

An Examination of Resourcing and Scheduling within the RCMP

by

Bruce Rout

B.Sc., University of Alberta, 1979.

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
MASTER OF SCIENCE
in the Department
of
Mathematics.

© Bruce Rout 2003

SIMON FRASER UNIVERSITY

February 2003

All rights reserved. This work may not be
reproduced in whole or in part, by photocopy
or other means, without the permission of the author.

APPROVAL

Name: Bruce Rout
Degree: Master of Science
Title of thesis: An Examination of Resourcing and Scheduling within the RCMP

Examining Committee: Dr. T. Brown
Chair

Dr. Luis Goddyn
Senior Supervisor

Dr. Keith Promislow

Dr. Patricia Brantingham
External Examiner

Date Approved: _____

Abstract

The problem of resourcing and staffing, or finding how much manpower is needed to meet demand, can be traced back to the times of the Roman Empire. We examine here the various means used by the Royal Canadian Mounted Police in trying to solve this problem efficiently. We also examine their latest attack in building a simulator to determine future demands on resources and we provide a solution to determine efficient staffing levels through an application of a scheduling algorithm using “rods”. This algorithm is characterized as a rod-scheduling method which can be reduced to a linear program. It has been found that the previous methods used by police departments in Canada and the United States are extremely cumbersome. The methods suggested here correct this. Although the methods are very similar to those used before in other industries they haven’t been applied to police work. What was previously done in policing is to optimize very simple constraints first and then try to fit the results to the needs of the user. In this work I have suggested first obtaining all legal inputs and then optimizing to obtain a final answer. The author uses this method to investigate different types of demand data. Furthermore, different integer programming techniques are investigated.

This document is meant for different users. It is hoped it can be read by various police departments as well as administrators and academics. In light of this the tone of the thesis is conversational. Much of the mathematical work is in sections three to seven.

Dedication

I would like to dedicate this thesis to my wife Roxanne who came up with the idea in the first place.

Acknowledgments

I would like firstly to acknowledge the help of Sgt. Wayne Mather of the RCMP without whose help most of this thesis would have been impossible. I would also like to thank Chief Supt. Lloyd Hickman for his support and assistance, Insp. Rick Noble, S.Sgt. Wolfgang Riemer, Sgt. Peter Fullbrandt, Sgt. Bill Shaw, and Sgt. Bill McDonnald, all of the RCMP and S.Sgt. Nora Skelding of the Ontario Provincial Police.

I would like to thank Dr. Patricia Brantingham for guiding me into using mathematics in the field of Criminology, Karen Meaghn and Misa Bilenky for their time and help, Alexander Kononov and Brett Stevens for giving me the idea of using “something like a one dimensional bin-packing algorithm - like rods” and most of all to Dr. Luis Goddyn for his time and effort in helping me with the intricacies of this thesis.

Contents

Abstract	iii
Dedication	iv
Acknowledgments	v
List of Figures	ix
1 Introduction and description of the problem	1
1.1 Background and description of approach to the problem	1
1.1.1 Steps that will be used to solve the problem	2
1.2 Some discussions with Sgt. Wayne Mather	3
1.2.1 Depth of the problem	3
1.2.2 Efficiency	5
1.2.3 Death Spiral	6
1.2.4 Effectiveness	6
1.2.5 Data Acquisition Problems and Forecasting	7
2 Brief background of attacks	9
2.1 Work in the 70s by Richard Larson	10
2.1.1 Hypercube	10
2.1.2 Reports by the Rand Corporation	11

2.2	No satisfactory software seems to exist	13
2.3	Appears to be an open problem	13
3	A new approach	15
3.1	Defining the problem through data modeling	17
3.1.1	The data model so far	17
3.2	Breaking the problem into three steps	18
3.2.1	Simulator to produce a demand table	19
3.2.2	What is a demand table	19
3.2.3	How these demand tables have been used	19
3.3	Hourly-shift method	21
3.4	A rod-schedule method	24
3.5	Applying the schedule to determine the number of officers needed	28
3.6	Reactive and proactive duties	29
3.6.1	Court and evening demands	29
4	Converting to Integers	30
4.1	Comparison of different Integer Programming methods	32
4.2	Integer programming for a packing problem	34
4.3	An Analysis and Motivation	35
4.3.1	A Comment on Usefulness	37
5	Rod Generation	39
5.1	Altering the Shift Schedules	39
6	Closeness of fit	41
6.1	Dual Solutions	41
6.2	Undermanning vs Overmanning	42

6.2.1	Choosing a convenient unit	44
7	Conclusions	51
Appendices		
A	Excerpts from a Report	53
A.1	Background	53
A.2	Current status	54
A.3	Conclusion	55
B	Description of data modeling	56
C	Table of data	61

List of Figures

3.1	Assignment Relationships	18
3.2	Typical Demand Data created by taking the number of calls in an hour for police service over a year in a typical large detachment. The number of calls were averaged for one week and adjusted to reflect an estimated demand for number of officers in each hour. This data was provided by Sgt. Mather and do not reflect the actual data of any particular detachment.	20
3.3	This is a graph showing the closeness of fit between the data of Figure 3.2 and the number of officers required by an optimal solution to the linear program described by the hourly-shift method.	23
3.4	This is a graph to show the closeness of fit between the data produced by using rods and the test data.	27
4.1	This graph looks at the number of officers needed to cover very simple constant demand data. It can be seen that the method involving rounding off deviates most from the line described by covering the data if the officers were allowed to be fractional. The process of rounding down and CPLEX integer solutions are identical. Also presented are results of solutions to the corresponding packing problems using CPLEX with integer flags turned on and the method described below. The solutions to the packing problems are identical.	33

6.1	Here we have taken the dual values produced by solving the problem with time moving forward and the dual values by solving the same problem but reversing time. We subtracted the two dual values and graphed them over the two-week period. Some values have remained the same while other have varied.	43
6.2	This is a graph depicting the size of the envelope around the data given. Notice the dramatic peak on Friday nights.	46
6.3	This is a graph of overmanning and undermanning over time as well as looking at the original test data. It appears that when the overmanning is small the undermanning is large. There is a good correlation between the values of the packing slack variables and the peaks of the original test data, showing greater difficulty in packing those areas. There is less correlation between the covering slack variables and the depth of the valleys of the test data.	47
6.4	This is a plot of the addition of the covering and packing slack variables. Resultant negative values were set to zero to make the graph more readable. We can see that Friday nights stand out most clearly. . . .	50
B.1	An entity	57
B.2	A many-to-many relationship	58
B.3	A one-to-many relationship	59
B.4	the resolved diagram	60

Chapter 1

Introduction and description of the problem

1.1 Background and description of approach to the problem

A background and description of the problem according to a report by Supt. Lloyd Hickman is included in Appendix A. In this report it is stated that the mandate of a resourcing RCMP task force is to develop a resourcing methodology that would determine the amount of human resources required to provide policing services at a given contractual location.

I show that this is an open problem in the policing industry and show that the construction of a simulator and the solution of a resulting scheduling problem will solve this problem. I will concentrate on the scheduling aspect.

In this thesis I propose to show that it is better to pre-specify constraints and work schedules and then apply a linear program rather than doing the reverse. This is a new scheduling process for the policing industry. Previously, a shift matrix has been specified which results in a nearly totally unimodular matrix. These shifts have then

been searched according to constraints to attempt to arrive at working schedules for personnel. I will show that this method is cumbersome and that a method I describe by using rods is superior.

I will use a linear program and then applying a heuristic to convert fractional values of police officers to integers. In the work presented here I demonstrate that the heuristic is equal to or better than that used by CPLEX.

I will examine dual and slack variables and show that the values of the slack variables give more information for this type of problem.

In examining the slack variables, I will associate packing and covering for the same demand data and constraint matrix. I will show that the slack variables for the packing problem correlate to peaks in demand and the depth of covering slack variables correlate to low points in demand data. I will also investigate using both covering and packing slack variables in finding times of greatest stress for the officers. Here I will apply a heuristic for obtaining an integer program (IP) for covering and a similar heuristic for an IP for packing. I show that the heuristic produces a result closer to the corresponding linear program (LP) in the covering problem than in the packing problem.

I also examine the difficulties in obtaining a good scheduling fit for data having varying raggedness and pointiness. I show that there is a linear relationship between less raggedness and the closeness of fit. This is to be expected. With an examination of the effect of pointiness of data to the closeness of fit I show there is a linear relationship up to a maximum and then a constant relationship occurs. This is also expected.

1.1.1 Steps that will be used to solve the problem

The overall problem can be broken down into steps. These steps are described as follows:

I Determine the input data from policy, geography, demographics and crime statistics.

- II Use a data model as described in Appendix B to determine relationships among these identities.
- III Write a simulator from the data model to determine a demand table for policing resources (i.e. number of police officers needed in each hour of the week).
- IV Use a linear program reduction to estimate work schedules that meet the demand for each hour.
- V Tailor the schedules for actual implementation.

1.2 Some discussions with Sgt. Wayne Mather

1.2.1 Depth of the problem

Sgt. Wayne Mather has been involved in Police Resourcing with the RCMP for about 27 years and assists detachments in their resource allocation planning in the Pacific Region. I know of very few other police officers who are also able to effectively assist detachments in their regions in planning resourcing requirements. There is no training, apprenticeship or mentoring program in place for future needs. Therefore we are faced with the situation that the “black art” of resource requirement prediction will be lost to our national police force.

We were referred to Supt. Hickman and eventually Sgt. Mather through the mentoring program at Simon Fraser University. We have had discussions starting in September 2000 and have attended workshops together to examine the possibility of writing software to assist in the prediction of resource requirements. (For the rest of this thesis we will call this prediction problem a resourcing problem.) Sgt. Mather says that this resourcing problem has plagued the RCMP since its inception and affects every aspect of RCMP operations. A literature search found a paper on resource allocations and its applications, albeit a simplified version, in the defense of the Roman Empire by Constantine the Great in 300 A.D. [1]. Effective and efficient

resourcing means that resources can be allocated and brought to bear without overly taxing resource supply.

Much of the resourcing problem lies in trying to define correctly the meaning of efficiency and what is effective. In physics, efficiency is defined as the amount of useful work produced by the total energy entered into a system. If police officers are sent to a crime scene and they are out-manned or out-gunned and unable to get a dangerous situation under control then they would be considered completely inefficient. Therefore, for certain types of calls for police services there is a minimum number of officers and a minimum level of equipment they must have in order to become somewhat efficient. An example may suffice to describe various inputs and outputs to a crime investigation. Suppose there was a burglary investigation and the burglar was no longer on the premises. A police officer is called to investigate the scene and takes a report from the victims of the crime and not much else can be done other than to see if the stolen items show up at a pawn shop and their serial numbers show up on a computer. Now let's say that when the police are called the burglar is still in the residence. This is a priority one call and could result in five or six cars being assigned to the scene and even a SWAT team. Here we have the same crime but the circumstances demand a far different input. Again, as before, let's suppose the burglar has left the premises and the police officer has been called in to investigate by the victims. And let's suppose the officer sees a window has been broken and there is blood on the broken glass or surroundings and perhaps even some fingerprints. The officer calls in the Identification Unit who may work for a half a day or a day collecting evidence. Here again, some slightly different circumstances will demand much more input to the investigation and there may still be no positive results. It can be seen that in some cases more police officers assigned to a situation will result in it's completion more quickly and in other cases it would make no difference how many more officers are assigned to a crime. Therefore, care has to be taken in defining efficiency.

1.2.2 Efficiency

Police work can be divided into four stages: reactive, proactive and special services, court and administrative. Let us concentrate on reactive policing and then see, in some future work, if the same principles will apply to the others. Reactive policing consists of a call for police services or a need for police service which is usually referred to as a call for service in the industry. Usually dispatchers receive a call for service and assign police to the scene where the call indicates police are needed. Police respond to the dispatcher and go the scene or various scenes involved and act according to the situation they find upon arrival. Then, when their work is completed at the scene, the call for service is cleared by the police by calling in a clear signal to the dispatcher. Please note that the police themselves can generate their own calls for service if they see something they feel requires their services while on patrol; but this can be considered as just a response for a call for service without a dispatcher even though the dispatcher would be notified of the officers' activities. There are many complications.

To simplify, let us say the call for service has been successfully handled if it results in a cleared call by the police and any suspects involved have been investigated. (Meaning arrested, ticketed, subpoenaed or whatever action is deemed necessary by the police officer). Basically, success means the call is cleared and the police got the bad guys.

It should be noted that police departments don't measure energy, they measure time. They measure the time it takes until a call is cleared and they use time in the form of Full Time Equivalent (an officer working full time for a year) for contracting purposes with municipalities. So any useful definition of efficiency should be temporal in nature. I would like to propose that if we knew the time police officers would take for a call if they were 100% efficient then we could define efficiency as this totally efficient time divided by the time the call actually takes to yield a successful completion. If a call is unsuccessful perhaps it could be considered 0% efficient. Note here that there may be circumstances in which officers may successfully complete a task faster than

the estimated totally efficient time. In this case there would be an efficiency greater than 100% and such circumstances would be interesting to operations departments.

1.2.3 Death Spiral

Finding these 100% efficiency times for various crimes or calls can end up being both a political and logistical nightmare. An examination of the logistics part of such a search brings to the fore the dreaded “death spiral”. The death spiral refers to the situation in which the time estimated to solving a crime is set as a standard and derived from historical data. There may be resourcing cutbacks and as a result the time for “solving” or clearing the crime is reduced because of the lack of resources to complete the task in what under normal circumstances would be a satisfactory conclusion. As a result, there is a lower estimate of the time to complete calls and a further reduction in resources. This process is repeated until crimes are considered that they can be solved and cleared in no time at all. This situation is referred to as the *death spiral* by resourcing experts in the RCMP.

The politics mentioned can destroy projects investigating resource methodologies. For example, the Metropolitan Toronto Police scrapped its resourcing project since its union felt that the project was going to result in the reduction in the number of already overworked police officers. The withdrawal of cooperation by police officers resulted in the two-year-old project being shelved.

1.2.4 Effectiveness

Furthermore, if more police officers are assigned there will probably be more arrests and it will appear that the crime rate will rise in an area’s statistics. Or it may be that with more time devoted to proactive duty more crimes will be reported, again resulting in a statistical increase in the crime rate. Therefore crime rate is not a good measure of an effective resourcing methodology, at least in the short term. The crime rate may well go down after a while but if resourcing is being effective, it will appear

to increase at the beginning.

Furthermore, there are deeper aspects of the resourcing problem, one of which is queuing. Queuing refers to the length of time a citizen's complaint is queued by either a dispatcher or a police officer. Queuing can be used as a measure of the effectiveness of resourcing. A shorter queue would mean there are more officers to respond to reactive calls. However there are geographical considerations to take into account. In an urban area it may be that a 15-minute queue is barely acceptable while in remote areas of the country it may take days to get a police officer to the scene of a complaint. The other side of queuing is the amount of time being spent on proactive work and special services. The philosophy apparently in place with the RCMP is that if there is any spare time it is attempted to be put to use in proactive policing, or in the prevention of crime. It would be nice if 30% of the total time of police officers on the job were spent in proactive work. At present the RCMP believe themselves to be a long way off. Unfortunately there is no measure being taken of this percentage. The length of the queue versus the percentage of proactive work can be balanced as a parameter in measuring resourcing effectiveness.

Another aspect of effective resourcing is acceptable scheduling. There are possible schedules that can be derived mathematically, which will fit the demands for police response very closely but which in turn are inhuman to the individual police officer. For example they may involve working 24 hours or working in one-hour shifts. Thus there is a tradeoff to be addressed. This is really where the "black art" starts to come in. Sgt. Mather jokingly said that to make a shift schedule you obtained a large white board and a bottle of whiskey, sat down at around eight o'clock at night and by about five in the morning you had yourself a shift schedule.

1.2.5 Data Acquisition Problems and Forecasting

Resourcing problems are not solved for situations that occur in the past; they are solved for situations that will occur in the future. So a resourcing methodology has to be able to predict the resourcing needs some time in the future. The Government of

Canada would like to have a twenty-year prediction of resourcing needs and municipal governments have indicated they would like to know what their policing needs will be for at least five years for budgeting purposes. In order to predict resourcing needs we need to be able to predict crime statistics. According to the RCMP resourcing experts, crime tends to be the same from year to year so long as you restrict yourself to only one year in the future. Two years in the future is considered really pushing things, and beyond that it is considered impossible to predict crime statistics. Five years in the future is completely unknown. There are also cases where the entire crime situation can change drastically on a national level without any premonition that such a thing was about to happen. Such is the case of the September 11, 2001 hijackings and destruction of the World Trade Center. The entire structure and expectations of policing has since completely changed. Furthermore court decisions such as Stinchcombe and policy changes such as the move to Restorative Justice and Community Policing have nearly doubled the demands on police resourcing for court proceedings. These incidences were unpredictable.

Nevertheless, some progress can be made in this area although it is limited. So long as our expectations are not too great we can make some qualitative analysis which may indicate whether the crime rate will increase or decrease. For example, demographics indicate that the usual crime committing age group has been less populous over the past ten years yet it is well known that the “echo” children of the baby boomers are now starting to come into this age group and we can expect crime to increase in the next few years. So a study of demographics can indicate and predict to some extent variation in the crime rate. Furthermore, if the effect of policy on the workload of police officers could be measured, then one would be able to budget or at least find the effect on budgets that policy changes could make. The long and short of it is that any detailed plan is due to fall into error in a very short time as unexpected events occur. If relationships are known budgets can be adjusted perhaps more quickly, but they cannot be accurately set very far into the future.

Chapter 2

Brief background of attacks

It seems that previous approaches to this problem have either been very complex in nature or overly simple. A rough description of a complex process is that the problem is turned into a linear or integer program and the results are then forced through many convolutions and iterations into some sort of final answer for the user. However, the iterations involve trial and error by the user to get software to obtain a satisfactory solution. An example is the PARR system, mentioned in the appendix, with which the user must first estimate a working schedule for officers and then have the program fit demand data. A *misallocation index* is then given which is a measure of the closeness of fit, and if the user is not satisfied with this, different parameters for a shift schedule of officers is attempted and the whole process is tried again. It can take days of frustrating effort to get this software to yield a satisfactory answer.

On the other hand, simplified approaches use linear approximations to the problem. The OPP use such a system. However, this system is overly simplified. It makes extensive use of averages in its calculations. An average of annual demand is guessed at by the user and a spreadsheet program multiplies this demand by some factor to produce the average number of officers the program predicts will be needed. The user is free to choose any demand average and can redo calculations until some number of police officers is produced that everyone will be happy with. But it is not clear how this system has anything to do with reality. Averages are very dangerous things

to work with in resourcing. For instance, let us say we have an average of six police officers needed on an evening shift according to the average demand for that period. But the demand fluctuates wildly during that shift because, let us say, the bars close late at night and maybe 12 officers are needed for that period. The six officers on duty would be unable to handle the demand. For any resourcing system to be successful it must handle the peaks of demand rather than averages. The problem is too complex to be handled by averaging programs and yet a simpler and more subtle process has to be found than that demanded by complicated programs. I propose that the idea of the complicated programs be retained but by reversing their processes. By doing things the other way around, an acceptable solution can be found. In other words, instead of doing a linear program first and then fitting resulting data, I propose obtaining all feasible possibilities to the problem and then applying a linear program to obtain a final solution.

First, however, let us examine in more detail the history of the problem.

2.1 Work in the 70s by Richard Larson

2.1.1 Hypercube

One of the first attempts at an approach to examining the distribution of police officers was that of Richard C. Larson and his Hypercube Queuing Model in 1973 [2]. This model basically describes a method of using the computer facilities available at the time to assist dispatchers in choosing the next available police car out of a queue. Administrators could then use these results in designing how many response units go into various districts for their communities and also in designing the layout of districts within a community. Measures of performance were computed for mean travel time, workload imbalance and fractions of dispatches that are interdistrict. Because of machine size restrictions at the time, a model of about ten cars was the maximum that could be used. Much of the literature in work done on police resourcing, dispatching and planning cites this early work by Larson.

The hypercube was an n -dimensional binary space with n being the number of cars in the model. Each vertex of the hypercube represented a car. The path (or adjacency) between points determines the order in which cars have been allocated. The binary value of the location on the hypercube is simply one if the car is busy and zero if the car is free. Larson then uses an exhaustive search to find the next available car to be dispatched. This model is then fed into a simulator which keeps track of the cars assigned according to an input data file that mimics or replicates calls demanding police service, or it can also act as a dispatching tool in real time. After the simulation, or real-time shift, data is analyzed and according to the expertise and experience of management, districts and communities can be determined to be under resourced or over resourced. This was a crude method but was the best methodology that was available at the time.

2.1.2 Reports by the Rand Corporation

A number of reports on emergency response systems were produced for the Rand Corporation in the 1970s. These include a report for computing performance characteristics of the previously described hypercube model by Larson [3]. Kolesar and Blum [4] derived useful square root laws for fire truck travel distances. This work shows that the average response distance in a region is inversely proportional to the square root of the number of locations per area from which fire companies are available to respond. They indicate that:

$$E_i = k_i \sqrt{\frac{A_r}{n - \lambda S}} \quad (2.1)$$

where E_i is the long run expected response distance of the i th closest fire company, A_r is the area of the region under consideration, λ is the expected number of alarms received per hour, S is the expected total service time spent by all fire companies that respond to and work at an alarm, k_i is a constant of proportionality that can be determined empirically and n is the number of fire companies assigned to the region. From this formula average travel times can be estimated. Further investigations were

undertaken by Hausner [5] in which he describes the mechanics of data gathering and computerized analysis to determine relationships among travel times, travel distances, time of day and weather conditions.

Chaiken [6] did an overview of mathematical models available in his day and stated that there were two basic types of models, those based on hazard formulas and those based on workload formulas. He concludes that hazard formula models have serious failings and are “not recommended for any purpose”. Chaiken says workload formula models have limited utility and that the best approach is to use computer programs that calculate a variety of performance measures and recommend allocations that meet the objectives desired. Unfortunately, Chaiken does not state what these two formulae are and how they are actually used so we just have to take his word that these two methods are unsatisfactory. I have seen a hazard formula in RCMP documentation but was unable to remember the reference. From what I remember a hazard formula is a simple one-line formula that estimates resources required. It would seem, from what is written above on the complexity of the problem, that it would be very unlikely that a simple algebraic formula would determine police resourcing. The January 17, 2002 front page of the Toronto Globe and Mail carried a story of a simple formula being used by London, England police in which 60% of the police were placed where there was the highest population density and 40% were placed where there was the heaviest density of calls for service. This formula was designed for the police by management consultants PA Consulting Group. No reason is given why there is a 60-40 split nor why the majority of police are not put where the demand is. Perhaps this formula would work under certain circumstances but it looks more like a rule of thumb under restrictive circumstances. Furthermore it does not shed any light on how many police officers are required to meet the demand stated.

In December of 1976 Kolesar et al [7] derived a set of guidelines for patrol car scheduling using linear programming. Their report includes an analysis of when mealtimes should be taken to ensure cars are available to respond. In the report, an ordinary linear program is used and the associated matrix looks totally unimodular and resulting solutions are all integers. From investigating the work of Cooper [8]

in 1977 it appears Kolesar and Cooper have come very close to discovering the use of rods which I describe later. However neither author uses or develops their methods so that they can be applied to resourcing.

2.2 No satisfactory software seems to exist

More recently, Berman et al [9] wrote a comprehensive report for the Metropolitan Toronto Police Department (MTP) and concluded that “none of the existing software packages and theoretical models is capable of capturing all the important features of the primary response system at MTP” [9, Page 2] They report on the fairly wide acceptance and use of Larson’s Hypercube Model, by now evolved to use a GUI interface, and earlier methods, namely the hazard and workload formulae which are rejected by the authors in favor of Larson’s model. They recommend a queuing model which would have the potential of capturing the essential characteristics needed by the MTP.

2.3 Appears to be an open problem

To quote Berman et al [9, Page 4], “Furthermore, no models dealing with community policing, a mandate for MTP, are available in the literature. Certainly, there are no existing studies that estimate the effect of community patrols on the rate of arrivals of service calls or in reducing crime. With regards to how community policing might be modeled, there are some suggestions but no attempts yet to make them concrete.” And also [9, Page 8], “The hypercube Model enables police to experiment with different designs before actually implementing the deployment plan. For example, with the click of a mouse the planner can change the basic locations of the patrol cars or the districts and re-plan the model. The Hypercube Model is used by about thirty police departments in North America. However, the Hypercube Model is certainly not a staffing model.” In their report, the authors state that they do not wish to deal

with shift scheduling and exclude it from their discussions. This is unfortunate. It can be shown that changes in shift schedules can affect not only how well demand for police response is covered but also how many police officers are needed.

Chapter 3

A new approach

Berman and his colleagues recommend the use of both a simulation model and optimization routines. Both the task force of the RCMP and myself agree with them. Much of the problem lies in trying to determine what is meant by the demand for police officers and how this can be expressed mathematically. Sometimes demand has been determined by studies on police officer time and at other times the number of calls for service has been used. The Chicago Police Department has done statistical studies on various methods of determining demand and James Thurmond of their computer services department reported that all of the methods tested gave proportionally very similar answers. He said they correlated to within 86%.

Because of these different methods, I would like to mathematically define police officer demand so it can be met through optimization routines.

Police officer demand (or just demand for the sake of this thesis) is defined as the minimum number of officers required each hour to adequately respond to and deal with police matters in a manner that is acceptable to the public, administrators, municipalities and police representatives.

The long and short of this definition is that police are involved in dealing with crime and that demand is defined as how many police officers are required each hour

to deal with the crime presented to them. Now, there exist extensive call for service data throughout police detachments in Canada. Since there is a correlation between calls for service in each hour and number of officers needed in each hour perhaps a relationship can be found. The municipalities of Surrey, a large municipality with many police officers, and White Rock, a small municipality, both in B.C. were examined in this regard. White Rock has an average of about one call for service in an hour. It was estimated that if a call for service took an average of one hour that 4 or 5 officers as an absolute minimum could handle the load. However there are about 18 officers in the White Rock detachment. Due to staggered call times a rough estimate is that on the average, for White Rock, each call for service requires about 3.5 hours, including response, investigation, arrest, interviews, court time, training, briefing and all other work which an officer has to do.

A similar very rough calculation for Surrey netted a similar result, approximately 3.5 to 4 hours for each call for service. Taking all of the above into consideration it would be tempting to simply take the number of calls for service a detachment receives, do a simple calculation to determine how many officers would be needed if each call took one hour, and multiply the result by 3.5 to get the number of officers required by the demand. Unfortunately there is a major problem here. This would base the calculation of demand on the staffing levels already in place, which may be very low and result in under-staffing throughout the country. Also it would not take into account major policy changes and changes in the world situation that would require major changes in staffing as described previously. It should be noted that both Surrey and White Rock use the same methodology to estimate resourcing needs which is based on calls for service. This methodology is contained in the PARR system which bases resourcing needs on setting the time required for completing a call for service at 3.5 hours.

The approach described here involves a simulator determining the demand followed by optimization routines to determine appropriate shift schedules and officers required. This can result in the estimation of a lower bound for the number of officers required on the roster.

3.1 Defining the problem through data modeling

The RCMP has effectively been in operation for more than 125 years and has become more and more sophisticated and complex, so that today the number and sheer size of its policies and modes of operation alone are quite staggering. To mathematically model the processes of a modern large police organization is impossible unless these processes are either simplified or organized into a comprehensible format. One way of doing this is through the process of data modeling. The task force took time and arduous effort over about a six-month period to produce a data model diagram. Data modeling is a way to take very complex systems and portray them so they can be understood by a group of people needing to work on problems associated with data generated by whatever the data model is about. A description of the data modeling process is included in Appendix B with references therein.

3.1.1 The data model so far

After much work, the RCMP Resourcing Task Force has come up with a preliminary data model diagram also known as an entity-relationship diagram. At the heart of the entity-relationship diagram are three entities: Activity, defined as something requiring the assignment of RCMP resources; Employee, defined as a human resource who may be assigned to an Activity; and an Assignment, defined as the allocation of an Employee to an Activity. Basically, the model is to find ways of assigning employees to activities. The relationships are presented in Figure 3.1.

At the top of the task force's entity-relationship diagram are four distinct flat data files which are inputs to the entire system. They are: Policy, Case File, Demographic Parameter and Geographic Area. These four input entities describe a very large input data stream into the proposed simulator. We will examine Policy as an example.

The entity Policy has huge ramifications within the RCMP. Policy can take the form of any directive to RCMP members from an Act of Parliament to a standing order at the detachment level and anything in between. None of the policy is in

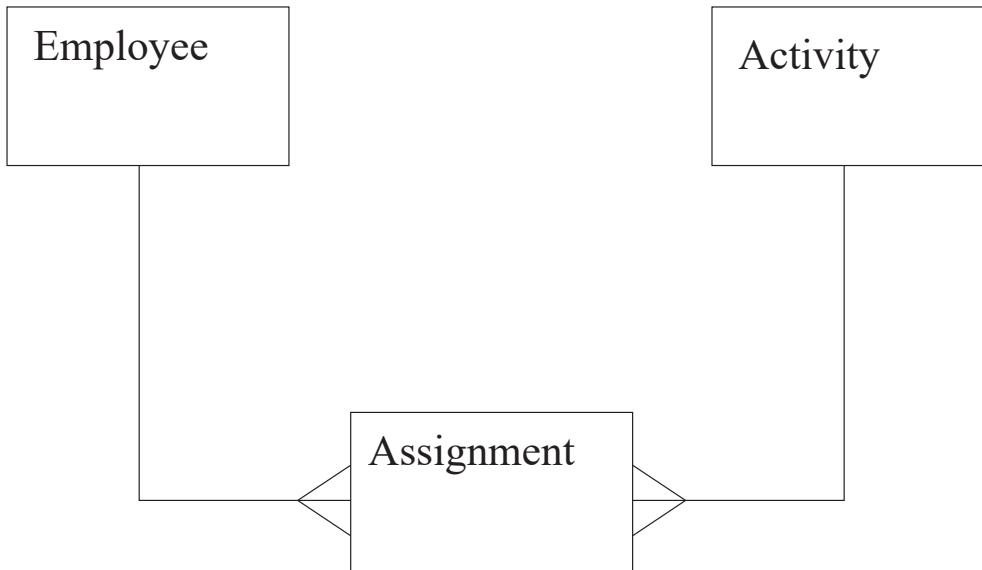


Figure 3.1: Assignment Relationships

a standard format. Somehow, all of this data has to be put into a standard form with the requirement that the form it takes shows how the policy affects resourcing in manpower, skill sets, training, equipment and time to do any particular activity. Though this may seem arduous there is a very big pay-back once accomplished. The model would then be able to predict the changes in resourcing requirements as a result of a change in policy. The same would be true of demographics, geographic parameters and crime rates once these data streams are put into a form which the model can use.

3.2 Breaking the problem into three steps

The overall resourcing problem can be broken down into three distinct steps. The data model diagram can be used to generate a simulator which will produce a table describing the demand. From that table a schedule can be produced to determine the most efficient way of using officers and from that schedule the minimum total number of officers can be determined under specific restrictions.

3.2.1 Simulator to produce a demand table

The four input data streams to the data model need to be arranged and formatted to take the form of usable data files. The complexity of this preprocessing of data depends on the morphology of data in its present form and what the data model determines is needed for a simulator to use. Within this simulator, the data is manipulated until the entities of Employee and Activity can be properly matched. These logical manipulations could take the form of look-up tables and if statements. Once the data has taken this form, the Assignment entity of the data model needs to know the number of officers assigned to an Activity and how long the officers will be at that activity. From this point it is fairly straightforward to create a demand table.

3.2.2 What is a demand table

The demand table produced by the described complex simulator can take the form of the number of officers required by a detachment or policing contractual area each hour for one week, using a week as an example time span. This can change. It may be that one week is not the same as another and a month could be used. The schedule could also be broken down into two-hour or four-hour segments. It has been suggested that the schedule be broken down into half-hour segments, which is possible mathematically but may create an unwieldy schedule. The demand table used here is simply a list of 168 numbers representing the number of officers needed on duty each hour of a 168-hour week.

3.2.3 How these demand tables have been used

Presently things have been done in a backwards fashion. Detachments are given the number of people they will be allowed and a schedule following the demand table is produced. Sometimes the demand table is ignored and a schedule is produced that just has the same number of officers on duty around the clock. This causes major problems during high demand times.

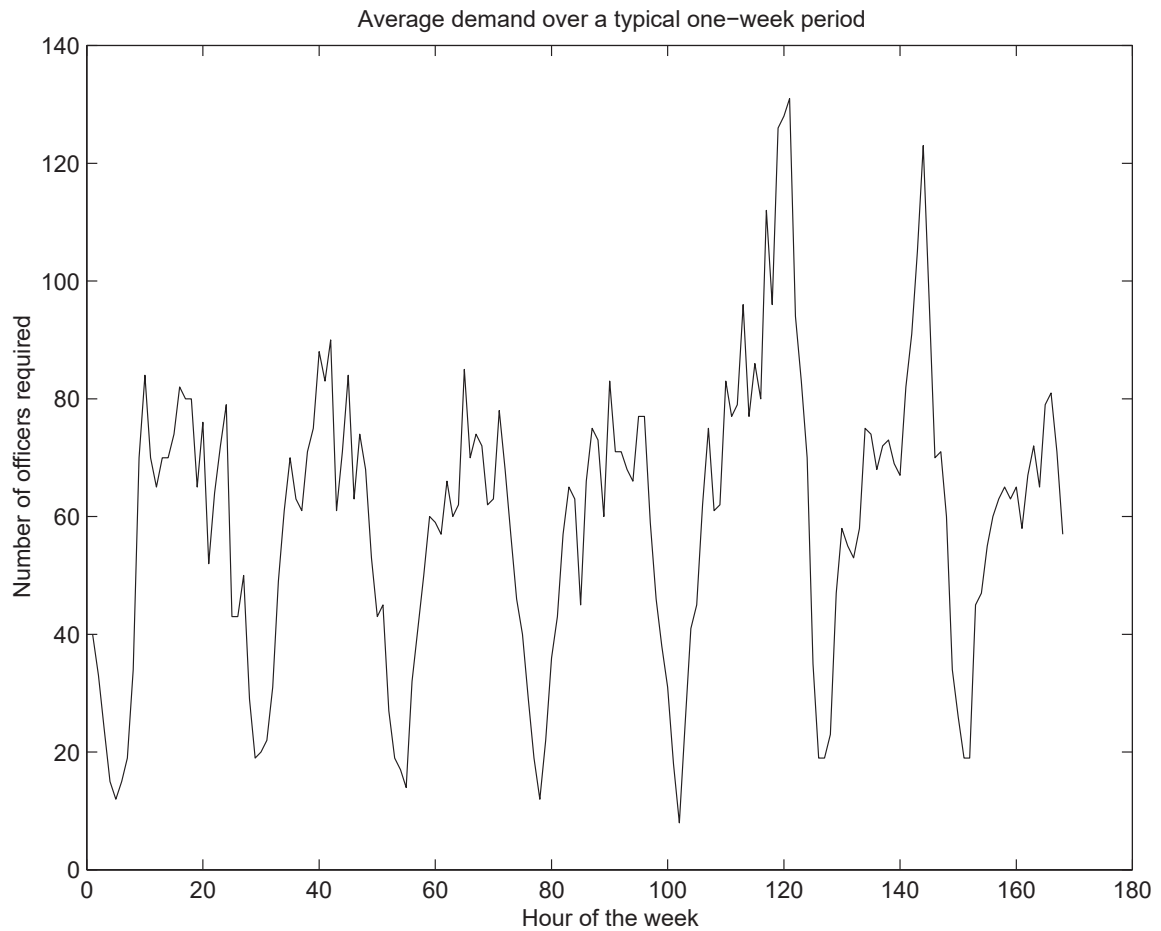


Figure 3.2: Typical Demand Data created by taking the number of calls in an hour for police service over a year in a typical large detachment. The number of calls were averaged for one week and adjusted to reflect an estimated demand for number of officers in each hour. This data was provided by Sgt. Mather and do not reflect the actual data of any particular detachment.

I will present two different LP approximations as ways of solving scheduling problems. The first is a method that has been used previously and is used by the PARR model which is referred to in Appendix A. I will call this previously used method the hourly-shift method. The second method I call the rod-schedule method. The hourly-shift method uses a linear program to determine how many officers should start at any particular hour. This fits the demand table fairly accurately but the shifts which in turn have to be produced by hand are too strenuous or stressful for the officers involved. For example they may result in officers working one-hour shifts. Even if someone produced a schedule through various means to come up with some sort of acceptable schedule, the mandate of the task force, to determine the total number of officers needed, would be missed.

We will use different durations of each shift, namely eight, nine, ten, eleven and twelve-hour shifts.

3.3 Hourly-shift method

We set $x_{l,t}$ equal to the number of police officers to start at time t and having shift length l . Our objective function then becomes to minimize the number of person-hours per week:

$$\min \sum_{l=8}^{12} \sum_{t=1}^{168} lx_{l,t} \quad (3.1)$$

subject to:

$$\sum_{l=8}^{12} \sum_{t=\tau-l}^{\tau} x_{l,t} \geq d_{\tau}, \tau = 1, 2, \dots, 168 \text{ and } x_{l,t} \geq 0 \quad (3.2)$$

Here d_{τ} is the demand at time τ . Inequality (3.2) ensures that the total number of people working at time τ must be at least as great as the number of police officers required at that time. We present the constraints (3.2) in the form:

$$\mathbf{Ax} \geq \mathbf{d} \quad (3.3)$$

where A is a binary-valued constraint matrix whose rows are indexed by the hours $1, 2, \dots, 168$ and columns indexed by the shifts $(l, t), l \in \{8, 9, 10, 11, 12\}, 1 \leq t \leq 168$. An entry in the constraint matrixes position $(\tau, (l, t))$ is one if and only if $\tau \leq t \leq \tau + l$. \mathbf{x} is a column vector containing the variables $x_{l,t}$ and \mathbf{d} is a vertical vector containing the demand values d_τ . That is, (3.3) can be expanded as:

$$\begin{pmatrix} 1000 & \cdots & 1111 \\ 1100 & \cdots & 1111 \\ \vdots & & 1111 \\ 0000 & \cdots 111 \cdots & 0000 \\ \vdots & & \\ 0000 & \cdots & 1111 \end{pmatrix} \begin{pmatrix} x_{8,1} \\ x_{8,2} \\ \vdots \\ x_{12,168} \end{pmatrix} \geq \begin{pmatrix} d_1 \\ d_2 \\ \vdots \\ d_{168} \end{pmatrix} \quad (3.4)$$

Here the constraint matrix is filled with the first row governing which shifts will have any officers working in the first hour; the next row governs which shifts will have any officers working in the second hour and so on.

Using the demand vector \mathbf{d} of Figure 3.2, and CPLEX, a linear programming package, an optimum integer-valued solution of \mathbf{x} was found. The random nature of \mathbf{d} leads one to suspect that A is unimodular or nearly unimodular in some sense, ie. relatively few full-rank square sub-matrices have determinants not in $\{-1, 0, 1\}$. It would be nice to do more research in this area.

Although this method gives a start table for when employees are to begin a shift and how many employees work on particular eight-hour, nine hour and so on shifts, it does not yield a complete shift schedule nor does it yield how many employees are needed. In particular, the shifts need to be grouped into a weekly work schedule for each officer. Please note that this method has been used in the past, and from this data, through very difficult calculations, an approximation for shift schedules has been produced. We may consider this an introduction to a much better method which is presented next.

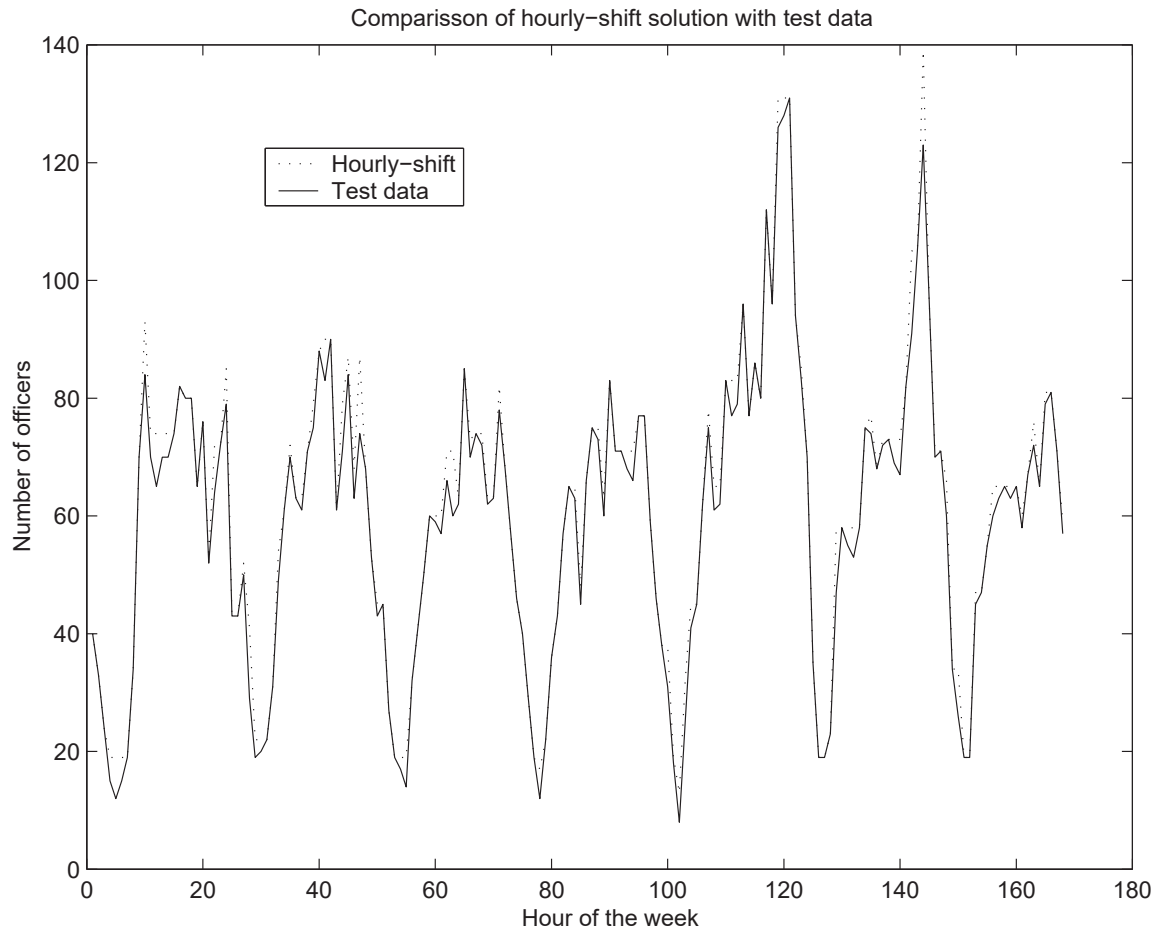


Figure 3.3: This is a graph showing the closeness of fit between the data of Figure 3.2 and the number of officers required by an optimal solution to the linear program described by the hourly-shift method.

3.4 A rod-schedule method

This method fortunately gives us the solution sought. Consider a rod representing a length of time of 168 hours. Then consider the rod sectioned into 168 sections, each section representing one hour of the week beginning from Monday 00:00 hours to the following Sunday 24:00 hours. Mark each of the sections with ones where an employee is working and with zeros where he or she is not. Consider a number of such rods marked so that each rod would represent a single weekly shift schedule. Also consider that each rod may have an unknown number of officers working on its particular weekly schedule. This method of using rods to represent weekly shift schedules was tried and it was quickly found that a one-week long rod was too restrictive in setting up different lengths of daily shifts, so I used two-week long rods having 336 hours each. I could also have used month-long rods if I had wished but two-week long schedules will suffice as an example here.

The reason for using two-week long rods is that a 12-hour shift pattern requires a 4-hour shift at the end of the working week, which is demoralizing. So a two-week schedule is used with an 8-hour shift at the end of two weeks to make up for time short of 80 hours. A similar situation arises for 9-hour and 11-hour patterns. (10-hour patterns fit nicely for four-day work-weeks however.) Let us say we can identify each rod by both the start time of each of many two-week work schedules and a particular identifiable pattern of lengths of shifts. A rod could be represented by a vector of ones and zeros like this:

$$A_\rho = \overbrace{\left(\underbrace{1, 1, 1, 1, 1, 1, \dots, 0, 0, 0, 0, 1, 1, 1, 1, 0, 0, 0, 0, 1, 1, 1, 0, 0, 1, 1, 1, 0, 0, 1, 1, 1, 1, 1, 0, \text{etc} \dots}_{\text{each binary number represents one hour}} \right)}^{\text{there are 336 coordinates of the vector}} \tag{3.5}$$

In the future we will consider A_ρ as a column vector. Here the positions of A_ρ each represent an hour of the week, the first position represents the first hour, the second position the second hour and so on. Furthermore, the ones and zeros can be in any pattern we like. It all depends on the shift schedule. Here I will use patterns

that have officers beginning their shifts at the same time each day with rest days in the middle of the two-week period. I am trying to choose rods to represent humane schedules which may be a result of union regulations or desires expressed by members of the police force. I also want the officers to work 80 hours in each two-week period so I will restrict the rods to contain only a total of 80 ones.

Let \mathcal{R} denote the set of allowable rods. In my example I have:

$$\rho \in \mathcal{R} = \{(l, t) : 8 \leq l \leq 12, 0 \leq t \leq 336\} \quad (3.6)$$

and the entries of A are:

$$A(\tau, \rho) = \begin{cases} 1 & \text{if } \tau \text{ is a work hour for rod } A_\rho \\ 0 & \text{otherwise} \end{cases} \quad (3.7)$$

I have used l to indicate the usual length of shifts in the schedule and t is the hour when this schedule would start. So a vector $A_{(8,7)}$ would indicate that people working on that rod or shift schedule would all work eight hours each day starting at 7 a.m. They would work five days followed by a weekend off and followed by another five days of work. This would require a pattern of six zeros, meaning from hour one to hour six the people on this shift schedule are off, then followed by eight ones, then 16 zeros, eight ones, 16 zeros and so on until the weekend where there would be all zeros until 7 a.m. the next Monday which would repeat the pattern of the first week. For eight-hour and ten-hour shift patterns the second week is the same as the first. For nine, eleven and twelve-hour shift patterns the second week is different than the first to ensure people working these rods all work 80 hours during the two weeks.

We now proceed in a fairly standard way to model the scheduling problem as a linear program. Let R_ρ be the number of officers working on a predefined two-week "rod". $R_{(l,t)}$ is a vertical vector 336×5 (or 1680) deep. Let d_τ be the number of officers demanded in hour τ . Let A be again a binary constraint matrix whose columns are merely the rod vectors described above. The rows of A are arranged as in equation 3.6. A is the matrix whose rows are indexed by $\tau \in \{1, 2, \dots, 336\}$ and whose columns are indexed by ρ .

Our objective function is to minimize the man-hours in a two-week period. Or:

$$\min \sum_{l=8}^{12} \sum_{t=1}^{336} 80R_{\rho} \quad (3.8)$$

since there are 80 working hours in a two-week span. The constraints are:

$$A\mathbf{r} \geq \mathbf{d} \text{ and } \mathbf{r} \geq 0 \quad (3.9)$$

where \mathbf{r} is a vertical vector whose members are R_{ρ} and \mathbf{d} is the demand consisting of members d_{τ} . We have:

$$\begin{pmatrix} 10000 & \cdots & 001111 & \cdots & 000 & \cdots & 11111 \\ 11000 & \cdots & 000111 & \cdots & 100 & \cdots & 11111 \\ 11100 & \cdots & 000011 & \cdots & 110 & \cdots & 11111 \\ \vdots & & & & & & \\ & & & & 000 & \cdots & 00001 \\ & & & & \vdots & & \\ & & & & & & \vdots \\ 00000 & \cdots & 111111 & \cdots & 000 & \cdots & 11111 \end{pmatrix} \begin{pmatrix} R_{8,1} \\ R_{8,2} \\ R_{8,3} \\ \vdots \\ R_{12,336} \end{pmatrix} \geq \begin{pmatrix} d_1 \\ d_2 \\ d_3 \\ \vdots \\ d_{336} \end{pmatrix} \quad (3.10)$$

Once again we can use the program CPLEX to find the values of the variables $R_{l,t}$. Most of the values of \mathbf{r} are zero as were the values for \mathbf{x} in the hourly-shift approach. In the hourly-shift approach, the constraint matrix A is in some sense close to being unimodular, hence the coordinates of the solution vector \mathbf{x} are almost certainly integer valued. This is not the case for the rod-schedule method. Here, in some sense, A is far from being unimodular. Thus optimized solutions of \mathbf{r} are generally non-integer valued.

I then rounded the values of \mathbf{r} to the nearest integer and examined the results. I found that by simply rounding the rational solution I could cover the demand with the exception of a few places. An iteration can be done using only those rods which had too few officers on duty as a result of simply rounding and a demand table consisting of the differences between this fitted data and given data. (This procedure converges

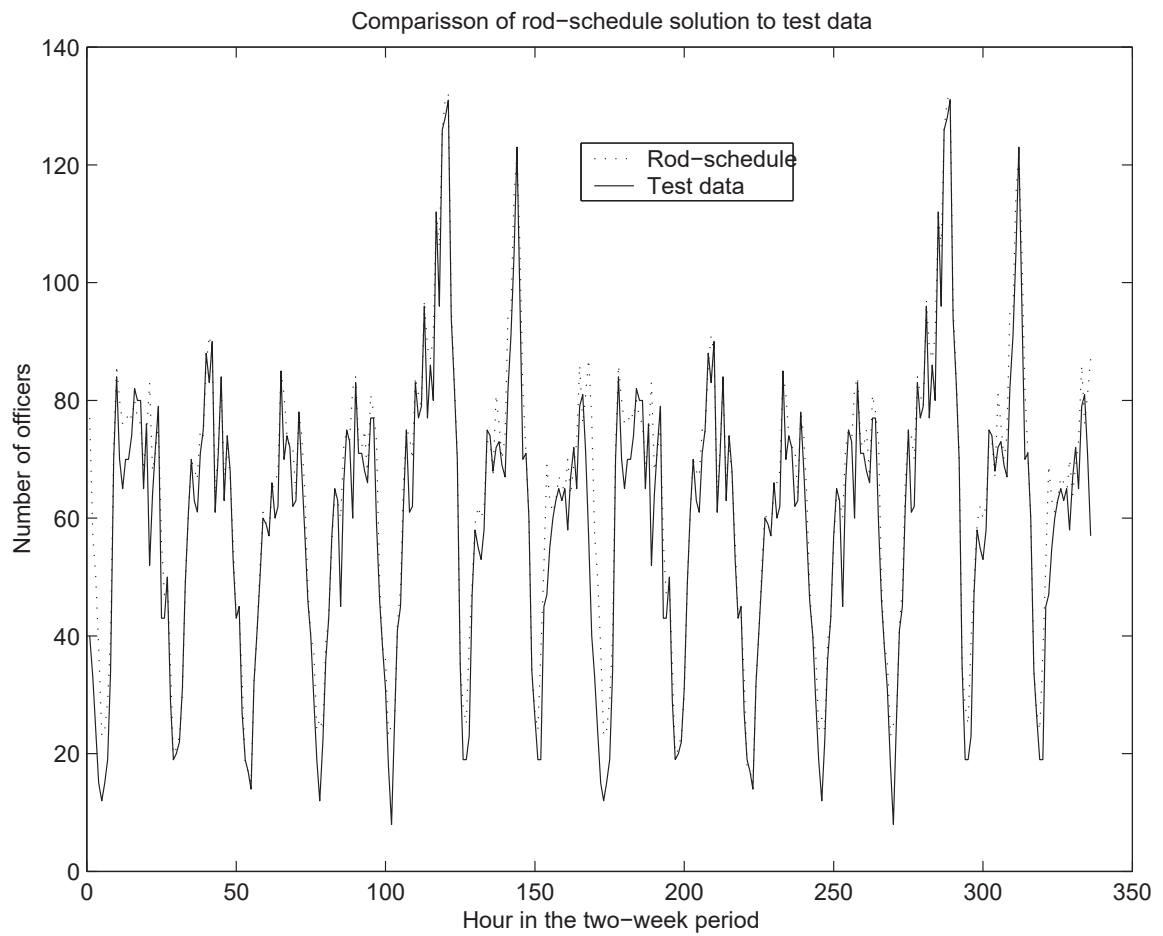


Figure 3.4: This is a graph to show the closeness of fit between the data produced by using rods and the test data.

and is discussed in more detail in the next chapter.) These new values can be added to the ones that were short in the previous iteration and the process repeated if necessary.

There is an important point to the rod-schedule method. Each of the rods can have any pattern and so long as the times indicated by the demand are covered, I am free to make up any shift pattern I like. In this case I have tried to put together as many combinations of shifts as possible having the officers start work at the same time each day and working 80 hours in the two-week period. But other patterns are possible.

3.5 Applying the schedule to determine the number of officers needed

The above method has netted a number of non-zero shifts and yielded the optimal number of officers to work on those shifts and cover the needs of the given demand. So we now have a roster of shift schedules. Since each officer works on only one shift schedule, (or each officer works on only one rod), we can determine how many officers are needed by totaling all of the officers on all of the rods. This is an absolute minimum number which does not take into account any extra work such as shift preparation and post-shift work. It also does not take into account activities that run over to the next shift and may involve some overtime. To now get a better picture of how many officers are truly needed it would be best to use this number as a seed for a second simulator which keeps track of officers, run them again through the matching up of employees and activities, take into consideration how long each activity takes and also taking into account sick leave etc. and see if the officers are being drastically overworked. This second simulator should be able to state how many more officers are required.

3.6 Reactive and proactive duties

It has been suggested at this point that to add in the number of officers needed to make up for proactive police assignments we just add an appropriate percentage of officers based on the number of reactive officers found by the previous method. There is a better way since most officer demand time is at night and most proactive time may be suited for the day. The times set aside for proactive work can be seen from the rod patterns. If there are a number of rods having predominantly daytime shifts then that is where daytime proactive police personnel can be added if the proactive duties are demanded in the day. The same can be done for evening. Sometimes a detachment may have personnel that just do proactive work and their schedules can be produced by having a separate demand table that reflects only proactive duties.

3.6.1 Court and evening demands

There are other demands in setting up the time schedules of the rods. Because of the nature of police work, the officers require a few evening shifts within about a week or so of an initial crime investigation in order to do at-home interviews with victims and witnesses. They also need court time during the day which must be a weekday that is not a holiday. There also has to be a spreading out of court times for the detachment personnel so that the entire detachment is not in court at the same time. It is suggested that a two-week schedule could be arranged by having each rod represent a first week of regular duty, then having some time off followed by a couple of evening shifts and then, after taking allowances for holidays and weekends, put in day shift time for court duty, then return to the regular duty portion of the rod. In this way there is a forward scheduling algorithm. It was also suggested that officers not begin work from midnight to seven or eight in the morning. This is achieved by not having rods with these shift start times.

Chapter 4

Converting to Integers

In the previous section we found a solution using rods. However the rods require fractional police officers. For example one of the rods could require 5.39627 police officers to work on a particular shift schedule. Since people only come as whole numbers, we have to find a way to convert these rod values into integers. The problem of deducing an acceptable integer-valued solution to the LP described by the rod-schedule method is the central problem in integer programming. If we round down we find the detachment is understaffed. If we round up we find the detachment is greatly overstaffed. For example if there are 40 rods, we would end up adding 40 officers across the board and this may be too many.

To solve this problem the following algorithm was tried:

- A) Solve for \mathbf{r} for the LP described by $A\mathbf{r} \geq \mathbf{d}$ with $\mathbf{r} \in \mathfrak{R}$.
- B) Round off the members of \mathbf{r} to the nearest integer. Let this rounded solution be $\hat{\mathbf{r}}$.
- C) Calculate how many more officers are needed as a result of this rounding. That is, calculate $\mathbf{d} - A\hat{\mathbf{r}}$ and let the result be denoted as $\hat{\mathbf{d}}$. If $\hat{\mathbf{d}} = 0$ then stop.
- D) Let $\tilde{\mathbf{d}} = \max(\hat{\mathbf{d}}, 0)$.

E) Solve the LP described by $A\mathbf{w} \geq \tilde{\mathbf{d}}$ for \mathbf{w} which will be the vector meeting values of how much each rod is undermanned as a result of rounding as in step B. If $\mathbf{w} = 0$ then stop.

F) Choose one of the following sub-steps:

- If any values of \mathbf{w} are ≥ 1.0 then round those values down and set all others to zero.
- If all values of $\mathbf{w} < 1.0$ then choose the greatest value of \mathbf{w} and set it to 1.0 and set all others to zero.
- If the values of $\mathbf{w} < 1.0$ and if more than one of the largest values of \mathbf{w} equal each other, then choose the last value in the vector \mathbf{w} of these values and set it to 1.0 and set all others to zero.

Then set $\mathbf{w} + \hat{\mathbf{r}} \rightarrow \mathbf{r}$.

G) Go to Step B.

This process tends to converge fairly rapidly. I would like to add some discussion on this process. It may have some shortcomings in that it may not be the optimal process. However, it is better than simply rounding everything upwards. The process described by the rod-schedule method produced a set of shift schedules (a total of 92 rods) that would require 269.05 officers to meet the demand data if fractional officers were allowed. The process described by steps A through G one produced integer numbers of police officers to be added to the rods which were sufficient to cover demand. The final number of officers if the rods were all integer values was 274, an addition of only 5 police officers. Out of curiosity I programmed the CPLEX package to treat the input data as an integer-only and it yielded a total of 81 rods requiring 270 officers to cover demand. So, the method described is less advantageous than CPLEX for this problem. This process is investigated further.

4.1 Comparison of different Integer Programming methods

To simplify data to investigate, I set demand data files which were “flat” or constant. These data files incorporated a demand of zero, one, two, three, four and five officers each hour over a two-week period. The rod-schedule method was used with rods having eight-hour daily shifts and a two-day rest period each week. The rods can start at any hour over the two-week period. The total number of officers to cover demand was found as integer solutions.

The CPLEX package with integer flags turned on and attempting to solve for a flat demand of three people each hour with 8-hour shifts and two-day weekends did not converge after 46 hours of CPU time and running for two days. The program was interrupted and a solution of 15 officers was obtained to cover a demand of three officers in each hour. The program reported a gap of 11.6%.

After viewing the results of using the steps described above for turning fractional officers into integers, I tried rounding down rather than rounding off in step B. All of these results are presented in Figure 4.1.

We can see from Figure 4.1 that if the officers were treated as fractions there would be a linear relationship between the number of officers required to cover demand and the number of officers required each hour. We also see that as soon as the constant demand jumps from two officers per hour to three, the process described by the steps above, that is by rounding off, deviates the most from the line described by fractional officers. We can also see that under the circumstances described, the results of the process of rounding down in Step B and the CPLEX integer method are the same.

Seeing that rounding down may produce a more efficient answer than rounding off the process of rounding down was tried on the original data and it was found that the data could be covered with 270 officers working on 83 rods. A further freeing up of 4 officers. CPLEX with integer flags turned on produced a solution having 270 officers working on 81 rods. Both methods produced the same values for the objective




Figure 4.1: This graph looks at the number of officers needed to cover very simple constant demand data. It can be seen that the method involving rounding off deviates most from the line described by covering the data if the officers were allowed to be fractional. The process of rounding down and CPLEX integer solutions are identical. Also presented are results of solutions to the corresponding packing problems using CPLEX with integer flags turned on and the method described below. The solutions to the packing problems are identical.

function while CPLEX produced a solution using two fewer rods.

4.2 Integer programming for a packing problem

In the next chapter I will be discussing overmanning and undermanning when trying to meet demand. As a result we need to find some way to turn fractional police officers into integers when doing a packing problem. I have used the following steps which are very similar to the steps used in the covering problem.

- a) Solve for \mathbf{r} for the LP with constraints $A\mathbf{r} \leq \mathbf{d}$ with $\mathbf{r} \in \mathfrak{R}$.
- b) Round up the members of \mathbf{r} to the nearest integer. Let this solution be $\hat{\mathbf{r}}$.
- c) Remove all rods having values ≤ 0 from further calculations.
- d) Calculate, as a result of this rounding, how many officers are over the demand. That is, calculate $A\hat{\mathbf{r}} - \mathbf{d}$ and let the result be denoted as $\hat{\mathbf{d}}$. If $\hat{\mathbf{d}} = 0$ then stop.
- e) Let $\tilde{\mathbf{d}} = \max(\hat{\mathbf{d}}, 0)$.
- f) Solve the LP described by $A\mathbf{w} \geq \tilde{\mathbf{d}}$ for \mathbf{w} . If $\mathbf{w} = 0$ then stop.
- g) Choose one of the following sub-steps:
 - If any values of \mathbf{w} are ≥ 1.0 then round those values down and set all others to zero.
 - If all values of $\mathbf{w} < 1.0$ then choose the greatest value of \mathbf{w} and set it to 1.0 and set all others to zero.
 - If the values of $\mathbf{w} < 1.0$ and more than one of the largest values of \mathbf{w} equal each other, then choose the last of these values and set it to 1.0 and set all others to zero.

Then set $\hat{\mathbf{r}} - \mathbf{w} \rightarrow \mathbf{r}$.

h) Go to Step b.

I present the results of comparing the CPLEX integer packing package with the method I have described above as an addition to Figure 4.1. The solutions are the same, however CPLEX was unable to converge for a constant demand of two officers each hour after 30 minutes of computer time. The program was stopped and results taken with a reported gap of 39%. The CPLEX program was also suspected of failing to converge or at least having problems converging for trying to solve the packing problem for a constant demand of 4 officers in each hour. The program was stopped with a reported gap of 12%.

To complete this experiment, the associated packing problem to the original data of Figure 3.2 over a two-week period was tried using both the procedure described by steps a) to h) and the CPLEX integer package.

A linear program solved the packing problem requiring 229.3 officers on 90 rods. The steps from a) to h) solved the packing problem with 228 officers on 73 rods and the CPLEX method netted 229 officers on 76 rods.

4.3 An Analysis and Motivation

Describing the previous algorithm to the covering problem we have . . .

$$\mathbf{s} = \mathbf{d} - A[\mathbf{x}] \tag{4.1}$$

where

\mathbf{s} = the slack of an associated problem which is made deliberately infeasible. \mathbf{d} = the original demand. A = the constraint matrix. \mathbf{x} = the solution of the fractional covering problem, $A\mathbf{x} \geq \mathbf{d}$

Following is the basis of my algorithm:

$$A(\lfloor \mathbf{x} \rfloor + \sum_k^n \Delta \mathbf{x}_k) \geq \mathbf{d} \quad (4.2)$$

where

$$\lfloor \mathbf{x} \rfloor + \sum_k^n \Delta \mathbf{x}_k = \mathbf{x}_{\text{optimum}} \quad (4.3)$$

and

$$\Delta \mathbf{x}_k \in \mathbf{Z} \quad (4.4)$$

Here $\mathbf{x}_{\text{optimum}}$ is the optimum final integer value to the problem. Also there are n iterations and each iteration adds $\Delta \mathbf{x}_k$ to $\lfloor \mathbf{x} \rfloor$.

As we go through each iterative loop set $\Delta \mathbf{x}_k$ as the amount to add to $\lfloor \mathbf{x} \rfloor$ to cover the infeasible slack. It can be shown that the optimum solution will be the solution with the least slack. Here we will let the least slack be defined as $\|\mathbf{s}\|_{L_1}$.

Consider a fractional vector \mathbf{F} such that:

$$\min \mathbf{F} \quad (4.5)$$

$$A\mathbf{F} \geq \mathbf{s}_k \quad (4.6)$$

and we can solve for \mathbf{F} . Here \mathbf{s}_k is the slack calculated for each iteration.

$$\mathbf{s}_{k+1} = \mathbf{d} - A(\lfloor \mathbf{x} \rfloor + \Delta \mathbf{x}_k) \quad (4.7)$$

The relationship between $\Delta \mathbf{x}_k$ and \mathbf{F} follows. Decompose \mathbf{F} in the following way:

$$\mathbf{F} = \left\{ \begin{array}{c} \begin{bmatrix} f_1 \\ 0 \\ 0 \\ 0 \\ \vdots \end{bmatrix} + \begin{bmatrix} 0 \\ f_2 \\ 0 \\ 0 \\ \vdots \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ f_3 \\ 0 \\ \vdots \end{bmatrix} + \cdots \end{array} \right\} \quad (4.8)$$

Or,

$$\mathbf{F} = \sum_{j=1}^N f_j \mathbf{e}_j \quad (4.9)$$

In the same way let

$$\Delta \mathbf{x}_k = \sum_{j=1}^N \delta x_{k,j} \mathbf{e}_j \quad (4.10)$$

In the cases where $0 \leq f_j \leq 1 \forall j$, I wish to choose some j such that the “error” created by choosing $\delta x_{k,j} = 1$ will be the least. Here I am defining the “error” = \mathbf{E} as the amount $\mathbf{x}_{\text{optimum}}$ will overshoot \mathbf{x} . Or,

$$\min \mathbf{E} = \min(\|\mathbf{x}_{\text{optimum}} - \mathbf{x}\|_{L_1}) \quad (4.11)$$

Since $A(\mathbf{x} + \mathbf{F})$ is the closest approach to \mathbf{d} then \mathbf{E} will be the least when $1 - f_j$ is the least. Therefore I am seeking the max f_j and using that index to set $\delta x_{j,k} = 1$.

4.3.1 A Comment on Usefulness

When the CPLEX package was tested against degenerate problems it ran into major difficulties trying to converge. The method I have come up with converged very quickly to correct answers. However the CPLEX package proved slightly better with problems that were not degenerate.

Please note that in the method I described, I am always increasing the values of the coordinates of $\lfloor \mathbf{x} \rfloor$ and this restricts the area of search for an optimal integer solution. In certain degenerate problems this seems to be an advantage.

I believe more work in this area may prove fruitful. For example, the objective function in the cases I have presented all have coordinates of unity. If the gradient of the objective function is altered would this make a difference? Furthermore, The coordinates of $\lfloor \mathbf{x} \rfloor$ are always increased; is there a way to decrease them according to some parameter? These are the questions that will have to be answered in the future.

Nevertheless, this is an approach that gives very similar objective function values in the problems presented here to that generated by CPLEX which uses branch and bound.

Chapter 5

Rod Generation

We now examine how rods can be generated as a result of the previous fit. We have used rods that let officers begin their shift at the same time each day with officers working any of eight to 12-hour shifts in their schedule. This results in a lower limit of officers to meet the demand table. This would also allow management to begin to set a budget. There are a couple of things we can try here. The first is to see if there is an effect on altering the shift schedules and to see if there is a resulting change in the number of officers required.

5.1 Altering the Shift Schedules

Note from Figure 3.4 that there is, on the whole, a tighter fit on the peaks than in the valleys of the data. Is it possible to make up new rods by sliding shifts along to include peaks? For example, let's say we have a rod of eight-hour shifts and the last shift before time off just misses including a heavy peak which has barely enough officers to cover it. Could we then take that one shift, and move it out of a valley to start an hour later and include the peak and see if this newly manufactured rod would help free up more police officers. In other words we make up a new rod all having eight-hour shifts and still beginning at the same time except for the last shift

before the weekend where we have the shift begin an hour later. Could we generate similar such rods and see if their addition would make a difference in the number of officers required to cover demand.

I took the following approach. I took all the differences between the negative of the amount of overmanning at the beginning of all shifts and the end of those shifts. I sorted them in order from least negative to most negative and noted the time of their beginning and the length of the shift. I then took the 20 greatest differences and calculated when these particular rods would have to begin the last shift before the weekend off. I then shifted these shifts one hour forward to include those areas where the overmanning would be the least and made up new rods. Briefly, I shifted the officers from times where the overmanning was the most and tried to put them where the overmanning was the least.

Out of 20 such manufactured rods an LP as described above stated two rods would be best used with 3 officers on each produced rod. Without generated rods there were 270 officers required to meet demand. With the generated rods 269 officers were required resulting in a saving of only one officer.

This seems to indicate that if there is a large enough and extensive enough vector of legal variables, or rods, then there is not much room for improvement by altering the rods or shift schedules. What this seems to mean is that complex shift schedules that may be hard or undesirable by the police officers may not be saving management very much in the way of efficiency of deployment of officers. It may be better in the long run to have shift schedules that are more acceptable to the officers and make sure they are extensive and variable enough to cover demand and optimize them than it is to try and make complex schedules that are disagreeable to the officers.

Chapter 6

Closeness of fit

6.1 Dual Solutions

From the solution to the linear problem it is possible to easily obtain a dual solution when solving the LP described by the rod-generation method. A dual solution will yield a relative worth of each hour of the two-week span and can, in a way, let the scheduler know the “strain” on officers at various times. The problem is so far from having a unique solution that little information is conveyed by a dual solution since random choices prevail in the method of determining a solution. To demonstrate, I took the problem outlined by the rod-generation method and found the dual through a simple command using CPLEX. This gave a relative worth of each hour in the demand table. Many of the dual values were zero. At these hours, due to the complementary slackness theory there are either just enough officers to meet demand or more than enough to meet demand. Where the dual has some positive value, there are just enough officers to meet demand. The dual, in business or commercial terms, gives a value for the relative worth of those hours. One might suspect that such dual variables, being greater than zero, would indicate a particularly troublesome hour. But such is not the case since random choices prevail.

I then took the same problem and solved it “upside-down”. I took the demand

table which was in terms of hours going from one to 336 and inverted it to go from hour 336 to one. I readjusted the coefficient matrix accordingly and left the rod vector \mathbf{r} alone. So in effect we have the same problem but we are solving it upside down. I again used CPLEX to solve the LP and yield the dual solutions to the same problem done this way. Again many of the dual values were zero and some were not. I reversed the order of these dual values to coincide with the dual values I had before and then subtracted the two vectors from each other hour by hour to see if there was a difference. There were significant differences in these values indicating that even though we have the same problem, by solving for the dual in a different order, I get different values for each dual. So I have found for this particular problem the values of the dual are not unique as seen in Figure 6.1. We therefore look to find another method of examining where officers may be hard pressed to cover demand.

6.2 Undermanning vs Overmanning

With Equation 3.9 we are assured that there are at least as many officers on duty at all times as that stated in the demand table. Since there is not an exact fit there are a few hours when there are more officers on duty than that demanded. So the schedule has been overmanned to meet demand. Examination shows this overmanning occurs mostly during the "valleys" or slow times of the demand graph. If we wish to add personnel at the busiest times perhaps we can find these times by turning the problem inside out and look at an undermanning situation. Here we look at maximizing our man-hours keeping manpower under the limit of the demand.

We have:

$$\max \sum_{l=8}^{12} \sum_{t=1}^{336} 80R_{l,t} \quad (6.1)$$

under the conditions:

$$\mathbf{Ar} \leq \mathbf{d} \quad (6.2)$$

A fit was produced using the above equations. The fit was then subtracted from the test data. The difference is given as the degree of undermanning. This is also

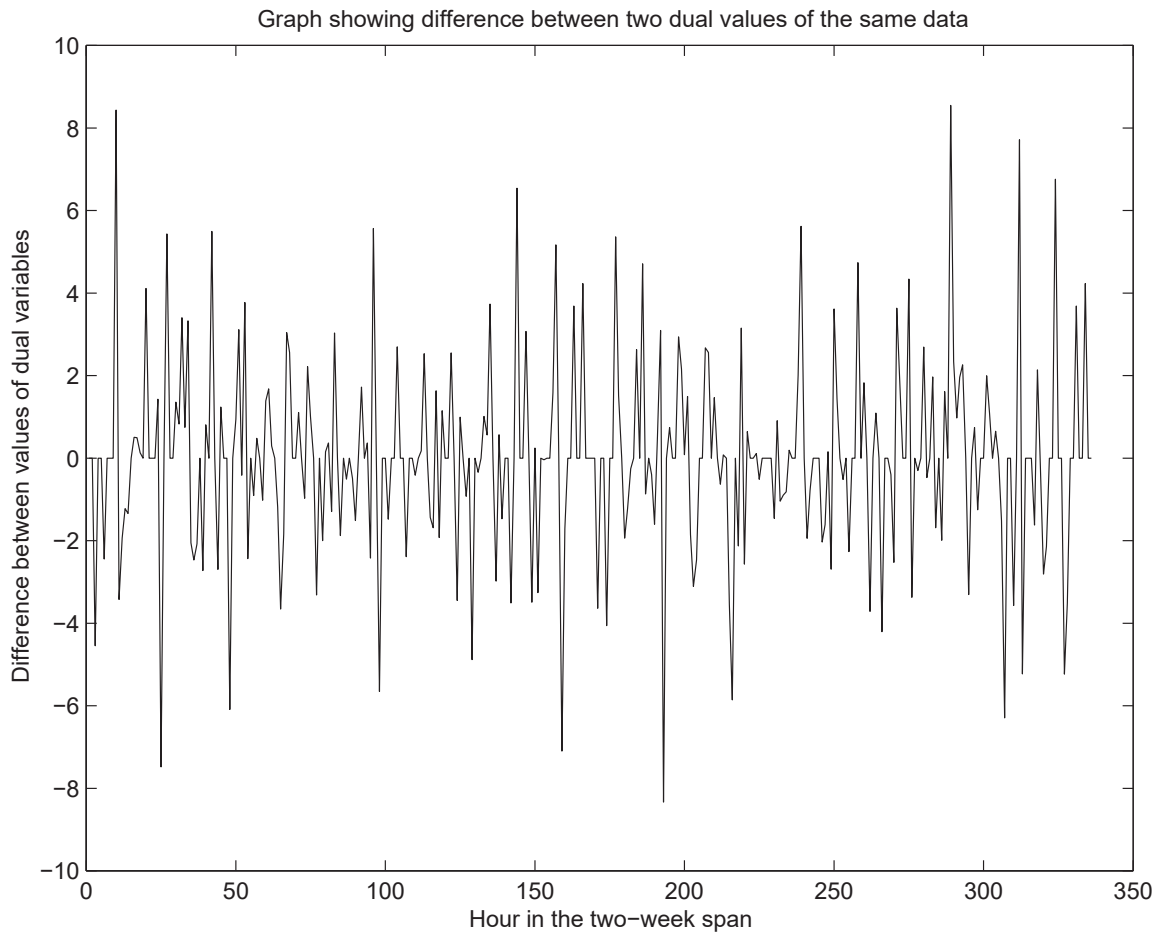


Figure 6.1: Here we have taken the dual values produced by solving the problem with time moving forward and the dual values by solving the same problem but reversing time. We subtracted the two dual values and graphed them over the two-week period. Some values have remained the same while other have varied.

known as the slack of the associated packing problem. We have also taken the fit produced by overmanning. I subtracted this fit from the test data and produced a difference I denote as the degree of overmanning, or the slack of the associated covering problem.

6.2.1 Choosing a convenient unit

This suggests the idea of an envelope surrounding the demand throughout the two-week period. Above the demand is overmanning and below is undermanning. I propose a useful way of measuring the size of this envelope. Overmanning is a covering problem while undermanning is a packing problem. We need to find some way to normalize the results we obtain using these techniques. Let n_c = the difference between the number of officers in the original data each hour and the number of officers predicted by covering (or the slack of the covering problem). Also let n_p = the difference between the number of officers in the original data each hour and the number of officers predicted by packing (or the slack of the packing problem).

In order to normalize the size of the envelope described I found its average size and divided by the total man hours which is $\sum_{\tau=1}^N d_{\tau}$.

Therefore I have:

$$F = \text{degree of fit} \tag{6.3}$$

where:

$$F = \frac{\overline{n_p - n_c}}{\sum_{\tau=1}^N d_{\tau}} \times 100 \tag{6.4}$$

where: $\overline{n_p - n_c}$ = the average difference between n_p and n_c over the time span under consideration.

$$N = \text{total number of time segments under consideration} \tag{6.5}$$

This gives F as a percentage.

We portray the percentage difference between overmanning and undermanning in Figure 6.2 throughout the two-week period. This data gave a degree of fit of .05%. Notice the peaks which stand out on Friday nights. If scheduling managers were given more officers they would probably put them on duty on Friday and Saturday nights simply from having experience that these are the busiest periods. Here I am attempting to show that there may be a relationship between a poor fit in packing and a higher stress in covering. We can see that using the idea of an envelope, Friday draws our attention. It would be nice if more investigation could be done in this area.

Also I have presented a graph of the slacks of packing and covering in Figure 6.3. It is presented to see if, on the whole, the slack for covering becomes less when the slack for packing becomes more negative. There are a couple of times when this seems to be true. Overall however I was unable to find a relationship in this regard. I have also placed the demand data on the same graph and it seems to show there is a good correlation between the values of the packing slack variables and the peaks of the demand data.

I will be using the idea of a covering/packing slack envelope in future experiments.

A proposal.

I am proposing that there is some use in looking at the associated packing and covering problems together. I have presented the fact that the dual of a covering problem in this instance may not be very useful. However, I believe that the slack of the associated packing problem can reveal quite a bit of information, particularly on where in the schedule to apply more officers should they become available. Or, could we use this information to investigate where officers are under the most stress? Perhaps defining the most stress as those times where there is the greatest demand compared to the least covering slack. An experienced officer would probably say Friday and Saturday nights are the most stressful. I am proposing that perhaps we can see this show up in a more mathematical way.

I have taken some of the values of the slack variables for the associated packing

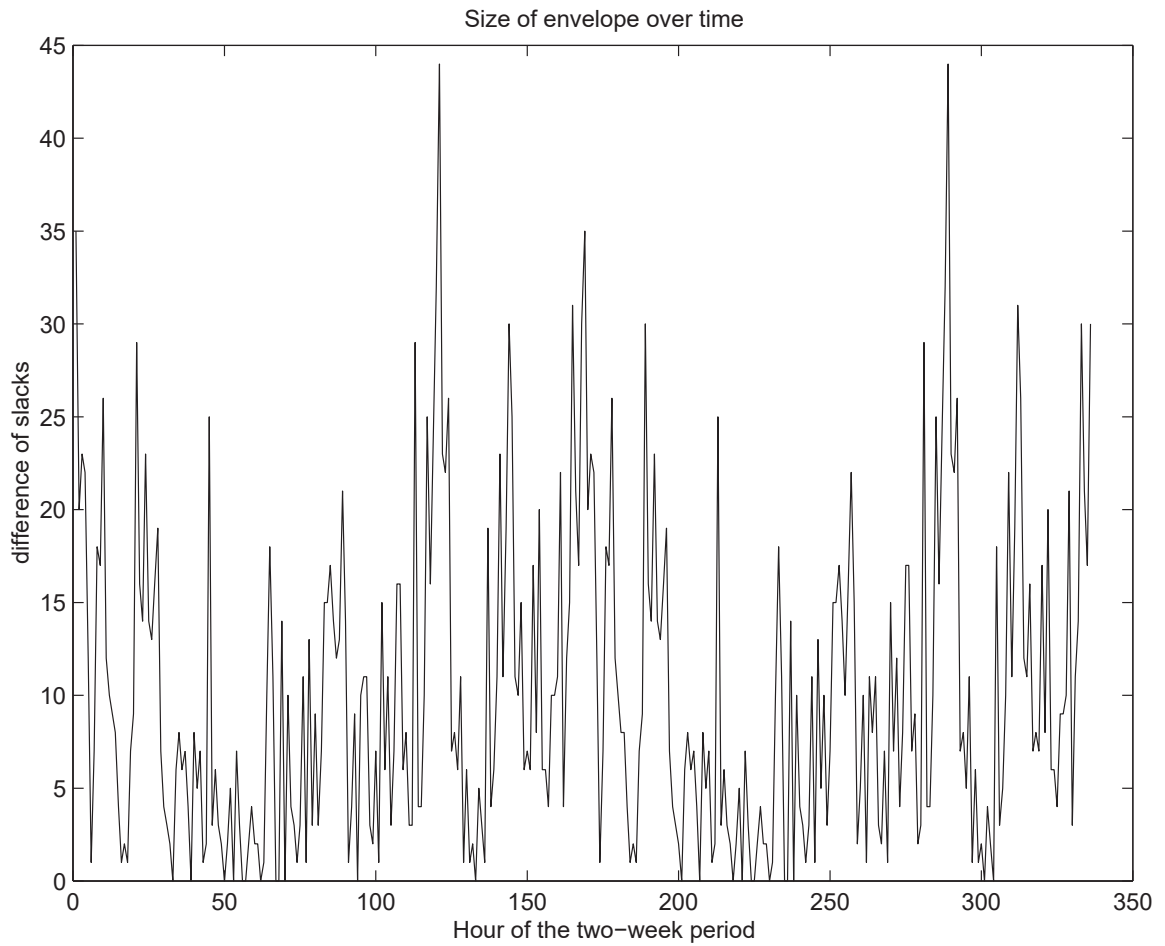


Figure 6.2: This is a graph depicting the size of the envelope around the data given. Notice the dramatic peak on Friday nights.

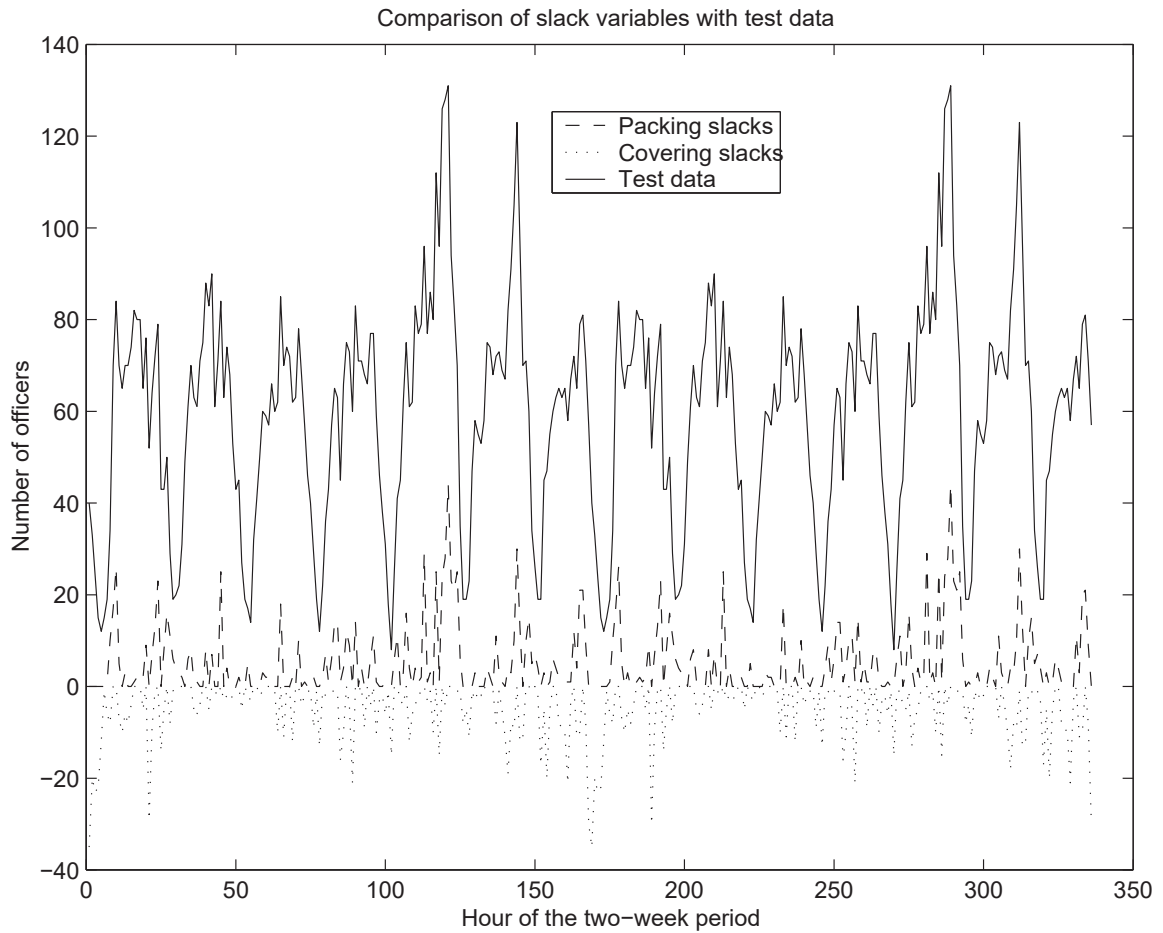


Figure 6.3: This is a graph of overmanning and undermanning over time as well as looking at the original test data. It appears that when the overmanning is small the undermanning is large. There is a good correlation between the values of the packing slack variables and the peaks of the original test data, showing greater difficulty in packing those areas. There is less correlation between the covering slack variables and the depth of the valleys of the test data.

problem and placed them in the following table.

Hour of the day	1	2	3	4	5	6	7	8	9	10	11	12
Friday	11	1	0	0	1	0	0	6	11	0	3	16
	4	0	4	1	2	29	1	3	0	25	1	24
Saturday	28	44	23	21	25	6	0	1	0	0	3	0
	0	0	4	2	0	11	3	1	0	4	3	10
Sunday	30	13	0	10	15	5	7	4	0	3	0	1

(6.6)

Here I have boldfaced the times which have a block of high values during Friday and Saturday nights. These values happened to be the highest of all the slack variables for the packing problem. In looking at this table it appears that two 8-hour periods, from 12 p.m. to 5 a.m. on Friday night-Saturday morning and later Saturday night and Sunday morning, are the times when the packing problem has the poorest fit and may correspond to the most stressful times in the schedule.

To further investigate this idea, I present the values of the slack variables for corresponding times to the covering problem in the following table.

Hour	1	2	3	4	5	6	7	8	9	10	11	12
Friday	0	-10	-3	-2	-6	-1	-15	0	0	-3	-4	0
	-12	-6	-4	-2	-1	0	-3	-1	-10	0	-15	0
Saturday	-4	0	0	-1	-1	-1	-8	-5	-11	-1	-3	-1
	-2	0	-1	-1	-1	-8	-1	-5	-11	-19	-8	-9
Sunday	0	-12	-11	0	0	-1	0	-2	-17	-5	-20	5

(6.7)

Here in table 6.7 we can see there is, overall, little negative slack on Friday night-Saturday morning. We can also see there are some times of little negative slack in the early hours of Sunday morning although 1 a.m. and 2 a.m. are very well staffed with negative slacks of -12 and -11 officers respectively.

It seems, from the data presented, that Friday night/Saturday morning stands out as the most stressful time in the schedule. It is still a judgment call on the part of the schedule manager as to where to put more officers however I feel the presentation of

the values of slack variables can give significant insight into the management of officer placement and perhaps give insight into matters concerning officer safety.

I present Figure 6.4. Here I have added the covering and packing slack variables and discarded negative values. We can see both Friday nights in the two-week period stand out fairly clearly.

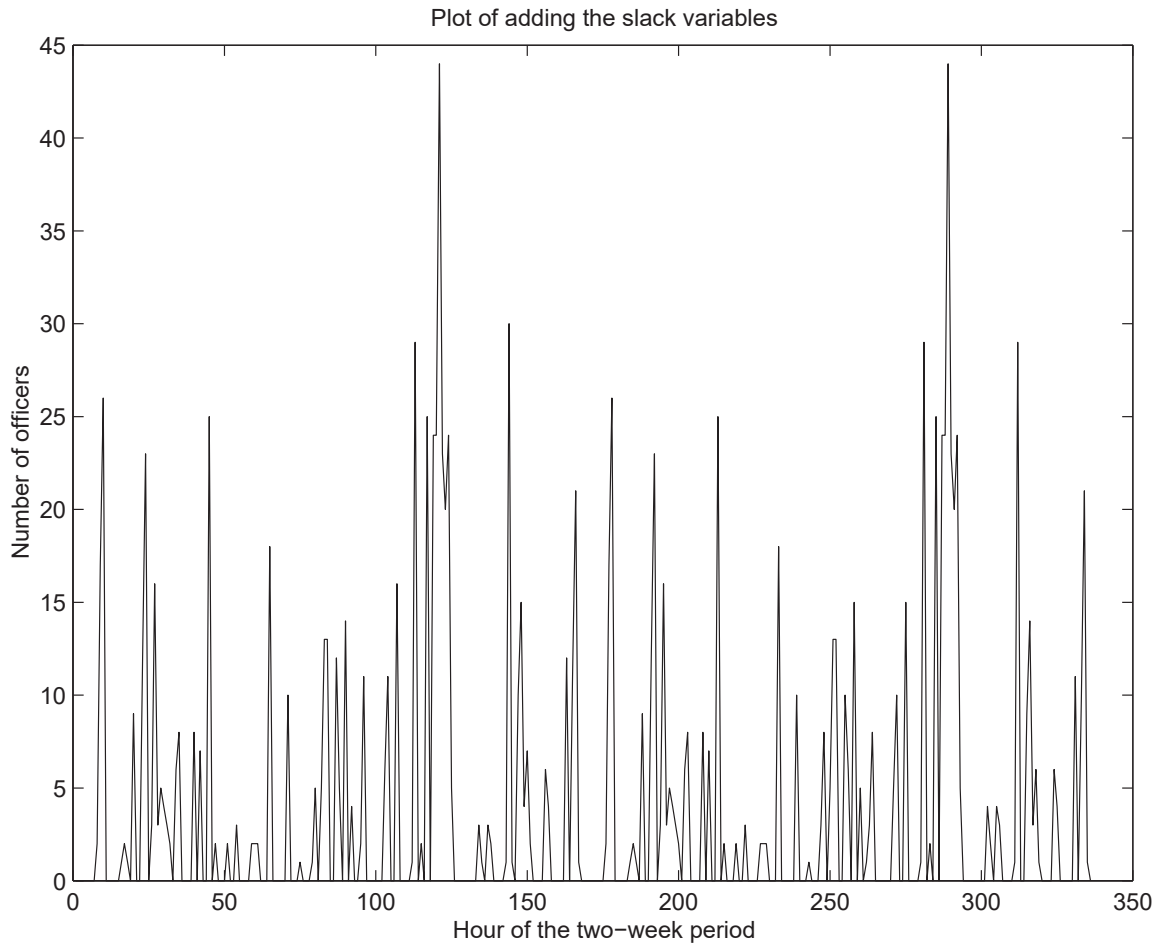


Figure 6.4: This is a plot of the addition of the covering and packing slack variables. Resultant negative values were set to zero to make the graph more readable. We can see that Friday nights stand out most clearly.

Chapter 7

Conclusions

From this thesis I conclude:

1. There exists a problem in police departments in accurately determining the number of resources required to meet demand
2. Past methods that have been used are inadequate in solving this problem
3. The manpower problem can be resolved using a simulator to determine future estimated demand followed by a scheduling routine
4. The way police officers are scheduled into their working shifts affects the quantity of resources needed
5. The hourly-shift method typically results in integer solutions due to the near unimodularity of the associated constraint matrix
6. The hourly-shift method has serious shortcomings; in particular, that it leaves the scheduler with a difficult, perhaps infeasible allocation problem
7. The rod-schedule method is an improvement in scheduling algorithms in the policing industry in that solutions provide a completely allocated schedule, despite having a more difficult IP problem

8. The integer conversion process described in the thesis for this type of problem is comparable to that used by the commercial package CPLEX.
9. The described rod technique can determine the number of officers required to cover demand, given a set of allowable rods
10. If there is a reasonable spread of allowable rods then there is not much to be gained by further rod generation as described in Chapter 5
11. Due to high degeneracy, dual solutions contain little information in this type of problem. Indeed, more information is contained in slack variables
12. Slack variables can be used to gain insight into how part-time employees can be deployed or in how adjustments can be made
13. The process of packing or covering tends to have more difficulty in filling the sharp peaks and valleys in demand typically encountered in police work. There is a possibility that this can be used to identify periods of high stress.

Appendix A

Excerpts from a Report

Printed with permission.

A.1 Background

The RCMP provides policing services to 8 provinces, 3 territories, over 200 municipalities and 62 First Nations, and through contract negotiations with the provinces and municipalities, has been expected to offer quality, client-centered police services for an agreed length of time. One of the most significant aspects of the contracting process is the determination of human resources required to offer an acceptable standard of operational policing services and the number of operational and organizational support resources who are required to support that complement of front line officers

Currently the human resource planners within the RCMP have developed and are currently using numerous different planning methodologies to allow them to determine human resource requirements and to allocate approved human resource levels. Pacific Region - Performance Analysis and Resource Review (PARR) and MALLOC. Northwest Region - Standards to Ensure Equitable Resourcing (STEER), Weighted Workload and Division Human Resource Model (DHRM). Atlantic Region and Headquarters - Standards Assisting Resource Planning in Law Enforcement (SARPLE). These systems have either been developed by H.Q. or individual Divisions and for the most part are out of date, rely on

unrealistic historical data, lack consistency and do not factor in current policing climate (FTE availability, supervision, officer safety, pro-active and community policing initiatives). In summary they do not meet the current needs of the RCMP or our contracting partners.

Furthermore, the rationale for additional resources of new contracts is often negotiated with each contracting partner, both at the provincial and municipal levels, in different formats and with different justifications. In current times, when our contracting partners discuss these types of issues, it is in the Forces best interest to have a national resourcing methodology in place so all of our contracting partners will be faced with a consistent methodology that is supportable, defensible and user friendly.

The ability or inability of the Force to determine, rationalize and provide accurate resourcing forecasts directly impacts front line service delivery, including the need to address members concerns for officer safety through appropriate backup and backfills, and the disparity in individual member workloads created by not utilizing a common method for allocating approved resources. As a result, the issue of resourcing has been on the Officer Safety Committee agenda since its inception, as well as being a floor item at conferences dating back to 1988.

During discussions with the Ontario Provincial Police, it became apparent they are facing many of the same problems the Force has encountered during negotiations with contracting partners on resource allocations. As a result of a series of negotiations, they agreed to participate in the project and have seconded one of their members to represent/liaise on behalf of the OPP. Their expertise and offer to share research material and statistics is enhancing the teams efforts and serve to support the projects final methodology formula.

The mandate of the project is: To develop a resourcing methodology that would determine the human resources required to provide policing services at a given contractual location.

A.2 Current status

Team members have been working on a number of initiatives addressing this complex task in a multi-pronged approach. These include the identification and assessment of any resourcing methodologies currently in use in other police agencies, both in North America and Europe,

as well as ongoing consultations, the identification of consultants and the preparation of building a new resourcing methodology. At the onset of this project, there was a belief that the resourcing model that would fit our needs was likely in place, either in the private sector, “off the shelf”, or in use by another police agency. Our efforts up to this point to identify any such model have been unsuccessful and indications are that we will have to design and build an entirely new model to meet both the Force needs and those of our contracting partners.

A.3 Conclusion

Policing has become a commodity, bought and sold like any other service. From a resourcing perspective, one of the most misunderstood aspects is that we buy and sell time. This appears to be a weakness with our current contractual processes in that we do not have any systems in place that tracks either the members availability or any systems available to accurately track time expended fulfilling our contractual role.

Additionally, any future process must be able to capture the resource costs in time, court decisions such as Stinchcombe and policy changes such as the move to Restorative Justice and Community Policing. All of these have a resource implication which in the past we have not been able to accurately track or forecast. Another policing component, adopted by the Force but not recognized in any of our present day processes, is pro-active policing. It is absolutely essential that this component be recognized in any new process.

Supt. L.T. HICKMAN, Project Leader, National Resourcing Methodology Project,
2001-01-2

Appendix B

Description of data modeling

Data modeling has existed for about 20 years and is used to generate databases and assist in creating computer programs that utilize Database Management Systems (DBMS) such as Oracle, Dataflex, Fox Pro, DB2 and so on. It is a process that is independent of computer jargon and is primarily aimed at business processes and can be used independently of computers in order to understand the workings of a large and complex organization. It is a map of associated data.

A data model's elements consist of entities and attributes which belong to entities. Its operations consist of relationships between entities. Through discussion and understanding of an organization's workings a diagram is produced showing entities of the organization and the relationships between them. The entity-relationship diagram produced is sometimes called a fork notation of a data model.

An entity is a person, place or thing and is denoted by a rectangle with the singular name of the entity placed inside the box. Let us consider creating an entity called Person. Its entity would therefore look like Figure B.1

There are four types of relationships: many-to-many, one-to-many, many-to-one and one to one. In fork notation, the entity having many relationships has a fork attached to it. Suppose we have Person who has many cars over time and that these cars were owned by many people over time. There would therefore exist a many-to-many relationship which could be represented by Figure B.2.



Figure B.1: An entity

Now, a many-to-many relationship as shown above, creates major problems in analysis. Computers, unless they are parallel processors, cannot handle many-to-many relationships and people have problems handling them too. Nevertheless, at higher levels of data model analysis it is good to put simple high-level relationships together to get a top-down view of an organization and resolve the many-to-many relationships later. A data modeler would take the high-level data model diagram which would probably contain many-to-many relationships and resolve it into one containing only one-to-many relationships. In our example above we have a relationship between Car and Person. The provincial government's Department of Motor Vehicles resolves this relationship using a vehicle registration. We note that each Vehicle Registration has one Person and one Car associated with it. We also note that each Person can have many Vehicle Registrations and each Car, over time, can also have many Vehicle Registrations. So, doing one relationship at a time we portray the one-to-many relationship between a Person and a Vehicle Registration as in Figure B.3.

Note that the fork in the diagram is attached to the entity which has a many relationship and there is no fork attached to the entity with only a singular relationship. Here there are many Vehicle Registrations to one Person.

To complete our diagram we would add in the Car entity to yield Figure B.4.




Figure B.2: A many-to-many relationship

The data model has now been resolved and only one-to-many relationships exist. There are also a number of one to one relationships, for example, Person would have a Name, an Address and a License Number which would be considered as attributes of the entity Person because they all describe or identify Person. As a database is being built the attributes form the fields in the record describing the data file of Person. Note that the entities Person and Car have no forks attached to them. They each relate to Vehicle Registration but Vehicle Registration is considered as not relating to either Car or Person. Because no entities are related to either Car or Person, they are at the top of the data model diagram and sometimes their associated data files are called flat data files. Sometimes data model diagrams become very large and cumbersome. I have personally come up with three rules in creating a data model diagram to assist in understanding what it represents and in its overall analysis. They are:

- All forks point downwards. The many entities are below the one entity.
- Lines connecting entities should not cross. This greatly simplifies the tangled appearance of many entity-relationship diagrams.
- There should be no repeated entities. Some data model diagrams repeat entities in

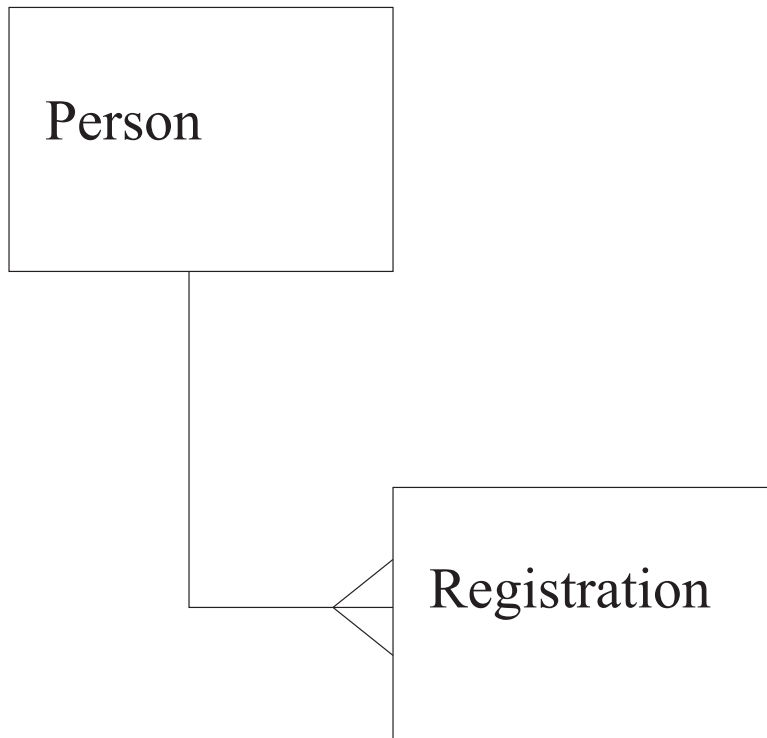


Figure B.3: A one-to-many relationship

different places by representing them as circles. This is confusing since there is no way to know which entities are being placed twice on the diagram.

In the resolved diagram B.4, where there are only one-to-many relationships, the data in Vehicle Registration needs only to contain data to link it to Person and to Car. (As well as any other information which is only pertinent to the Vehicle Registration itself like, dates etc.) If Person has an identifying index and Car has an identifying index, then, to access these two parent files, the child, Vehicle Registration, only needs to contain these indices. All in depth data an enquirer may wish to access can be found by opening the Vehicle Registration file, obtaining indexes of parent files and then very quickly accessing data in the parent files which relate to the child. There is a similarity here between matrices and data file structure in which the indices act as addresses to a cell or record. There is room for more exploration in this direction.

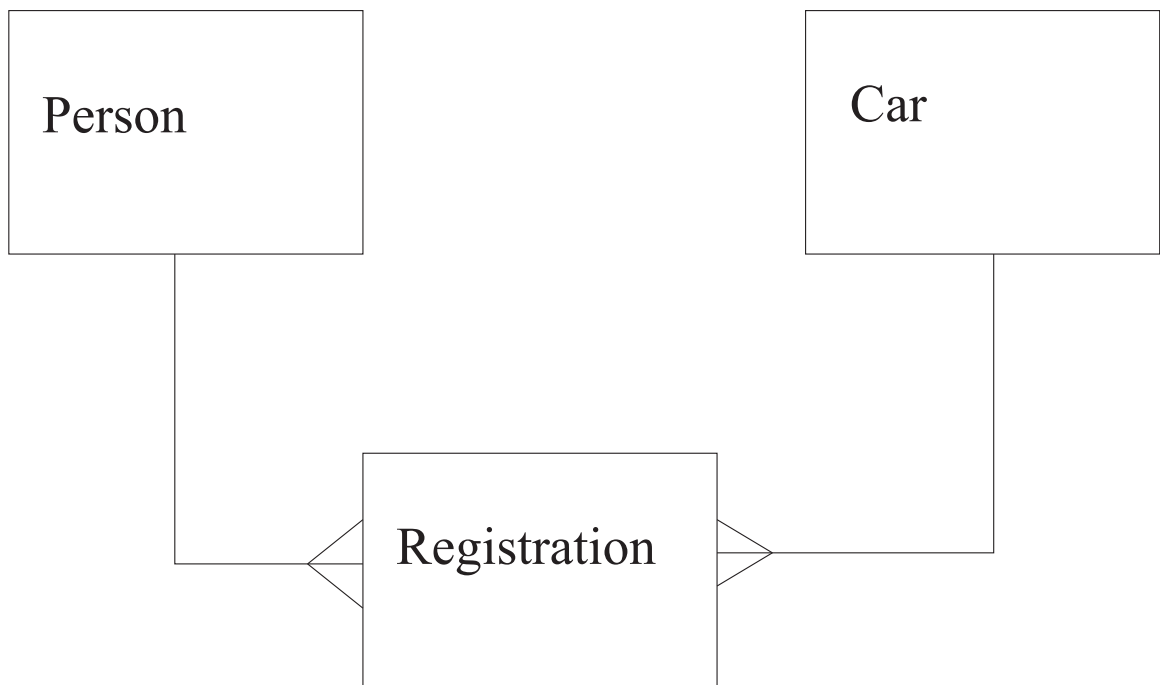


Figure B.4: the resolved diagram

Appendix C

Table of data

Following is a table of the data presented in figure 3.2:

Hour of the day	1	2	3	4	5	6	7	8	9	10	11	12
Monday	40	33	24	15	12	15	19	34	70	84	70	65
	70	70	74	82	80	80	65	76	52	64	72	79
Tuesday	43	43	50	29	19	20	22	31	49	61	70	63
	61	71	75	88	83	90	61	71	84	63	74	68
Wednesday	53	43	45	27	19	17	14	32	41	50	60	59
	57	66	60	62	85	70	74	72	62	63	78	68
Thursday	57	46	40	29	19	12	22	36	43	57	65	63
	45	66	75	73	60	83	71	71	68	66	77	77
Friday	59	46	38	31	18	8	25	41	45	62	75	61
	62	83	77	79	96	77	86	80	112	96	126	128
Saturday	131	94	83	70	35	19	19	23	47	58	55	53
	58	75	74	68	72	73	69	67	82	91	105	123
Sunday	97	70	71	60	34	26	19	19	45	47	55	60
	63	65	63	65	58	67	72	65	79	81	71	57

(C.1)

Bibliography

- [1] ReVelle, Charles S. Rosing, Kenneth E., *Defendens Imperium Romanum - A Classical Problem in Military Strategy*, Amer Math Monthly 107 (2000) No 7 - 585-594
- [2] Larson, Richard C., *A Hypercube Queueing Model for Facility Location and Redistricting in Urban Emergency Services*, Prepared for the Department of Housing and Urban Development, The New York City Rand Institute, New York, March, 1973.
- [3] Larson, Richard C., *Urban Emergency Service Systems: An Iterative Procedure for Approximating Performance Characteristics*, The New York City Rand Institute, New York, January, 1974.
- [4] Kolesar, Peter and Blum, Edward H., *Square Root Laws for Fire Company Travel Distances*, The New York City Rand Institute, New York, June, 1975.
- [5] Hausner, Jack, *Determining the Travel Characteristics of Emergency Service Vehicles*, The New York City Rand Institute, New York, April 1975.
- [6] Chaiken, Jan, M., *Patrol Allocation Methodology for Police Departments*, The New York City Rand Institute, New York, September 1975.
- [7] Kolesar, Peter J.; Rider, Kenneth L; Crabill, Thomas B., *Guidelines for Scheduling Police Patrol Cars*, The New York City Rand Institute, New York, December, 1976.
- [8] Cooper, Leon; LeBlanc, Larry J.; Bhat, U. Narayan, *Introduction to Operation Research Models*, Saunders, Philadelphia, 1977.
- [9] Berman, Oded; Krass, Dmitry; Sethi, Suresh, *Proposal to Model the Staffing Requirement for Primary Response System at the Metropolitan Toronto Police Department*,

Unpublished Report to Metropolitan Toronto Police Department, Toronto, November 9, 1994.