OPERATIONS RESEARCH CENTER

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

The Potential Impact of Expert Systems in Urban Police Services

by

Jacqueline Kee

Richard C. Larson

OR 140-85 June 1985

The research for this report was funded in part by the National Institute of Justice on grant no. 83-IJ-CX-0065 and by the National Science Foundation on grant no. 8411871-SES.

The opinions expressed in this report are those of the authors and do not necessarily reflect the views of the sponsors.

ABSTRACT

We give a short overview of *expert systems* and discuss the potential applications of these systems in urban police services. Two applications in the prioritization of incoming 911 emergency calls and the dispatch of police units to service these calls are discussed in detail.

CONTENTS

1. INTRODUCTION

What are *expert systems* and why is this term suddenly catching the attention of so many people? How are these systems any different from the computer programs created in the past? And, more importantly, how will they help enhance our daily lives? In this paper, we will try to answer some of these questions and provide some references which the reader may consult for additional information. Also, we will study applications in the area of urban police services where we feel that *expert systems* can make a significant impact.

In urban police services, the resources are limited and bad decisions may result in the damage of property or loss of lives. Any overworked urban police department would welcome some assistance in the management of its resources. The existing computer technology in most police departments serve as electronic file drawers where information can be entered and retrieved easily and also as expensive calculators where certain statistics can be computed from the data entered into the system. Rarely do these systems provide any strategies or advice on how to carry out a service, such as the dispatching of a patrol car to answer a 911 emergency call. Frequently, overworked and underpaid police personnel make crucial decisions based only on their intuition and experience. It is evident that there is a need for a more advanced technology, like that of expert systems, to help in decision-making processes of the police department so that the quality of police services to the public can be improved.

The paper is organized as follows. In section 2, we present an overview of expert systems. This section will provide the reader with some information on the structure of expert systems and also describe an expert system called MYCIN, which is used in medical diagnosis. Section 3 examines the 911 call categorization process and identifies two categories of calls where expert systems technology can be of use. Section 4 investigates the addition of an expert systems factor into an Operations Research (OR) algorithm called the Exact Hypercube Queueing Model, which is used to analyze problems of vehicle location and response district design in urban emergency services.

2. AN OVERVIEW OF EXPERT SYSTEMS

To help us define expert systems, we first need to discuss the motivation behind building such systems. Expert systems evolved from research into developing computer programs to solve very difficult tasks. Medical diagnosis and mineral exploration are two examples of such tasks. Researchers realized that most complex tasks are not solved by following straightforward algorithms and hence can not be solved by conventional programming techniques. What the researchers needed instead were computer programs that were similar in design to the human reasoning process and that could capture the knowledge of human experts. They needed *intelligent* computer programs; programs that could represent knowledge about complicated tasks and that could make inferences from this knowledge. Another motivation was that conventional programs are understood by those that create them but not always by their users. It is rare for a user to be able to make changes or additions to these programs. If a program could be replaced by sets of rules that express some fundamental knowledge about the domain of the problem to be solved, then making modifications would only involve adding new rules or changing old ones.

Inevitably, these researchers turned to the field of artificial intelligence in their quest for *intelligent* computer programs. The result was some problem-solving computer programs called expert systems. Expert Systems are an attempt to identify, formalize, encode, and use the knowledge of human experts as the basis for a high-performance program [Davis, 1984]. These systems are used to solve very complicated tasks that require a high level of expertise and are designed in a radically different way from previous problem-solving programs. Conventional computer programs were mere coded instructions to the computer on how to carry out a task. Just like amateurs carrying out an assignment, these codes knew what to do but had no knowledge of why they were doing it. The new systems, on the other hand, possess the knowledge about the problems they are given to solve and are able to make inferences from their knowledge to solve specific tasks just like human experts. Because of their ability to understand what they are doing, like the experts, these systems are able to explain the solutions they find and are flexible enough to learn from their mistakes. Therefore, it is of no surprise that these programs were called expert systems.

2.1 STRUCTURE OF EXPERT SYSTEMS

Expert systems can be made up of two or more components; the number of components varies according to the type and function of each system. The following are some basic components present in most systems:

- the *knowledge base*

  This contains all the necessary knowledge about the expert system's domain. The domain is essentially the problem space which contains all possible types of problems that the expert system might be given to solve.

- the *inference engine*

  This component is the interpreter that selects the relevant information from the knowledge base in order to solve a particular task. It only uses the knowledge about the specific task in order to solve the task.

- the *database*

  This component stores all the data about the domain. When the inference engine is working on a specific task, it might occasionally refer to the database for some data about that task.

An important feature of an expert system is that the knowledge base is separated from the inference engine. The reason why the knowledge base is kept apart from the inference engine is to make the knowledge more easily identified, more explicit, and more accessible [Davis, 1982]. Knowledge about a particular domain is stored in the knowledge base; so whenever changes or additions to the domain knowledge are required, it is easy to know exactly where to do it. Similarly, if an error in the inference engine needs to be traced, the problem can be located easily. Moreover, since the inference engine is not dependent on the domain knowledge, it can be used with different knowledge bases. An example of a domain-independent system is EMYCIN [van Melle, 1979]. EMYCIN's knowledge base can be replaced with a knowledge base from another expert system, and it will function in the same way as the system where the knowledge base was taken from. MYCIN, an expert system that will be discussed later in the chapter, has a knowledge base that is compatible with EMYCIN's inference engine; the combination of MYCIN's knowledge base and EMYCIN's inference engine performs just as well as the original MYCIN system.

The first step in building an expert system in any domain is to obtain the knowledge. Acquiring the knowledge from the experts is not an easy task since it requires that experts formalize their knowledge and express intricate relationships between the basic concepts in their domain [Hayes-Roth, et al. 1983]. This knowledge-acquiring process is called *knowledge engineering*. Recent developments have produced some knowledge engineering tools that can assist in the building of a knowledge base. EMYCIN, which was introduced earlier, possesses this ability to help transfer expertise from the expert to the knowledge base [Buchanan, Shortliffe, 1984]. Another knowledge engineering tool is EXPERT which was developed for building consultation models based on classification problems and is used primarily in medical diagnosis [Weiss and Kulikowski,1979].

The next step in the process of creating an expert system is to find a suitable form of representation for the knowledge acquired. The method of representation chosen depends largely on the domain and the results to be achieved; the method should be able to model all the intricate relationships in the system's domain. It is important that the selected representation make the knowledge as transparent as possible for future changes and improvements. Then after ten years, any user of the system will still be able to understand what is in the system and can update the knowledge. The most common method of representing knowledge is by using rules. The rules are usually *if-then* statements, i.e., *if* certain conditions are satisfied, *then* the action will be carried out. The action carried out by the rule usually involves the calling of another rule, and this sequential calling of rules continues until a solution has been found for the given problem. Examples of expert systems that use rules to represent the knowledge base are MYCIN which is used in the diagnosis of infectious diseases [Shortliffe, 1976] and R1, an expert computer configuration system at Digital Equipment Corporation [Mcdermott, 1982].

2.2 MYCIN - AN EXAMPLE OF AN EXPERT SYSTEM

MYCIN is an expert system developed by E.H. Shortliffe in 1974 at Stanford to assist in antimicrobial selection. An antimicrobial agent is any drug designed to kill bacteria or to arrest their growth. Only specific antimicrobial agents are effective against particular types of bacteria. Drugs which are useful against certain types of bacteria are often not effective against other types and in some cases, they have adverse effects. In a study conducted in a community hospital, about 50% of the patients treated with antimicrobial agents were victims of the misuse of the drugs [Roberts, Visconti, 1972]. Misuse of antimicrobial agents includes wrong selection of the drug, wrong dosage, allergic reactions to the drugs and inappropriate combination of the drugs [Shortliffe, 1976]. Because of the misuse and the overuse of antimicrobial agents and the lack of human experts, assistance was needed to help select the right drug for a specific patient.

MYCIN, like the physician, follows four steps in the process of antimicrobial selection [Shortliffe, 1976]. It first decides whether the patient has significant infection and if the answer is positive, it then tries to determine the likely identity of the offending organism. Next, it decides what drugs are effective against the organism and chooses the most appropriate drug given the patient's condition . These steps are encoded in 500 decision rules which are in the knowledge base of the MYCIN system. These decision rules, which are obtained by interviewing expert physicians, are in the form of if-then statements where each if-then statement represents a small *chunk* of knowledge. Besides identifying the bacteria causing the infection and prescribing a drug for the infection, MYCIN is also able to explain how it arrived at these conclusions. In addition, it permits the user to change a rule or to add a rule to its knowledge base. A typical rule in MYCIN looks like:

IF: 1) the stain of the organism is gram-negative, and

2) the morphology of the organism is rod, and

3) the aerobicity of the organism is aerobic

THEN:

there is strongly suggestive evidence (0.8) that the class of organism is enterobacteriaceae.

The number in the round brackets, i.e., (0.8), is called the certainty factor (CF). A special feature of MYCIN is that the uncertainty in medical diagnosis, due to the lack of perfect information, can be reflected by the assignment of a certainty factor to each MYCIN rule. The CF measures the belief that the decision rule will be true given the premises of the rule are satisfied and it is initially assigned by expert physicians during the development of the knowledge base. The CF has a range from -1 (the rule is false) to + 1 (the rule is true). When the CF is in between the values of -0.2 to +0.2, no conclusion can be drawn. At each step in the MYCIN's deduction process, the certainty factors are combined so that the final conclusion has a certainty factor assigned to it.

In MYCIN's inference engine, the rules are invoked in a backward-chaining deduction process. All the rules which state the identity of the bacteria are evaluated first. Then the process uses other rules to infer backwards so as to check if the clinical information available agrees with the premises of the initial identification rule. After the bacteria is identified, MYCIN uses other rules to select the appropriate antimicrobial agent. This method of deduction allows MYCIN to justify every step of its deduction process to the user. For example, if MYCIN asks the user if his patient has a particular symptom, it can justify this question by saying that the presence of that symptom is one of the premises of the rule which would identify the bacteria causing the infection. The following excerpt from Buchanan and Shortliffe's Rule-Based Expert Systems shows how MYCIN justifies its questions:

** indicates user input

```
** WHY DO YOU ASK WHETHER THE PATIENT HAS A FEVER OF UNKNOWN
   ORIGIN?
The rules listed below use:
        whether the patient has a true fever of unknown origin

RULE180

IF:   1) The site of the culture is blood, and
      2) The number of cultures in the series including the
         culture is greater than or equal to 1, and
      3) The number of cultures in this series which were
         positive for the organism is less than or equal to 1,
         and
      4) The patient has a true fever of unknown origin, and
      5) Cardiac-surgery is a relevant item from the history of
         the patient
Then:    There is suggestive evidence (.6) that the infection is
         infective-endocarditis

** WHAT DOES GU MANIPULATION TELL YOU ABOUT THE IDENTITY OF AN
   ORGANISM?
The rules listed below use:
        whether the patient has had a genito-urinary manipulative
            procedure
to conclude about:
        the identity of the organism
156, 163, 190
Which do you wish to see?
        ** 156

RULE156

IF:   1) The site of the culture is blood, and
      2) The gram stain of the organism is gramneg, and
      3) The morphology of the organism is rod, and
      4) The portal of entry of the organism is urine, and
      5) The patient has not had a genito-urinary manipulative
         procedure, and
      6) Cystitis is not a problem for which the patient has
         been treated
THEN:    There is suggestive evidence (.6) that the identity of
         the organism is e.coli
```

The following diagram from Shortliffe's Computer-Based Medical Consultations: MYCIN presents an overview of the structure of MYCIN:

STRUCTURE OF THE MYCIN EXPERT SYSTEM

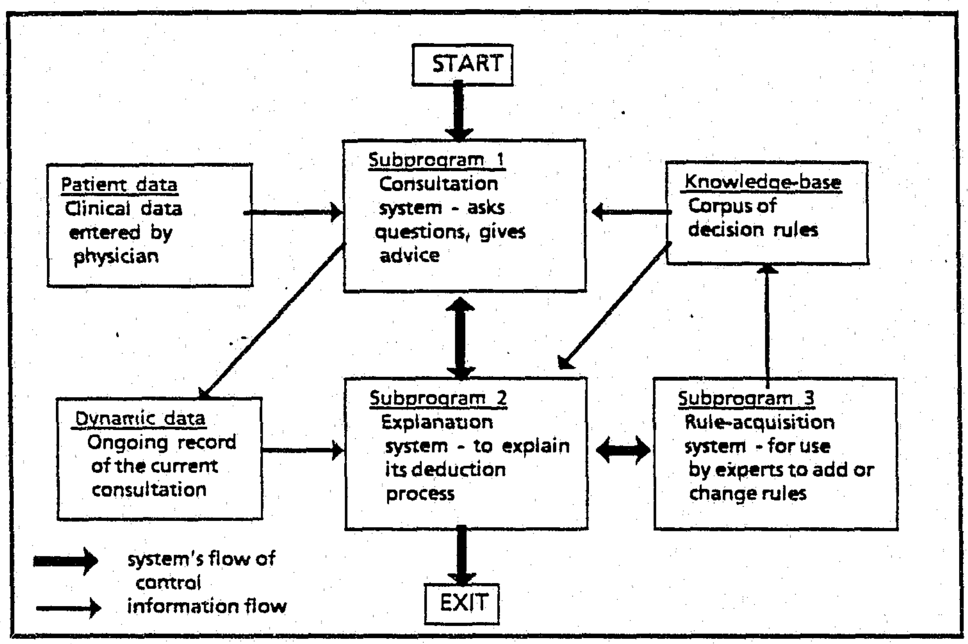

2.3 THE POTENTIAL OF EXPERT SYSTEMS

The rate of emergence of expert systems in a wide range of fields has been growing rapidly in the past years. What started out as a research project in an academic environment has ended up in numerous industrial facilities and commercial offices. Today, expert systems have been developed in such diverse fields as medical diagnosis, chemistry, mathematics, mineral exploration and management. This rise in popularity of expert systems is not only due to their ability to capture the knowledge of human experts. The systems have other appealing features such as their portability, i.e., they are independent of computer architecture and hardware. They are not restricted to run only on the machines where they were initially designed; it is conceivable that expert systems could be build to run on personal computers in the future. Another attractive feature is the transparency of the programs. Since the programs are usually written in LISP, which resembles English text to some extent, any user can understand how the program works, and can easily make changes to suit his goals.

The future ahead for expert systems looks very promising and challenging indeed. Just as changes were brought on by the industrial revolution and later by the computer age, so also will there be major improvements in our society from the emergence of expert systems. These systems will have tremendous social and economic impact on our daily lives; expert consultations will become very accessible, inexpensive and reliable since expertise will be available wherever an expert system can be set up.

3. IMPROVING THE 911 CALL CATEGORIZATION PROCESS

In most cities, there are emergency numbers for citizens to call the police, the fire stations or the ambulance service twenty-four hours a day. The 911 number in Boston is one such emergency number. Incoming 911 calls are first screened at a central police communications center by phone operators whose job is to determine what types of services are needed. If police service is required, the necessary information is channeled to the police dispatchers who then assign patrol cars to service the call.

The 911 operator plays an important role in determining the outcome of life-or-death situations. For example, in the Village of Kenmore in Buffalo, New York, a woman was brutally murdered by a burglar because the 911 operator mistook the street address which the woman gave as another street in the city of Buffalo. The police officers were sent to the wrong location and the error was not detected until fourteen minutes later when it was much too late. Yet another example occurred when the misinterpretation of the urgency of a call resulted in the death of a woman in Dallas, Texas. The nurse who was handling the 911 calls repeatedly asked to speak with the woman who was having difficulty breathing although the caller said that the woman was unable to talk. Finally, an ambulance was dispatched but the woman could not be saved.

More importantly, policemen rely heavily on the information obtained by the 911 operator to determine what types of incidents they are responding to and how defensive they should be upon entering the locations of the incidents. From interviews with patrol policemen, it was pointed out that there were certain occasions when the policemen were not given any other information except the address of where they were to go to check out a disturbance report. These policemen were not told the cause of the disturbance; for all they knew, it could just have been a group of youngsters hanging out on a sidewalk. However, as it turned out in some serious cases, the disturbances were caused by a person with a dangerous weapon.

The 911 operator's role is clearly a crucial one and demands accurate diagnosis of the 911 emergency call. Expertise in the categorization process of 911 calls is acquired through the experience of dealing with many different types of cases as in the medical diagnosis process. Therefore, employing expert systems in the categorization of calls is natural. Such systems would diminish the reliance placed on the experience of the 911 call-taker. This is particularly important given the high turnover rate of civilian operators. For instance, the turnover rate of 911 operators in Boston's Police Department went up to as high as 200% in 1984.

3.1 The 911 Call Categorization Process

The response time for each incoming 911 call depends on the initial categorization and prioritization of the call by the operator. The number of categories for the 911 calls can range from fifty to over a hundred and these codes cover criminal incidents, minor disturbances, traffic offenses and many other incidents that are frequently reported to the 911 system. Most police departments operate on a three-priority system. Priority 1 calls involve life-threatening incidents or crimes in progress and require immediate response. Priority 2 calls include property-threatening incidents or crimes that have already been committed and do not need an immediate response but will be serviced immediately if there are no priority 1 calls awaiting service. Priority 3 calls include minor violations and will be serviced only when there are no longer any higher priority calls left to be serviced.

In larger police departments, the 911 operator types the information given by callers into a terminal attached to a computer-aided dispatch (CAD) system. When a new call comes in, the 911 operator calls up an information tableau on the terminal screen and proceeds to enter the required information in the tableau. Most tableaus require the caller's name, address and telephone number and the exact location of the incident being reported. In emergency cases, the 911 operator only needs to get the address of the incident before sending the information to the dispatcher. Depending on her experience, the 911 operator may ask for the time of the incident, whether there were any injuries, whether any weapons were present, and other information relevant to the incident being reported. This additional information is important in helping the operator to correctly categorize and prioritize the call. On most CAD systems there are no specific information tableaus to help the 911 operator to elicit the information needed for different types of incidents. The same basic tableau is usually used for all the different types of incidents. It is usually left up to the operator's discretion to ask for more information to help her in the decision making process.

Most 911 operators hold low paying jobs with small pay increases over time. Their jobs can be very demanding and stressful on busy weekend nights. They are also very easy targets for abuse by irate citizens and by children playing pranks on the telephone. Because of this working environment, it is no wonder that the turnover rate of civilian operators is quite high. Also, the formal training that 911 operators receive is minimal; a 175-city study by the Police Executive Research

Forum revealed that less than 70% of the surveyed police departments provided any training and the training that was provided usually only involved 1 to 2 weeks of basic training [Sumrall, et al. 1981]. In addition, the 911 operators are under little supervision as their role is not considered as crucial as that of the patrol policemen [Percy, Scott, 1983].

The high turnover rate of civilian 911 operators coupled with the minimal training and supervision of the operators can lead to relatively ineffective call-taking procedures. The 911 operators have to make quick decisions on every call they receive and if they are unqualified to do so, there is a higher chance of making fatal errors. Incoming 911 calls which are obvious emergencies will get assigned high priorities by the operators. So, instead of looking at these emergency calls, we will focus on calls where it is not evident that there are emergencies but where there is a possibility that they might in fact be or might escalate into priority 1 calls. We will set up some protocols which will help the 911 operators evaluate two examples of the types of calls that tend to be ambiguous, i.e., the family trouble incidents and the barking dog incidents.

These protocols are illustrative of a broader expert system's methodology that has yet to be tested on the CAD system. A possible scenario is the following: when the 911 operator receives a call, he first channels the call to fire station or the ambulance service or the police or decides that the call does not need any service. If the call needs police service, he types in the category of the call into the CAD expert system. Depending on what category it is, the appropriate protocol is called up and the CAD expert system leads the 911 operator through asking the caller a sequence of questions. The aim of these protocols is to help the 911 operators correctly prioritize the calls and obtain sufficient information to help the police service the call.

Like the MYCIN system, the CAD expert system should be able to explain why it chose a particular sequence of questions and how it assigned a priority to a certain call. It should also allow the user to change or to add rules to existing protocols. Besides being used in the actual 911 categorization process, the CAD expert system can also prove to be a useful tool for the training of new 911 operators.

An important distinction between a medical diagnosis expert system like MYCIN and a CAD expert system is that there is no real-time constraint in the medical consultation session whereas in the 911 categorization process, there is. In the CAD system, if the caller faces an emergency, there is no time to go through an extended sequence of questions. The 911 operator should be able to call up an alternative protocol to handle emergency calls so that the minimum amount of information can be obtained from the caller as fast as possible. Also, some callers are not co-operative and might not be willing to answer the questions asked by the 911 operator. The CAD expert system should be flexible enough to handle these types of situations.

3.2 FAMILY TROUBLE CALLS

Family trouble incidents reported to the 911 system can range from simple verbal abuses to homicides. Because of their ambiguous nature, the 911 operators need more information before they can determine the correct priorities of the family trouble calls. In a study conducted by Eva Buzawa in 1972, more often than not, after having screened the calls, the 911 operators delayed or prevented the dispatch of officers to domestic violence incidents. Questions about the extent of the injuries or the presence of a weapon or the imminence of the offender's return do not appear to have been asked in most cases by the operator unless the caller explicitly volunteers the information [Parnas, 1971]. It is evident then that there is a need for some protocols for the 911 operators to follow when they receive family trouble calls.

To further emphasize the importance of protocols for domestic violence incidents, we note that the FBI's Uniform Crime Report (UCR) states that approximately two-third's of all murders were committed by relatives or friends of the victims and that family trouble calls were the third highest cause of death among police officers killed in the line of duty. Also, according to the statistics released by the Boston Police Department, family disturbance calls were one of the ten most-reported incidents in Boston.

The first step in setting up useful protocols for the 911 phone operators is the identification of those characteristics common to most family trouble calls. In a study conducted by the Police Foundation in Kansas City, the three main predictors of physical violence in family trouble incidents were the presence of a weapon, a history of previous disturbances and the presence of alcohol. The presence of a gun was the strongest indicator that the family dispute would involve physical violence and it accounted for 29% of the difference between disputes involving physical violence and those that did not. The second highest indicator was a history of previous disturbances with a 11.4% difference. This same study also revealed that over a two-year period, in 85% of the family trouble calls that resulted in homicides, the police had been called in to intervene in previous disturbances before the homicides occurred. In fact, in about 50% of the calls the police were called in *five* *or more times* to previous disturbances before the homicides occurred. There was a case in Boston where the police were called to intervene in a family disturbance about eight different times and on the last call the police found that the man had killed his wife and seven children. The third highest indicator, with a difference of 11.3%, was the presence of alcohol. Alcohol when abused, presents a startling correlation with crime and violence. One possible explanation, as noted in a staff report to the National Commission on the Causes and Prevention of Violence, is that alcohol has the effect of removing a person's inhibitions and reducing control of his actions.

From interviews with experienced 911 operators and from tapes of 911 calls, it is apparent that the expert 911 operators do rely on the three above-mentioned indicators of physical violence to help them determine the urgency of the family trouble calls. Our research also showed that the callers reporting family trouble calls fall into two categories: observers or participants. The category of observers also includes partial observers and listeners. Thus, we need two different protocols for the family trouble calls as the caller who is the participant has a different perspective of the incident than the observer.

Based on the above results, we have suggest the following protocols for family trouble calls:

PROTOCOL FOR FAMILY TROUBLE CALLS

(caller is a victim)

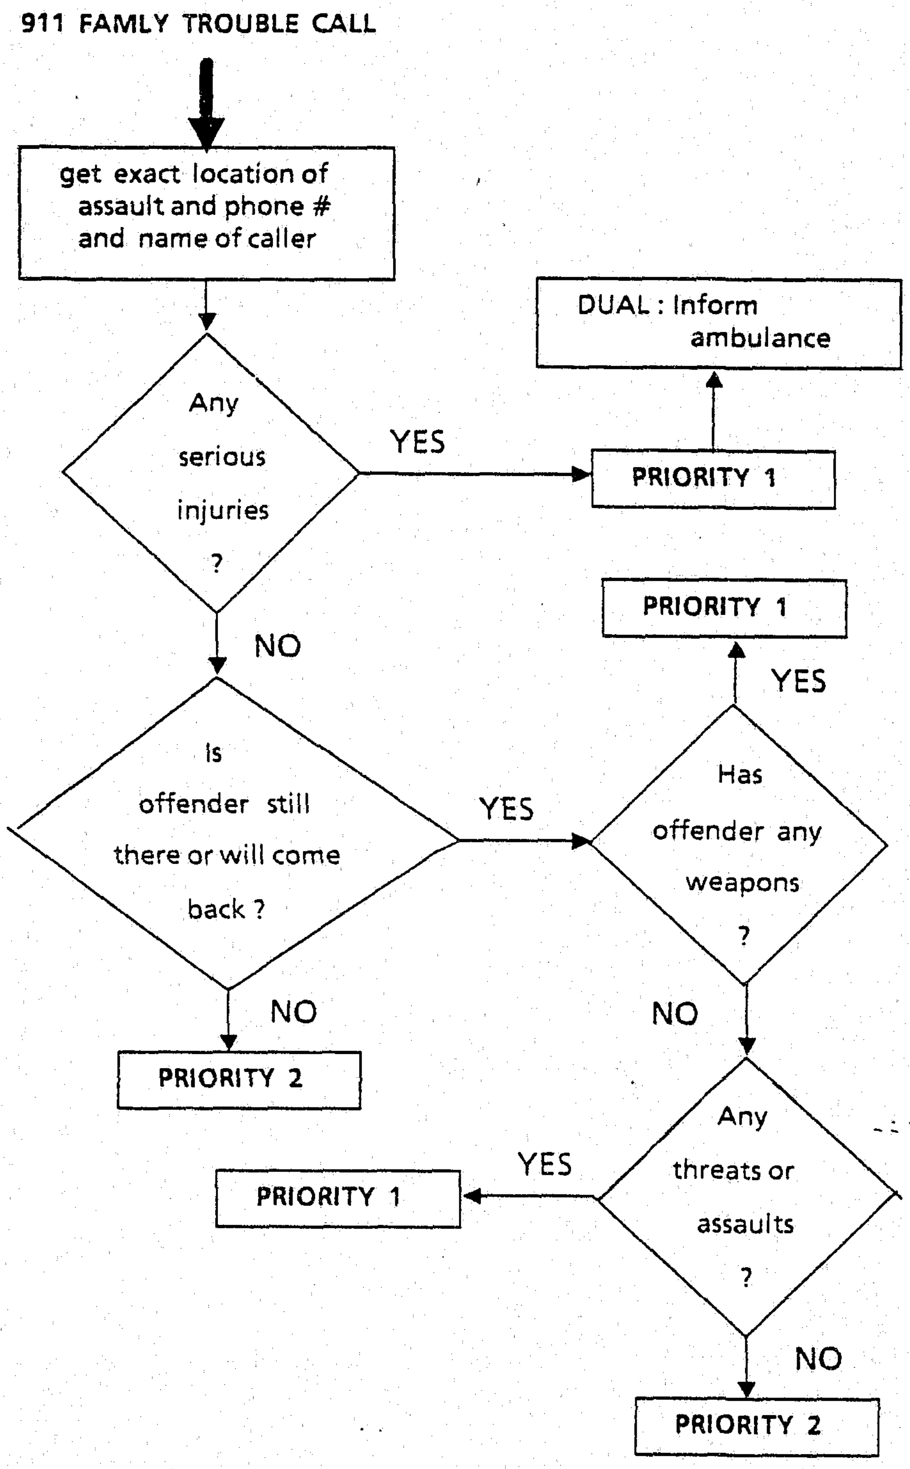

PROTOCOL FOR FAMILY TROUBLE CALLS

(caller is an observer)

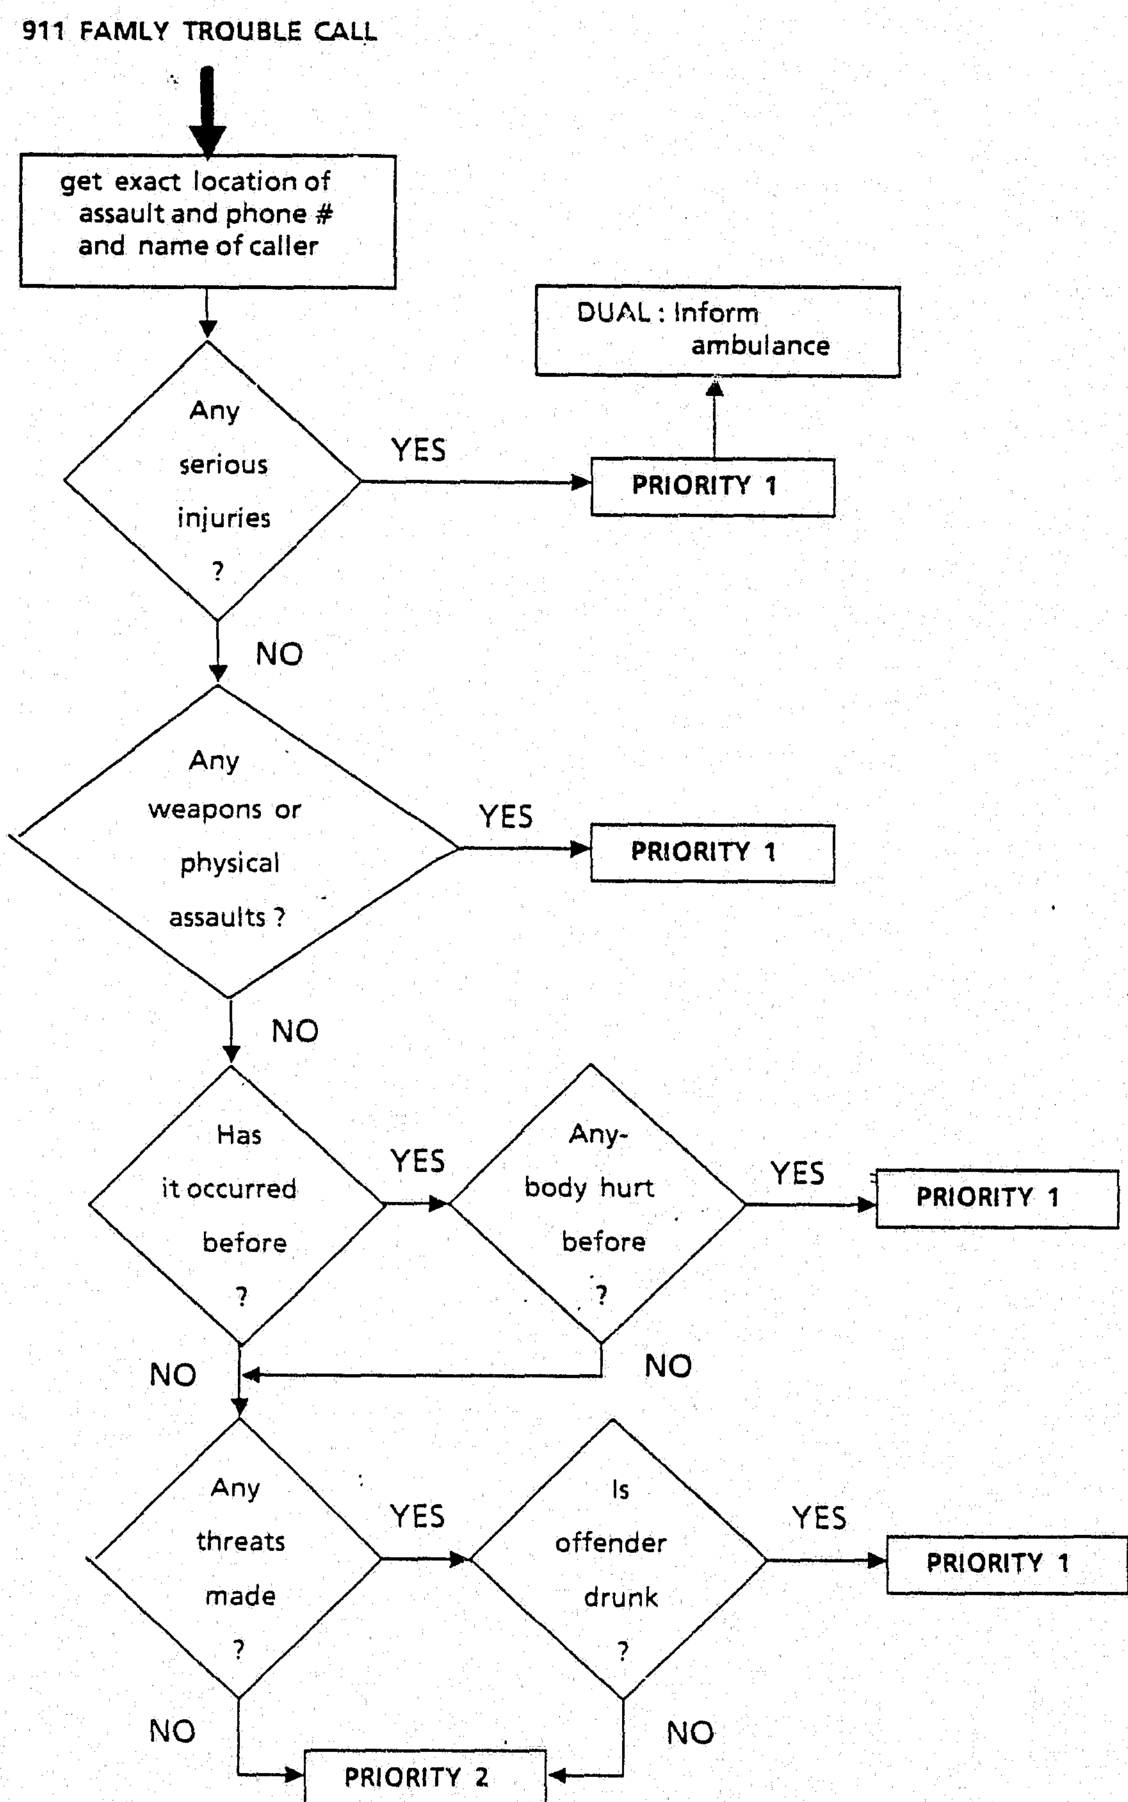

3.3 BARKING DOG INCIDENTS

Barking dog reports are rarely regarded as emergency calls and are usually designated as priority 3 calls by the 911 operators. However, there have been some cases in the past where barking dog incidents have escalated into priority 1 incidents. For example, a complaint about a barking dog in the early hours of the morning actually turned out to be a burglary attempt that resulted in the murder of the dog's owner. In another case, when police responded to a barking dog complaint, they discovered that a rape had occurred. Our aim then is to set up a protocol which will make the 911 operators aware of the possibility that some barking dog incidents might escalate into more urgent calls.

Based on our interviews with experienced 911 operators and police patrolmen, we suggest the following protocol to handle barking dog incidents:

PROTOCOL FOR BARKING DOG CALLS

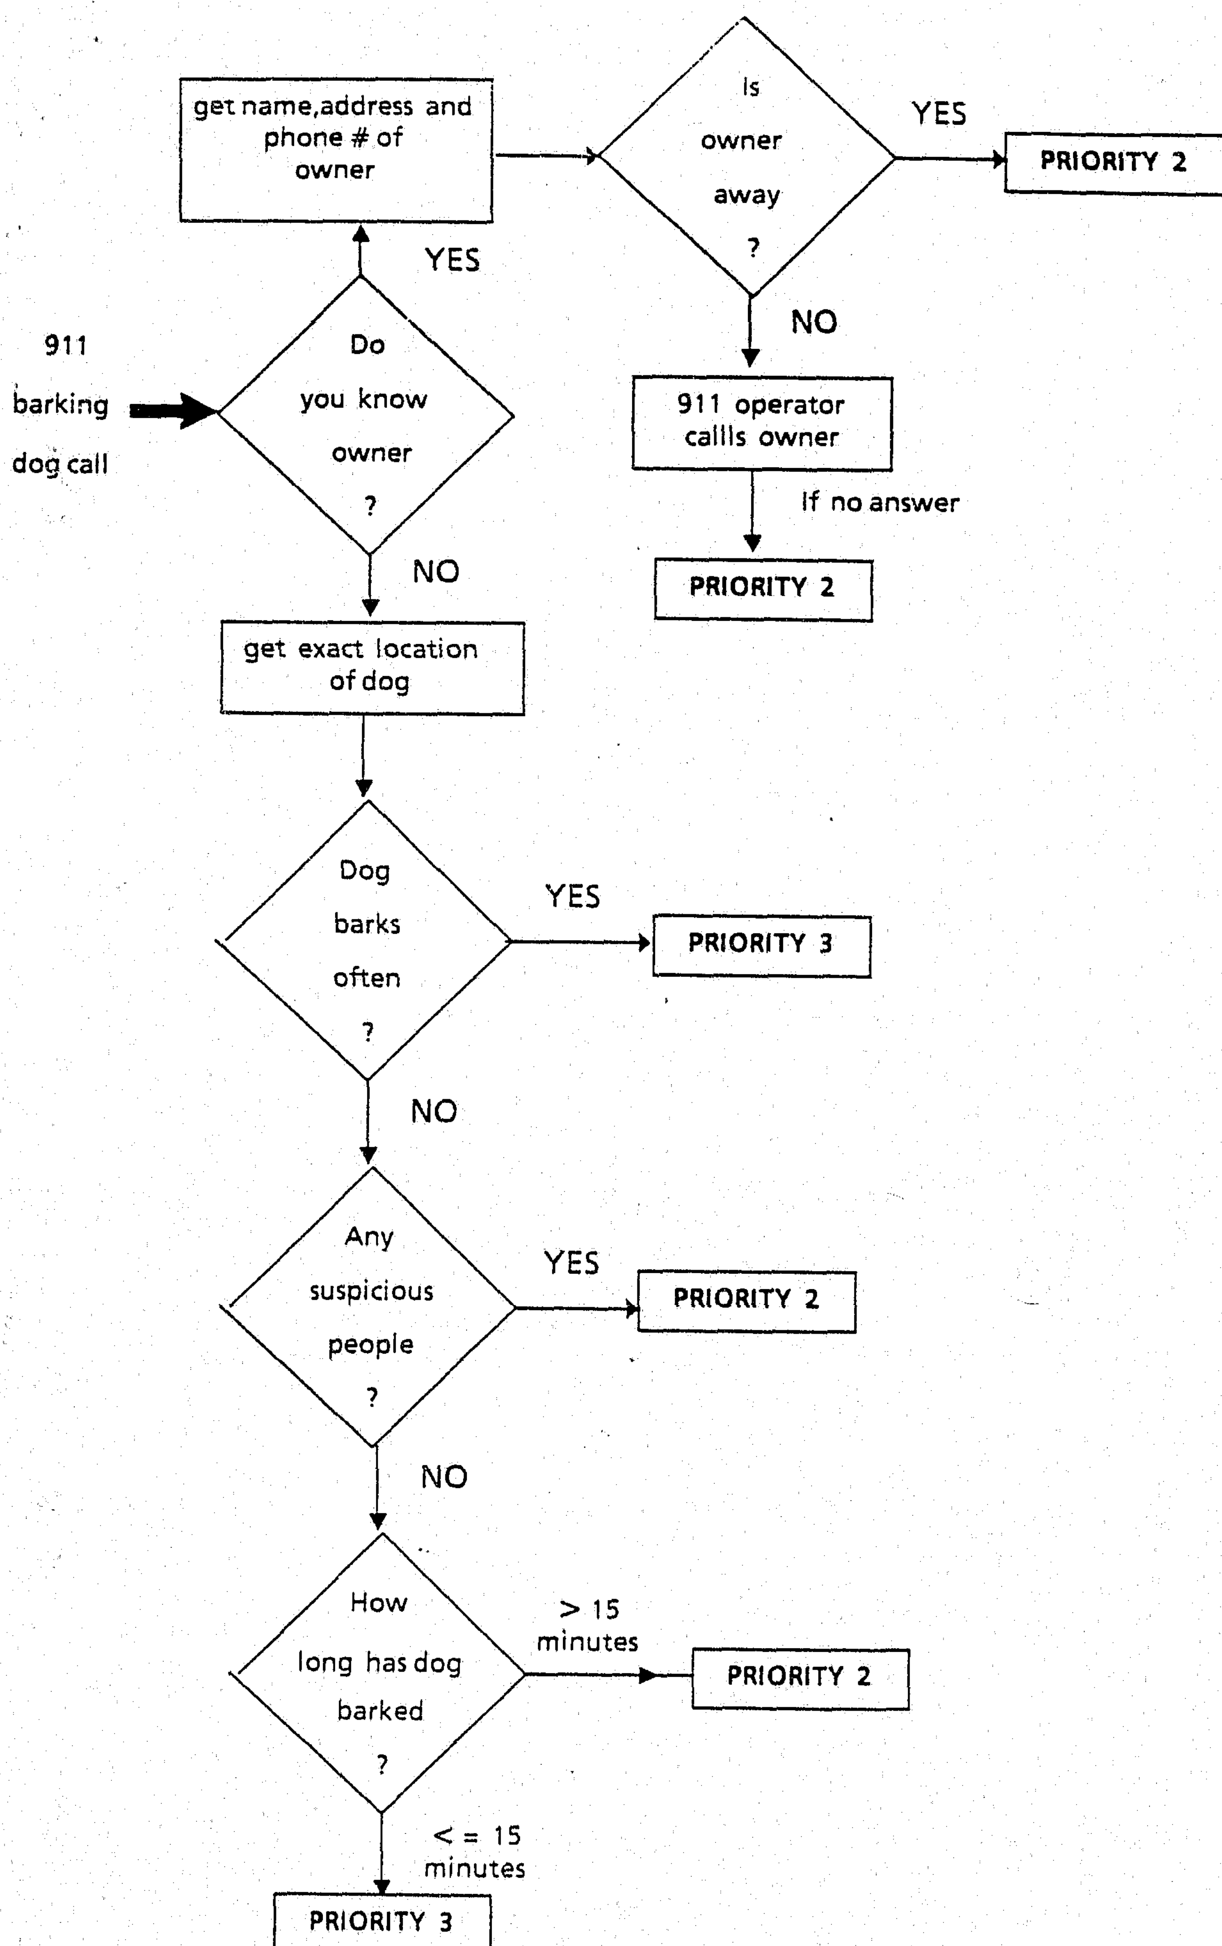

3.4 FUTURE RESEARCH

In this paper, we have only studied two categories of the many types of 911 calls. Protocols similar to the protocols we have presented can also be set up for the other types of calls, in particular, for calls which are usually ambiguous such as disturbances and suspicious persons. Devising these protocols takes a lot of time since interviews need to be performed with experienced 911 operators, other police personnel, and sometimes with psychologists, and other social workers. These interviews with the experts in the various related fields will help determine what questions the 911 caller should be asked and in what sequence these questions should be presented. Listening to tape recordings of 911 operators handling incoming 911 calls can also help in the design of suitable protocols.

After the protocols have been set up, careful and vigorous testing should be performed on these protocols. Tests can be performed by reenacting 911 calls which were recorded on tape. In these reenactments, the protocols will guide the 911 operator through the call categorization process. The advice of these protocols will then be compared to the actual decisions of the 911 operators who handled the calls and the eventual outcomes of the calls. If the protocols are found to perform badly, then corrections should be made and the tests should be run again. Needless to say, the protocols should not be implemented until they perform as well as any experienced 911 operator .

4. EXTENSIONS TO THE EXACT HYPERCUBE QUEUING MODEL

The Exact Hypercube Queueing Model is a model developed by Professor Richard C. Larson to analyze facility location, and redistricting problems in urban emergency services [Larson, 1978]. The hypercube queueing model is essentially an analytical OR tool which explores the operational behavior of an urban emergency service system under various resource-allocation strategies.

The main criticism of OR-type programs, such as the hypercube model, is that the programs do not take into account the human factor present in most decision-making processes. OR Techniques and algorithms are usually too rigid and often disregard idiosyncrasies in the system that a human decision-maker would otherwise consider. One solution to this drawback is to incorporate expert systems techniques into the OR algorithms. In this section, we will study the addition of the expert systems factor to the hypercube model as an example of the *marriage* between expert systems and OR.

4.1 DESCRIPTION OF THE EXACT HYPERCUBE MODEL

A detailed account of the structure of the hypercube model can be found in the second chapter of Larson's Police Deployment: New Tools for Planners. Briefly, the region in which the model provides service is divided into small geographical atoms which can be assigned to a patrol unit's beat. Over a long interval of time, it is possible to obtain the data for the arrival rates which is the average number of calls for service originating from each atom and the service rates which includes the travel time of each patrol unit to the atom and the on-scene service time .

Exactly one unit will be dispatched to service a call from any atom. If the first preferred unit is busy, then the next available unit will be sent. We will use the hypercube model to help us analyze this queuing system. There are two possible states a patrol unit can be in: the unit is either free (which is denoted by 0) or the unit is busy (which is denoted by 1). Each state in the state space of the hypercube model is a combination of the states of all the patrol units and can be pictured as a corner of a regular cube. For example, the states in a two-server system can be represented by the corners of a square:

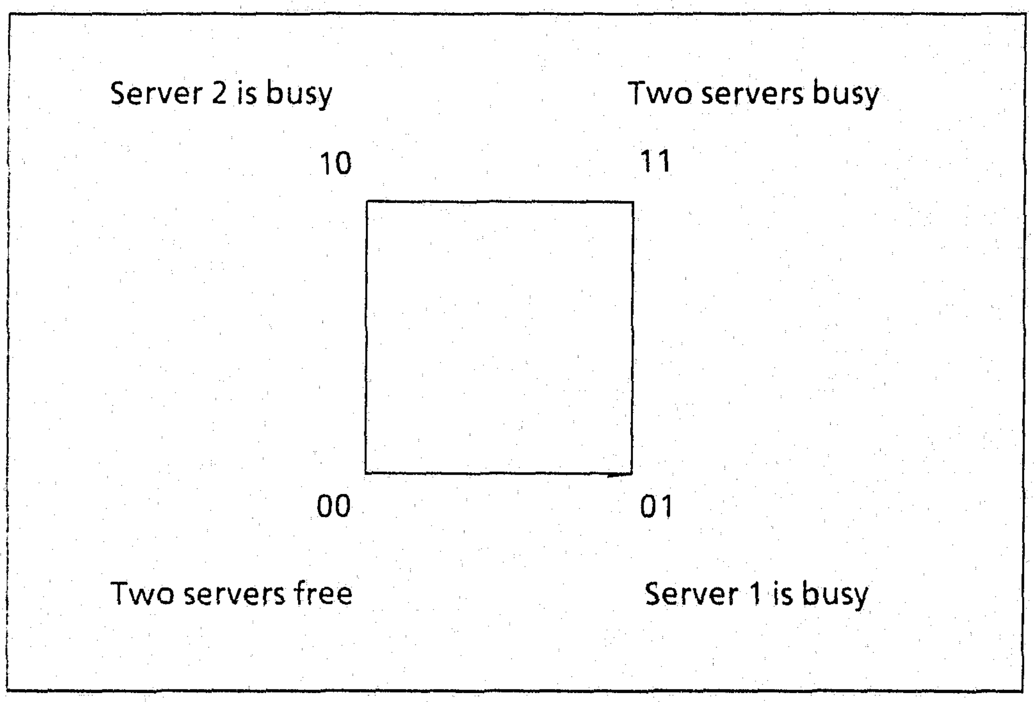

Only one-step transitions are allowed and the transition from one state to another occurs along the edges of the cube. The upward transitions are transitions that occur when a unit is dispatched and their rates can be calculated given the arrival rates of calls for service and the dispatch preferences for each atom. On the other hand, downward transitions are transitions that occur when a unit becomes free and their rates are just the service rates.

Over a long period of time, the rate of transitions into a state equals the rate of transitions out of a state. The system is then said to be in a steady state. By writing down the steady-state equations for each state and then by solving for them, we obtain the steady-state probabilities, i.e., the average proportion of time spent in each state. These steady-state probabilities can be used to calculate the performance measures of the system such as the workload rates for each patrol unit, the travel times, the workload imbalance and the interatom dispatch frequencies. These performance measures will allow the user to explore different beat designs in order to find the best design suited for their goals. For example, one user might like to design her beats to minimize the travel times of the patrol units, while another might prefer to have equal workloads in all his beats. The latter criteria is usually chosen when police beats have to be redesigned so as to reflect the changing distribution of criminal incidents and beat population. This is because if the distribution of workloads becomes the same in all the beats, then the police response to each call could be much faster than for an uneven distribution of workloads.

4.2 ADDITION OF THE EXPERT SYSTEM FACTOR

The current hypercube queueing model only recommends that the first free and preferred server be sent to service calls from a given atom. However, in most real-life situations, the police dispatcher also considers other patrol units as possibilities to be dispatched. Although the travel times of these other units might be longer, there are other factors that the dispatcher might consider. For example, an experienced police dispatcher takes into account the need for a bilingual patrol unit and the expertise of each unit in dealing with certain types of calls. Also, the dispatcher might know that the most preferred unit is overworked and fatigued, and would not dispatch that unit to service a call.

We propose making an extension to the exact hypercube model which will allow the possibility that the next three free and preferred servers will be dispatched. In our extended version of the exact hypercube queueing model, the probability that we will dispatch the second preferred server instead of the first preferred server is expressed as a function of the difference in travel time between the two servers. The rationale for making this probability a function of the difference in travel times is to reflect to some extent the costs involved when we choose a longer travel time. The probability that we will use the third preferred server is also a function of the difference in travel time between the first preferred and third preferred servers multiplied by the probability that we did not use the second preferred server. These probabilities are then used to calculate the new upward transition rates and consequently, the steady-state probabilities in the hypercube model.

We will consider the general case of N servers in the system. The following notation is used:

λ_j^k = Arrival rate of calls in atom j where the car dispatched is the k^{th} preferred and available car for atom j

p_j^k = Probability that the k^{th} preferred and available car will be dispatched to atom j where

$k = 1,2,..........,N$ and $j = 1,2,................, \#$ of atoms

There are N types of equations for the upward transition rates in our N-server system. For the general case when i servers are busy, the upward transition rate is given by the following equation:

$$\sum_j \lambda_j^1(p_j^1) + \sum_j \lambda_j^2(p_j^2) + \dots\dots + \sum_j \lambda_j^{N-i} p_j^{N-i}$$

Note that in this example, we will use a simple exponential function, $f(t) = e^{-t}$, to obtain the p_j^k's. The simple exponential function is a reasonble choice because there is a higher probability that we will use the next preferred car if the difference in travel times is small and we would expect this probability to decrease rapidly when the difference in travel times increases. The p_j^k's are obtained by the following equations:

$p_j^2 = f(t_1)$ where t_1 is the difference in average response time between the first and second preferred servers

$p_j^3 = f(t_2)(1 - p_j^2)$ where t_2 is the difference in average response time between the first and third preferred servers

$p_j^1 = 1 - p_j^2 - p_j^3$

Then with these revised upward transition rates, we can set up the steady-state equations to find the steady-state probabilities for the 2^N states and consequently, obtain the performance measures for the system.

4.3 NUMERICAL EXAMPLE FOR A THREE-SERVER SYSTEM

We will now consider the three-server system example in Section 5.4.3 of Larson and Odoni's Urban Operations Research. There are 10 geographical atoms in the system with the following arrival rates of calls for service:

$$\lambda_1 = 0.25, \lambda_2 = 0.25, \lambda_3 = 0.10, \lambda_4 = 0.25, \lambda_5 = 0.15, \lambda_6 = 0.10, \lambda_7 = 0.10, \lambda_8 = 0.10, \lambda_9 = 0.10, \lambda_{10} = 0.10$$

The dispatch preference for the three-server system is given by the table below which was extracted from Urban Operations Research.

Dispatch preferences for three-server city.

Atom Number	*First Preference Unit*	*Second Preference Unit*	*Third Preference Unit*
1	1	2	3
2	1	2	3
3	2	1	3
4	1	3	2
5	2	1	3
6	2	1	3
7	3	1	2
8	3	1	2
9	3	1	2
10	3	1	2

The graphic illustration of the hypercube model for the three-server example with the new transition rates is:

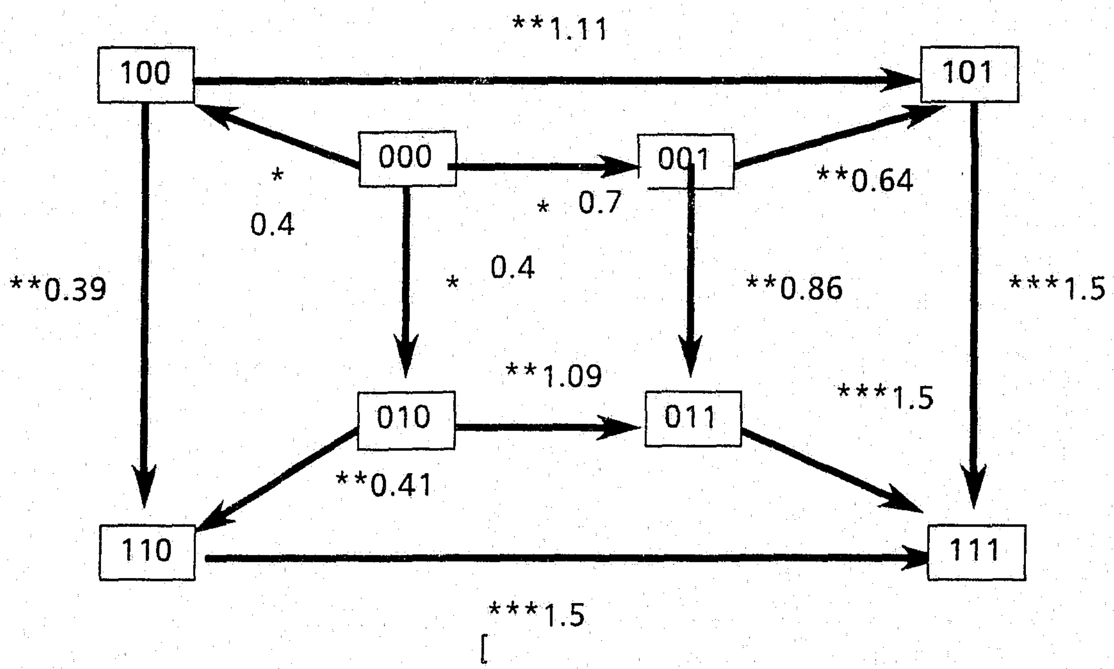

(The * indicates which equation was used to calculate the upward transition rate)

The upward transition rates are obtained by using the following equations:

* For the case when all servers are free, the upward transition rate is:

$$\sum_j \lambda_j^1(p_j^1) + \sum_j \lambda_j^2(p_j^2) + \sum_j \lambda_j^3(p_j^3)$$

**For the case when one server is busy, the upward transition rate is:

$$\sum_j \lambda_j^1(p_j^1) + \sum_j \lambda_j^2(p_j^2)$$

*** For the case when two servers are busy, the upward transition rate is:

$$\sum_j \lambda_j^1$$

After solving the steady-state balance equations, we have the steady-state probabilities:

P_{000} = 0.21053 [0.21053]
P_{001} = 0.13173 [0.13669]
P_{010} = 0.09380 [0.08863]
P_{100} = 0.09025 [0.09047]
P_{110} = 0.05488 [0.05031]
P_{101} = 0.08655 [0.08894]
P_{011} = 0.09542 [0.09489]
P_{111} = 0.11842 [0.11842]

Note: the values from the original hypercube model are in square brackets

The performance measures obtained from the steady-state probabilities are:

(1)The workload rates for servers #1, 2, 3 :

ρ_1 = 0.55054 [0.5574]
ρ_2 = 0.48094 [0.4734]
ρ_3 = 0.46852 [0.4693]

(2) Interatom dispatch frequencies:

Atom Number (j)	Unit Number (n) 1	2	3
1	0.07491[0.07376]	0.03638[0.03760]	0.01590[0.01581]
	0.08807[0.08692]	0.04954[0.05076]	0.02906[0.02897]
2	0.07491[0.07376]	0.03638[0.03760]	0.01590[0.01581]
	0.08807[0.08692]	0.04954[0.05076]	0.02906[0.02897]
3	0.00991[0.00994]	0.03461[0.03511]	0.00636[0.00633]
	0.01517[0.01470]	0.03987[0.04037]	0.01162[0.01159]
4	0.07491[0.07376]	0.01442[0.01482]	0.03786[0.03860]
	0.08807[0.08692]	0.02758[0.02798]	0.05102[0.05176]
5	0.01487[0.01416]	0.05192[0.05286]	0.00954[0.00949]
	0.02276[0.02205	0.05982[0.06056]	0.01743[0.01738]
6	0.00991[0.00994]	0.03461[0.03511]	0.00636[0.00633]
	0.01517[0.01470]	0.03987[0.04037]	0.01162[0.01159]
7	0.00968[0.00957]	0.00577[0.00593]	0.03543[0.03538]
	0.01494[0.01483]	0.01103[0.01119]	0.04069[0.04064]
8	0.00968[0.00957]	0.00577[0.00593]	0.03543[0.03538]
	0.01494[0.01483]	0.01103[0.01119]	0.04069[0.04064]
9	0.00968[0.00957]	0.00577[0.00593]	0.03543[0.03538]
	0.01494[0.01483]	0.01103[0.01119]	0.04069[0.04064]
10	0.00968[0.00957]	0.00577[0.00593]	0.03543[0.03538]
	0.01494[0.01483]	0.01103[0.01119]	0.04069[0.04064]

For the table on the previous page:

Upper figure is $f_{nj}^{[1]}$ = fraction of all dispatches that send unit n to atom j and i incur no queue delay

Lower figure is f_{nj} = fraction of all dispatches that send unit to geographical atom j

(3) Mean Travel Times:

j	Average Travel Time to Atom j, T_j (minutes)	Average Travel Time to Response Area n, TRA_n (minutes)	Average Travel Time of unit n, TU_n (minutes)
1	3.35[3.37]	3.14[3.13]	3.14[3.14]
2	4.31[4.29]	3.41[3.42]	3.18[3.18]
3	5.29[5.26]	3.29[3.30]	3.45[3.45]
4	1.73[1.75]		
5	1.79[1.77]		
6	3.60[3.64]		
7	2.67[2.64]		
8	2.31[2.31]		
9	4.54[4.55]		
10	4.18[4.19]		

The changes in the upward transition rates resulted in changes in the steady-state probabilities which ranged from 2.9% decrease for state 101 to 9.1% increase for state 110. The differences in the new workload rates reflected to some extent the possibility that the second and third prefered cars could be considered for dispatch. The workload rate of patrol unit 1 decreased by 1.2% since it no longer had as much responsibility as it did in the old model where it was the first preferred unit for 50% of the city's workload and the first backup unit for the rest of the city. On the other hand, units 2 and 3 had more likelihood of being dispatched in the new model and so their workload rates increased by 1.6% and 0.2% respectively. The range of changes in f_{nj} were from a 2.4% decrease to a 3.2% increase. The changes for the mean travel times also were not very significant.

For this three-server example, we observed only minor changes in the performance measures. However, we expect to find in future research, other cases where the differences will be more significant.

5. CONCLUSION

In this paper, we have presented two possible applications of expert systems in urban police services. We have proposed expert systems protocols for assisting 911 operators to prioritize two categories of ambiguous calls and also added an expert systems factor to the hypercube queueing model, which is an OR planning tool for police dispatchers.

It should be emphasized that these two areas are just a small sample of the countless number of applications of expert systems in urban police services. For example, a potential use is in the analysis of the frequencies of the different types of crime in each neighborhood. This analysis is helpful in predicting future occurrences of certain types of crime in the neighborhood. An expert system can be designed to monitor current levels of crime in each neighborhood and subsequently, to project future trends in criminal incidents. The expert system can also help the police department plan its allocation of resources such as manpower and equipment, based on its projection of future occurrences of crime.

Expert systems can also serve as training tools for police patrolmen. The expert system can create simulations of criminal incidents and analyze the responses of the patrolmen. From its analysis, the expert system can then identify the weaknesses of each patrolman and propose suitable remedies. Yet another larger and more crucial application is during major emergencies caused by floods, fires or other types of catastrophes. An expert system can be used to effectively mobilize police resources and to deploy rescue teams. The argument for using expert systems in such disasters is that the disasters rarely occur and very few people know what to do when they do occur. Storing the knowledge of a person experienced in dealing with such critical situations, in the expert system, can prove to be very valuable in times of sudden emergencies.

On a broader scope, expert systems can also be used in areas that are related to police work such as criminal justice. One such possible application in criminal justice is the assignment of prisoners to suitable cell blocks. Incompatibility of prisoners in the same block can lead to unnecessary abuse and violence. An expert system can be used to assess each prisoner's preferences and based on this assessment, it can make suitable assignments. Another application is the selection of appropriate sentences for offenders of the law. There are no hard-and-fast rules for choosing sentences for specific types of crime; the judge has to consider many issues before deciding on the appropriate sentence. Expert systems can help in the sentence-selection process by evaluating the evidence against the offender and then suggesting suitable sentences.

It is evident that there are many interesting applications of expert systems in urban police services and other related areas. Although these applications have yet to be explored, the recognition of their existence is a significant step towards their implementation in the future.

REFERENCES

Buchanan, Bruce G., and Edward H. Shortliffe. 1984. Rule-Based Expert Systems, Reading, Massachusetts: Addison-Wesley Publishing Company, Inc.

Davis, Randall. 1984. " Amplifying Expertise with Expert Systems." The AI Business: Commercial Uses of Artificial Intelligence, Cambridge, Massachusetts: The MIT Press.

Davis, Randall. 1982. " Expert Systems: Where Are We? And Where Do We Go From Here?" The AI Magazine, Spring issue.

Hayes-Roth, Frederick and Donald A. Waterman, and Douglas B. Lenat. 1983. Building Expert Systems, Reading, Massachusetts: Addison-Wesley Publishing Company, Inc.

Larson, Richard C. 1978. Police Deployment: New Tools for Planners, Lexington, Massachusetts: Lexington Books.

Larson, Richard C. and Amedeo R. Odoni. 1981. Urban Operations Research, Englewood Cliffs, New Jersey: Prentice-Hall, Inc.

McDermott, John. 1982. "R1: A Rule-Based Configuration of Computer Systems." Artificial Intelligence, Vol.19, no.1.

Parnas, R. 1971. " Police Discretion and Diversion of Incidents of Intra-Family Violence." Law and Contemporary Problems, Vol. 36.

Percy, Stephen L. and Eric J. Scott. 1983. Demand Processing and Performance in Public Services Agencies, Bloomington, Indiana: Indiana University, Workshop in Political Theory and Policy Analysis.

Roberts, A.W., and J.A. Visconti. 1972. "The Rational and Irrational Use of Systemic Microbial Drugs." American Journal of Hospital Planning, Vol. 29.

Shortliffe, Edward H. 1976. Computer - Based Medical Consultations: MYCIN, New York, New York: Elsevier.

Sumrall, Raymond O., and Jane Roberts, and Michael T. Farmer. 1981. Differential Police Response Strategies, Washington D.C. : Police Executive Research Forum.

van Melle, William. 1979. "A Domain-Independent Production-Rule System for Consultation Programs." Proceedings of the Sixth International Joint Conference on Artificial Intelligence, Tokyo, Japan.

Weiss, S.M. and C.A. Kulikowski. 1977. "EXPERT: A System for Developing Consultation Models." Proceedings of the Fifth International Joint Conference on Artificial Intelligence, Cambridge, Massachusetts.