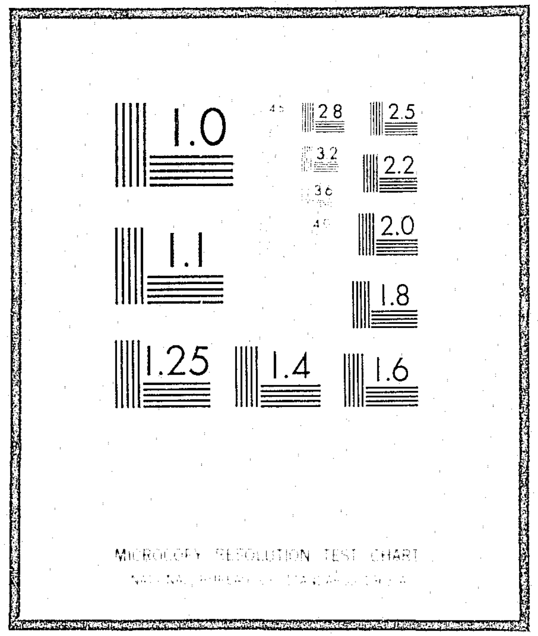

Date filmed 6/28/76

HYPERCUBE QUEUING MODEL: USER'S MANUAL

PREPARED FOR THE DEPARTMENT OF HOUSING AND URBAN DEVELOPMENT

RICHARD C. LARSON R-1688/2-HUD JULY 1975

29826
SNI

This report was prepared as an account of work sponsored by the United States government, Office of Policy Development and Research, Department of Housing and Urban Development (HUD), under Contract H-2164. Neither the United States nor HUD, nor any person acting on behalf of HUD, makes any warranty, expressed or implied, or assumes responsibility for the accuracy or completeness of the information herein, or assumes liability with respect to the use of any information or method herein.

Published by The Rand Corporation

HYPERCUBE QUEUING MODEL: USER'S MANUAL

PREPARED FOR THE OFFICE OF POLICY
DEVELOPMENT AND RESEARCH, DEPARTMENT
OF HOUSING AND URBAN DEVELOPMENT

RICHARD C. LARSON

R-1688/2-HUD
JULY 1975

THE
NEW YORK CITY
RAND
INSTITUTE

Rand
SANTA MONICA, CA. 90406

PREFACE

The Hypercube Queuing Model is a computer program that calculates selected performance measures of emergency service systems (police, fire, and medical). It is especially useful for assisting police and emergency medical agencies in designing response districts for their mobile vehicles.

The model is completely described in five reports, available from The Rand Corporation. There are two versions of the model--an *exact* model and an *approximate* model--both of which are incorporated in a single computer program. The mathematical formulations of the two models are presented in the following reports:

- 11696 Richard C. Larson, *A Hypercube Queuing Model for Facility Location and Redistricting in Urban Emergency Services*, R-1238-HUD, and
- 13979 Richard C. Larson, *Urban Emergency Service Systems: An Iterative Procedure for Approximating Performance Characteristics*, R-1493-HUD.

Nontechnical descriptions of potential applications of the Hypercube Queuing Model, how it works, when it should be used in preference to other models, and the resources needed to use it are given in the summary of this report (published separately):

- 29825 Jan M. Chaiken, *Hypercube Queuing Model: Executive Summary*, R-1688/1-HUD.

29826

The present report (R-1688/2-HUD) is a manual for users of the model. It describes and gives examples of applications, describes the procedures to operate the computer program once it has been installed in the user's computer system, and discusses the decisions to be made (such as the dispatching strategy employed), the results, and the costs and requirements for operation. The fifth report

- Richard C. Larson, *Hypercube Queuing Model: Program Description*, R-1688/3-HUD

gives a listing of the computer program and provides information for programmers who wish to install the program.

The author, a consultant to The Rand Corporation, is Associate Professor of Electrical Engineering and Urban Studies at the Massachusetts Institute of Technology (MIT). All work on this model, from design through documentation, has been supported jointly by grants to the Massachusetts Institute of Technology from the National Science Foundation (NSF) and by contracts between The New York City-Rand Institute and the U.S. Department of Housing and Urban Development (HUD). The most recent work has been funded by the Division of Social Systems and Human Resources (Research Applied to National Needs) at NSF and the Office of Policy Development and Research at HUD. Reports describing the Hypercube Model and copies of the computer program are available from both Rand and MIT by writing to the addresses shown in Appendix B.

The project funded by HUD has resulted in the development, field testing, and documentation of a number of models for improving the deployment of municipal emergency services. Further information about the models themselves and case studies of applications of the models in several cities can be obtained from the Rand or HUD address in Appendix B.

ACKNOWLEDGMENTS

The theoretical development and computer program design of the models and procedures described in this report began in 1971. The work has been supported by the National Science Foundation, Research Applied to National Needs (Division of Social Systems and Human Resources); and by the U.S. Department of Housing and Urban Development (HUD) under contract to the New York City-Rand Institute.

Many people have contributed to the evolution of the model and its computer implementations. The author thanks especially the Police Advisory Panel of the Innovative Resource Planning Project at MIT, the NSF/RANN Site Visit team for the IRP Project, the students in the 1974 offering of "Analysis of Urban Service Systems" at MIT (July 1974), and the participants in the November Workshop at MIT on police sector design. In addition, the author thanks David Fahrland, Hugh Findley (Director of Planning and Research), and Officer Joe Molloy, all of the Quincy (MA) Police Department, and Larry Deetjen (formerly associated with the Arlington (MA) Police Department), for being willing to use a preliminary version of the model in redesigning police sectors in their own communities. For helpful suggestions from time to time, thanks are due Kenneth Chelst, Richard Jarvis, Mark McKnew, Keith Stevenson, and Richard Weissberg (all of MIT), Jan Chaiken (The Rand Corporation), and Robert Baumgardner (HUD).

All computer programming and computer runs were performed at the MIT Information Processing Center.

CONTENTS

FIGURES

TABLES

GLOSSARY[*]

Atom

See *geographical atom.*

Call for service

A communication to an emergency service originating from a citizen, an alarm system, a police officer, or other detector, reporting an incident that requires on-scene assistance by a *response unit.*

Command

An area or region comprising several *districts* that is administratively distinct, usually having a station-house or garage used as a base of operations. Often called precincts or zones. *Dispatch assignments* are nearly always intra-command assignments.

Dispatch assignment

A directive by the *dispatcher* to a *response unit* assigning the unit to respond to the scene of a reported incident or *call for service.*

Dispatcher

An individual who has responsibility for assigning available radio-dispatchable *response units* to reported incidents.

District

A collection of *geographical atoms* that are primarily associated with a particular *response unit.* For certain dispatch strategies, the district's response unit always receives first preference in dispatching decisions. In police applications, a district (often called a beat or sector) is the area in which the patrol unit can perform *preventive patrol.* Over the entire *region,* the set of districts need not be mutually exclusive nor collectively exhaustive.

Effective travel speed

That speed which, if constantly maintained over the path of a response journey, would result in the same *travel time* as that actually experienced by the dispatched *response unit.*

EMCM: Expected modified center of mass

A dispatch strategy that calculates the probabilistic location of units, representing the best that any dispatcher could do without

---

[*] Italicized words in definitions are themselves defined elsewhere in the Glossary.

knowing the exact real-time position of available units. See Sec. 4.4.3.

ESCM: Expected strict center of mass

A dispatch strategy that estimates the statistically average travel distance from each of the unit's geographical atoms (weighted by the likelihood of the unit being located in that atom) to the atom of the incident (weighted by the likelihood of the incident being located in that atom). See Sec. 4.4.4.

Geographical atom

A subarea within a *command*, typically no more than a few city blocks in size, that is used as the smallest geographical unit for aggregating statistics on the spatial distributions of *calls for service* and positions of the *response units*.

Interdistrict (or cross-district) assignment

A *dispatch assignment* to a *district* other than the unit's district.

MCM: Modified center of mass

A dispatch strategy in which the exact location of the incident is used to make travel time (distance) estimates. See Sec. 4.4.2.

Overlapping districts

Districts that have at least some areas in common; partially shared districts.

Preventive patrol

An activity undertaken by a police *response unit*, in which the unit tours an area, with the officer(s) checking for crime hazards (for example, open doors and windows) and attempting to intercept any crimes that are in progress.

Region

The entire collection of *geographical atoms* included in a particular set of runs of the model. Can be an entire city or part of a city (e.g., a *command*).

Response unit

A patrol car, scooter, or wagon, and its assigned police officer(s); a radio-dispatchable footpatrolman; an ambulance; a fire truck.

SCM: Strict center of mass

A dispatch strategy in which the dispatcher makes travel time estimates acting as if the unit were located at the statistical center

of its district and the incident were at the statistical center of its district. See Sec. 4.4.1.

Service time

The total "off the air" time per *call for service* for a *response unit*. Includes *travel time*, on-scene time, and possibly related off-scene time.

Travel time

The time required for the dispatched *response unit* to travel to the scene of the reported incident.

Utilization factor

The fraction of time a *response unit* is unavailable to respond to dispatch requests. In this model, it is assumed that a unit can only be unavailable because of call-servicing duties. Sometimes called utilization rate.

Workload

Same as *utilization factor*.

0. INTRODUCTION

In two earlier papers,(1,2) both precise and approximate mathematical models were described for the numerical evaluation of certain performance characteristics of urban emergency service systems (e.g., police, emergency medical, and fire services). This report details the use of the computer program, written in the PL/I language, that implements these two models. Although it is recommended that a user of the models be familiar with the mathematics described in the earlier papers, detailed knowledge is not required. In fact, in order to develop an intuitive understanding for the application of the models, including a set of operational rules of thumb, it is recommended that the potential user (including nonpolice users) first consult the nontechnical illustrative case study,(3) which applies the model to a police district in Boston. Additional case studies are described in Refs. 4 and 5. An overview of the model, its potential uses, and its data requirements is found in the Executive Summary.(6)

This user's manual is organized as follows: Section 1 briefly reviews the model's assumptions, data requirements, and outputs. Section 2 outlines some illustrative applications of the models. Section 3 illustrates a simple sample run without using any of the advanced features of the computer program. Section 4 then details all of the options that are included to facilitate implementation in complex urban environments. Section 5 discusses definitions and conventions that are used in certain complicated situations (e.g., involving overlapping districts). Section 6 provides a concise technical summary of program use. The program listing is given in a separate volume:

- Richard C. Larson, *Hypercube Queuing Model: Program Description*, The Rand Corporation, R-1688/3-HUD, 1975.

1. MODEL ASSUMPTIONS, DATA REQUIREMENTS, AND OUTPUTS

Each of the models--the "exact" hypercube model and the "approximate" model--requires the same data and produces the same outputs. We will thus treat the two models as one until the appropriate time in describing program options.

Briefly, the model can be used to estimate certain performance measures of any spatially distributed emergency service system which can be modeled as follows:

1. The area in which the system provides service (called the region) can be broken down into a number of "reporting areas" or "geographical atoms." Typically no more than a few city blocks in size, the atom is the smallest geographical unit for aggregating statistics on the spatial distributions of calls for service and positions of the response units.
2. Calls for emergency assistance ("calls for service") are generated independently from each of the reporting areas. Although the exact time and location of any particular call cannot be predicted in advance, long-term statistical averages are available to predict the relative workload generated from each of the reporting areas.
3. Data are available to estimate the travel time from each reporting area to every other reporting area.
4. There are M spatially distributed response units, each of which may travel to any of the reporting areas in the serviced region.
5. The location of each response unit not servicing a call is known (at least statistically). For instance, a patrolling police car may allocate 50 percent of its patrol time to reporting area 7 and 25 percent each to reporting areas 8 and 11. A fixed-position unit, such as an ambulance, would always be located in one particular reporting area when not providing emergency medical service.

6. Geographical atoms are collected together to form "districts." For mobile units, any atom in which the unit spends some of its available time must be included in its district; in addition, other atoms (in which no available time is spent) may be assigned to a unit's district. Districts may overlap. In police applications these districts are usually called beats, sectors, or patrol areas. For fixed-position units, the atom containing the fixed position must be included in in the unit's district; in addition, any other atoms may be specified to be within the unit's district. Often in this case (e.g., ambulance and fire department applications), a unit's district is defined to be all points closer to that unit than to any other unit. In addition to preventive patrol assignments, the district is used in determining dispatch strategies and in computing output performance measures.
7. In reponse to each call for service, exactly one response unit is dispatched to the scene of the call, provided at least one unit is available within the service region. If no unit is available, the call either enters a queue with other backlogged calls or it is serviced by some backup system (e.g., police providing backup to an ambulance service or a neighboring community dispatching units into the temporarily saturated community). If the call enters a queue, it is later dispatched on a first-come, first-served basis. (The assumption of dispatching exactly one unit to a call indicates that the model does not accurately portray the performance of those fire departments that send many units to the scene of a fire alarm.)
8. The service time for a call, including travel time and on-scene time, has a known average value. In general, each response unit may have its own average service time. Moreover, reflecting the unpredictability of service times in actual systems, there is considerable variability about the

average value(s). As one measure of variability, the standard deviation of the service time is assumed to be approximately equal to the mean.[*]

9. Variations in the service time that are due solely to variations in travel time are assumed to be of minor significance compared to variations of on-scene service time. This assumption, which limits the applicability of the model, is most nearly satisfied by urban police departments and least nearly satisfied by rural emergency services (especially rural ambulance services).

In practice, no emergency service system will ever conform to all of the model's assumptions exactly. In applying the model, the user must weigh the extent to which the actual system does not fit the rigidities of the model (and the associated loss in predictive accuracy) against alternative methods (with their own limitations) to choose that method which best suits the resource allocation purposes at hand.

Given the required data (whose precise input formats will be specified in later sections), the model computes numerical values for the following performance characteristics:

1. Region-wide mean travel time.
2. Region-wide workload and workload imbalance.
3. Region-wide fraction of dispatches that remove a unit from its district.
4. Workload of each response unit (measured in fraction of time unit is busy servicing calls).
5. Mean travel time to each geographical atom.
6. Mean travel time to each district.
7. Mean travel time of each response unit.
8. Fraction of responses in each unit's district that are handled by other units.

---

[*] The exact model assumes negative exponential service times. Slight deviations in this assumption do not markedly alter the predictive accuracy of the model.

9. Fraction of responses of each response unit that dispatch the unit outside of its district.
10. Fraction of responses within each geographical atom that are handled by each of the response units.

In addition, if the emergency service being analyzed is a police patrol force, then the user can calculate

11. The frequency of preventive patrol passings in each of the reporting areas.

2. SOME USES OF THE MODEL

The model is a powerful planning tool that can be used in a variety of applications by planners and administrators of emergency service agencies (i.e., police, ambulance, fire, emergency repair services). In this section are briefly described several applications that are likely to be important in many cities and towns.

2.1 POLICE BEAT DESIGN

Suppose a city's police department has not redesigned its beats (sometimes known as sectors) for many years. Then, due to changing population patterns and other factors in the evolution of the city, the distribution of crimes and other incidents that give rise to calls for police service is likely to have changed significantly from the time of the last beat design. This could result in an intolerable situation in which some patrol officers are working considerably longer hours responding to calls than are others. Compounding the problem, crime preventive patrol is probably least prevalent in the high workload areas, since the high call-for-service workload in these areas sharply reduces the time available for patrol; yet, it is probably these areas that most need such patrol.

In the case described above, the model can be used to assist the police planner in redesigning beats to correct imbalances. The model provides outputs on travel times, workloads of each police vehicle, preventive patrol frequencies, and other factors that allow the simultaneous consideration of response time reduction, workload balancing, preventive patrol strength, and so forth. The model reveals the tradeoffs one must accept in attempting to reach acceptable performance in each of these categories.

In using the model the police planner must specify the beat configuration that he desires. Then the model computes numerical values for each of the performance measures (e.g., travel times, patrol car utilizations, etc.). Undoubtedly each police planner will have his own set of issues--some quantitatively oriented and some not--that will be

important in the beat design process. For instance, the well-known police planner O. W. Wilson focused heavily on workload balancing in his recommended beat design procedures.(7) In most cases, however, regardless of the planner's particular issues and their relative priorities, the model described here should be useful in his thinking--primarily because it computes rapidly and effectively many operationally oriented performance measures that come into play in the beat design process.

A word of caution: The model does not "optimize" any performance measures to find the "best" beat design. The philosophy behind the construction of the model is that in a public service agency as complex and multifaceted as an urban police department, the word "optimize" has little meaning. Rather, it is felt that a police planner with an intimate knowledge of his own city can be an excellent judge of the qualitative as well as quantitative factors that are relevant.

Using the computer to calculate the important performance measures, beat design can be viewed as an iterative process. First the planner proposes a particular design of beats and has the computer calculate the resulting values of the performance measures. He then incorporates this evidence, including possible workload imbalances and/or inequities in accessibility to police service, with his additional knowledge of the area under consideration, and decides whether to accept the proposed beat plan or to devise an altered one. In the latter case, the entire process is repeated one, two, or several times until a satisfactory beat design is obtained. In this way, good use is made of the planner's talents and the computer's computational power.

2.2 AMBULANCE DISTRICT DESIGN

Suppose that a city disperses its ambulances throughout the city, prepositioning them in a way that best anticipates likely calls for emergency medical service. Then, the ambulance system planner needs assistance in determining good locations for the ambulances and reasonable areas of primary responsibility for each. The model can be used for this purpose in much the same way as a police planner would use it to design police beats. Here, however, the ambulances (when not responding to medical emergencies) are fixed at preselected sites, whereas the police cars are likely to patrol throughout their beats. Also, the time for an ambulance to service a medical emergency usually includes travel time to and from a hospital (to transport the patient), a time not experienced in the police example (except when transporting arrestees to a police station-house). So in the ambulance case it is much more likely that travel times (time to the scene, time from the scene to the hospital, time from the hospital back to the prepositioning site) will play a dominant role in the overall time required per incident. (In the police case, on-scene service time is usually significantly greater than travel time.)

The ambulance system planner can use the model to explore the consequences of alternative prepositioning sites for his ambulances and alternative districts of primary responsibility for each. Since travel times play such an important role in ambulance services, it is likely that the planner will have to adjust the service time of each ambulance separately to reflect the different geographical travel time factors affecting each one.

The final ambulance site selection and district design could include factors of workload balance, travel time reduction, neighborhood integrity, etc. Again, analogous to the police beat example, the exact tradeoff among the various factors must be determined by the user of the model, not by the model itself.

2.3 ASSIGNING BILINGUAL PERSONNEL

If police officers in a particular car or attendants in a particular ambulance are bilingual--fluent in English and a second language (say Spanish, Chinese, Portugese, Italian, or some other language predominant in one or more sections of the city), then the planner using the model (for beat or district design or any other purpose) would want to be careful to give first preference to this bilingual unit in responding to calls for service from a neighborhood having the second language as its primary language. This consideration of matching the service capabilities of the unit to the needs of the neighborhood would probably outweigh narrow efficiency considerations such as minimizing travel times.

The planner using the model can specify that the dispatching process be such that this specialized unit will be assigned to any call from these neighborhoods, provided that the unit is available when the call arrives. If not, then another (less preferred) unit would be assigned to the call. The model will provide the planner with information about the consequences of this policy.

The planner may wish to explore various prepositioning sites or patrol areas for the bilingual unit as well as for other units within the area under study.

2.4 BACKUP UNITS

Sergeant's cars are an example of backup units. Suppose a police command area (often called district, precinct, or area) is divided into beats with one or two sergeant's cars assigned command-wide responsibility. These cars might patrol the entire region, or they might divide the region among them. In any case, each would patrol several beats, with primary responsibility being supervision of the regular (beat) patrol cars not responding to calls for service. However, should all the beat cars be busy, then the sergeant's cars may be used by the dispatcher as backup cars to assign to calls for service in order to avoid delay and congestion at the dispatcher's position.

The user of the model can easily take account of this situation by adjusting the dispatching strategy to assign last preference to the sergeant's cars. In this way, the model assigns the beat cars to all calls for service in the area as long as there are beat cars available. However, whenever they are all simultaneously busy, the model (imitating what the actual dispatcher would do) will assign the sergeant's cars to calls for service. Given this added complication, the user can still address issues of workload balance, response time reduction, prepositioning, and so forth, including a calculation of the call-for-service workloads of the sergeant's cars and a calculation of how frequently they respond to each of the neighborhoods in the area.

2.5 OVERLAPPING BEATS

Most police departments, when considering the beat design process, view beats as separate nonoverlapping areas where primary responsibility can be assigned to one patrol unit. However, we are seeing more and more variations on a new theme--overlapping beats, where the same neighborhood(s) will be patrolled by two or more police cars.

Example 1: One large U.S. city has "umbrella cars," each of which is assigned to patrol two regular contiguous beats. Thus, if beats A and B are side-by-side, they will be patrolled by car A (patrolling beat A), car B (patrolling beat B), and an umbrella car (patrolling both beats A and B). The dispatching strategy for calls from beat A is usually to assign car A, if available, and otherwise to try (in sequence) car B, the umbrella car, and other cars in the command area.

Example 2: Another large U.S. city divides each command area into two or more sergeant's zones. Each sergeant's zone is assigned one sergeant's car and several regular patrol cars. The patrol cars share patrol responsibility for the entire zone, and there are no regular beats within the zone.

Example 3: Several smaller communities divide their area into regular beats, but assign two patrol cars to each beat. This is perhaps the simplest form of overlapping beat structure.

The model will allow exploration of a wide variety of such overlapping beat plans.

These and other overlapping beat structures can be studied--in terms of the performance measures computed in the model--by the police planner. With overlapping beats, it is especially difficult to predict ahead of time the call-for-service workloads or preventive patrol levels of each of the units. The model performs this task, aiding the planner in considering the many different factors that come into play.

2.6 PRIORITIES

Although many emergency services place priorities on the types of calls they receive--either explicitly or implicitly--it is usually not necessary to consider these priorities in beat or district design, positioning, etc. For those cases in which priorities must be included in the model, the user can adapt the model to reflect a certain limited class of priority call-handling and dispatching procedures. However, complicated priority-oriented procedures cannot be treated by the model--and this perhaps is the model's largest single limitation at its current state of development. (See Appendix A for future modifications to the program.)

As an example of a situation that can be treated by the model, consider a police command that has a patrol unit specializing in family disputes. If a family dispute is reported anywhere within the command and the special unit is available for dispatch, then it is assigned to the call. Otherwise another (nonspecialty) unit is assigned. To model this situation, the user essentially splits each reporting area into two reporting areas--one generating family dispute calls and the other generating all other police calls. The family dispute unit is then given first dispatch preference for all "family dispute reporting areas." This way of adapting the model to a special type of call and a special type of unit will allow the user to compute separately the travel times to each type of call, the workloads of each of the units for each type of call, and the fraction of family dispute calls handled by the family dispute car.

If additional types of specialty units and/or calls are considered, such as one-man versus two-man cars, this procedure of splitting reporting areas by type of call can be continued into three, four, or more splits. We call this process "layering" of reporting areas. However, our experience is that the volume of data produced by the model can quickly overwhelm the user and obscure simple relationships that he could more readily see if there were fewer priority levels. Thus, there exists an important tradeoff between the fine-grained detail of the system replicated by the model and the ability of the user to comprehend, interpret, and act intelligently upon the output of the model.*

The model *cannot* be used to study the following types of priority dispatching schemes:

---

*As currently structured, the user must work hard to use the layering process since the inputs and the outputs are not geared to the concept of layering. However, we hope to make the layering process a standard advanced feature of the model in the next version.

- Selective stacking (queuing) of low priority calls to await the attention of the beat car once it becomes available.
- Preemption (interruption) of a unit on a low priority call to send it to a high priority call.
- Priority-oriented schemes for dispatching calls in queue.

If the system planner wishes to analyze these types of operation (which may be very important in a particular application), he should probably use a simulation model.*

2.7 PREVENTIVE PATROL DESIGN

Although the model is generally considered to be of primary importance for designing the geometry of beats or the locations of units, the user can examine the consequences of alternative preventive patrol strategies, without altering the geometry of beats. This is because the model allows the user to specify the fraction of time (while on patrol) that a patrol unit will spend in each of its reporting areas (within its beat). For instance, higher crime rate reporting areas can be given greater patrol attention than lower crime rate areas. At one extreme, the user could examine the effect of having the patrol unit spend all of its patrol time in the highest crime rate reporting area. At another extreme, the patrol unit could be directed to allocate its patrol time equally among all the reporting areas in its beat.

The user may ask, "What performance measures will be affected by altering these patrol time assignments (assuming fixed-beat boundaries)?" Essentially all performance measures will change in value as patrol time assignments within a beat are changed, even if there is no change in beat boundaries. While it is easy to see that frequency of patrol passings will change, it is not so easy to see why workloads, travel times, and frequency of cross-beat dispatches will change. The reason for their change lies in the fact that by altering the fractions of time spent by a unit in each part of its beat, the likely position of the unit at times of dispatch is also changed.

---

*See, for example, Chap. 6 in Ref. 8, or Ref. 9.

As a simple example, a unit currently may be spending its patrol time in the western part of its beat. Because it has a high probability of being near the western boundary of the beat while on patrol, the unit is probably the second preferred unit to dispatch to calls for service in the adjacent beat to the west. If the unit is moved to the eastern part of its beat (by changing the patrol time assignments), then some other unit would most likely become the second preferred unit to dispatch to the adjacent beat to the west; similarly, the adjacent beat to the east will probably now assign second dispatch preference to this unit due to its (newly acquired) geographical proximity. *Thus, changing patrol time assignments changes dispatch preferences and thus changes nearly all computed performance measures.*

The above seven subsections have described a variety of applications of the model. There are many others, and these become apparent by using the model, first for the simpler applications and then--as confidence and model familiarity build in the model user--in more advanced applications. In this spirit, the next section indicates the data requirements (including computer card formating) for using the model in one of its simpler modes. Subsequent sections describe the more advanced features.

3. A SIMPLE EXAMPLE

In this section is demonstrated the simplest possible use of the model--to analyze a system with three response units and seven geographical atoms. The locations of the atoms are shown in Fig. 1; the travel directions are shown as north, south, east, and west, but of course in actual applications the coordinate grid can be rotated in any arbitrary manner. The resulting district configuration is shown later in Fig. 11 (p. 51).

3.1 DATA CARDS

We now demonstrate how to set up the deck of input data cards to run the model.

Card Type 1: Basic Program Specifications

The first data card is prepared as follows:

```
M = 3     R = 7     NUM = 2     ESTSTAT = 1;
```

In the computer the number of response units is given the symbol M, and thus the equation M = 3 indicates that this run of the model will use exactly three response units. The total number of geographical atoms is denoted by R, which in this case equals seven. The symbol NUM indicates the number of different workload levels (or call rates) the user wishes to examine with the current set of input data (see card type 10). Since NUM = 2, the user wishes numerical output for two different workload levels. Finally, the symbol ESTSTAT is used to indicate whether the user wishes output from the exact hypercube model (ESTSTAT = 0), from the approximate model (ESTSTAT = 1), or from both models (ESTSTAT = 2). In this case, since ESTSTAT = 1, the user wishes to see output only from the approximate model. To summarize,

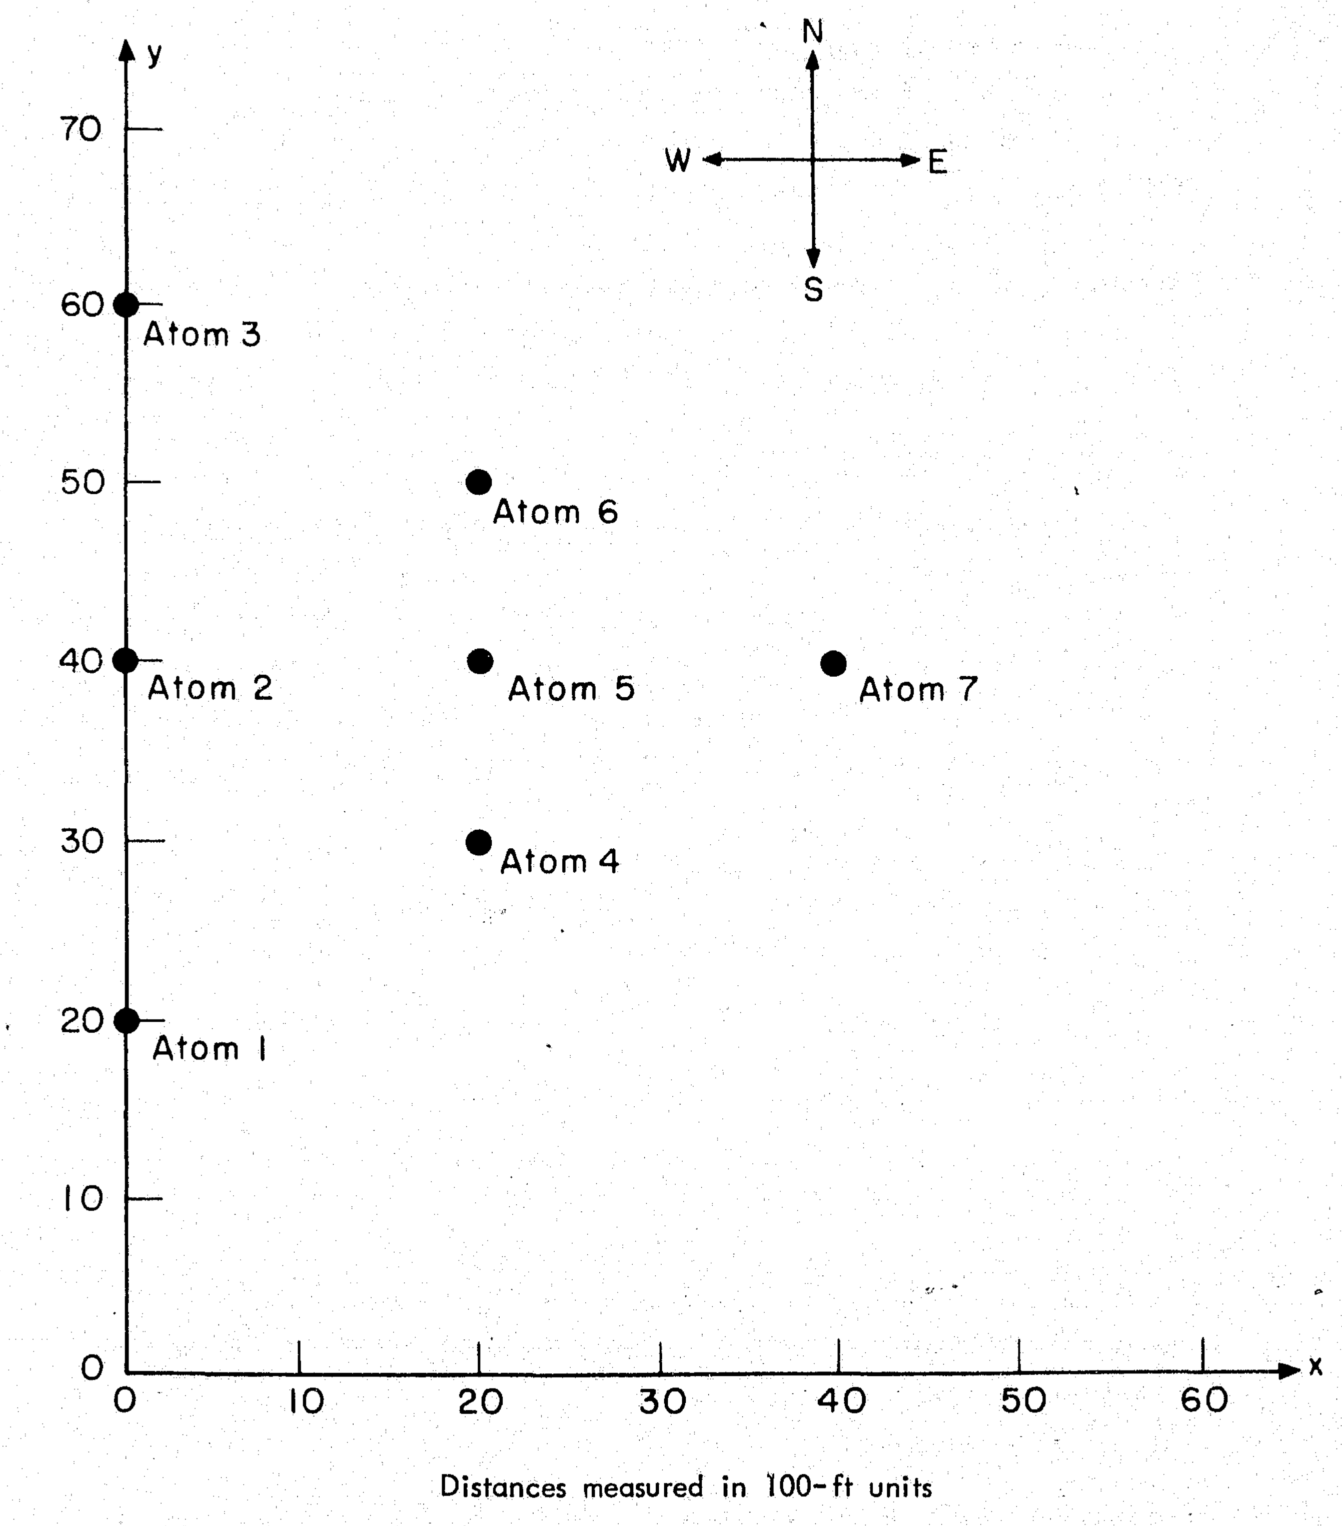

Fig. 1 — Hypothetical region containing seven geographical atoms

M = number of response units
R = number of geographical atoms (reporting areas)
NUM = number of different workload levels
ESTSTAT = 0,1,2, indicating type of model(s) to be used (exact, approximate, or both)

IMPORTANT: Note the required semicolon after the last data entry on the card. The reader will observe that most data cards do not require a semicolon. Those that do will be flagged.

The first three symbols (M, R, and NUM) have maximum values. When using the exact model, M cannot exceed 15 (i.e., no more than 15 response units), whereas M can take on any reasonable value (say less than 50) when using the approximate model (setting ESTSTAT = 1). The number of reporting areas R cannot exceed 200 and the number of workload levels per run (NUM) cannot exceed 10. The debug timing option may be added to this card; it is explained in Sec. 6.3.

Card Type 2: Title Card

The second card is particularly easy to fill out. In this simple example case it reads

```
'TITLE'     'SAMPLE 3-CAR RUN WITH 7 ATOMS'
```

This title card is used to uniquely label the run that the user is submitting. The word TITLE is typed within single quotes, followed by at least one space, and then by the title of the run (not to exceed 50 characters), placed within single quotes. The title is one of the first pieces of information to be printed by the output component of the program.

Card Type 3: Workload Distribution

The third card is printed as follows:

```
          1      2      3      4      5      6      7
'LAM'   1000.  1000.  1000.  1000.  2000.  1000.  1000.
```

This card, denoted by 'LAM,' indicates the relative workload distribution (in numbers of calls for service) among the geographical atoms. The entries are ordered in the same way that the geographical atoms are numbered in Fig. 1. In general, the atoms in the region can be ordered in any arbitrary way, but this ordering must be used consistently in preparing the input data card. In this example, atom 5 incurs twice the call-for-service workload of any of the other atoms. It is *not* necessary to normalize the sum of the workloads (so that they add to 1 or 100 or 1000, for instance). The relatively large entries in this case could correspond to the actual number of calls for service recorded from each atom during a six-month or one-year data collection period.

Format requires that at least one space follow 'LAM' and each entry is separated by at least one space. It is allowable to continue printing the entries on successive data cards, should one card be insufficient. However, 'LAM' is typed only once--as the first entry on the first type 3 data card. No special entry on a continuation card is required.

Card Type 4: Spatial Allocation of Response Units

Perhaps the key type of data card in any particular run, the 'SS' card specifies the location of a response unit when it is not servicing calls. The location may be fixed or mobile. One card must be submitted for each response unit. For this example the three cards read as follows:

```
'SS'   1    3    1   1.0   2   0.0   3   1.0      1st card in deck
'SS'   2    3    4   1.0   6   1.0   7   2.0      2nd card in deck
'SS'   3    1    5   1.0                          3rd card in deck
  ↑    ↑    ↑    ↑    ↑    ↑    ↑    ↑    ↑
 (a)  (b)  (c)  (d)  (e)  (f)  (g)  (h)  (i)
```

Each card contains the following information, in order from left to right:

a. card type identifier (this field will always contain 'SS')
b. response unit's identification number
c. total number of geographical atoms in that unit's district
d. identity number of first such atom
e. relative amount of nonbusy time spent in first atom (decimal) automatically normalized by the computer
f. identity number of second such atom
g. relative amount of nonbusy time spent in second atom (decimal)...

(This continues for as many atoms as necessary. In the example items (h) and (i) refer to the third atom.) Note, in the example, that units 1 and 2 have mobile locations, whereas unit 3 is fixed in atom 5.

Points to note:

- The decimal quantities indicating the relative amounts of time spent in each atom need not be normalized to 1 or 100 or any other number. The computer performs the normalization automatically.
- The atoms in each unit's district may be listed in any arbitrary order.
- Districts may overlap. That is, we could have one or more geographical atoms contained in two or more districts. Example:

```
'SS'  1  3  1  1.0  2  0.0  3  1.0    1st card in deck
'SS'  2  3  4  1.0  6  2.0  7  1.0    2nd card in deck
'SS'  3  2  5  1.0  3  4.0            3rd card in deck
```

The quantities on the third card indicate that unit number 3 allocates 80 percent of its nonbusy time in atom 3 (and 20

percent in atom 5). Similarly, from the first card, unit number 1 allocates 50 percent of its nonbusy time in atom number 3 (and equal time in atom 1). Thus atom 3 "belongs" to both districts 1 and 3. (See Sec. 5.1 for further discussion.)

- In police applications, every geographical atom must be contained in at least one unit's district.* This raises the following question: "What does one do with atoms in which no nonbusy time is spent by any response unit?" Each such atom must still be "in" a district, and so must appear on an 'SS' card. The only change is that the user sets the relative amount of nonbusy time in that atom equal to 0.0, as is done in the first 'SS' card for atom 2 in the above example.

Card Type 5: Response Speed (mph)

This simple card indicates the effective speed of response of the response units, in miles per hour:

```
'SPEED'    10.0
```

In this case, the effective response speed (a decimal) is 10.0 mph.

Card Type 6: Locations of Atoms (100-ft units)

This type of card specifies the location, in X-Y coordinates, of each of the atoms listed in the same order as they were on card type 3. For the example, the card reads as follows:

```
'TX'  0.0  20.0  0.0  40.0  0.0  60.0  20.0  30.0   1st card in deck
 20.0  40.0  20.0  50.0  40.0  40.0                 2nd card in deck
```

---

*This requirement actually applies only when using two of the four preprogrammed dispatch strategies. See Sec. 4.4.

The card identifier, 'TX,' is followed by the X-Y coordinates (in decimal, measured in units of 100 ft) of each of the seven atoms in our example. (There is a continuation card because all of the information could not be typed on the first card.) For instance, atom 1 is located at $X_1 = 0.0$, $Y_1 = 20.0$ (see Fig. 1), and these are the first two numbers printed on the 'TX' card. As another example, atom 6 is located at $X_6 = 20.0$, $Y_6 = 50.0$, and these figures are displayed as the sixth coordinate pair on the two cards. Since the coordinates are given in units of 100 ft, we see that atom 6 is (50.0 - 20.0) × 100 = 3000 ft "north" of atom 1 and (20.0 - 0.0) × 100 = 2000 ft "east" of atom 1.

The computer model uses this location information when estimating travel time between atoms. Assuming a street grid structure for the city, the computer model requires that the responding unit travel the sum of the east-west distance and the north-south distance between its initial location and the location of the incident.* For example, the travel distance between atoms 1 and 7 is computed as follows:

$$\text{East-west distance} = X_7 - X_1 = 40.0 - 0.0 = 40.0$$
$$\text{North-south distance} = Y_7 - Y_1 = 40.0 - 20.0 = 20.0$$

Total travel distance = 60.0 100-ft units or 6000 ft or about 1.13 mi. The travel time is equal to the travel distance divided by the effective travel speed, which in this case equals [1.13 mi/10 mi/hr] = 0.113 hr = 6.82 min.

Administrators in some cities may be fortunate enough to have an empirically devised table of travel times from point to point. In that case the 'TX' card is unnecessary and the empirical values can be read in directly; the details are discussed in Sec. 4.3.

Card Type 7: Dispatch Procedure

This card indicates the type of position and response time estimation procedure the dispatcher employs when making dispatch assignments. The sets of options are brought about by units whose (mobile) positions are not known with certainty at the time of dispatch. The detailed discussion of the four options is rather technical and is therefore deferred to Sec. 4.4. For the example we fill out the card as follows:

'SCM'

The identifier 'SCM' indicates that a "strict center of mass" dispatching policy is used. This policy, which closely models the behavior of many police dispatchers, assumes that the dispatcher acts as if each available response unit were positioned at the statistical center of its district and as if the incident were positioned at the statistical center of its district. Travel time estimates, and therefore dispatch decisions, are made on the basis of these assumptions. There are three other (preferable) policies that the dispatcher (or more precisely, the dispatching algorithm) can employ; they are discussed in detail in Sec. 4.4.

Card Type 8: Queue Capacity

This optional card, a default override card, specifies how the response system functions when all units are simultaneously busy and additional calls for service arrive. The card contains no data other than the card identifier:

'CAP'

It signifies that there is limitless (or unlimited) queuing capacity and that backlogged calls should be held in queue and dispatched to units in a first-come, first-served manner. The absence of this card indicates that there is zero queue capacity, that is, no queuing is to occur and calls which arrive when all units are busy are to be handled by a backup response system. For example, a police department

---

[*] The user can input exceptions to this rule (see Sec. 4.3).

will often provide such backup capability for an ambulance service. The model does not compute performance measures for the backup system other than the total fraction of calls that are handled by that system.

Central to the use of the model is the idea of a "default" in the computer program, and the 'CAP' card illustrates a default override card. In certain cases in which the user does not specify an action or a value for a parameter or variable, the computer program assumes a default value. It is important to recognize these defaults, since they play a critical role in defining the assumptions and operating rules of the model. In this example, the default assumption is that there exists a backup system to handle calls that arrive when all units are busy. The default override indicates that, to the contrary, such calls are to be held in queue and dispatched on a first-come, first-served basis.

Card Type 9: Service Time Card (min)

This type of card gives the average total time (including travel time as well as on-scene time*) required for a unit to service a call. This total time is called the *service time*. The card reads as follows:

```
'SERVTM'                    25.
```

Service time is measured in minutes; thus the card states that the average service time is 25 min. If card type 9 is not inserted into the input deck, the *default* value for average service time is 30 min.

Card Type 10: The 'RUN' Card (calls/hour)

The final input card, called the 'RUN' card, provides the final two pieces of information necessary to run the model:

---

[*] See Sec. 1 (particularly point 9) and Sec. 2.2 for a further discussion of service time, travel time, and on-scene time.

```
'RUN'          2.88          2.0
```

The first entry following 'RUN' (the card identifier) tells the computer that an average workload of 2.88 calls per hour is experienced by the emergency response system under study. So, the first run will be executed with this workload.

Recalling card type 1, which gives basic program specifications, the fact that NUM = 2 indicated that two runs are to be made with this input data deck. The second run will be executed with a higher workload of (2.88 + 2.00) = 4.88 calls per hour. The 2.0 is the increment to be added to the workload for each additional run. If we had had NUM = 3, then a third run would be executed with a workload of (2.88 + 2(2.00)) = 6.88 calls per hour.

To summarize the entries on the 'RUN' card, the first is the workload (in calls per hour) generated from the region being modeled; the second is the increment to be added to the workload to compute the workload for each successive run.

It must be noted that for unlimited queue capacity systems, the workload cannot be so great that there are not enough units (and time) to service all calls. For such systems all call-for-service workloads (including added increments) must be less than the number of servers (M) divided by the average service time in hours (SERVTM/60).

3.2 PRINTED RESULTS OF THE RUN

The program outputs from the sample run set up in Sec. 3.1 are displayed in Figs. 2-8. Interpretation of these results is given in Sec. 3.3 below.

3.3 INTERPRETATION OF THE OUTPUT

The first page of the output (Fig. 2) simply verifies some of the input data. In particular, for this example it states that there are three response units and consequently three districts, and seven geographical atoms; it gives the probability distribution of calls for service, by atom. In examining the distribution, it is seen that 25

```
NSF/RANN-HUD SPATIALLY DISTRIBUTED QUEUING MODEL OF AN
URBAN EMERGENCY SERVICE SYSTEM
RESPONSE UNIT=          RESPONSE_UNIT     TOTAL NUMBER=  3
DISTRICT=               DISTRICT          TOTAL NUMBER=  3
GEOGRAPHICAL ATOM=      ATOM              TOTAL NUMBER=  7


    CALLS FOR SERVICE  DISTRIBUTION, BY ATOM
        1         0.12500
        2         0.12500
        3         0.12500
        4         0.12500
        5         0.25000
        6         0.12500
        7         0.12500
```

Fig. 2—Certain input data (first page of program output)

```
               MEAN TRAVEL TIMES FOR EACH RESPONSE_UNIT
               TO EACH ATOM

    ATOM            ID OF RESPONSE_UNIT
    ID
    NO     UNIT        UNIT        UNIT
             1           2           3
      1    2.27        5.68        4.55
      2    2.27        3.98        2.27
      3    2.27        5.68        4.55
      4    4.55        2.27        1.14
      5    4.55        1.70        0.00
      6    4.55        2.27        1.14
      7    6.82        1.70        2.27
```

Fig. 3—Mean travel times for each unit to each atom
(second page of program output)

```
    STRICT CENTER-OF-MASS DISPATCHING

    ESTIMATED "COST" OF DISPATCHING I_TH RESPONSE_UNIT
         TO J_TH ATOM

    ATOM            ID OF RESPONSE_UNIT
    ID
    NO     UNIT        UNIT        UNIT
             1           2           3
      1    0.00        3.41        2.27
      2    0.00        3.41        2.27
      3    0.00        3.41        2.27
      4    3.03        0.38        0.76
      5    2.27        1.14        0.00
      6    3.03        0.38        0.76
      7    3.03        0.38        0.76
```

Fig. 4—Estimated dispatch costs (third page of program output)

```
    RESPONSE_UNIT         SPATIAL ALLOCATION, WHILE AVAILABLE
ATOM              ID OF RESPONSE_UNIT
NO.        UNIT         UNIT         UNIT
             1            2            3
   1       0.500        0.000        0.000
   2      -0.500        0.000        0.000
   3       0.500        0.000        0.000
   4       0.000        0.250        0.000
   5       0.000        0.000        1.000
   6       0.000        0.250        0.000
   7       0.000        0.500        0.000


THE FOLLOWING ATOM      GROUPS HAVE BEEN FORMED BECAUSE
OF IDENTICAL DISPATCH PREFERENCES
      1      2      3
      4      6      7
```

Fig. 5—Spatial allocation of response units, while available
(fourth page of program output)

```
NSF/RANN - HUD SPATIALLY DISTRIBUTED QUEUING MODEL OF AN
URBAN SERVICE SYSTEM:  COMPUTED PERFORMANCE MEASURES
PROBLEM TITLE: SAMPLE 3-CAR RUN WITH 7 ATOMS

    ***ITERATIVE APPROXIMATION METHOD USED***
    NUMBER OF ITERATIONS REQUIRED:          3
UNLIMITED CAPACITY QUEUE WITH 1_ST-COME 1_ST-SERVED QUEUE DISCIPLINE
RUN NUMBER:  1
RESPONSE_UNIT        ...TOTAL NUMBER OF =  3
ATOM       ...TOTAL NUMBER OF =   7
AVERAGE SERVICE TIME=  25.00 MINUTES
AVERAGE NUMBER PER HOUR OF CALLS FOR SERVICE  =   2.880
AVERAGE NUMBER PER  25.00 MINUTES OF CALLS FOR SERVICE  =   1.200
AVERAGE UTILIZATION FACTOR (IN THE CASE OF UNLIMITED LINE CAPACITY)=   0.400


REGION-WIDE AVERAGE TRAVEL TIME=   2.202 MINUTES


AVERAGE TRAVEL TIME FOR QUEUED CALLS=   2.912 MINUTES
PROBABILITY OF SATURATION= 0.14118
REGION-WIDE AVERAGE WORKLOAD (% TIME BUSY)= 0.40000
STANDARD DEVIATION OF WORKLOAD=   0.010
MAXIMUM WORKLOAD IMBALANCE= 0.02019


FRACTION OF DISPATCHES THAT ARE INTER-DISTRICT = 0.35255
```

Fig. 6— Region-wide performance measures (fifth page of program output)

```
PERFORMANCE MEASURES THAT ARE SPECIFIC TO EACH RESPONSE_UNIT

   ID OF
   RESPONSE_UNIT                      FRACTION OF
                WORKLOAD     % OF  DISPATCHES       % OF     AVERAGE
NAME         NO  OF UNIT     MEAN  OUT OF DISTRICT MEAN  TRAVEL TIME
 UNIT         1   0.388      97.1     .2351         66.7       2.759
 UNIT         2   0.409     102.1     .2975         84.4       2.370
 UNIT         3   0.403     100.8     .5209        147.8       1.496


PERFORMANCE MEASURES THAT ARE SPECIFIC TO EACH DISTRICT

   ID OF
  DISTRICT                            FRACTION OF
                WORKLOAD     % OF  DISPATCHES       % OF     AVERAGE
NAME         NO  OF DISTRICT MEAN  INTER-DISTRICT   MEAN  TRAVEL TIME
DIST          1   0.450     112.5     .3418         97.0       2.911
DIST          2   0.450     112.5     .3609        102.4       2.335
DIST          3   0.300      75.0     .3556        100.9       0.937
```

Fig. 7—Unit-specific and area-specific performance measures (sixth page of program output)

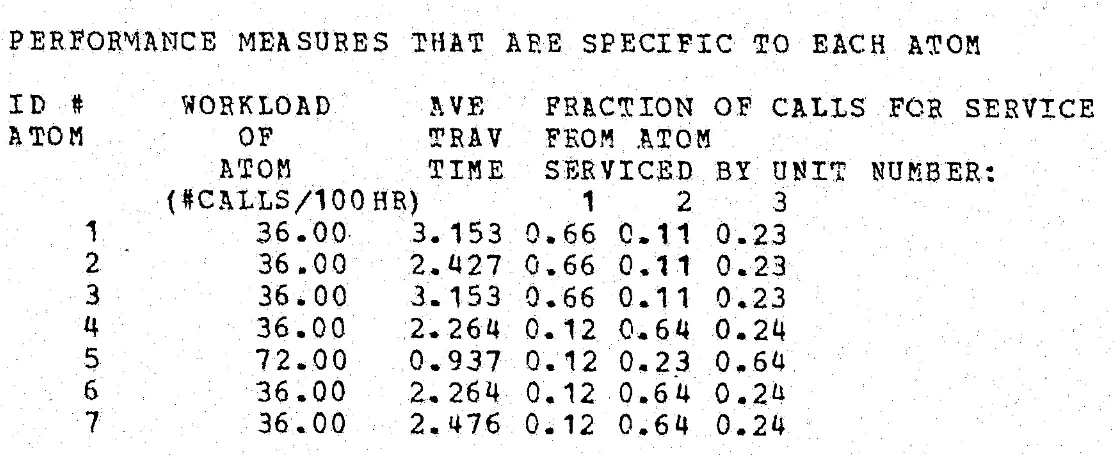

```
PERFORMANCE MEASURES THAT ARE SPECIFIC TO EACH ATOM

ID #      WORKLOAD       AVE   FRACTION OF CALLS FOR SERVICE
ATOM        OF           TRAV  FROM ATOM
           ATOM          TIME  SERVICED BY UNIT NUMBER:
       (#CALLS/100HR)             1    2    3
    1        36.00      3.153 0.66 0.11 0.23
    2        36.00      2.427 0.66 0.11 0.23
    3        36.00      3.153 0.66 0.11 0.23
    4        36.00      2.264 0.12 0.64 0.24
    5        72.00      0.937 0.12 0.23 0.64
    6        36.00      2.264 0.12 0.64 0.24
    7        36.00      2.476 0.12 0.64 0.24
```

Fig. 8—Atom-specific performance measures (seventh page of program output)

percent of the calls are generated from atom 5 and 12.5 percent are generated from each of the other atoms. This is consistent with the data read in on card type 3 (workload distribution).

The second page of the output (Fig. 3) gives the calculated mean travel times for each unit (while positioned in its own district) to travel to each of the geographical atoms. For instance, the time for unit 2 to travel to atom 3 is 5.68 min.

The third page (Fig. 4) gives the estimated "cost"--the travel time as estimated by the dispatcher (or dispatching algorithm)--of dispatching each of the units to each of the atoms, assuming in this case strict center of mass dispatching. Dispatching strategies are discussed in Sec. 4.4; it will be useful when reading Sec. 4.4 to compare the third page of the printout (Fig. 4) to the entries in Tables 1-3 and Fig. 9 (p. 44).

The fourth page indicates what fraction of available time each unit spends in each of the atoms. Any entry with a minus (-) allocation means that the corresponding unit spends no available time there but that the atom in question is "contained" in that unit's district. (The question of atoms belonging to districts is important when considering dispatch policies--see Sec. 4.4.) From Fig. 5, it is seen that unit 2 spends 50 percent of its available time in atom 7 and 25 percent in each of atoms 4 and 6. Unit 3 spends all of its time in atom 5. Unit 1 splits its efforts 50-50 between atoms 1 and 3, but considers atom 2 to be in its district. These figures are consistent with the data read in from card type 4.

Also on the fourth page the computer prints out the results of some internal computations. It states that for many computations, atoms 1, 2, and 3 will be considered to be equivalent (as will atoms 4, 6, and 7) because of identical dispatch preferences. That means that the ordered rankings of preferred units are identical for atoms 1, 2, and 3 and for atoms 4, 6, and 7. In actual police or ambulance operations this information may be useful in itself, with or without consideration of the Hypercube Model.*

---

*See Refs. 1 and 2 for a more complete discussion of fixed-preference dispatch strategies.

The major output segment, containing the majority of the computed performance measures, is given on the fifth page of the output (Fig. 6). Each of the printed output lines will be interpreted in turn:

```
NSF/RANN-HUD SPATIALLY DISTRIBUTED QUEUING MODEL OF AN URBAN
SERVICE SYSTEM:  COMPUTED PERFORMANCE MEASURES
PROBLEM TITLE:  SAMPLE 3-CAR RUN WITH 7 ATOMS
```

These lines are self-explanatory. Note that the title is consistent with that read in from card type 2.

```
***ITERATIVE APPROXIMATION METHOD USED***
```

This means that the values of the performance measures were computed from the approximate model rather than from the exact Hypercube Model. (This line would not appear if the exact model were used.)

```
NUMBER OF ITERATIONS REQUIRED:  3
```

This tells the user the number of times that the computer iterated through the approximation equations (Ref. 2) to obtain a satisfactory solution. In this case three iterations were required.

```
UNLIMITED CAPACITY QUEUE WITH 1-ST COME 1-ST SERVED QUEUE DISCIPLINE
```

This says that calls for service that arrive when all units are simultaneously busy are entered into a queue or waiting line, which is depleted or serviced in a first-come, first-served (FCFS) manner. This line is consistent with our use of the 'CAP' card (Sec. 3.1).

```
RUN NUMBER:  1
```

This states that the current output (on the fifth page, in this case) represents run number 1 (as defined on card type 10).

```
RESPONSE_UNIT    ...TOTAL NUMBER OF=3
ATOM     ...TOTAL NUMBER OF = 7
```

Self-explanatory.

```
AVERAGE SERVICE TIME = 25.00 MINUTES
```

This confirms the input value for average service time obtained from card type 9.

```
AVERAGE NUMBER PER HOUR OF CALLS FOR SERVICE = 2.880
AVERAGE NUMBER PER 25.00 MINUTES OF CALLS FOR SERVICE=1.200
```

Since in run 1 there are 2.88 calls for service generated per hour (on average), there are 2.88 × (25/60) = 1.200 calls for service generated each 25 min (on average). With unlimited queue capacity, this is important for the following reason: If there were only one response unit to handle this workload, then an average of 1.200 calls would arrive during the time required (on average) to service one call. Thus, one response unit would not be able to handle this workload. However, two (or more) could do the job without having backlogs of calls build up indefinitely. If this figure had been 2.200 instead of 1.200, then at least three response units would have been required to handle the workload. In general, whatever this figure is, the next highest integer is the minimum number of response units required to do the job. If the user attempts to run the model with too few response units (assuming unlimited queue capacity), then the run stops and the following error message is printed: "QUEUE SATURATED."

With zero queue capacity, the user can operate the model with any number of response units. However, if this number is too small (in comparison to the workload generated per hour), then a large fraction of calls for service will be handled by the backup response system (see Sec. 3.1, description of the 'CAP' card).

AVERAGE UTILIZATION FACTOR (IN THE CASE OF UNLIMITED LINE CAPACITY) = 0.400

Since 1.200 calls for service arrive (on average) each 25 min and all (eventually) are assigned to a unit, then on average 1.200 × (1/3) = 0.400 call is assigned to any particular response unit every 25 min. But each such assignment requires (on average) 25 min to service, thus the "average response unit" is busy servicing calls 40.0 percent of the time. This figure is called the average utilization factor (referring to the fraction or percent of time that response units are servicing calls). Reflecting the discussion above, this utilization factor must be less than 1.0 (or 100.0 percent) for the case of an unlimited capacity queue.

For the case of a zero line capacity queue, the units would incur an average utilization factor somewhat less than 40.0 percent, due to overflow calls being sent to a backup emergency response system.

All of the printed output lines to this point have simply restated various input data. The remaining lines give values for the various performance measures, as computed from the model.

REGION-WIDE AVERAGE TRAVEL TIME=2.202 MINUTES

This indicates that the average travel time to a call for service, averaged over all the atoms in the region, is 2.202 min. Since the travel speed is 10 mph (or 1/6 mi/min), this implies that the average distance traveled per response is (1/6) × (2.202) ≈ 0.367 mi, a figure that is intuitively reasonable. (The entire region is just 4000 ft or 0.76 mi east-west and north-south.)

AVERAGE TRAVEL TIME FOR QUEUED CALLS=2.912 MINUTES

Here the program is showing the somewhat larger average travel time incurred by calls that are delayed in queue (averaged over all calls delayed in queue, regardless of geographical atom). Of course, in application, a travel time of 2, 3, or 5 min may be insignificant compared to a queue delay of 15, 30, and 60 min.*

PROBABILITY OF SATURATION = 0.14118

Saturation is said to occur when all units are simultaneously busy. If this occurs X percent of the time, then (due to the random arrival patterns of calls for service) X percent of the calls reach a saturated system and thus must be held in the dispatcher queue. In this case, 14.118 percent of all calls for service are held in queue. (In the case of a zero line capacity system, if X percent of the time all units are busy, then X percent of all calls for service are transferred to the backup system.)

REGION-WIDE AVERAGE WORKLOAD (% TIME BUSY) = 0.400

This is the average fraction of time that units are computed to be busy. In the case of the infinite line capacity system, this figure will equal (within acceptable round-off errors) the average utilization factor discussed earlier.

STANDARD DEVIATION OF WORKLOAD = 0.01

This is the standard deviation of the workload distribution, which is one measure of the imbalance in workloads among units. The larger this quantity, the greater the imbalance would be. If this quantity were zero, then the workloads of all units would be equal.

MAXIMUM WORKLOAD IMBALANCE = 0.02019

Subtracting the workload of the least busy unit (unit 1 in this case) from the workload of the busiest unit (unit 2) gives the *maximum*

---

[*] Later versions of the model will print out the estimated average delay in queue (see Appendix A).

workload imbalance. (In this case the maximum workload imbalance is 0.409 - 0.388 = 0.021, which, within error tolerances, is the same as 0.02019. The least busy unit has only 5 percent less work than the busiest unit, a very small imbalance.)

```
FRACTION OF DISPATCHES THAT ARE INTER-DISTRICT=0.35255
```

This says that 35.255 percent of all dispatch assignments (including those from a queue of calls) cause the assigned unit to travel to a reporting area not in its own district. Thus, for a randomly selected call for service, there is a 35.255 chance that the unit which responds to that call will not be the unit whose district contains the call.

The above completes the summary of region-wide performance measures. The remaining performance measures are based first on the response unit, second on the district, and third on the atom.

<u>Unit-Specific Performance Measures</u>

Examining the performance measures for unit 1 (identified in the far left columns in the top portion of the sixth page of output, Fig. 7), we see the following (reading from left to right):

- Unit 1 spends 38.8 percent of its time handling calls for service.
- That workload is 97.1 percent of the average workload of all three units (which is 40.0 percent in this case).
- Only 23.51 percent of the dispatch assignments to unit 1 cause the unit to leave its district.
- That figure for interdistrict dispatch frequency is only 66.7 percent of the mean for the three units (which is 35.255 percent in this case).
- The average time it takes for unit 1 to travel to the scene of an incident is 2.759 min.

Similar interpretations apply to units 2 and 3.

<u>District-Specific Performance Measures</u>

The lower portion of Fig. 7, sixth page of output, gives district-specific performance measures. For district 1, we see the following (reading from left to right):

- The district's workload is enough to cause one response unit to remain busy servicing calls 45.0 percent of the time (if that response unit handled all of district 1's calls and no others).
- This figure is 12.5 percent above the mean for the three districts.
- The fraction of the district's dispatches that require an out-of-district unit (either unit 2 or 3) because unit 1 is unavailable is 34.18 percent.
- The average travel time to incidents in district 1 is 2.911 min.

Similar interpretations apply to districts 2 and 3.

This completes our discussion of the sixth page of output, which is the major output page of the model.

<u>Atom-Specific Performance Measures</u>

The seventh page of output (Fig. 8) contains values of several atom-specific performance measures. For atom 1, for instance, we see the following (reading from left to right):

- That atom generates an average of 36.00 calls per 100 hours.
- The average travel time to the atom is 3.153 min.
- Sixty-six percent of all calls from the atom are handled by unit 1, 11 percent by unit 2, and 23 percent by unit 3.

Judicious examination of the results in this figure will allow the user to spot inequities in the distribution of service accessibility to neighborhoods (service accessibility being measured by average travel time).

Run Number 2

The results of run number 2, in which the call rate is increased to 4.88 calls/hour, are provided in the same format as run number 1 and thus are not shown.

4. OPTIONS

The previous section presented a simple example of use of the model; this section will discuss each of the several input data options in detail. While it is helpful to read this section in its entirety, most users will not find it necessary to use each of the options discussed. To the extent possible, each of the options is discussed independently of the others, so users can treat this section simply as a reference handbook.

4.1 GLOSSARY

Every city or town has its own labels for response units, neighborhoods, districts, etc. The intent of the glossary option is to allow the user to input these names into the program so that the final output describes the response patterns in terms familiar to residents of the community being modeled.

If a user wishes to use the glossary option, he inserts the following card in the input deck (immediately after card type 1, the card containing the basic program specifications):

```
'GLOSSARY'
```

Then, immediately following the 'GLOSSARY' card, the user inserts data giving the names for the various components of the emergency response system being modeled that he would like to change.

The following are examples of changes to each of the terms that can be made using the 'GLOSSARY' option:

```
R_DIST='BEAT' NM_UNIT(3)='SGT_CAR' NO_UNIT(3)=2NM_DIST(3)='DOWNTOWN' NO_DIST(3)=2
    R_UNIT='POLICE_CAR' T_COST='TRAVEL TIME'  CFS='CALLS FOR SERVICE'
          ATOM='ATOM';
```

Note that a semicolon must follow the last entry.

As the above data cards indicate, there are nine types of definitions that can be included in the glossary. Each of these is treated individually below. The information on each is summarized in Table 1.

R_DIST

The variable R_DIST is the generic name for district in the region being modeled. In a police context, R_DIST may be 'BEAT', 'SECTOR', 'ZONE', 'DISTRICT', etc. Whatever the choice, the definition for R_DIST cannot exceed eight characters in length. (Here, as elsewhere in the manual, the underline (_) in R_DIST represents a *typed* underline, not a blank.) In this case, the output of the program will refer to beats rather than districts.

If the user does not choose to define R_DIST, then as a default the computer sets R_DIST = 'DISTRICT'. All of the other eight glossary terms have associated default definitions, as discussed below and summarized in Table 1.

NM_UNIT(I), NO_UNIT(I)

The variable NM_UNIT(I) is the name (usually a type of unit) given to the Ith response unit. In the example, the third response unit is given the name 'SGT_CAR'. Since there are M response units, in general the user may provide up to M names for response units. The definition for NM_UNIT(I) cannot exceed eight characters in length. If no definition is given for unit I, then the default sets UNIT(I) = 'UNIT'.

By using the variable NO_UNIT(I), the user has the option of giving the Ith response unit a number (a nonnegative integer less than 1000). In the example, the third response unit, which has been named 'SGT_CAR', is assigned the number 2. (The output from the model will list this car as SGT_CAR 2.) This suggests that the third response unit is the second sergeant's car that is being fielded.

Some cities like to use three digits to identify their response units, the first being a prefix indicating the command (precinct, hospital zone, etc.) of the unit, and the second two uniquely identifying the unit. These numbers can be assigned by using this option.

Table 1

SUMMARY OF GLOSSARY OPTIONS

Glossary Variable	Illustrative Definitions	Maximum No. of Characters (including spaces)	Default Definition
R_DIST	'BEAT' 'SECTOR' 'DISTRICT' 'ZONE' 'RES_AREA' 'HOSP_ZN'	8	'DISTRICT'
NM_UNIT(I)	'CAR' 'SGT CAR' 'AMBUL' 'TRUCK'	8	'UNIT'
NO_UNIT(I)	I,901,367	3 digits	I
NM_DIST(I)	'BEAT' 'SECTOR' 'ZONE' 'DOWNTOWN' 'BACK BAY' 'CEN PARK'	8	'DIST'
NO_DIST(I)	I,427,67	3 digits	I
R_UNIT	'AMBULANCE' 'POLICE CAR' 'LADDER TRUCK' 'WAGON'	18	'RESPONSE UNIT'
T_COST	'TRAVEL TIME' 'TRAVEL DISTANCE'	18	'TRAVEL TIME'
CFS	'CALLS FOR SERVICE' 'AMBULANCE CALLS' 'FIRE ALARMS' 'NO EMER REPAIRS'	18	'CALLS FOR SERVICE'
ATOM	'REP AREA' 'ZONE'	8	'ATOM'

The default is NO_UNIT(I) = I. The two sets of defaults--for NM_UNIT(I) and NO_UNIT(I)--may produce strange-looking printouts if only a few values of I are used with this option (thereby relying on the default for the other values). The user is cautioned to think carefully about the names and numbers of each of the units if he uses this option for some (but not all) of the units.

NM_DIST(I), NO_DIST(I)

The variable NM_DIST(I) is the name given to the Ith district. In the example, the third district (which has been defined to be a beat by the definition of R_DIST) is given the name 'DOWNTOWN'. The definition for NM_DIST(I) cannot exceed eight characters in length. The default is NM_DIST(I) = 'DIST', no matter what definition is used for R_DIST.

In a manner similar to numbering response units, the user can assign numbers to districts by using the variable NO_DIST(I). In the example, the DOWNTOWN beat is assigned the number 2, and will be referred to in the output as DOWNTOWN 2. Again, the number may be any nonnegative integer less than 1000.

The default is NO_DIST(I) = I. The same comments regarding defaults that apply to naming and numbering response units apply to naming and numbering districts.

R_UNIT

The variable R_UNIT is the generic name for all response units within the region. The definition for R_UNIT may contain up to 18 characters, with a default of R_UNIT = 'RESPONSE_UNIT'. In the example, R_UNIT = 'POLICE_CAR', so the output from the model will refer to police cars instead of response units.

T_COST

The variable T_COST is the name given to the cost of travel for units to reach the scenes of incidents. Usually (as in the example) T_COST = 'TRAVEL TIME', but T_COST could be some other measure of cost of delay such as T_COST = 'TRAVEL DISTANCE'. The maximum length of the definition is 18 characters, with a default of T_COST = 'TRAVEL_TIME'.

CFS

The variable CFS is the name given to calls for service. Due to convention, CFS is stated as a plural. Thus, the user could have CFS = 'CALLS FOR SERVICE' (the default), or CFS = 'AMBULANCE CALLS', or CFS = 'FIRE ALARMS', or CFS = 'NO EMER REPAIRS'. The maximum length of the definition is 18 characters.

ATOM

The variable ATOM is the name for the small reporting areas upon which the geographical data are based. The default is ATOM = 'ATOM', and a maximum of eight characters is permitted. Note that unless the 'ATOM_NO' option is used (see Sec. 4.8), the program will assume that atoms are numbered consecutively starting with 1.

4.2 PRINTOUT

Currently the print options of the program are very simple (and limited). Eventually, based on user experience, the printout options will be made more flexible, perhaps automatically suppressing detailed printout of obviously poor runs.

Printing the Travel Time Matrix

The inter-atom travel time matrix is denoted by TR, where

TR(I,J) = travel time from atom I to atom J

Since this matrix is often very large, having R^2 entries (where R is the number of atoms), it is not usually printed out. However, the user can have it printed out (perhaps to verify the values contained therein) by inserting the following card between card type 2 (the 'TITLE' card) and card type 3 (the 'LAM' card, indicating workload distribution):

'PRNT_TR'

The corresponding printout for the seven-atom example of Sec. 2 is given in Fig. 9.

```
        TRAVEL TIME MATRIX:   INTFR-ATOM

ATOM    NUMBER: ORIGIN       ATOM    NUMBER:DESTINATION
                               1       2       3       4       5       5       7
              1              0.00    2.27    4.55    3.41    4.55    5.68    6.82
              2              2.27    0.00    2.27    3.41    2.27    3.41    4.55
              3              4.55    2.27    0.00    5.68    4.55    3.41    6.82
              4              3.41    3.41    5.68    0.00    1.14    2.27    3.41
              5              4.55    2.27    4.55    1.14    0.00    1.14    2.27
              6              5.68    3.41    3.41    2.27    1.14    0.00    3.41
              7              6.82    4.55    6.82    3.41    2.27    3.41    0.00
```

Fig. 9—Printout of inter-atom travel times (photoreduced)

Suppress Printing of Atom-Specific Performance Measures

If the user does not wish to see atom-specific data (e.g., travel times, fraction of incidents handled by each of the units, preventive patrol frequencies[*]), then he inserts the following card between card type 2 and card type 3:

'NO_PRNT_AT'

4.3 TRAVEL TIME DATA

The travel time assumption used as the standard default option is that travel times between geographical atoms are proportional to "right-angle" or "Manhattan" travel distances between the atoms. This subsection describes several ways in which to modify this travel time assumption. Three options are outlined, and their logical dependence is indicated in Fig. 10, at the end of Sec. 4.3.2.

4.3.1 Travel Times: Allowing Nonzero Intra-Atom Travel Times

Using the assumption that travel distances can be modeled with the Manhattan distance metric (reflecting a mutually perpendicular grid of streets), the computer calculates that all intra-atom travel times are zero. Thus, an ambulance could be positioned at the center of an atom, and travel times to incidents within that atom would be set equal to 0.0 (min).

To correct for this situation, the user can insert the following card immediately following card type 6 (the 'TX' card):

'CORTM'   0.667

The 'CORTM' card indicates that intra-atom travel times are to be corrected. The number following 'CORTM' reflects the constant of proportionality to be used in a "square-root law" that estimates intra-atom travel time to be proportional to the square root of the area of the atom. That is, when the 'CORTM' option is used, the computer estimates the mean intra-atom travel time in atom I to be

$$TR(I,I) \approx C\sqrt{\frac{A_i}{XSPEED \cdot YSPEED}}$$

where C is the constant of proportionality (0.667 in this case), A_i is the area (in sq mi), and XSPEED and YSPEED are the directional travel velocities, in mph (see Sec. 4.9). For patrolling police units, it is reasonable to set C = 0.667, whereas for units stationary near the center of the atom, the smaller value of C = 0.50 is very often appropriate.[*] Of course, some cities may have their own empirically measured value of C, and this can also be used.

---

[*] See Sec. 4.6.

[*] See Ref. 8, Chap. 3.

Immediately following the 'CORTM' card (and the accompanying proportionality constant), the user must provide as input the area of each of the geographical atoms, sequentially from the first to the last atom. An example might read as follows for a seven-atom example:

```
0.02    0.10    1.2    0.0    0.50    1.9    10.1
```

In this case, atom 3, for instance, has an area of 1.2 sq mi, while atom 4 has zero area.

4.3.2 Travel Times: Override of X-Y Coordinates Assumption

In many applications the assumption that travel distance follows the Manhattan metric, which reflects a mutually perpendicular grid of streets, is not always satisfactory. Especially in areas of the city having barriers to travel such as rivers, parks, and cemeteries, or other irregularities, the user would wish to include the effects of these situations in the model. Or, in the fortunate circumstance that the user has an empirically measured set of travel times, he may wish to by-pass the Manhattan metric assumption entirely.

Selective Override of X-Y Coordinates Assumption

For situations in which the assumption of a mutually perpendicular street grid is a reasonable one "most of the time" (say for 90 percent or more of all possible inter-atom travel times), then the user would wish to provide as input the exceptions to this assumption. To do this he inserts immediately following the 'CORTM' card or, if no 'CORTM' card is used, card type 6 (the 'TX' card) the following card:

```
'TX_OV'
```

This card signifies that data constituting an override of the 'TX' card are about to follow. Immediately following the 'TX_OV' card, the user supplies the override travel time data in minutes. For example:

```
TR(2,3) = 1.6  TR(3,2) = 2.1  TR(1,7) = 10.9  TR(6,6) = 1.0;
```

Note that a semicolon follows the completion of the override data.

The travel time from atom I to atom J is stored in the computer as TR(I,J). Thus the sample card above indicates that, regardless of the X-Y coordinates assumption, the travel time from atom 2 to atom 3 is 1.6 min, the travel time from atom 3 to atom 2 is 2.1 min (in general, T(I,J) can be different from T(J,I)), the travel time from atom 1 to atom 7 is 10.9 min, and the mean travel time within atom 6 (TR(6,6)) is 1.0 min.

As many entries as desired may be included following the 'TX_OV' card. Several data cards may be used if required. However, following the last entry on the last card, a semicolon must be inserted, signifying completion of the 'TX' override.

Empirically Derived Travel Time Matrix

If the model is being applied in an area for which inter-atom travel times have been empirically measured, then no 'TX' card (or 'TX_OV' card) is required. Instead of the 'TX' card, the following card is inserted as card type 6:

```
'TR'
```

The 'TR' card indicates that a complete inter-atom travel time matrix is to follow. *If this card is used, the user must use the 'EMCM' card as the preprogrammed dispatch selection procedure.* (See Sec. 4.4, particularly subsection 4.4.3.)

The following is an example for a region with R = 4 atoms:

Card 1:
```
1.1    2.1    1.9    3.2
```

Card 2:	2.1	0.6	2.4	4.8
Card 3:	1.86	2.4	1.4	6.1
Card 4:	4.9	9.2	6.2	0.2

As examples from these cards, the travel time from atom 1 to atom 2 is TR(1,2) = 2.1 min; the travel time from atom 3 to atom 2 is TR(3,2) = 2.4 min; and the travel time within atom 3 is TR(3,3) = 1.4 min. The program will read in these data assuming that the first R entries correspond to travel times from atom 1 to each of the other atoms (numbered sequentially), the next R entries correspond to travel times from atom 2 to each of the other atoms, etc. Note that in general R^2 entries are required, a sizable data requirement for many regions.

Figure 10 indicates the logical interdependencies of the travel time options available with the computer program.

4.4 DISPATCH SELECTION: PREPROGRAMMED PROCEDURES[*]

The model assumes a "fixed-preference" dispatch selection procedure, described as follows: Suppose a call for service arrives from atom i. There is an ordered list of units, say (3,1,7,5,6,4,2) for a seven-unit problem, that specifies the dispatcher's preference for units to assign to atom i. The dispatcher starts with the first entry in the list, unit 3 in this case, and assigns that unit, if available. If the first preferred unit is not available, then the dispatcher assigns the second preferred unit (unit 1 in this case), if available. The dispatcher continues down the list until the first available unit is found (if there is one), and assigns that unit. This procedure is called "fixed-preference" because the ordered list of preferences does not change with the state of the system. However, the list of preferences may change by atom; for instance, the list for atom i + 1 may be (5,6,3,1,7,2,4). The model can operate with any fixed-preference dispatch policy. If

---

[*] This subsection is substantially more technical than the rest of the manual. The reader is referred to the summary of dispatch strategy definitions at the end of subsection 4.4.

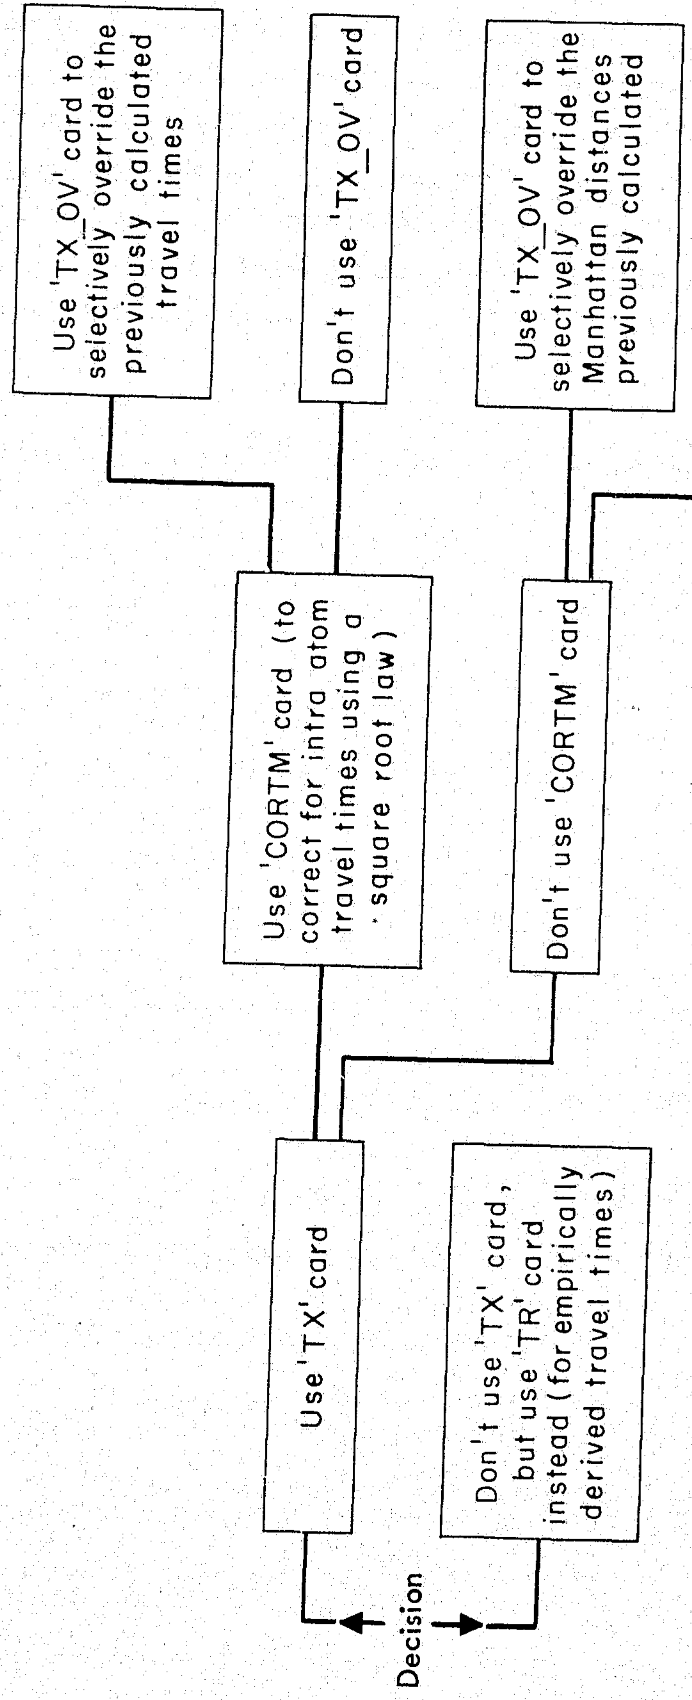

Fig. 10 — Logical dependence among travel time options

there are R geographical atoms and M response units, then there are $(M!)^R$ possible different dispatch policies. If the possibility of ties is included, the number of different policies becomes even more enormous.

Thus it is desirable to select a small number of dispatch policies that have a certain intuitive appeal as a starting basis for operating the model. As will be seen, it is possible to alter these arbitrarily so that, in fact, any one of the possible dispatch policies could be investigated. We start with the four basic "center of mass" strategies discussed in Ref. 8, Chap. 3. To define these strategies precisely, first let

(X_j,Y_j) = random variables indicating the position coordinates of response unit j.

(X,Y) = random variables indicating position coordinates of the incident (conditioned only on the district of the incident).*

(x_j^o,y_j^o) = center of mass coordinates of response unit j.

$(E[X|k],E[Y|k])$ = center of mass coordinates of incidents in district k, where E[ ] is the mathematical expectation of the random variable contained within the brackets.

(x,y) = exact position coordinates of the incident.

The meaning of these definitions can be illustrated by reconsidering the simple three-unit example of Sec. 3 (see particularly Fig. 1 and card type 4: Spatial Allocation of Response Units). The random variables indicating the position coordinates of response unit 1, (X_1,Y_1), can take on the values (0,20), (0,40), (0,60); the probability that (X_1,Y_1) will assume either of the pairs (0,20) or (0,60) is 1/2 (as indicated on the first 'SS' card). The center of mass coordinates of response unit 1, (x_1^o,y_1^o), are the respective statistical averages of the three possible coordinate pairs. Thus,

---

*Subscripts could be appended to (X,Y) to indicate explicitly the district of occurrence. It is not necessary here since the meaning is clear.

$$x_1^o = 1/2\cdot 0 + 0\cdot 0 + 1/2\cdot 0 = 0$$

$$y_1^o = 1/2\cdot 20 + 0\cdot 40 + 1/2\cdot 60 = 120/3 = 40$$

Similarly, using the second and third 'SS' cards,

$$(x_2^o,y_2^o) = (30,40)$$

$$(x_3^o,y_3^o) = (20,40)$$

Now suppose an incident occurs in atom 1. Then the exact coordinates of the incident are (x,y) = (0,20). However, because of the way geographical information is often coded for the dispatcher, he may only be aware that the incident is located "somewhere" in district 1. In that case, the location of the incident is random, conditioned only on the fact that it is located in district 1. The question then becomes, "What is a reasonable probability assignment for the incident over the atoms in district 1?" It is assumed that these probabilities are proportional to the call-for-service workloads generated from each of the atoms. Since the workload of each of atoms 1, 2, and 3 is the same (see card type 3: Workload Distribution), it can be reasonably assumed that the incident is equally likely to be in any one of the three atoms, given only that it is in district 1. In that case the center of mass coordinates of incidents in district 1 are

$$E[X|1] = 0\cdot 1/3 + 0\cdot 1/3 + 0\cdot 1/3 = 0$$

$$E[Y|1] = 20\cdot 1/3 + 40\cdot 1/3 + 60\cdot 1/3 = 40$$

Similarly, for the other two districts

$$(E[X|2],E[Y|2]) = (26.7,40.0)$$

$$(E[X|3],E[Y|3]) = (20,40)$$

These calculations are displayed on a map in Fig. 11.

The four major dispatch strategies can now be defined and interpreted. The table below shows a formula for each of the four strategies for estimating the travel time for unit j (which is assumed to be available) to travel to an incident at (x,y) located in district k:

Strategy	Travel Time Estimation Procedure
1. Strict center of mass (SCM)	$\frac{\lvert x_j^o - E[X\lvert k] \rvert}{XSPEED} + \frac{\lvert y_j^o - E[Y\lvert k]\rvert}{YSPEED}$
2. Modified center of mass (MCM)	$\frac{\lvert x_j^o - x\rvert}{XSPEED} + \frac{\lvert y_j^o - y\rvert}{YSPEED}$
3. Expected strict center of mass (ESCM)	$\frac{E[\lvert X_j - X\rvert]}{XSPEED} + \frac{E[\lvert Y_j - Y\rvert]}{YSPEED}$
4. Expected modified center of mass (EMCM)	$\frac{E[\lvert X_j - x\rvert]}{XSPEED} + \frac{E[\lvert Y_j - y\rvert]}{YSPEED}$

In these formulas, XSPEED and YSPEED are the effective speeds of response in the north-south and east-west directions, respectively.

Each of the four strategies is considered in turn. For simplicity in the discussion, we set XSPEED = YSPEED = 1 and thus consider travel time and response distance interchangably.

4.4.1 Strict Center of Mass (SCM)

With a strict center of mass strategy, the dispatcher makes travel time estimates acting as if the unit were located at the statistical center of its district and the incident were at the statistical center of its district.* In each case the underlying probability distributions of incidents and units may be different because, in general, the spatial distribution of incidents and the unit within a district are

---

*To use the SCM strategy, each geographical atom must be included in at least one district (via the 'SS' card or the 'S' card).

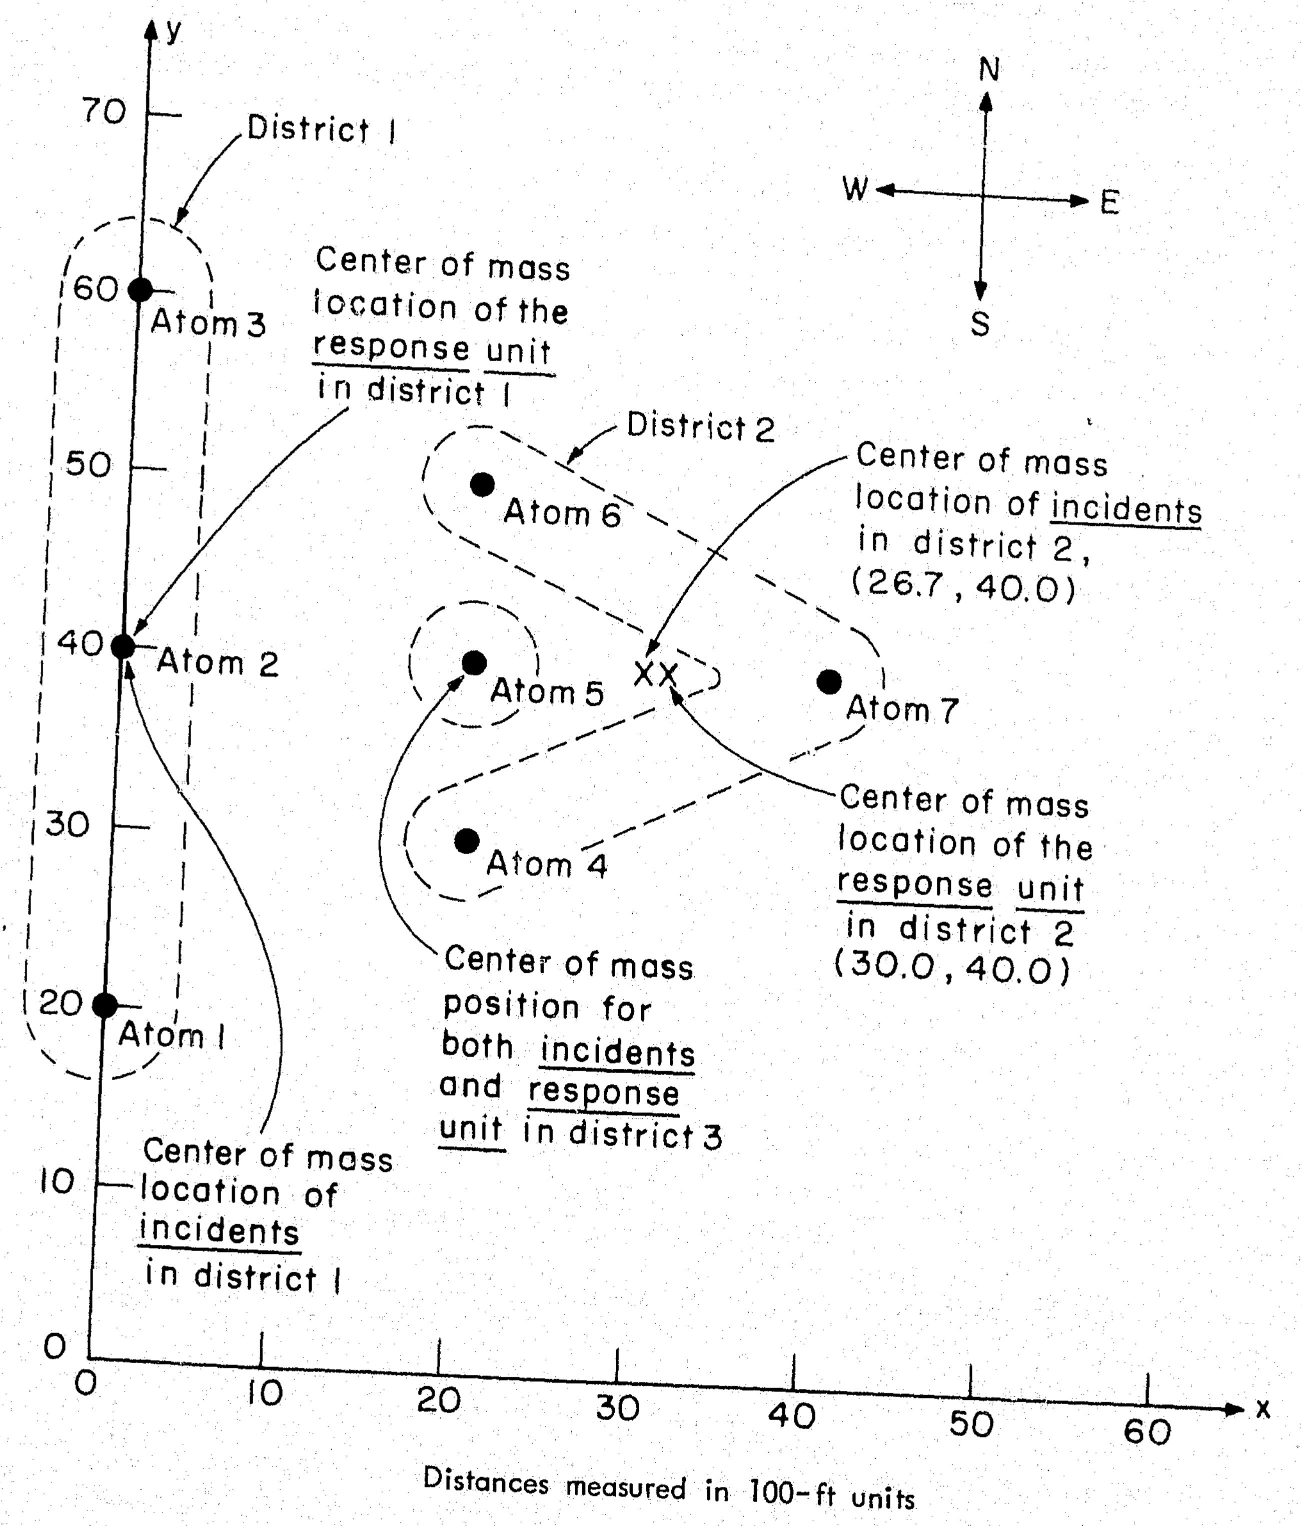

Fig. 11 — Hypothetical region: Response unit locations and districts

different. Continuing with the same example as before, suppose we reexamine the case in which the incident is in atom 1. The estimated position of the incident is the center of mass of incidents in district 1:

$$(E[X|1],E[Y|1]) = (0,40)$$

Thus, the estimated travel distances for each of the units, recognizing that the units are assumed to be at the centers of mass of their respective districts, are as follows:

Unit 1: Estimated travel distance = $|0 - 0| + |40 - 40| = 0$
Unit 2: Estimated travel distance = $|30 - 0| + |40 - 40| = 30$
Unit 3: Estimated travel distance = $|20 - 0| + |40 - 40| = 20$

Using an SCM policy, unit 1 would be the first preferred unit to dispatch to a call from atom 1, unit 2 would be the second preferred, and unit 3 would be the third preferred. The entire set of SCM dispatch preferences, for each of the atoms, is shown in Table 2.

Table 2

SCM DISPATCH POLICY: 3-UNIT, 7-ATOM EXAMPLE

	First Preferred Unit		Second Preferred Unit		Third Preferred Unit	
Atom No.	Unit No.	Estimated Dispatch Cost	Unit No.	Estimated Dispatch Cost	Unit No.	Estimated Dispatch Cost
1	1	0.0	3	20.0	2	30.0
2	1	0.0	3	20.0	2	30.0
3	1	0.0	3	20.0	2	30.0
4	2	3.33	3	6.67	1	26.67
5	3	0.0	2	10.0	1	20.0
6	2	3.33	3	6.67	1	26.67
7	3	3.33	3	6.67	1	26.67

The SCM strategy may be summarized as follows:

SCM (Strict Center of Mass)

Incident Location: Assumed at statistical center of its district, based based on statistics describing the historical distribution of incidents.

Unit Location: Assumed at statistical center of its district, based on statistics describing the location (mobile or fixed) of the unit when not servicing incidents.

Estimated Travel Distance: From district center (for unit) to district center (for incidents).

4.4.2 Modified Center of Mass (MCM)

The modified center of mass (MCM) strategy is the same as the SCM strategy, except that the exact position of the incident is used in making travel time (distance) estimates. The MCM strategy may be summarized as follows:

MCM (Modified Center of Mass)

Incident Location: Assumed at center of incident's geographical atom.

Unit Location: Assumed at statistical center of its district, based on statistics describing the location (mobile or fixed) of the unit when not servicing incidents (same as SCM).

Estimated Travel Distance: From district center (for unit) to center of incident's geographical atom.

An MCM strategy provides a good model for a dispatcher who is intimately familiar with the various neighborhoods of "his" part of the city. Thus, even though a dispatch card may only stipulate that the reported incident is in district 1, he knows from the street address the reporting area in which the incident is located.

Continuing with the example in which an incident is located in atom 1, we obtain the following estimated travel distances using an MCM strategy:

Unit 1: Estimated travel distance = |0 - 0| + |40 - 20| = 20
Unit 2: Estimated travel distance = |30 - 0| + |40 - 20| = 50
Unit 3: Estimated travel distance = |20 - 0| + |40 - 20| = 40

Compared to SCM estimates, note how these larger travel distance estimates more nearly reflect true travel distances to atom 1 (a relatively isolated atom). The entire set of MCM dispatch preferences is shown in Table 3, in which it may be seen that it is possible to have a unit other than the one assigned to the district of the incident be the first preferred unit (atoms 4 and 6).

Table 3

MCM DISPATCH POLICY: 3-UNIT, 7-ATOM EXAMPLE

	First Preferred Unit		Second Preferred Unit		Third Preferred Unit	
Atom No.	Unit No.	Estimated Dispatch Cost	Unit No.	Estimated Dispatch Cost	Unit No.	Estimated Dispatch Cost
1	1	20.0	3	40.0	2	50.0
2	1	0.0	3	20.0	2	30.0
3	1	20.0	3	40.0	2	50.0
4	3	10.0	2	20.0	1	30.0
5	3	0.0	2	10.0	1	20.0
6	3	10.0	2	20.0	1	30.0
7	2	10.0	3	20.0	1	40.0

4.4.3 Expected Modified Center of Mass (EMCM)

This and the next preprogrammed dispatch strategy assume more sophistication on the part of the dispatcher. The EMCM strategy assumes the most sophistication, and represents the best that any dispatcher (human or computer) can do, assuming that the exact real-time positions of available units with mobile positions are not known. (The dispatcher could, of course, do better with real-time position information, such as that provided by AVL--automatic vehicle locator--systems.)

It is important to note that EMCM is the *only* preprogrammed dispatch strategy that is allowed if the user uses the 'TR' card (Sec. 4.3.2).

Simply stated, the EMCM strategy calculates the statistically correct average travel time (distance), assuming that the reporting area of the incident is known. The statistical distribution in the calculation depicts the probable locations of mobile units. In our example, the following EMCM distances are calculated for an incident from reporting area 1:

Unit 1: Estimated travel distance = 1/2(0) + 1/2(40) = 20.0
Unit 2: Estimated travel distance = 1/4(30) + 1/4(50) + 1/2(60) = 50.0
Unit 3: Estimated travel distance = 40.0

Note that these distance estimates are identical to those of the MCM strategy. Consider now, however, an incident from reporting area 2. The MCM strategy, since it assumes that the unit is located at the statistical center of its district, yields estimated travel distances of 0.0, 30.0, and 20.0 for units 1, 2, and 3, respectively. The EMCM strategy, on the other hand, gives the correct statistical weights to each of the possible locations of a mobile unit, yielding for an incident from reporting area 2 the following:

Unit 1: Estimated travel distance = 20
Unit 2: Estimated travel distance = 1/2(30) + 1/2(40) = 35.0
Unit 3: Estimated travel distance = 20

The complete set of calculations for this example is shown in Table 4. It may be seen that since units 1 and 3 are both equally preferable to assign to atom 2, there is a tie for first preference. This presents no problem for the computer algorithm which, in effect, flips a fair coin each time that both units are available to determine which one to assign. (The issue of ties is discussed at greater length in Sec. 5.2.)

Table 4

EMCM DISPATCH POLICY: 3-UNIT, 7-ATOM EXAMPLE

	First Preferred Unit		Second Preferred Unit		Third Preferred Unit	
Atom No.	Unit No.	Estimated Dispatch Cost	Unit No.	Estimated Dispatch Cost	Unit No.	Estimated Dispatch Cost
1	1	20.0	3	40.0	2	50.0
2	1,3 tie	20.0	1,3 tie	20.0	2	35.0
3	1	20.0	3	40.0	2	50.0
4	3	10.0	2	20.0	1	40.0
5	3	0.0	2	15.0	1	40.0
6	3	10.0	2	20.0	1	40.0
7	2	15.0	3	20.0	1	60.0

The EMCM strategy can be summarized as follows:

EMCM (Expected Modified Center of Mass)

Incident Location: Assumed at center of incident's geographical atom.

Unit Location: Correctly distributed statistically over atoms in its district, based on statistics describing the location (mobile or fixed) of the unit when not servicing incidents.

Estimated Travel Distance: Statistically average travel distance from each of the unit's geographical atoms (weighted by the likelihood of the unit being located in that atom) to the atom of the incident.

4.4.4. Expected Strict Center of Mass (ESCM)

This strategy is the most difficult computationally, although it does not yield distance estimates that are "as good" in general as either the MCM or EMCM strategies. It is, however, an improvement over the SCM strategy. Basically, analogous to the simple SCM strategy, the dispatcher does not consider the reporting area of the incident but only its district.* The dispatcher is too sophisticated, however, to act as if incidents and response units were located precisely at appropriate

---

*To use the ESCM strategy, each geographical atom must be included in at least one district (via the 'SS' card or the 'S' card). The same restriction is applied to SCM dispatching.

statistical centers. The dispatcher thus assigns weights to both the locations of incidents within districts and response units within districts to arrive at estimated travel distances.

As an example, consider an incident from reporting area 1. First, the dispatcher recognizes only that the incident is in district 1, with equal likelihood that it is located in reporting areas 1, 2, or 3, respectively. Second, the locations of the response units are correctly weighted by their respective likelihoods of being located in the various reporting areas. Thus, the following distance estimates are arrived at for an incident from reporting area 1:

Unit 1: Estimated travel distance = 1/2[1/3(0 + 20 + 40)]2 = 20.0

Unit 2: Estimated travel distance = 2/3[1/2(30) + 1/2(60)] + 1/3[1/2(30) + 1/2(40)] = 41.7

Unit 3: Estimated travel distance = 1/3 · 40 + 1/3 · 20 + 1/3 · 40 = 33.3

For unit 1, the factor of 1/2 is the probability that unit 1 is located in reporting area 1. The 0, 20, and 40 within the inner parentheses correspond to the three possible travel distances, reflecting an incident in reporting areas 1, 2, and 3, respectively; these are each weighted by 1/3 since each possibility is equally likely. Similar explanations apply to the other two calculations. The entire set of estimated travel distances is shown for this example in Table 5.

Table 5

ESCM DISPATCH POLICY: 3-UNIT, 7-ATOM EXAMPLE

	First Preferred Unit		Second Preferred Unit		Third Preferred Unit	
Atom No.	Unit No.	Estimated Dispatch Cost	Unit No.	Estimated Dispatch Cost	Unit No.	Estimated Dispatch Cost
1	1	20.0	3	33.3	2	41.7
2	1	20.0	3	33.3	2	41.7
3	1	20.0	3	33.3	2	41.7
4	3	13.3	2	18.3	1	46.7
5	3	0.0	2	15.0	1	40.0
6	3	13.3	2	18.3	1	46.7
7	3	13.3	2	18.3	1	46.7

The ESCM strategy may be summarized as follows:

ESCM (Expected Strict Center of Mass)

Incident Location: Assumed distributed over geographical atoms in its district, based on statistics describing the historical distribution of incidents.

Unit Location: Correctly distributed statistically over atoms in its district, based on statistics describing the location (mobile or fixed) of the unit when not servicing incidents.

Estimated Travel Distance: Statistical average travel distance from each of the unit's geographical atoms (weighted by the likelihood of the unit being located in that atom) to the atom of the incident (weighted by the likelihood of the incident being located in that atom).

The four preprogrammed dispatch strategies are summarized below.

SUMMARY OF DISPATCH STRATEGY DEFINITIONS

EMCM	Expected modified center of mass. A dispatch strategy that calculates the probablistic location of units, representing the best that any dispatcher could do without knowing the exact real-time position of available units; see Sec. 4.4.3.
ESCM	Expected strict center of mass. A dispatch strategy that calculates the estimated statistical average travel distance from each of the unit's geographical atoms (weighted by the likelihood of the unit being located in that atom) to the atom of the incident (weighted by the likelihood of the incident being located in that atom); see Sec. 4.4.4.
MCM	Modified center of mass. A dispatch strategy in which the exact location of the incident is used to make travel time (distance) estimates; see Sec. 4.4.2.
SCM	Strict center of mass. A dispatch strategy in which the dispatcher makes travel time estimates acting as if the unit were located at the statistical center of its district and the incident were at the statistical center of its district; see Sec. 4.4.1.

Input Formats for Dispatch Procedures

Selection of one of the four preprogrammed dispatch procedures is particularly easy. On card type 7, the user puts 'SCM,' 'MCM,' 'EMCM,' or 'ESCM,' depending on which of the procedures is selected. If no procedure is selected, the default is 'EMCM.'

4.5 DISPATCH SELECTION: SELECTIVE OVERRIDE OF PREPROGRAMMED PROCEDURES

The four dispatch procedures described above form the basis for analyzing a rich variety of alternative dispatch philosophies. Next are described the fairly simple input card formats that are required to generate alternative procedures.

4.5.1 District's Unit Gets First Preference

Often there are situations in which it is desired to assign the unit associated with the district to the incident, even though there may be a closer available unit in travel time. For instance, in police applications, arguments based on sector identity[*] say that a radio-dispatched patrol unit should be assigned to as many of the calls originating in its "own" sector as possible. Even in ambulance applications, situations can be imagined in which it would be preferable to assign a unit that is slightly further away than the closest unit, provided the former unit has intimate familiarity with the street and traffic patterns of the neighborhoods through which it will be traveling.

It is very simple to give the district's unit first preference. Following the dispatch procedure card (card type 7), the user simply inserts an additional card:

---

[*] Chapter 8 in Ref. 8, or Ref. 10.

'FRST'

This card sets the estimated "cost" of assigning a unit to an incident in its own district equal to *zero*, thereby guaranteeing first preference for that unit.*

As an example, consider the three-unit, seven-atom example of Sec. 4.4 using an EMCM dispatching procedure. The entries in Table 4 are now transformed, assuming the 'FRST' card is used, to the entries in Table 6. Note from this table that all cost entries in the column under the first preferred unit are zero, as dictated by the 'FRST' card. Comparing Table 6 to Table 4, it is seen that the 'FRST' card has caused several important changes in the dispatch preference orderings. First, units 1 and 3 are no longer tied for first preference for incidents from atom 2, which by definition (see the first 'SS' card in Sec. 3) is contained in district 1 (even though no available time is spent there by response unit 1). Since the "cost" of assigning unit 1 to any incident within district 1 is now zero, response unit 1 is given first preference. Second, the first two preferences for units for incidents from atoms 4 and 6 are reversed; now it is preferable to dispatch unit 2 to these incidents since atoms 4 and 6 are contained in district 2. The rank-ordered dispatch preferences for the other four atoms remain unchanged in this case, even though the cost of dispatching the district's unit in each case is dropped to zero.

4.5.2 More General Modifications of Dispatch Procedures

In many applications the user will wish to have a dispatch policy which is "basically EMCM, say, with a few exceptions." Or the policy may be "most like MCM with the district's unit getting first preference except for several units." So, it is desirable to have some way to build on the four (or eight, depending how it's counted) preprogrammed dispatch procedures and to modify them to suit the situation at hand.

---

[*] Or, at least a tie for the first preference in the unlikely event that another unit is also estimated to have zero cost of assignment.

Table 6

EMCM DISPATCH POLICY, WHERE DISTRICT'S UNIT GETS FIRST PREFERENCE: 3-UNIT, 7-ATOM EXAMPLE

Atom No.	First Preferred Unit		Second Preferred Unit		Third Preferred Unit	
	Unit No.	Estimated Dispatch Cost	Unit No.	Estimated Dispatch Cost	Unit No.	Estimated Dispatch Cost
1	1	0.0	3	40.0	2	50.0
2	1	0.0	3	20.0	2	35.0
3	1	0.0	3	40.0	2	50.0
4	2	0.0	3	10.0	1	40.0
5	3	0.0	2	15.0	1	40.0
6	2	0.0	3	10.0	1	40.0
7	2	0.0	3	20.0	1	60.0

If the user wishes to do this, he must first insert the following card after card type 7 (and after the 'FRST' card, if one is used), but before card type 8:

'DISP_OV_RD'

This card signals that the user wishes to use the selective *dispatch override* feature of the program.

Cards are inserted that explicitly indicate the details of the dispatch override immediately following the 'RUN' card (card type 10). There are three types of cards that may be used. Immediately following these cards the user must submit a 'END_OV_RD' card, signifying the end of the dispatch override.

'FRONT'

If the user wishes to assign first preference to a particular unit, he uses a 'FRONT' card. The following is an example:

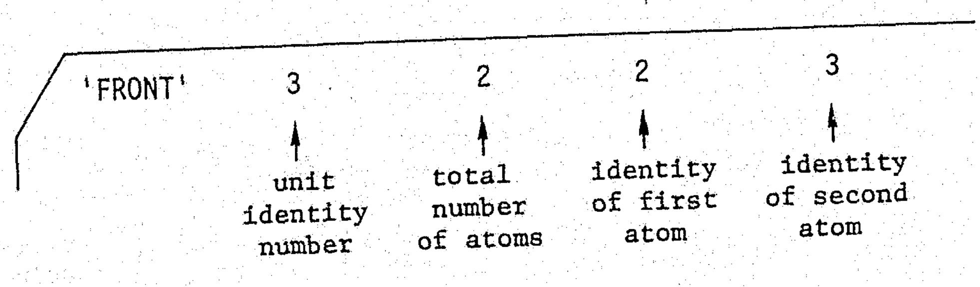

Use of this card makes unit 3 the first preference for calls from atoms 2 and 3, regardless of the previous dispatch policy. In the computer program, this is accomplished by setting the cost of dispatch of unit 3 to incidents in atom 2 or atom 3 equal to zero. Note that if the 'FRST' card is used and if atom 2 and/or atom 3 is not within district 3, then we would have a tie* for first preference, since two units would have zero dispatch cost for responding to incidents in atoms 2 and 3. The general format of this card is 'FRONT', followed by the identity number of the response unit, followed by the total number of atoms, and followed finally by the successive identity numbers of each of the affected atoms. As usual, at least one space must appear between successive entries.

If two or more 'FRONT' cards are applied with different units to the same atom, then there would be two or more ties for first preference for that atom.

An example in the police area might involve a unit with bilingual (say Spanish-speaking as well as English-speaking) officers. If the residents in one reporting area not in that unit's patrol sector are predominantly Spanish-speaking, then it may be preferable to dispatch that unit to incidents from that atom, even though travel time is increased above the minimum possible. This is accomplished by using the 'FRONT' card, where the first digit is the identity number of the car with bilingual officers, the second is equal to 1 (indicating one out-of-sector Spanish-speaking neighborhood [atom]), and the third is the identity number of the atom in question.

---

*See Sec. 5.2 for a further discussion of ties.

'BACK'

If the user wishes to assign last preference to a particular unit, he uses a 'BACK' card. The following is an example:

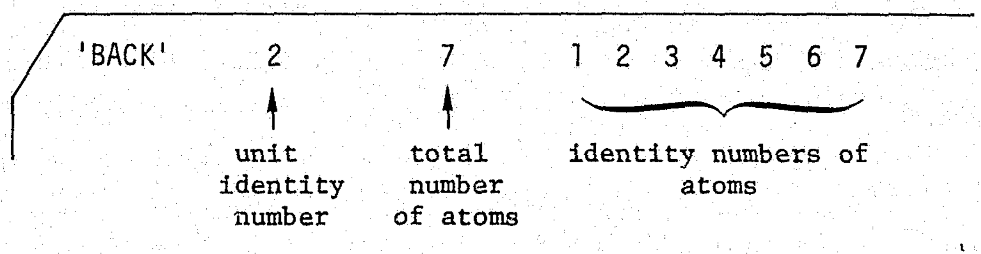

Use of this card makes unit 2 the last preference for calls from atoms 1 through 7, regardless of the previous dispatch policy. In the computer program, this is accomplished by setting the cost of dispatch of unit 2 to incidents in these atoms equal to a very large number (999). The general format of this card is 'BACK,' followed by the identity number of the response unit, followed by the total number of atoms, and followed finally by the successive identity numbers of each of the affected atoms. (Of course, at least one space must separate successive entries.)

If two or more 'BACK' cards are applied with different units to the same atom, then there would be two or more ties for last preference for that atom.

Consider another example from the police area. Often a sergeant's car will be assigned to patrol all or part of the entire region being modeled, and this car is to be used only as a last resort for dispatch purposes. This policy can be accomplished in the model by using the 'BACK' card, followed by the identity number of the sergeant's car, followed by the total number of atoms in the region being modeled, and followed finally by the successive list of atom numbers. For instance, the above card containing seven atoms would accomplish this purpose for the example in Sec. 3, where car 2 is considered a sergeant's car. Other cities use "umbrella" cars and/or "backup" cars that are to be dispatched only as a last resort; again, use of the 'BACK' card is called for in these cases.

'MIDDLE'

If a user wishes to modify dispatch preferences in a more general way than is represented by the 'FRONT' and 'BACK' cards, he can use the 'MIDDLE' card, an example of which is the following:

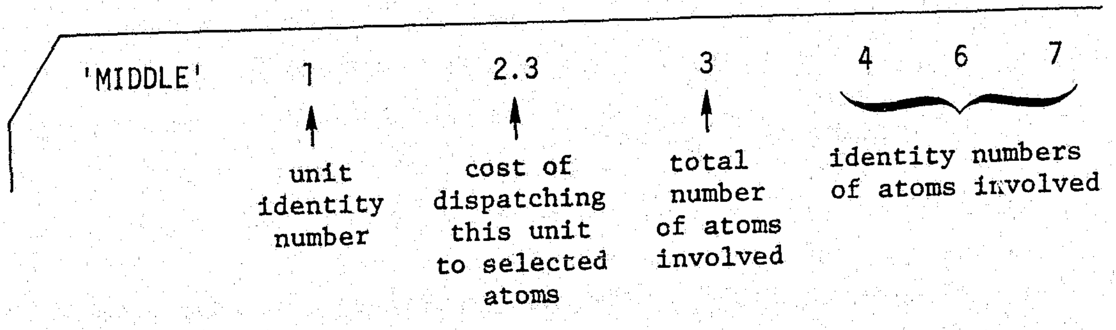

Use of this card places a cost of 2.3 to assigning unit 1 to calls from atoms 4, 6, and 7. Note that the nature of the data contained on the 'MIDDLE' card is different from that of the 'FRONT' and 'BACK' cards. This is because the user must specify one additional piece of information, namely the cost of dispatching the unit to each of a set of comparable atoms. In fact, the user could accomplish the equivalent of the 'FRONT' and 'BACK' cards by using only the 'MIDDLE' card. This is not recommended, however, except for users who feel very comfortable with the concept of dispatch cost and how it is used to arrive at dispatch preferences.

If the user has a large problem (i.e., many units and reporting areas), he may not be able to use the 'MIDDLE' card effectively until he has submitted one run without it. As a result of the first run, he could determine the value of the dispatch cost for the 'MIDDLE' card required to place the unit anywhere in the rank-ordered list of dispatch preferences.

The three types of dispatch override cards have now been discussed. In any particular run, any combination of these cards may be used (and they are processed by the computer program in the order inserted in the deck). All such cards *must* follow immediately after the 'RUN' card, and an 'END_OV_RD' card must follow the last of these cards.

4.6 POLICE PREVENTIVE PATROL

In the cases for which a police patrol force is being modeled, the user may wish to calculate the frequency of passings of units on preventive patrol in each of the geographical atoms. Thus, for instance, he could compute that, on the average, a patrolling unit passes by a randomly selected point in atom 4, 2.1 times every hour and in atom 12, 0.5 times every hour. To perform the computations, the computer utilizes a well-known formula for the average frequency of patrol passings of a randomly selected point* in atom J:

$$PF(J) = \sum_{I=1}^{M} \frac{(SP) f_{IJ}}{L(J)} (1 - WORKLOAD(I))$$

where PF(J) = average frequency of patrol passings of a randomly selected point in atom J.

SP = effective speed of a unit performing preventive patrol (in mph or other standard unit).

L(J) = number of patrollable street miles in atom J.

f_{IJ} = the relative amount of time that unit I spends in atom J while on patrol. (These factors are obtained directly from the 'SS' or 'S' cards.)

WORKLOAD(I) = fraction of time that the unit is busy answering calls for service (this quantity is computed from the model).

Note that the formula allows for overlapping police patrol sectors (corresponding to more than one f_{IJ} nonzero for a given value of J), as well as nonoverlapping sectors.

To use this option, the user inserts the following card immediately following the response speed card (card type 5):

'PATROL' 6.5

---

*Reference 8, Chap. 4.

The 'PATROL' card indicates that the user wants patrol frequencies calculated. The constant following 'PATROL' (6.5 in this case) is the effective speed (in mph) of the patrolling unit.

Immediately following this card, the user provides as input the patrollable street miles of each of the atoms (sequentially, from the first to the last atom). An example for a seven-atom case may read as follows:*

```
0.5     2.5     6.0     0.0     0.6     5.0     48.3
```

An extra output column is printed for each atom, the entry in the column indicating the average patrol frequency for that atom. An example for a seven-unit, seven-atom run is shown in Fig. 12.

4.7 DEFAULT FOR LOCATIONS OF MOBILE UNITS

In Sec. 2 it was shown that mobile or fixed position units could be positioned arbitrarily by using the 'SS' card as card type 4. However, there may be circumstances utilizing mobile units in which the user does not know the relative amounts of time spent by units in each of "its" atoms and therefore may be willing to settle for a reasonable default procedure. This default for locations of mobile units assumes that the fraction of available time that a unit spends in a particular atom contained in that unit's district is *proportional* to the workload (in incidents) generated from that atom. The default is invoked by replacing the 'SS' card with the 'S' card. For the example of Sec. 2, the new data appear as follows:

```
'S'     1     3     1     2     3
'S'     2     3     4     6     7
```

---

*An atom with zero patrollable miles is defined to have zero patrol frequency.

```
PERFORMANCE MEASURES THAT ARE SPECIFIC TO EACH ATOM
ID #      WORKLOAD       AVE    FRACTION OF CALLS FOR SERVICE
ATOM         OF          TRAV   FROM ATOM
            ATOM         TIME   SERVICED BY UNIT NUMBER:
         (#CALLS/100HR)           1    2    3
   1         72.00      3.774 0.42 0.27 0.31
   2         72.00      2.546 0.42 0.27 0.31
   3         72.00      3.774 0.42 0.27 0.31
   4         72.00      2.472 0.27 0.41 0.31
   5        144.00      1.712 0.27 0.32 0.41
   6         72.00      2.472 0.27 0.41 0.31
   7         72.00      3.333 0.27 0.41 0.31
```

Fig. 12—Illustrative output of preventive patrol frequencies

'S'	3	1	5
card type identifier	response unit number	total number of geographical atoms in that unit's district	identity of atoms in each unit's district.

Note that the format is similar to that of the 'SS' cards, except that there is no information regarding the relative amount of available time spent in each of the atoms. Since this default now assumes that the likelihood of a unit being in a particular atom (while available) is proportional to the incident workload from the atom, the above three 'S' cards are equivalent to the following three 'SS' cards:

'SS'	1	3	1	0.333	2	0.333	3	0.333
'SS'	2	3	4	0.333	6	0.333	7	0.333
'SS'	3	1	5	1.000				

The uniform distribution of positions over atoms in districts 1 and 2 is due to the fact that incidents are distributed uniformly over the atoms in each district (see card type 3--Workload Distribution--in Sec. 3).

If the user desires, he can intermix 'S' and 'SS' cards. With this option, it is particularly attractive to use the 'S' card for region-wide roaming units--such as sergeant's cars--for which it may may make sense to approximate the likelihood of its being located in a particular atom to be proportional to that atom's incident workload.

4.8 NONSEQUENTIAL NUMBERING OF ATOMS

Throughout this manual it is assumed that the geographical atoms are numbered consecutively from 1 to the highest number, R. Cities may number their atoms differently, or in one command (in which the model may be applied), numbering may start at some arbitrary point (say 302).

To correct at least partially for this we allow the user to input the following optional card (between card type 3--Workload Distribution--and card type 4--Spatial Allocation of Response Units):

'ATOM_NO'	7	10	3	101	999	27	302

Here 'ATOM_NO' specifies that the following entries (there must be R of them) indicate the number assigned by the city (or by the emergency service agency) to the atoms that are numbered (when inputting data) consecutively from 1 to R. Thus, for the seven-atom example of Sec. 3, the above 'ATOM_NO' card says that the first atom considered by the computer is actually atom 7 as numbered by the city, the second is actually atom 10, and so forth. (These numbers cannot exceed 999.)

If this option is used, then the city's atom numbers will be used on printouts of data and performance measures. This greatly increases the ability of the user to relate the output figures to his own city's situation.

A word of caution about inputting data: It is assumed in examples and discussions in this manual that the atoms are numbered consecutively from 1 to R when inputting data. To avoid confusion, the user is advised to prepare a simple list showing side-by-side the computer's internal (input data) number for the atom and the city's actual number.

4.9 DIRECTION-SPECIFIC TRAVEL SPEEDS

Many communities are characterized by different obtainable travel speeds in each of the two directions of travel (say east-west and north-south). For instance, in downtown Manhattan during working hours, the effective average uptown-downtown speed may be 15 or 20 mph, whereas the crosstown speed may be less than 5 mph.

To model such a situation, the user can input different speeds for the X direction (XSPEED) and the Y direction (YSPEED). This is done by replacing card type 5 (the 'SPEED' card) by two cards, as follows:

```
'XSPEED'          5.0
```

```
'YSPEED'          15.0
```

In this case, reflecting perhaps the situation in Manhattan, the variable XSPEED is set to 5 mph (crosstown) and YSPEED is set to 15 mph (uptown-downtown).

4.10 'RERUN'

On occasion the user may wish to run the program under one set of assumptions regarding dispatch policy, geography, workloads, and so forth and then immediately rerun the program with a minimal number of relatively minor modifications. Sometimes this can be done most efficiently by putting all the data changes on data cards following the 'RUN' card (which was used for the first set of runs) and then by using the 'RERUN' card instead of the 'RUN' card to initiate the second set of runs.

Referring to the example discussed in Sec. 3, suppose the user places the following cards immediately after the 'RUN' card:

```
'TITLE'              'SAMPLE RUN - 2ND PASS'
'SPEED'              10.0
'SERVTM'             25.
'RERUN'              2.88      2.0
```

In this case, the computer will have saved all the information previously obtained from the "missing" data cards, that is, card types 1, 3, 4, 6, and 7. In essence, then, the new cards tell the computer to rerun the previous runs but with *zero queue capacity*. This change is evident since there is no 'CAP' card in the above set of new cards.

In using the 'RERUN' option, the following conventions must be observed:

1. There must be no changes in card types 1, 3, 4, 6, or 7 or in any optional cards that relate to data contained on those cards.
2. All variables that have default values (e.g., travel speeds, service times, print options, etc.) will be reset to their default values after the earlier run. They must be reset (or otherwise changed) to their proper values by the correct data cards before the 'RERUN' card.

There are two major purposes for the 'RERUN' card: (1) it saves the user time since he does not have to retype essentially repetitive (and sometimes lengthy) data; (2) it saves computer execution time by reducing the number of computer operations that have to be performed during the second run.

4.11 SERVICE TIMES THAT VARY BY RESPONSE UNIT

In some applications each response unit may have a unique average service time. This may be due to characteristics of the personnel assigned to the unit, to special capabilities of the unit, or to other factors. If the differences in average service times are known and are sufficiently large so that they cannot be ignored in a planning model, then the user should input these average service times for inclusion in the model. At the time of this writing, this capability is available for the exact Hypercube Model only. (It is hoped that in the near future the approximate model will also be able to handle response unit-specific service times.)

To use this capability, the user replaces card type 9 (the 'SERVTM' card) with the following card:

```
'VAR_SER_TM'     30.      17.5      36.2
```

Following the card identifier ('VAR_SER_TM'), the user types (in decimal) the average service times of each of the units in order. The above numbers for the three-unit (seven-atom) example indicate that the average service time of unit 1 is 30.0 min, that of unit 2 is 17.5 min, and that of unit 3 is 36.2 min.

5. COMPLICATIONS

There are several items that may arise in use of the program that are grouped here in the category of "complications." Each of them may arise in any particular application, and do not in fact represent complications at all, provided that the definitions and conventions used here are understood. The items are:

1. Overlapping Districts
2. Ties for Dispatch Preference
3. Use of Center of Mass Dispatching Strategies with Arbritary Inter-Atom Travel Times.

If any of these items is relevant to a particular application, the user should read the appropriate section before running the program.

5.1 OVERLAPPING DISTRICTS

If overlapping districts are used (relevant only for mobile units), the following issues must be resolved:

- o With center of mass dispatching strategies, to which district does an incident "belong?"
- o How are cross-district dispatches to be counted?
- o If the 'S' card is used to provide a default for mobile locations, in what ways (if any) do overlapping districts affect the default?
- o If preventive patrol frequencies are to be computed, do overlapping districts affect the computation in any way? Each of these questions is addressed in the following paragraphs.

District Identity of an Incident

Very simply stated, the convention used in the program is that *an incident "belongs" to the lowest numbered district containing the* *geographical atom of the incident*.[*] Thus, if the incident is located in atom 13, which is included in four (different) districts, 3, 5, 11, and 12, then the incident "belongs" to district 3. As an example, if an SCM dispatching strategy were used with the 'S' card option for positioning units, unit 3 would get first preference for incidents from atom 13 (since the center of mass locations of incidents in district 3, including those in atom 13, would coincide with the center of mass location of the unit in district 3, yielding an estimated travel cost of 0.0).

This convention is quite reasonable to follow and to interpret in practice. It suggests, for instance, that region-wide roving units should be given the larger identity numbers, and the units responsible for smaller territories should be given the smaller identity numbers.

Counting Cross-District Dispatches

There are two possible conventions here. First, a dispatch could be counted as crossing district boundaries only if the unit responding is traveling to an atom not contained in "its" district. Second, following the ideas above for district identity of an incident, a dispatch could be counted as crossing district boundaries if the unit is traveling to an atom whose incidents do not "belong" to the unit's district. Clearly, the second convention will count more dispatches as cross-district dispatches, and that is the one which is chosen here for the computer program. (Based on users' comments, we may make either of the conventions possible by insertion of a control card in the input deck.)

The 'S' Card

Briefly stated, use of the 'S' card does not cause difficulties when overlapping response areas are employed. The program treats each 'S' card individually and computes the mobile location likelihood of

---

[*] A district "contains a geographical atom" if that atom is entered on the corresponding unit's 'SS' or 'S' card. Thus, even an atom which receives zero attention (i.e., time spent there) while the unit is available is contained in that unit's district if it is listed (with an accompanying 0.0) on the unit's 'SS' card.

each unit over atoms in its district, regardless of the possible existence of districts which may overlap the unit's district.

Preventive Patrol Frequencies

Again, there are no problems encountered with computing and interpreting preventive patrol frequencies in atoms belonging to more than one district. This is basically because the total frequency of patrol passings (in passes per hour) is additive over the units that are patrolling the atom (see Sec. 4.6).

5.2 TIES FOR DISPATCH PREFERENCE

On occasion there may be equal travel costs associated with assigning two or more units to incidents in a particular geographical atom. This situation is called a "tie for dispatch preference" and is particularly likely to occur if the user employs the 'FRONT' or 'BACK' dispatch override option with two or more units (see Sec. 4.5) or the 'FRONT' option in conjunction with the 'FRST' option.

Ties present no problem for the exact Hypercube Model (which is invoked by setting ESTSTAT = 0, as discussed in Sec. 3). Each time there is a tie between two units, the program effectively flips a fair coin to choose the unit to be assigned. Since the model is not a simulation model, but rather an analytical model, the fine-grained process of flipping a coin for each incident cannot be actually duplicated because incidents are not treated individually. Rather, situations involving ties result in the workload from the geographical atom in question being split 50-50 between the two units (given an availability pattern for units in which both units are available and all other available units require higher travel costs). Similar procedures are employed if three or more units tie for dispatch preference.

In the approximation procedure, ties are not treated in the same precise way as in the exact hypercube model procedure. This is not likely to cause problems in practice since the error caused by imprecise consideration of ties is not likely to exceed the level of error introduced by the approximation procedure itself. The only situation in which the treatment of ties may give intuitively unsatisfactory results is one in which two (or more) units have identical dispatch preferences associated with all atoms in the region. Then, for instance, we would expect the performance measure computed for each to be identical, whereas the heuristic procedure for handling ties may yield slightly different values for these performance measures.

5.3 CENTER OF MASS DISPATCHING STRATEGIES WITH ARBITRARY INTER-ATOM TRAVEL TIMES

When the computer calculates the center of mass position of a unit or an incident, it is unlikely that the resulting coordinates will coincide with the coordinates of any particular atom. This causes no problems if Manhattan distances are assumed throughout, since the Manhattan distance to a statistical center is a well-defined quantity. However, complications arise when other than Manhattan distances are employed (indicating use of the 'TX_OV' card, as discussed in Sec. 4.3.2). In these cases it is unclear how to compute the expected travel distance to a statistical center since perhaps one or more barriers or other complications to travel (as reflected in the values inserted in the inter-atom travel time matrix) may be encountered. It certainly does not make sense to employ blindly the Manhattan distance metric to the statistical center, which in effect would ignore perturbations read in by the user.

As a not unreasonable solution to this dilemma, if either 'TX_OV' or 'TR' is used in conjunction with either the SCM or the MCM preprogrammed dispatching strategy, then the statistical center of units and/or incidents is shifted to the nearest* geographical atom center; and all travel times are computed from (to) this atom center, using the exact values in the matrix (TR) of inter-atom travel times. In effect, this admittedly approximate solution to the problem states that the center of mass location is now the center of a geographical atom, and that the dispatcher utilizing either an SCM or MCM policy will act as if the unit (incident) is located at that point; the dispatcher will still be quite "intelligent," since he will use empirically measured (or otherwise verified) travel times from (to) that point.

---

*Here "nearness" is measured by right-angle travel time (using XSPEED and YSPEED).

6. TECHNICAL SUMMARY

For the convenience of the user, this final section summarizes in concise form the operating procedures for the program and discusses several other technical points.

6.1 INDEX AND ORDERING OF INSTRUCTIONS

Table 7 contains an alphabetized list of input instructions, summarizing the following for each:

- The function of the instruction, concisely stated (column 2).
- Whether or not the instruction is optional (column 3).
- The default that is associated with an optional instruction (column 4)
- The section(s) in which the instruction is described.

Table 8 lists the instructions (mandatory and optional) in the order in which they must appear in the input deck. The three columns contain the following information:

Column 1: The card type, including the card type number (if aplicable) and the name of the card in single quotes.

Column 2: The card or cards that are replaced by the card in question.

Column 3: Special conditions that are associated with the card in question.

It is believed that Tables 7 and 8 will provide a convenient and compact reference source for frequent users of the program.

6.2 SUMMARY OF UNITS OF MEASUREMENT

When using the program it has been our experience that a common source of data error results from using units of measurement inconsistent with those of the program. This can be avoided simply by checking the review definitions in Table 9. These are the most common:

CONTINUED

1 OF 2

6. TECHNICAL SUMMARY

For the convenience of the user, this final section summarizes in concise form the operating procedures for the program and discusses several other technical points.

6.1 INDEX AND ORDERING OF INSTRUCTIONS

Table 7 contains an alphabetized list of input instructions, summarizing the following for each:

- o The function of the instruction, concisely stated (column 2).
- o Whether or not the instruction is optional (column 3).
- o The default that is associated with an optional instruction (column 4)
- o The section(s) in which the instruction is described.

Table 8 lists the instructions (mandatory and optional) in the order in which they must appear in the input deck. The three columns contain the following information:

Column 1: The card type, including the card type number (if aplicable) and the name of the card in single quotes.

Column 2: The card or cards that are replaced by the card in question.

Column 3: Special conditions that are associated with the card in question.

It is believed that Tables 7 and 8 will provide a convenient and compact reference source for frequent users of the program.

6.2 SUMMARY OF UNITS OF MEASUREMENT

When using the program it has been our experience that a common source of data error results from using units of measurement inconsistent with those of the program. This can be avoided simply by checking the review definitions in Table 9. These are the most common:

Table 7

INDEX OF INSTRUCTIONS

Instruction	Function	Optional	Relevant Default	Described in Section
'ATOM_NO'	Assigns nonsequential numbers (no larger than 999) to atoms	Yes	Atoms numbered sequentially starting at 1	4.8
'BACK'	Places unit as last preference for dispatch	Yes	No override	4.5.2
'CAP'	Specifies unlimited queue capacity	Yes	No queue capacity, backup system assumed	3.1
'CORTM'	Correction factor for intra-atom travel times	Yes	Intra-atom travel times all zero (unless empirical travel times are used)	4.3.1 (see also Sec. 4.3.2)
'DISP_OV_RD'	Allows for up three different types of override of the preprogrammed dispatch procedures	Yes	No override	4.5.2
'EMCM'	One of four preprogrammed dispatch strategies	Yes	'EMCM'	4.4
'END_OV_RD'	Terminates dispatch overrides	Yes	No override	4.5.2
'ESCM'	One of four preprogrammed dispatch strategies	Yes	'EMCM'	4.4
'FRONT'	Places unit as first preference for dispatch	Yes	No override	4.5.2
'FRST'	Indicates that dispatch unit gets first preference	Yes	District's unit does not automatically get first preference	4.5.1
'GLOSSARY'	Inputs city-specific names for response units, geographical regions, etc.	Yes	See Table 1	4.1
'LAM'	Specifies distribution of calls for service over atoms	No	(a)	3.1
'MCM'	One of four programmed dispatch strategies	Yes	'EMCM'	4.4
'MIDDLE'	Places unit at specified position in dispatch preference list	Yes	No override	4.5.2
'NO_PRNT_AT'	Suppresses printing of atom-specific performance measures	Yes	Performance measures are printed	4.2
'PATROL'	Specifies that frequency of preventive patrol is to be computed	Yes	No preventive patrol	4.6
'PRNT_TR'	Specifies that inter-atom travel time matrix is to be printed	Yes	Matrix not printed	4.2

Table 7--(continued)

Instruction	Function	Optional	Relevant Default	Described in Section
'RERUN'	Reruns the program with minimal (nongeographical) changes to the input data	Yes	Use 'RUN' card and input entire set of data cards	4.10
'RUN'	Provides call rates for current runs	No (unless rerun is used)	(a)	3.1 (see also 'RERUN' in Sec. 4.10)
'S'	An 'SS' card substitute which assumes a simple default for locations of units	Yes	Use 'SS' card	4.7 (see also Sec. 2.1)
'SCM'	One of four preprogrammed dispatch strategies	Yes	'ESCM'	4.4
'SERVTM'	Specifies mean service time (in minutes)	Yes	30 minutes	3.1 (see also Sec. 4.11) for server dependent service times
'SPEED'	Specify speed of responding units	Yes	Travel speed = 10 mph	3.1 (see also 'XSPEED' and 'YSPEED' in Sec. 4.4)
'SS'	Specifies spatial distribution of response units	No (unless 'S' card used)	(a)	3.1
'TITLE'	Specifies title of run(s)	No	(a)	3.1
'TR'	Specifies that *all* travel times are empirically measured	Yes	Right-angle travel times (with possible selective override) are assumed	4.3.2 (see also Sec. 4.3.1)
'TX'	Specifies coordinates of centers of atoms	No (unless empirical travel times are used)	(a)	3.1 (see also 'TR' in Sec. 4.3.2)
'TX_OV'	Allows selective override of right-angle travel times	Yes	No override	4.3.2 (see also Sec. 4.3.1)
'VAR_SER_TM'	Reads in average service times that vary by response unit	Yes	Use 'SERVTM' card or default value of 30 minutes	4.11
'XSPEED'	Inputs speed of response in X direction	Yes	Use 'SPEED' card	4.9
'YSPEED'	Inputs speed of response in Y direction	Yes	Use 'SPEED' card	4.9

[a] Instruction not optional.

Table 8

ORDERING OF INSTRUCTIONS

Card No. and Type	Replaces Card Type	Special Conditions
1. Basic program specifications		Semicolon after last data entry; DEBUG = 1 is optional last entry (see Sec. 6.3).
'GLOSSARY'		Optional; free-form input; semicolon at end.
2. Title card, 'TITLE'		Title cannot exceed 50 characters.
'PRNT_TR'		Optional.
'NO_PRNT_AT'		Optional.
3. Workload distribution, 'LAM'		Workloads inserted sequentially by atom in decimal; workloads do not have to be normalized.
'ATOM_NO'		Optional.
4. Spatial allocation of response units, 'SS'	'SS' is standard; can use 'S' instead.	First data entry is unit number (integer), second is total number of atoms in response area (integer), and then (in pairs) the atom numbers in district with the relative amount of available time spent in each.
'S'	'SS'	Can be mixed with 'SS' cards, if desired.
5. Response speed, 'SPEED'	'SPEED' is standard if 'TX' is used. Can be replaced with 'XSPEED' and 'YSPEED'.	Data entry is speed in mph (decimal).
'XSPEED'	'SPEED'	Optional; if 'XSPEED' is used, 'YSPEED' must also be used.
'YSPEED'	'SPEED'	Optional; if 'YSPEED' is used, 'XSPEED' must also be used.
'PATROL'		Optional.
6. Locations of atoms, 'TX'	'TX' not required if 'TR' used.	Data entries are coordinate pairs (decimal) of centers of successive atoms.
'CORTM'	Overrides assumptions of 'TX' card.	Optional; data entry is a constant of proportionality in square-root response-time law.
'TX_OV'	Overrides assumptions of 'TX' card and 'CORTM' card (if 'CORTM' is used).	Optional; free-form data entries (in minutes, decimal); semicolon at end.
'TR'	'TX', 'CORTM', 'TX_OV'.	Optional; requires R^2 data entries. Must use 'EMCM' card as dispatch procedure.
7. Dispatch procedure, 'SCM'	Any of other 3 dispatch procedure cards.	Requires each atom to be in at least one district. Use of 'SCM' with 'TX_OV' causes statistical centers to be "moved" to closest atom center.
'MCM'	Any of other 3 dispatch procedure cards.	Use of 'MCM' with 'TX_OV' causes statistical centers of response units to be "moved" to closest atom center.
'ESCM'	Any of other 3 dispatch procedure cards.	Requires each atom to be in at least one district.
'EMCM'	Any of other 3 dispatch procedure cards.	Default; must be used if 'TR' is used.
'FRST'		Optional.
'DISP_OV_RD'		Optional.
8. Queue capacity, 'CAP'		Optional.
9. Average service time, 'SERVTM'	Can use 'VAR_SER_TM' instead.	Data entry is average service time per call (in minutes, decimal).
'VAR_SER_TM'	'SERVTM'	Optional; can use only with exact Hypercube Model.
10. Run card, 'RUN'	Can use 'RERUN' under special circumstances.	First data entry is smallest region-wide rate of calls for service (per hour) to be considered; second data entry is increment to be added on each successive run.
'RERUN'	'RUN'	Can use only on second or higher order-runs and only under certain conditions (see Sec. 4.10).
'FRONT'		Optional.
'BACK'		Optional.
'MIDDLE'		Optional.
'END_OV_RD'		Last card in sequence of 'FRONT', 'BACK', and 'MIDDLE' cards.

Table 9

UNITS OF MEASUREMENT

Card Label	Definition of Relevant Variable	Variable Name	Units of Measurement	Default Value
'CORTM'	Areas of the geographical atoms	A_i	Sq mi	None
'CORTM'	Constant of proportionality for square-root law	C	None	None
'LAM'	Distribution of calls for service over atoms		Arbitrary (computer automatically normalizes)	None
'MIDDLE'	Cost of dispatching unit I to atom J	C(I,J)	Usually min or mi	None
'PATROL'	Speed of patrolling unit	SP	Mph	None
'RERUN'	See 'RUN' (same format)			
'RUN'	Average number of calls for service per hour from the region being modeled		Calls per hour	None
'RUN'	Increment in average number of calls for service per hour from the region being modeled		Calls per hour	None
'SERVTM'	Mean service time	SERVTM	Min	30 min
'SPEED'	Speed of responding unit	SPEED	Mph	10 mph
'TR'	Inter-atom travel times	TR(I,J) = travel time from atom I to atom J	Min	None
'TX'	Coordinates of centers of atoms	(X(J),Y(J)) = center coordinates of atom J	100-ft units	None
'TX_OV'	Inter-atom travel times (input via override)	TR(I,J) = travel time from atom I to atom J	Min	None (except right-angle metric, if 'TX_OV' not used)
'VAR_SER_TM'	Average service times of each of the units	SERVTM(I) = average service time of unit I	Min	30.0 min (assuming absence of 'VAR_SER_TM' card)
'XSPEED'	Speed of response in X direction	XSPEED	Mph	10 mph
'YSPEED'	Speed of response in Y direction	YSPEED	Mph	10 mph

- Coordinates of centers of atoms are specified in 100-ft units (on an X-Y grid)
- Service times are in minutes
- Areas of atoms are in square miles
- Inter-atom travel times are in minutes
- Speed of patrol is in miles per hour
- Response speed is in miles per hour
- The call rate from the region is in calls per hour.

6.3 DEBUG TIMER

The user, when running the program, may wish to trace its progress through the various stages of execution. If this is the case, the user types DEBUG = 1 *before* the semicolon on card type 1. Thus, using the debug timing option, the first data card for the example described in Sec. 3 would read as follows:

```
M = 3     R = 7     NUM = 2     ESTSTAT = 1     DEBUG = 1;
```

Given this option, the moment after the first data card is read in (during execution) the following is printed:

START

CURRENT TIME = x_1x_2 HR y_1y_2 MIN z_1z_2 SEC $w_1w_2w_3$ MILLISEC.

This tells the user that the execution has started at the time (on a twenty-four hour clock) indicated. For instance, the time could read

CURRENT TIME = 13 HR 34 MIN 36 SEC 454 MILLISEC.

At prespecified points throughout the program, additional time and location indicators are given. These points include completion of data read-in, initiation of iterations for solving equations, completion of iterations, start of performance measures printout, and completion of run. Detailed understanding of each of these location signals requires intimate knowledge of the program structure--which is beyond the scope of this user's manual.

At the completion of the run, the following is printed (assuming the debug timer is used):

TIME AT COMPLETION OF RUN
CURRENT TIME = $x'_1 x'_2$ HR $y'_1 y'_2$ MIN $z'_1 z'_2$ SEC $w'_1 w'_2 w'_3$ MILLISEC.

Thus, the user who is not familiar with the detailed workings of the program can still determine the exact amount of computer time (to the nearest millisecond) required to execute his run.

6.4 COSTS AND CORE STORAGE REQUIREMENTS

In running the model, the execution time and core storage requirement (both of which determine cost) depend largely on the number of response units and the number of geographical atoms in the model. Regarding core storage requirements, virtually all arrays in the computer program have variable dimensions, their values depending on M (the number of response units), R (the number of atoms), and NUM (the number of runs in a set of runs). Most runs of the model (either approximate or exact) require 300K bytes or less of core storage. (Soon we hope to have a version of the approximate model that requires less than 150K for most applications.) For large problems (10 or more response units, 50 or more atoms), the exact model may require up to 500K bytes of core storage.

Once the program is compiled, the cost per run for runs having less than 10 response units and less than 50 atoms has usually been less than $5 on MIT's IBM Model 370/168 computer. For runs having more than 10 response units, the user can save considerably on costs by using the approximate model rather than the exact model. As a rough rule of thumb, the marginal cost per run of the exact model (for runs with more than 10 response units) doubles for each additional unit included in the model. For the approximate model, the cost per run increases only slightly faster than linearly with each additional unit included in the model.

Since each set of runs (executed from one set of data cards) involves considerable fixed set-up costs (e.g., data read-in, storage allocation, variable initialization, etc.), the user is advised to run all workload levels of interest (specified by 'NUM' on card type 1 and by the two data entries on card type 10, and 'RUN' card) in one set of runs. In this way, the cost per individual run is considerably reduced.

Regarding the number of atoms in the model, the cost per run grows approximately linearly with the number of atoms; this applies to both the approximate and the exact model. Due to the fixed set-up costs, however, this does not mean that a 100-atom run will cost precisely twice that of a comparable 50-atom run; in almost all instances, it will cost less than twice the original amount.

Thus, in "typical" sets of runs employing the approximate model, the cost per run is usually less than $5 and rarely exceeds $10. The exact model is usually within these limits for small and moderate-sized problems, but could cost as much as $50 or $100 per run for runs with nearly 15 response units and more than 100 geographical atoms. All of these costs--even at the high extreme--compare remarkably favorably with those of a simulation model, which is the only other type of model currently known to the author that computes the same or similar performance measures.

6.5 SAMPLE DATA DECKS

Figures 13 through 17 contain sets of data cards that might be used in various applications. These particular examples have been selected to illustrate various capabilities of the computer program; in actual applications, it is likely that a larger number of reporting areas and/or response units in the region being modeled will result in the data decks containing more cards than those illustrated here.

```
M=3   R=7   NUM=2   ESTSTAT=1     ;
'TITLE'  'SAMPLE 3-CAR RUN, 7 ATOMS, PRNT TR'
'PRNT_TR'
'LAM'  1000. 1000. 1000. 1000. 2000. 1000. 1000.
'SS' 1 3 1 1.0 2 0.0 3 1.0
'SS' 2 3 4 1.0 6 1.0 7 2.0
'SS' 3 1 5 1.0
'SPEED'  10.0
'PATROL'  6.5
.5 2.5 6. 0. .6 5. 48.3
'TX' 0.  20.  .0  40.  .0  60.  20.  30.  20.  40.  20.  50.  40.  40.
'SCM'
'CAP'
'SERVTM' 25.
'RUN' 2.88 2.88
```

Fig. 13—Illustrative data deck: 7 atoms, 3 cars, print TR matrix

---

```
M=5   R=7   NUM=2   ESTSTAT=2     ;
'TITLE'   'SAMPLE 5-CAR RUN, 7 ATOMS, 2 REGION-WIDE CARS'
'PRNT_TR'
'LAM'  1000. 1000. 1000. 1000. 2000. 1000. 1000.
'SS' 1 3 1 1.0 2 0.0 3 1.0
'SS' 2 3 4 1.0 6 1.0 7 2.0
'SS' 3 1 5 1.0
'S' 4 7 1 2 3 4 5 6 7
'S' 5 7 1 2 3 4 5 6 7
'SPEED'  13.72
'PATROL'  7.1
.5 2.5 6. 0. .6 5. 48.3
'TX' 0.  20.  .0  40.  .0  60.  20.  30.  20.  40.  20.  50.  40.  40.
'EMCM'
'CAP'
'SERVTM' 35.76
'RUN' 3. 1.
```

Fig. 14—Illustrative data deck: 7 atoms, 5 cars, (2 region-wide)

```
M=5   R=7   NUM=2   ESTSTAT=2    ;
'TITLE'  '5 CARS,2 REGION-WIDE,ATOM 1 ISOLATED'
'PRNT_TR'
'LAM'  1000. 1000. 1000. 1000. 2000. 1000. 1000.
'SS' 1 3 1 1.0 2 0.0 3 1.0
'SS' 2 3 4 1.0 6 1.0 7 2.0
'SS' 3 1 5 1.0
'S' 4 7 1 2 3 4 5 6 7
'S' 5 7 1 2 3 4 5 6 7
'SPEED'  13.72
'PATROL'  7.1
.5 2.5 6. 0. .6 5. 48.3
'TX' 0.  20.  .0  40.  .0  60.  20.  30.  20.  40.  20.  50.  40.  40.
'TX_OV'
TR(1,2)=12.27 TR(1,3)=14.55 TR(1,5)=14.55 TR(1,6)=15.68 TR(1,7)=16.82
TR(2,1)=12.27 TR(3,1)=14.55 TR(5,1)=14.55 TR(6,1)=15.58 TR(7,1)=16.82;
'EMCM'
'DISP_OV_RD'
'CAP'
'SERVTM' 35.76
'RUN' 3. 1.
'BACK' 5 7 1 2 3 4 5 6 7
'MIDDLE' 4 998. 7 1 2 3 4 5 6 7
'END_OV_RD'
```

Fig. 15—Illustrative data deck: 7 atoms, 5 cars, one atom isolated

---

```
M=4   R=7   NUM=2   ESTSTAT=2     ;
'GLOSSARY'
NM_UNIT(1)='AMBULNCE' NM_UNIT(2)='AMBULNCE' NM_UNIT(3)='AMBULNCE'
NM_UNIT(4)='ROVER'
R_DIST='AMB_ZONE'   R_UNIT='AMBULANCE'
CFS='CALLS FOR AMBULNCE'   ATOM='REP AREA'
NM_DIST(1)='WEST ZN'   NO_DIST(1)=101
NM_DIST(2)='EAST ZN'   NO_DIST(2)=102
NM_DIST(3)='CENTRAL'   NO_DIST(3)=91
NM_DIST(4)='CITYWIDE'  NO_DIST(4)=500    ;
'TITLE'  'AMBULANCE SERVICE WITH ROVING UNIT'
'PRNT_TR'
'LAM'  1000. 1000. 1000. 1000. 2000. 1000. 1000.
'S' 1 1 1
'S' 2 1 5
'S' 3 1 7
'S' 4 7 1 2 3 4 5 6 7
'XSPEED'   10
'YSPEED'   20.
'TX' 0.  20.  .0  40.  .0  60.  20.  30.  20.  40.  20.  50.  40.  40.
'CORTM'   0.5
1. 0. 0. 0. 2.3 0. 7.1
'EMCM'
'CAP'
'VAR_SER_TM'  15.  15.  15.  35.
'RUN' 4. .1
```

Fig. 16—Illustrative data deck: Ambulance service with roving unit

```
M=5   R=14   NUM=2   ESTSTAT=2     ;
'TITLE' '5-CAR RUN, 14 ATOMS, EMPIRICAL TRAV TIME MATRIX'
'PRNT_TR'
'LAM'  1245. 1673. 5002. 993. 3351. 4448. 128. 444. 2001. 987.
       1704. 1442. 1899. 2006.
'S' 1 1 1
'S' 2 1 10
'SS' 3 3 2 1. 3 1. 4 2.
'SS' 4 3 13 1. 14 3. 12 1.
'SS' 5 6 2 1. 3 1. 4 2. 13 1. 14 3. 12 1.
'SPEED'  9.67
'PATROL'  5.9
20.1 2.9 5. 9. 1.9 4.8 90. 33. 20. 5.
   33. 22. 11. 8.99
'TR'
1. 2. 3. 4. 5. 6. 7. 8. 9. 9. 9. 9. 9. 9.
2. 1. 2. 3. 4. 5. 6. 7. 8. 9. 9. 9. 9. 9.
3. 2. 1. 2. 3. 4. 5. 6. 7. 8. 9. 9. 9. 9.
4. 3. 2. 1. 0. 2. 1. 3. 4. 5. 6. 7. 8. 9.
5. 3. 0. 0. 2. 2. 3. 4. 5. 6. 7. 8. 9. 4.
6. 5. 4. 3. 2. 1. 0. 1. 2. 3. 4. 5. 6. 7.
7. 6. 5. 4. 3. 2. 1. 0. 1. 4. 6. 7. 9. 9.
8. 9. 9. 9. 6. 6. 6. 4. 1. 1. 3. 5. 6. 7.
9. 8. 7. 6. 5. 4. 3. 2. 1. 2. 3. 4. 5. 6.
10. 2. 1. 0. 4. 5. 6. 7. 8. 9. 5. 4. 7. 9.
11. 4. 4. 5. 6. 7. 8. 9. 4. 3. 2. 1. 3. 7.
3.2 2. 4. 6. 8. 9. 8. 6. 5. 4. 3. 2. 1. 5.
5.1 3.2 8.4 3. 6. 2. 9. 8. 7. 6. 5. 4. 3. 2.
3. 4.67 2. 5. 9. 7. 9. 6. 4. 1. 0. 8. 7. 6.
'EMCM'
'SERVTM' 29.67
'RUN' 3.89 2.55
```

Fig. 17—Illustrative data deck: 5 cars with empirical travel time matrix

---

Appendix A

FUTURE MODIFICATIONS TO THE PROGRAM

The program described in this user's manual is undergoing continual modification and improvement, based primarily on feedback from users.[*] To the extent possible, we plan to make these changes in an "upward compatible" fashion. Thus, additional features will be invisible to the user who is accustomed to working with an earlier version of the program, and the instructions in this manual will still be applicable for later versions.

Some of the changes that have already been suggested and are being considered for the next version are as follows:

1. Print out the mean dispatcher queue delay (in the case of an unlimited line capacity system).
2. Add an option to suppress printout of much of the input data and related initialization matrices.
3. Allow unit-specific mean service times in the hypercube approximation procedure.
4. Improve the glossary option.
5. Add options to facilitate the "layering" process for different types or priorities of calls.
6. Add option to print out *distributions* of travel times (not just mean values).
7. Make the inter-atom travel time matrix easier to modify when there are barriers or other obstructions to travel.
8. Add an option permitting the user to specify when a response will be counted as intra-district, in the case of overlapping districts.

Readers are invited to submit other suggestions and to comment on the possibilities listed above.

---

[*] These modifications are being made at the Massachusetts Institute of Technology under the grant from the National Science Foundation mentioned in the Preface.

Appendix B

ADDRESSES FOR FURTHER INFORMATION

1. For copies of the Hypercube Model on card or tape, additional documentation of the model, information about related emergency service deployment models, and output listings generated from each of the five sample data decks:

Professor Richard Larson
Room 4-209
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139
(617) 253-1358

Dr. Jan Chaiken
The Rand Corporation
1700 Main Street
Santa Monica, California 90406
(213) 393-0411

2. Research sponsor:

U.S. Department of Housing and Urban Development
Alan Siegel, Director
Hartley Campbell Fitts, Program Manager
Office of Policy Development and Research
Community Development and Management Research Division
451 Seventh Street, S.W.
Washington, D.C. 20410
(202) 755-5580

REFERENCES

1. Larson, R. C., "A Hypercube Queuing Model for Facility Location and Redistricting in Urban Emergency Services," *Computers and Operations Research*, Vol. 1, No. 1, March 1974, pp. 67-95. (Also available from The Rand Corporation as R-1238-HUD).

2. Larson, R. C., *Urban Emergency Service Systems: An Iterative Procedure for Approximating Performance Characteristics*, The Rand Corporation, R-1493-HUD, 1974. (Revised version to appear in *Operations Research*.)

3. Larson, R. C., "Illustrative Police Sector Redesign in District 4 in Boston," *Urban Analysis*, Vol. 2, 1974, pp. 51-91.

4. Chelst, Kenneth, *Implementing the Hypercube Queueing Model in the New Haven Department of Police Services: A Case Study in Technology Transfer*, The Rand Corporation, R-1566/6-HUD, 1975.

5. Jarvis, James, Mark McKnew, and Larry Deetjen, *Data Collection and Computer Analysis for Police Manpower Allocation in Arlington, Massachusetts*, Innovative Resource Planning in Urban Public Safety Systems, Massachusetts Institute of Technology, Cambridge, Massachusetts (to appear).

6. Chaiken, Jan M., *Hypercube Queuing Model: Executive Summary*, The Rand Corporation, R-1688/1-HUD, July 1975.

7. Wilson, O. W., *Police Administration*, McGraw-Hill, New York, 1963.

8. Larson, R. C., *Urban Police Patrol Analysis*, MIT Press, Cambridge, Massachusetts, 1972.

9. Kolesar, Peter, and Warren Walker, *A Simulation Model of Police Patrol Operations*, The Rand Corporation, R-1625/1-HUD (*Executive Summary*), R-1625/2-HUD/NYC (*Program Description*), 1975.

10. Chaiken, Jan M., *Patrol Allocation Methodology for Police Departments*, The Rand Corporation, R-1852-HUD, 1975.

END