HOME OFFICE
Fire Research and Development Group

FIRE COVER MODELLING FOR BRIGADES

by

Cath Reynolds

Jim Pedroza, Systems Options Ltd.

FRDG Publication Number 6/98

ISBN 1840821493

Home Office
Fire Research and Development Group
Horseferry House
Dean Ryle Street
LONDON SW1P 2AW

ABSTRACT

The Fire Cover Model has been developed by the Home Office to assist fire brigades with the planning and disposition of resources to achieve optimum fire cover. This report outlines the principles and development of the model.

MANAGEMENT SUMMARY

The Home Office Fire Cover Model has been developed over a number of years to assist brigades with the task of fire cover planning.

Fire cover principally involves siting fire appliances in locations which best serve the local conditions and meet the national standards of fire cover. The national standards state that a given number appliances should reach an area in a given time, depending upon its risk. The task faced by brigades is therefore to find optimum sites for resources, taking the local geography, incident patterns and road network into account.

The fire cover model allows users in brigades to view maps of their brigade area and then to overlay information relevant to fire cover planning, such as the locations of stations and incidents. The user can build a model of the road network in the area using simple road links and junctions and then use the model to calculate how long it takes for appliances to reach each junction. Once validated for actual travel times, the model can then be used as a predictive tool, to give quantitative analyses of hypothetical fire cover scenarios, for example, what happens to the overall fire cover within a brigade if a station is moved from one site to another.

The model is currently in 25 brigades throughout the UK and is likely to be updated shortly, as a result of initiatives within the Home Office to review the national standards of fire cover.

CONTENTS

1 INTRODUCTION TO FIRE COVER

Traditionally the term ‘fire cover’ has been used describe the disposition and deployment of firefighting equipment and personnel.

Whilst the primary role of the fire service is firefighting, they also attend and deal with many other sorts of incidents, such as flooding, road traffic accidents, people stuck in lifts or caught in machinery. These are collectively known as special services and account for almost half of the operational calls.

The fire service in the UK usually respond to calls in fire appliances, which carry 4 or 5 crew and appropriate firefighting and rescue equipment. Most appliances are general purpose vehicles which are sent to all types of incidents. Some appliances are more specialised and carry specialist equipment. For example most brigades have specialised vehicles for dealing with road traffic accidents and these carry specialised equipment for cutting people from cars.

The UK currently has national guidelines for the weight and timeliness of responses to fire calls. These guidelines were first formulated by the Riverdale Committee in 1936 and have been modified a number of times, most recently by the Joint Committee on Standards of Fire Cover in 1985. The guidelines depend largely upon the risk of fire spread within an area. The risk of fire spread in an area is ascribed to one of four broad categories; ‘A’, ‘B’, ‘C’ and ‘D’. In each risk category, the required response is specified as a number of ‘pumps’ (first line pumping appliances) which must arrive at an incident within the time limits shown below.

Risk	1st appliance	2nd appliance	3rd appliance
A	5 mins	5 mins	8 mins
B	5 mins	8 mins	
C	10 mins		
D	20 mins		

Table 1-1: Fire Risk Categories and Attendance Standards

These standards apply only to calls to fires (and only those which do not occur whilst another fire is being attended by the brigade). At present there are no standards for special service calls, although draft guidelines have recently been introduced. They are guidance only, but fire services are required to provide performance indicators based upon them.

2 ORIGIN AND BACKGROUND OF THE MODEL

The problem of planning the best location and distribution of resources to provide the an optimum level of fire cover is quite complex. It involves not only the size and position of fire stations, but also the number and type of incidents attended, the standards which fire brigades strive to meet, the geographic constraints of the area and, of course, the financial constraints of the fire authority.

In principle, however, it should be possible to construct a model which balances the cost of providing a fire service against the cost of fire losses, in terms of both property and life.

In 1970, the Departmental Committee on the Fire Service, chaired by Sir Ronald Holroyd produced a report(1) which suggested an optimum size for a fire brigade in terms of the number of stations. In the same year, Hogg produced a report(2) which detailed a distribution model for an emergency service. This model attempted to balance the cost of fire losses against the cost of provision of fire cover and hence to find an optimal solution. Unfortunately, much of the data required as input for the model was not available then and, even now, is difficult to obtain.

In 1975 Rutstein published a report(3) which progressed the original study (with the help of some additional data gathered during 1970 and 1971 on fire damage), to establish a relationship between the losses caused by fire and the attendance time of the fire brigade. The state of development the model was reviewed by Harwood and Taylor in 1982(4) to provide information for the review of fire cover standards being undertaken by the Joint Committee on the Standards of Fire Cover. Subsequently, the model was used (without the loss attendance relationships) to assess the impact of potential changes to standards of fire cover in brigades(5).

In 1992, with the emergence of desktop computer technology the model was rewritten for use with personal computers so that brigades could use the software themselves rather than relying on consultants to run it for them. In 1994 the model was updated further by incorporation into a geographical information system which made interaction with the model easier.

3 THE PRINCIPLES OF THE MODEL

3.1 Introduction

The Fire Cover Model is a computer representation of the operational aspects of a brigade which are relevant to fire cover planning. The model carries out the same planning process that brigade's officers have traditionally undertaken, but it enables planning to be done much more quickly and allows quantifiable evaluation of alternative, hypothetical fire cover scenarios.

3.2 The calculation of minimum attendance times

The first step in modelling fire cover is to design a model which can predict the time it takes for an appliance to reach an incident, This first step makes some fairly crude assumptions about the nature of incidents, such as: an appliance is always available at the nearest station, that incidents happen one at a time and are evenly spaced throughout the 24 hour day. (The effect of these assumptions can be modelled later).

The main aspects of fire cover which need to be taken into account when calculating the minimum time taken for appliances to reach an incident are:

- geography of the area
- incident patterns
- travel times of appliances on the run
- turn-out times of appliances
- risk categories
- brigade resources

3.2.1 The geography of the area

Each brigade is unique in its road network and other geographical features. These geographical aspects are represented in the model by a network of the roads most frequently used by the fire appliances. Road junctions are represented by "nodes" in the network, where road links join.

Nodes can be sited at various types of location :

- major road junctions,
- station locations, and possible future sites,
- incident clusters which are not near to other nodes,
- motorway accesses, on sections which give a significant number of incidents.

3.2.2 Incident Patterns

Since a brigade attends many thousands of incidents each year, it is not practical to consider the attendance time to each incident. Incidents are therefore grouped together into zones. Each zone is a compact area, of one risk category, whenever possible. The number of incidents per year in each zone is estimated from past records of incidents.

The way in which zones are defined is usually chosen to fit the way that incidents' locations have been recorded. The number of incidents in each zone must ultimately be allocated to one or more of the nodes in the network.

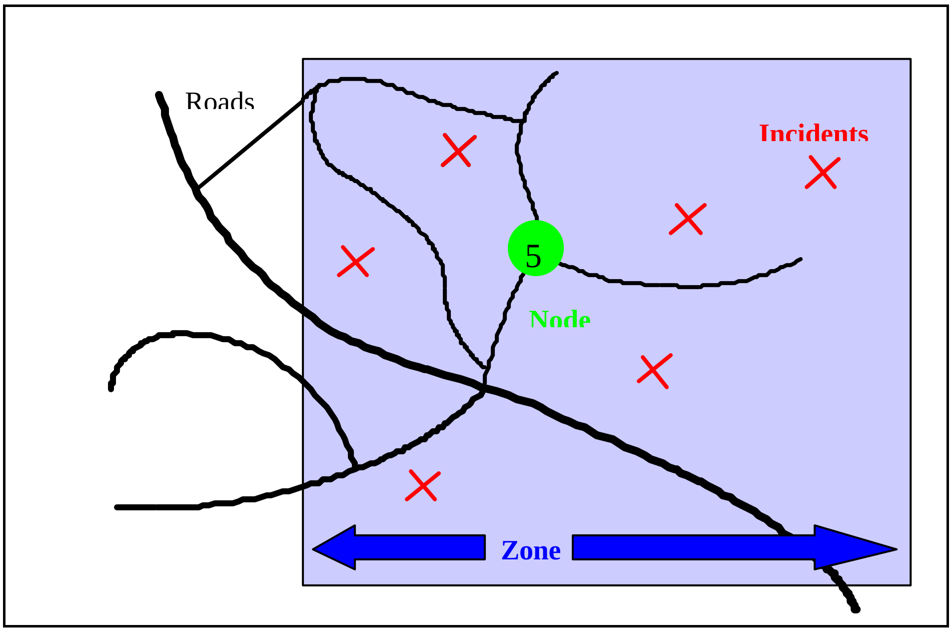

Figure 3-1: The allocation of incidents to nodes

Each node is allocated the number of incidents (per year) which could be reached from that node more quickly than from any other node. In this way, the distribution of the brigade's incidents over its area is represented.

3.2.3 Travel times

A nodal network of the brigade area is developed. Each node has a number of neighbours. A neighbour is a node which would normally be reached without going via any third node.

In this example, Node 1 has neighbours 2, 3 and 5. Node 4 can best be reached via nodes 2 or 3 and therefore is not a neighbour of Node 1.

For each pair of neighbour nodes, the average time needed (under real conditions of an emergency call) to travel along that stretch of road, is input to the model or can be calculated from road speeds. This is called the inter-nodal time. When this has been done for all neighbour nodes, the model calculates the time needed to travel from every node to every other node, by the

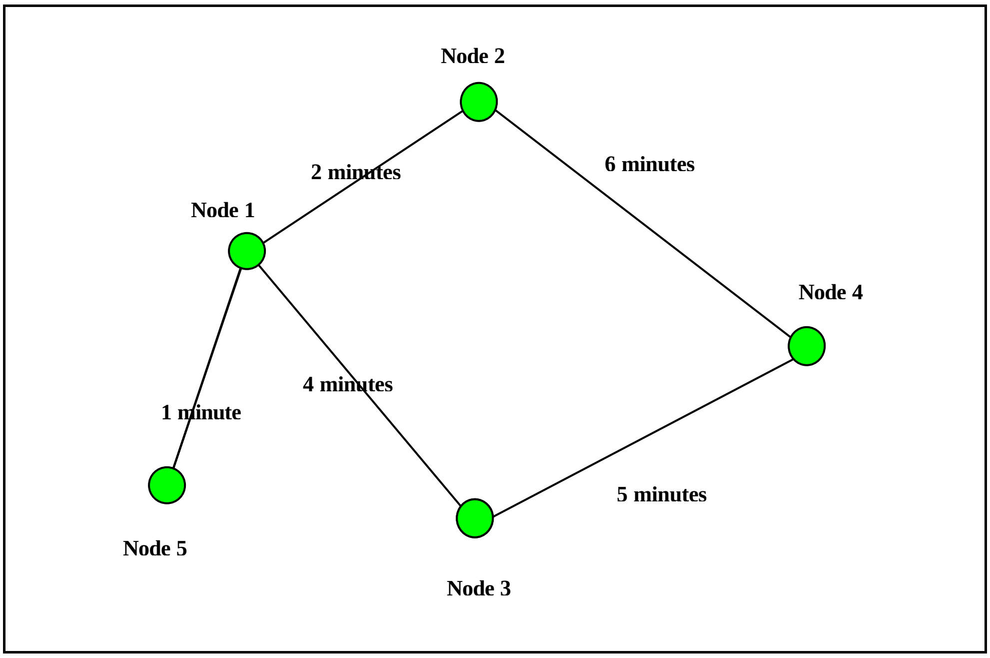

Figure 3-2: Travel times between nodes

quickest path.

The result of this calculation is the time taken to travel from one node to another, but no account has yet been taken of the time needed for the final part of the journey, namely the extra time taken to travel on minor roads to the location of the fire. The majority of incidents will not occur exactly at a node, so a small extra time is added for travelling from the node to the incident. This part of the journey lies entirely within one zone - that is, an incident zone - and

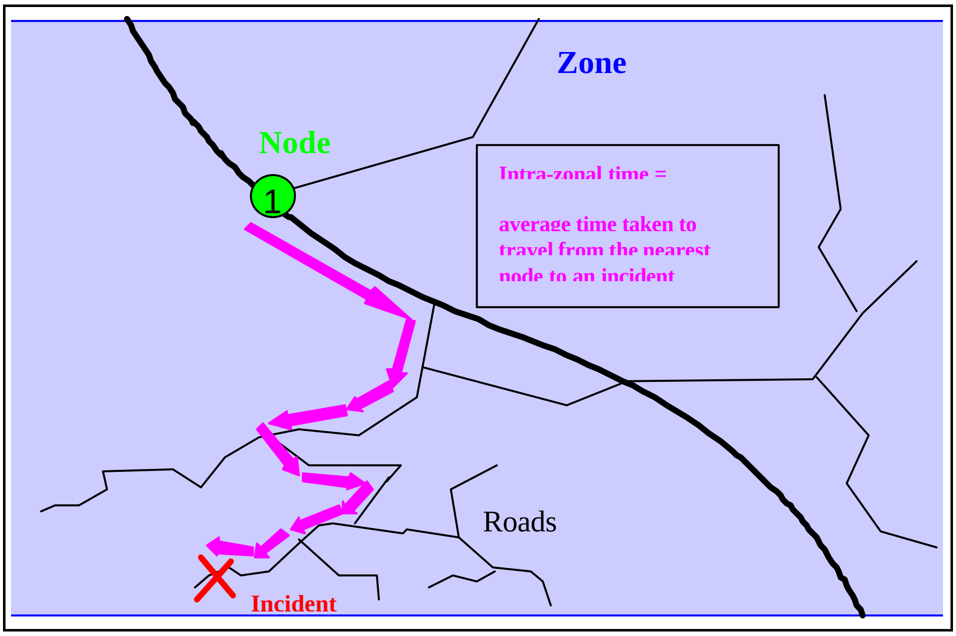

Figure 3-3: Intra-zonal travel time

is therefore called the "within-zone", or "intra-zone" time. It is supplied to the model as an average value for the zone concerned. In urban areas with small zones, it may be no more than a fraction of a minute, but in rural areas with scattered incidents it might be several minutes.

The total travel time from a station to a call in zone X is the sum of the inter-nodal times, by the quickest route to the node serving zone X, plus the within-zone time for zone X.

3.2.4 Turn-out times

In addition to the local geographical information items described above, the model requires figures for the turn-out times for stations. The turn-out time of a station is the time taken for an appliance to be mobilised from the time the emergency call was received at the station.

The default turn out times in the model are :

whole-time crewing :	1 minute
day crewing :	1 minute during day shift, and 2.5 minutes at other times
retained crewing :	5 minutes.

although these can be easily changed globally, or for individual stations.

The response time to a node is therefore the sum of:

(a) the turn-out time

(b) all the link times on the shortest-time path

(c) the intra-zonal time.

The model does this calculation for every combination of station and node. For example, if there are 400 nodes and 30 stations, this produces 12,000 response times, which are stored for later use.

3.2.5 Risk categories

Each node is given a risk category for each incident type, which reflects the area it serves. Usually the risk category reflects the fire risk categorisation suggested by the Home Office, which depends upon a number of factors including the type of buildings present. However, for incidents such as special services, where no standard risk categorisation exists, users can define their own risks and associated pre-determined attendances.

3.2.6 Brigade resources

Each station is represented as a node on the network. The number of pumps (or special appliances) at each station is specified, together with the type of crewing of each appliance.

3.2.7 Level of cover required

The level of cover required is mainly determined by the guidelines issued by the Home Office. However, local policies may also have an impact. For example, it is common practice in many brigades to send two appliances to house fires in 'C' risk areas. The model uses the Home Office guidelines as default for the level of cover required, but allows modification of these to take account of local policy.

3.3 The calculation of average attendance times

In the discussion so far, the model has not done anything beyond normal planning practice in brigades - estimating the total response times to areas within the brigade. However, the model then goes on to consider the effects of various complications which occur in practice. These include :

- the effect of workload on response times;
- multi-pump stations and incidents;
- the amount of time needed to deal with incidents;
- the variations in call-rate throughout the day.
- risk categories.

3.3.1 The effect of workload on response times

In the response time calculation described above, the model calculates the time needed to respond to a call from a particular node, by appliances at each station. Normally, a zone will be served by the station which is nearest in time. But there will be occasions when the nearest station's pumps are already attending another call, so that a station further away may need to be called out. The response time of this more distant station will be greater than that of the nearest station, and so the average response time to the node being considered will be greater than the response time from the nearest station. The extent to which the response time is increased depends on the proportion of time that the nearest
station is attending calls - that is, its workload.

For example, consider the case of an area with two stations, A and B, each with one pump, in which calls never require more than one pump to respond. Consider a single node X in the area, which has a response time from the nearer station, A, of 4 minutes, and of 8 minutes from the other station.

If station A has a workload which keeps it busy attending calls for 25% of its day to nodes other than the nearest to it, then the proportion of calls at node X which are served by station B will be 25%, with 75% served by the nearer station A.

The average response time to node X will be the weighted average of the response time from station A (4 minutes) and that from B (8 minutes), that is:

average response time at node X = 75% x 4 + 25% x 8

= 3 + 2 = 5 minutes

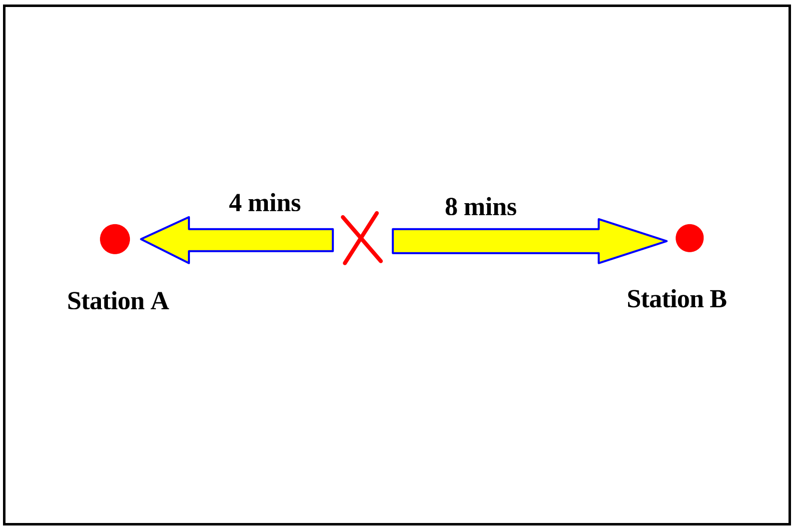

Figure 3-4: Workload and response times

This response time compares with the 4 minutes expected on the simple assumption that the nearer station is always available.

If the workload on station A is decreased, say to 10% instead of 25%, then the average response time becomes:

average response time at node X = 90% x 4 + 10% x 8

= 4.4 minutes.

This illustrates that when stations have a light workload, the response time is very little more than the minimum response time.

This "simplest case" illustration contains two difficulties which need to be examined. First, the next-best station's own workload has been ignored. Second, the workload calculation needs explanation.

Dealing first with the next-best station's workload, it is apparent that if station A's workload is 25% and station B's is 10%, then on average for 10% of 25% of calls, that is 2.5% of calls, both pumps will be busy and no response is possible. But this is because our example is too simplistic. In practice there are more than two stations, and the model simply deals with all that have been included in the study area. That is to say that the model calculates the proportion of calls from node X that will be attended by the nearest station, the second nearest, third nearest, ... and so on. In practice, the probability that a one-pump call from node X will need to be serviced by the station which is the fourth nearest to node X is virtually zero. But the model can, and does, include this. In this way, the workload on every station is allowed to have its influence on response times.

The second omission in the simple calculation carried out above is how the workloads were found. This example assumed that station A had a known workload which kept its pump busy for 25% of the time. The model calculates this proportion by finding which nodes are nearest in time to station A, and then multiplying the number of incidents at each node by the number of pumps needed. This calculation gives, to a first approximation, the proportion of hours in each day that the pump at station A will be busy.

However, this is only a first approximation. Whenever station A is busy and a call comes in from a node nearest to A, this is serviced by another station. This effectively reduces the workload of station A, whilst increasing that of the station which serves the call. This problem is dealt with by repeating the calculation in a series of stages, each of which uses the estimate of workload on each station produced by the previous stage in the following manner.

If we begin by assuming that the probability that an appliance is available is 1 (i.e. it has no initial workload). As soon as the appliance attends an incident, then it is unavailable to attend another call in the area, and therefore the second call must be answered by the next nearest appliance. This obviously has a knock on effect on the availability of the next nearest appliance to answer a call in its area. Similar workload consideration can affect all the appliances in the brigade area.

To model this, consider first the expected attendance time calculation (step 1).

Let

P_i = the current estimate of the probability that the ith nearest pump is available

$IncidentRate_n$ = the number of incidents at the node n

AttendTime = the average length of attendance at each incident

The probability that the nearest appliance is available is P_1. Therefore, the initial estimate of the workload for the first appliance in attending single appliance incidents at node 1 is

$$P_1 * AttendTime * IncidentRate_1$$

However, there is a probability $(1-P_1)$ that the incidents at this node are not answered by the nearest appliance because it is already allocated. The probability that the appliance will have to come from the second nearest location is

$$(1-P_1) * P_2$$

and the corresponding additional workload for the second nearest appliance is

$$(1-P_1) * P_2 * AttendTime * IncidentRate_1$$

But again, there is a probability that an incident cannot be attended by the two nearest appliances, which is given by $(1-P_1) * (1-P_2)$. Hence, the corresponding increase in workload for the third nearest appliance will be

$$(1 - P_1) * (1 - P_2) * P_3 * AttendTime * IncidentRate_1$$

This allocation of work at each node continues in this way until the 10th nearest appliance has been considered, although in practice the probability that an appliance has not been allocated usually becomes negligible before the 10th nearest appliance is reached.

Where the requirement is for two appliances to each incident, then the allocation can be generalised such that $(P_1 * P_2)$ is the probability that two appliances will come from the nearest and next nearest locations.

So

$$P_1 * P_2 * AttendTime * IncidentRate_1$$

should be added to the workload for these appliances. However, there is a probability

$$(1 - (P_1 * P_2))$$

that the incidents at this node cannot be serviced by this combination of appliances, when the first appliance is not available then the 2nd and 3rd nearest must attend, so the next alternative is to consider attendance from the nearest and third nearest appliances.

The probability that the second appliance will have to come from the third nearest location is, with the first appliance still coming from the nearest

$$(1 - (P_1 * P_2)) * P_1 * P_3$$

so

$$(1 - (P_1 * P_2)) * P_1 * P_3 * \text{AttendTime} * \text{IncidentRate}_1$$

should be added to the workload of the first and third nearest appliances.

This process continues for all appliance pairs in the area under study until steady state availability probabilities for all appliances are reached. For example, for a model with ten appliances, when considering the availability of the 10th appliance, there is a possibility of passing incidents to the 1st appliance. This gives rise to an iterative consideration of all appliance availabilities, which stops when the changes in availability with each iteration are sufficiently small to be negligible. In practice, the answer settles down to a virtually unchanging figure after only two repetitions.

These “steady state” availabilities can then be used to calculate which appliances attend incidents and hence average attendance times to incidents.

To summarise, the calculation of the workload effect proceeds through the following stages:

1. For each node, its nearest station is identified, which defines the "catchment area" of each station;

2. The workload of each station is calculated from the number of calls at each node in its catchment area, multiplied by the time needed to deal with them;

3. For each station the workload is recalculated by apportioning the calls at each node to all the stations serving it, using the workload factors calculated in the previous stage to estimate the proportion of time that pumps will be available at each station;

4. The calculation in stage 3 is repeated until the station workloads have arrived at their stable values;

5. For each node, the average response time is calculated, by the weighted average of the response times from all the stations, using the workloads at each station arrived at in stage 4.

This completes the explanation of how the average response times are calculated, but a number of details have been deferred, in order to simplify the explanation.

3.3.2 Multi-pump stations and incidents

The model is provided with the number of pumps at each station, and with the type of crewing of each of these pumps. When calls at a particular node are being considered in the calculation of response times, the model does not assume that each incident needs only one pump. Instead, it uses the distribution of the numbers of pumps needed which is defined by the user in the model.

By comparing these numbers with the number of pumps at the nearest station, the model calculates the proportion of occasions when the call will require 1, 2, 3 or more pumps from other stations. The calculations are lengthy, but no more complex than for one-pump calls. For example, if 10% of incidents produce a call for 4 pumps and the nearest station has at most two pumps, then the model will calculate the proportion of calls when 4 pumps are required from other stations - because both pumps at the nearest station are busy - and the corresponding proportions of calls for three pumps, two pumps, one pump and no pumps. If 3% of calls require 5 pumps, then a similar calculation is carried out for five pumps, and so on, up to a maximum of about ten pumps.

3.3.3 The amount of time needed to deal with incidents

The workload of each station depends on the number of calls that it responds to, and on the time taken by each call. False alarms, and all other calls on the pumps are included. The time needed to deal with a call is very variable, but the model simplifies this to an average figure. However, this average varies from one time of day to another. The 24 hours are divided into six periods of four hours, and each of these periods has its own average time for dealing with an incident. The values currently used by the model in each time period are shown below. They were derived from national statistics and are constant in the model, they may not be exactly the figures found in some brigades. These figures can be changed to suit the local conditions within the model.

Time of day	Average time - minutes
noon to 4 pm	97.0
4 pm to 8 pm	64.0
8 pm to midnight	61.0
midnight to 4 am	48.0
4 am to 8 am	53.0
8 am to noon	74.0

3.3.4 The variation in call-rate throughout the day

The call rate for each node is calculated as calls per year from the incident records of the brigade. In the model, this is then converted to calls per day, and then to calls at various times of day. It is obvious that average call rates vary greatly according to the time of day, so the 24 hours are divided into six four-hour periods, as for the time per incident. The variation in call rates, about their 24-hour average, has been found by analysis of national figures, and is built into the model. The figures currently used in the model are given below, but these can be changed within the model to reflect a brigade's activity.

Time of Day	% of incidents
noon to 4 pm	6.0
4 pm to 8 pm	18.2
8 pm to midnight	36.4
midnight to 4 am	12.1
4 am to 8 am	15.2
8 am to noon	12.1

3.3.5 Risk categories

Risk categories affect response times through their effect on the number of pumps required for calls. This obviously affects the workload of stations and through this the response times and is therefore also considered in the model.

3.4 Results from the Model

3.4.1 Overall measures of performance

At this stage, the model has calculated the average response time, at every node, under the conditions being tested. This gives the response time for the first, second, third.....down to the tenth pump, to arrive. This mass of information can be queried and analysed in detail, but for ease of use it is also condensed into two fire cover indices.

The first of these, the fire cover failure index, measures the importance of failures of the fire cover standards. To take the example of a B-risk node, having an average of 20 incidents a year, which has a response time of 6.5 minutes for the first pump and 10.0 minutes for the second pump, the Failure Index at this node would be -

$$20 \times (6.5 - 5) + 20 \times (10.0 - 8) = 30.0 + 40.0 = 70.0$$

The 5 (minutes) in the first bracket, and the 8 (minutes) in the second corresponds to the Fire Cover Standards for a B-risk area - 5 minutes for the first pump, and 8 minutes for the second.

If this node had the same response times, but only half as many incidents, the value of the fire cover failure index would be halved - to 35.0 .

The fire cover failure index only takes account of the degree of failure of a response time - no "credit" is given for a pump which arrives in less time than the Standards require. If, in the above case, the first pump had arrived well within the 5 minute standard, say in 3 minutes, then the value of the failure index would have been:

$$20 \times (10.0 - 8) = 40.0$$

This means that a change in station location which worsens the average response times of pumps can lead to an improvement in the failure index, under circumstances where it would be hard to justify this change. An example is described below.

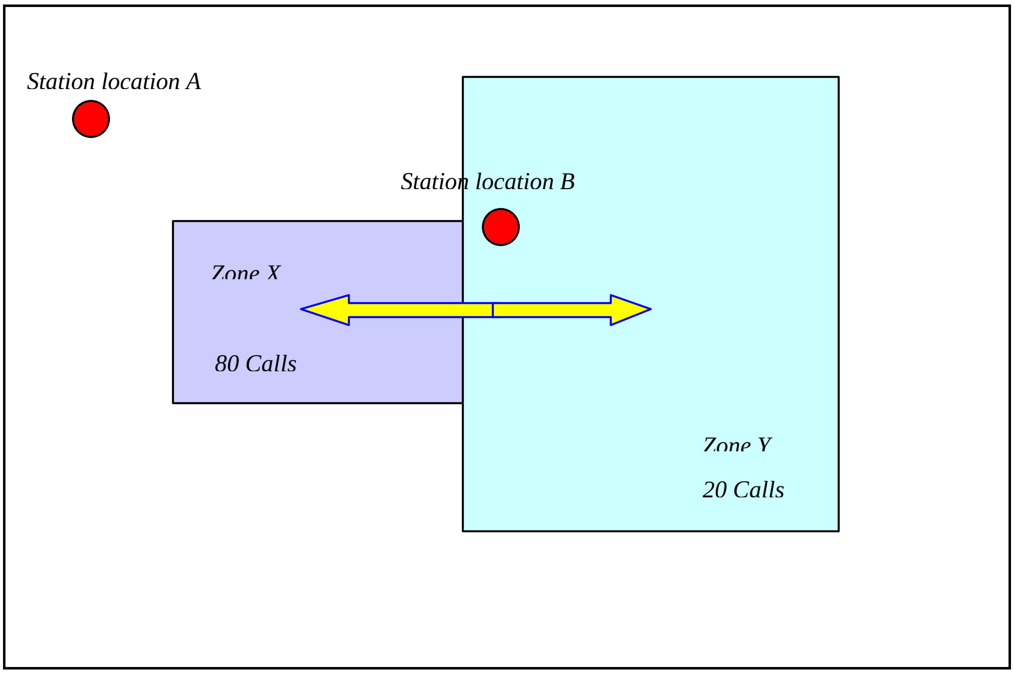

Figure 3-5: Calculation of failure index

Consider an area with just two zones X and Y, both of C risk - taken as an 8 minute standard. Zone X has 80 calls a year, and Zone Y has 20 calls. The station serving them is located near to the centre of Zone X, giving a response time of 4 minutes in Zone X. However, Zone Y has a response time of 10 minutes, which is 2 minutes outside the Standard, so that the Failure Index is:

20 calls x (10-8) minutes, which is 40. The weighted average of response times to calls in Zones X and Y is :

$$\frac{80 \times 4 + 20 \times 10}{100} = \frac{320 + 200}{100} = 5.2 \text{ minutes}$$

If the station is relocated to a point which serves Zone Y better, reducing this zone's response time from say 10 to 8 minutes, but serves Zone X worse, say increasing its response time from 4 to 6 minutes, then the Failure Index is reduced, from 40 to zero, since both zones are within the 8 minute Standard. However, the average response time is now:

$$\frac{80 \times 6 + 20 \times 8}{100} = \frac{480 + 160}{100} = 6.4 \text{ minutes}$$

A change which has increased average response time from 5.2 to 6.4 minutes has reduced the failure index from 40 to zero. This shows the effect of increasing the response time to the busier zone, in order to get the less-busy, but outside-standards zone within the Standards.

In order to draw attention to this kind of effect, the fire cover failure index is supplemented by a second measure, the overall performance measure, which reflects equally both increases and reductions in response times. Taking the previous example, with the station's original position, the overall performance measure would be:

Overall performance measure = 80 incidents x 4 mins + 20 incidents x 10 mins

= 520

After the station is moved nearer to zone Y,

Overall performance measure

= 80 x 6 + 20 x 8

= 480 + 160 = 640

This shows that the overall performance measure is proportional to the weighted average response time at all the nodes in the area, regardless of whether or not they meet the defined standards.

These two indices, taken together, provide good indicators of how fire cover in the study area would be affected by changes in the provision of resources, and their location.

4 THE CURRENT SOFTWARE

4.1 The GIS Environment

The fire cover model is currently incorporated into a Geographic Information System (GIS) which allows users in brigades to view and manipulate maps of the brigade area and then to add features, which relate to the fire cover model, such as fire cover nodes, roads and fire stations, in overlays over the maps.

For ease of reference, this chapter has been subdivided into sections which correspond as far as possible with those in the previous chapter.

4.2 The calculation of minimum attendance times

4.2.1 The geography of the area

A model of the brigade road network, taking into account the geography of the area, is built up by the users using road link and node features, as shown in Figure 4-1.

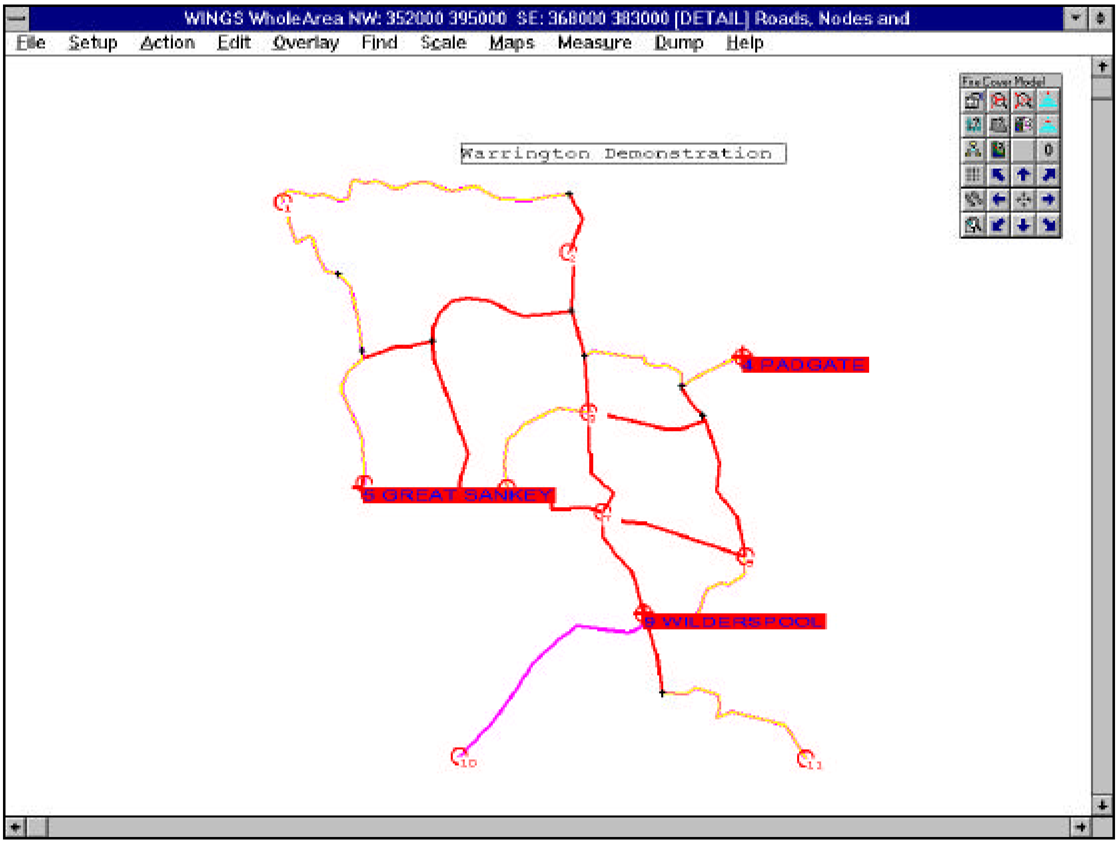

Figure 4-1: A road network

4.2.2 Incident Patterns

Fire cover zones can be defined as areas around each node and incidents input and displayed where they occur within the zones. Incident records can be loaded directly from the brigade mobilising computer or can be input manually as single, or groups of incidents.

Figure 4-2: Fire cover zones and incidents

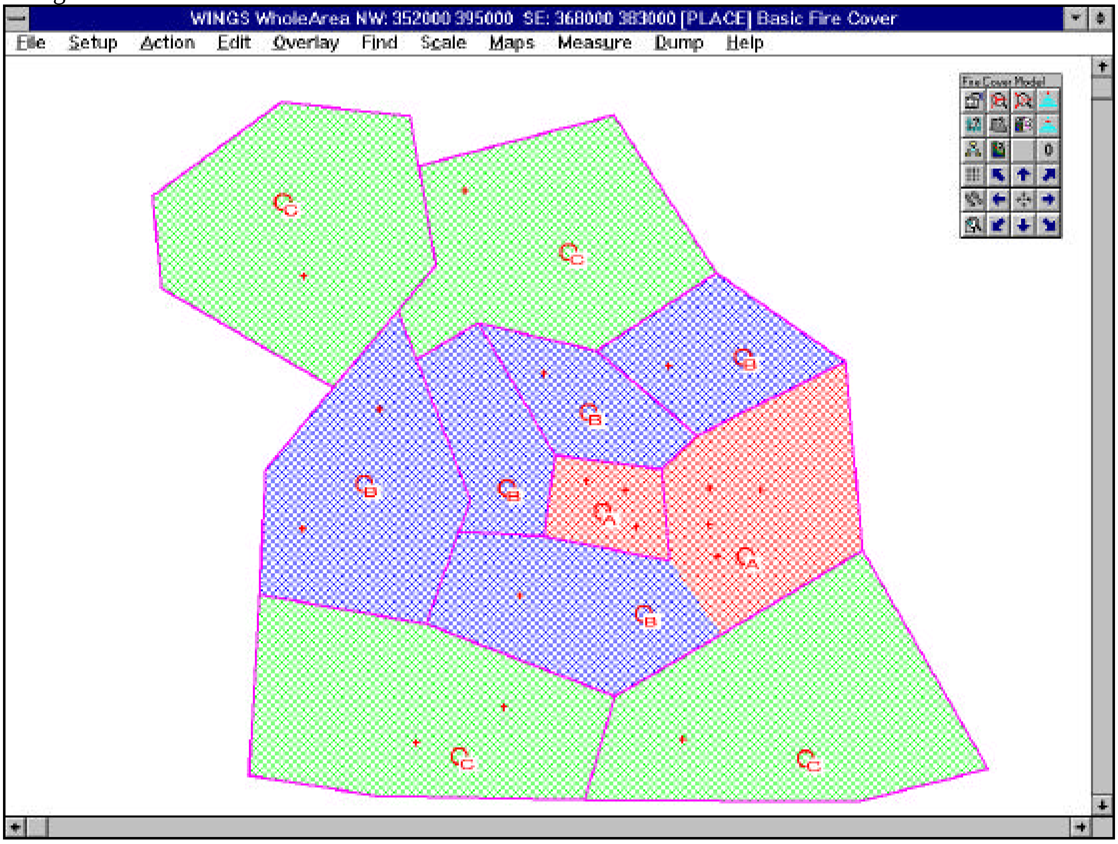

Consolidate incidents to nodes

This automatically counts the number of incidents of each type within a zone and then adds them to the node. Incidents are usually grouped into similar types, for example, all property fires might be grouped together (Figure 4-3).

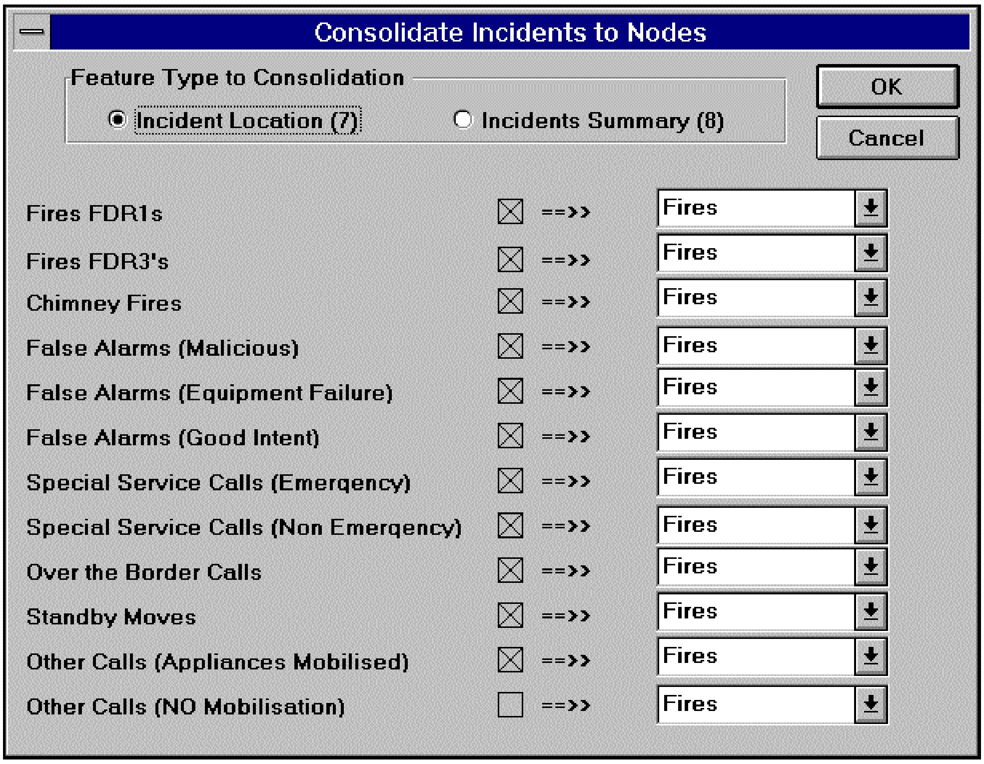

Figure 4-3: Consolidate Incidents to Nodes Function

4.2.3 Travel Times

Travel times are calculated in two steps within the current software:

a) Determine Nearest Neighbours

This option allows the user to automatically determine the nearest neighbours of nodes in the area. The option allows the results of the calculation to be stored in a number of places, so that results from several runs of the model can be compared without having to rerun the calculations.

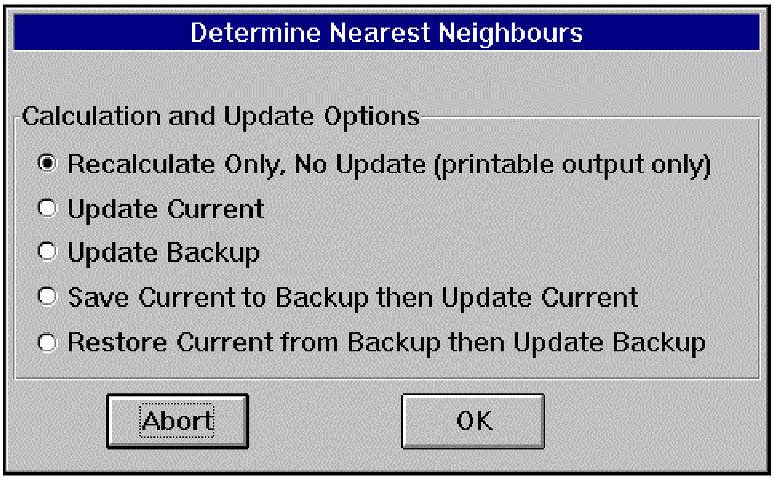

Figure 4-4: Determine Nearest Neighbours Screen

The node record, Figure 4-5, shows two sets of nearest neighbour information. The nearest neighbour numbers and times shown in the boxes are those used by the model in its calculations. There is also a backup area, beneath each of the boxes which shows the backup nearest neighbours. This backup area can be used for storing the results of nearest neighbour recalculations without overwriting the current neighbour information.

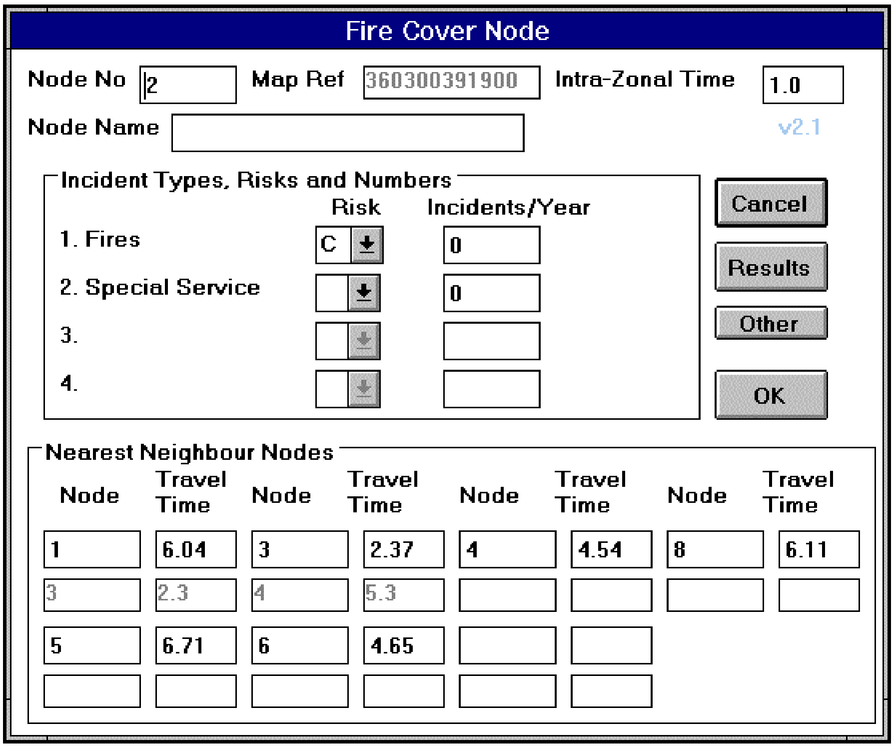

Figure 4-5: Fire Cover Node Showing Calculated Nearest Neighbours

This example shows the node record for node number 2. The model has calculated that this node has 6 nearest neighbours: 1, 2, 3, 4, 8, 5 and 6. To reach each of theses nodes from node number 2 takes 6.04, 2.37, 4.54, 6.11, 6.71 and 4.65 minutes respectively. A secondary calculation of another node configuration is shown in the backup area where nodes 3 and 4 are nearest neighbours and can be reached in 2.3 and 5.3 minutes respectively.

b) <u>Generate Travel Time Matrix</u>

This generates the matrix which contains the travel times between each of the stations to all the nodes in the area. The travel time matrix uses the shortest route in time to get to nodes. It does not include station turn out times or the effect of workload on appliance availability.

The results of the travel time calculation are written back to the node record and are displayed in the minimum arrival time box.

4.2.4 Turn Out Times

The user can change the default turn out times which are applied to all stations. The default turn out times are shown on the screen here. Changing these times means that the new turn out time is applied to all appliances which have not had a turn out time specified on the station record.

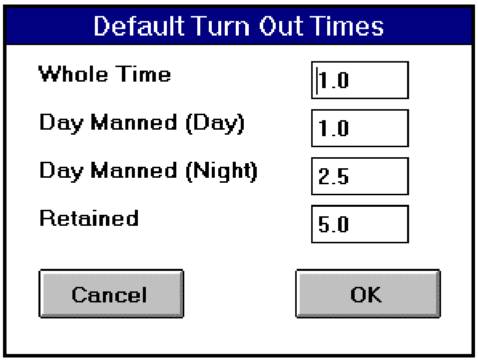

Figure 4-6: Default Turn Out Times Screen

Individual stations and appliances can have their own turn out times by changing the station record.

4.2.5 Risk Categories

This option allows the user change the current fire cover standards and also to create new standards using the following screen:

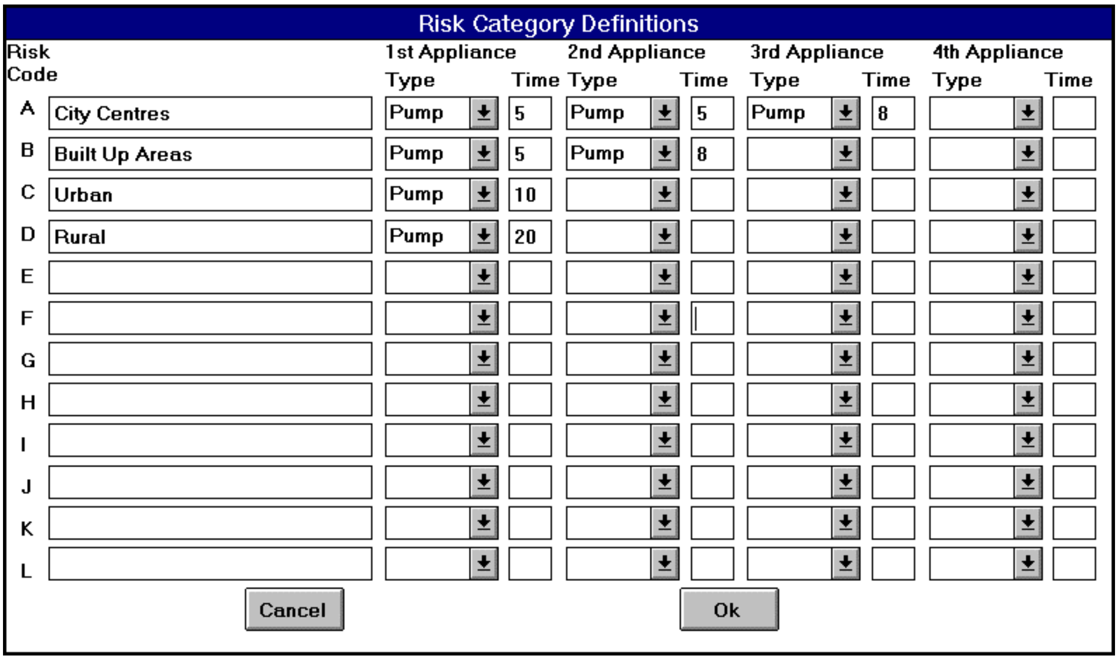

Figure 4-7: Risk Category Definition Screen

The default standards are those of the national guidelines - 3 pumping appliances attend A risk areas, the first two to arrive within 5 minutes and the third to arrive within 8 minutes, and similarly for the other standard risks. There is also space for the user to define their own risk categories, using up to four different appliances.

4.2.6 Brigade resources

The software allows the user to place and interrogate station information using the screen shown below. Standard or special appliances can be modelled using different crewing types and turnout times.

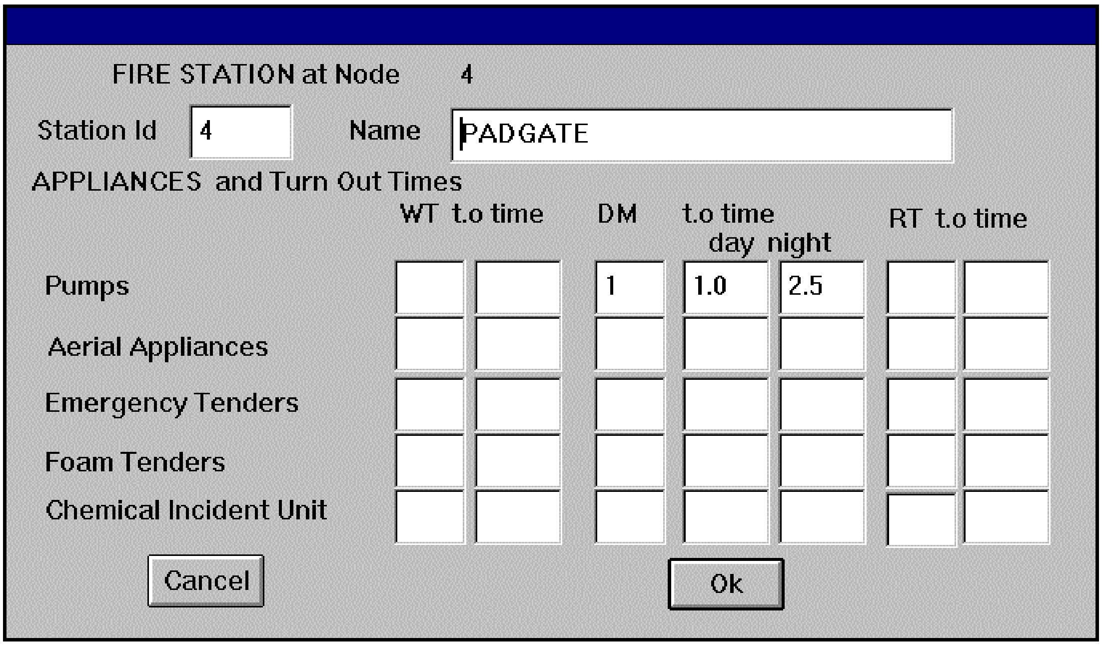

Figure 4-8: Station record

Figure 4-8 shows the station record for station number 4 Padgate. This station has one day crewed appliance which turns out in 1 minute during the day and 2.5 minutes during the night.

4.3 The calculation of average attendance times

The software uses a function called the Study Record which allows users to change factors which may affect the fire cover within an area, such as the default road speeds and turnout times, the duration and distribution of incidents throughout a 24 hour period and the proportion of calls which require more than the standard number of appliances.

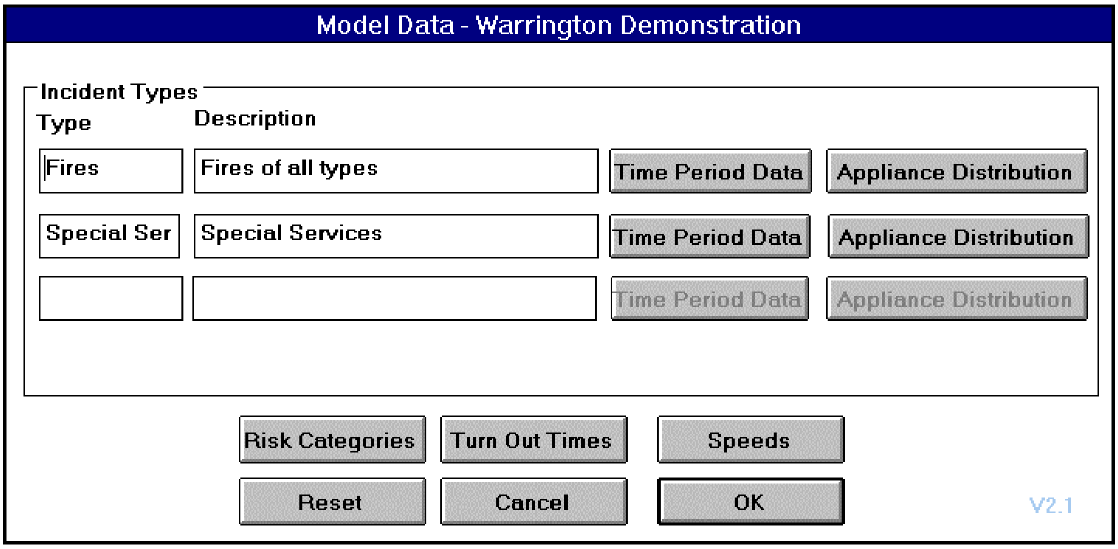

Figure 4-9: The Study Record

The screen contains three areas for defining types of incidents. These incidents categories once defined upon this screen are used throughout the model.

For the demonstration area shown here, two incident types have been defined: fires and RTA's.

4.3.1 The effect of workload on response times

Time period data option in the software allows the user to define

(a) the distribution of incidents throughout a 24 period, and

(b) the average duration (attendance time) of an incident.

These are used by the model in establishing the workloads of appliances as described in Sections 3.3.3 and 3.3.4

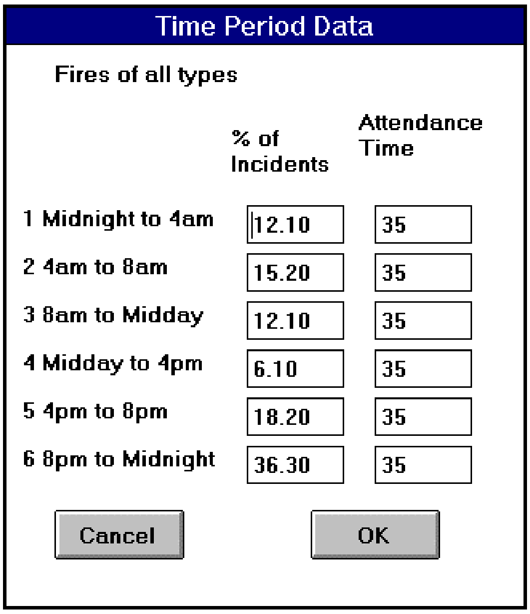

Figure 4-10: Distribution of Incidents and Duration Throughout the Day

4.3.2 Multi-pump Stations and Incidents

The software allows users to change the proportion of incidents which require 1, 2, 3 etc. appliances, (as described in Section 3.3.2). The total proportion for 10 appliances must total 1.

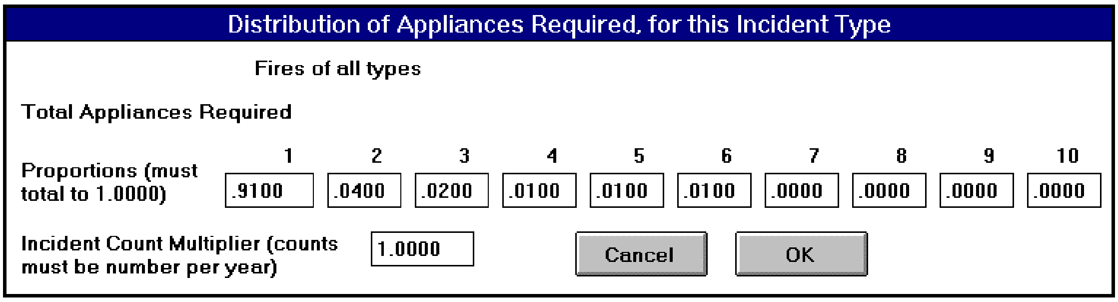

Figure 4-11: Appliance Distribution Screen

The Incident Count Multiplier allows the user to enter a multiplying factor, which will be applied to all incidents of this type. This is a useful predictive tool to, for example, examine fire cover if the number of fire incidents increases by 10% over the next 5 years or if fires become more serious and require greater attendance.

4.3.3 Run model

This option runs the fire cover model using the information on the travel time matrix. The results now include the effects of station turn out times and also the number and location of incidents. The results from running the model are written back into the node record and are recorded under average arrival time.

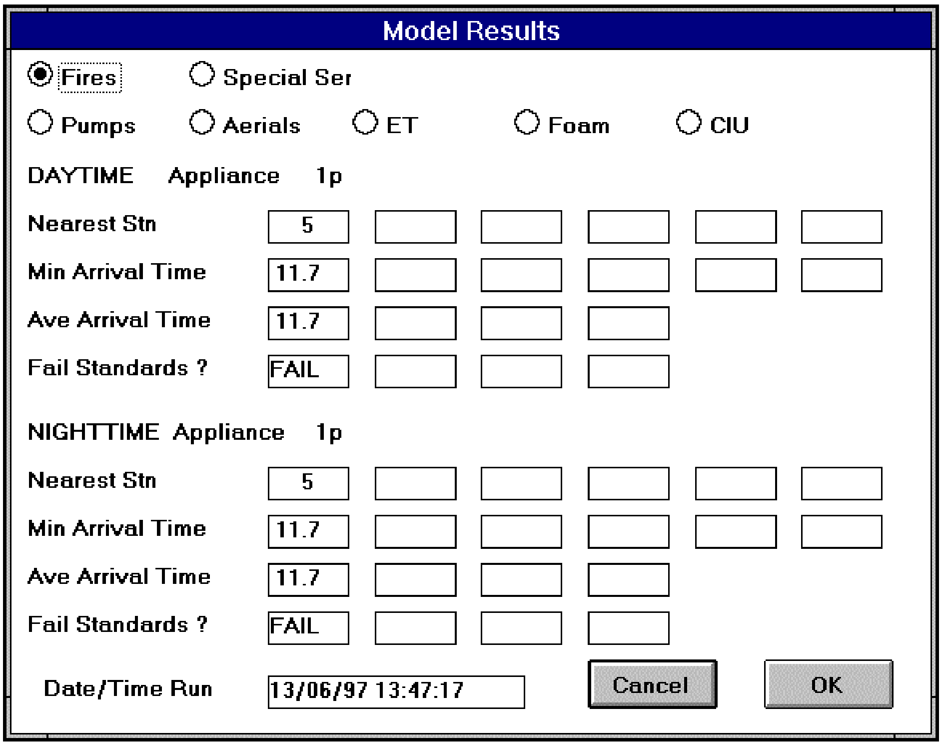

Figure 4-12: Nodes Results Screen

In this example, the minimum and average arrival times are the same, indicating that the workload is light.

4.4 Results from the Model

The software allows users to view and print tabulated results for each run of the model.

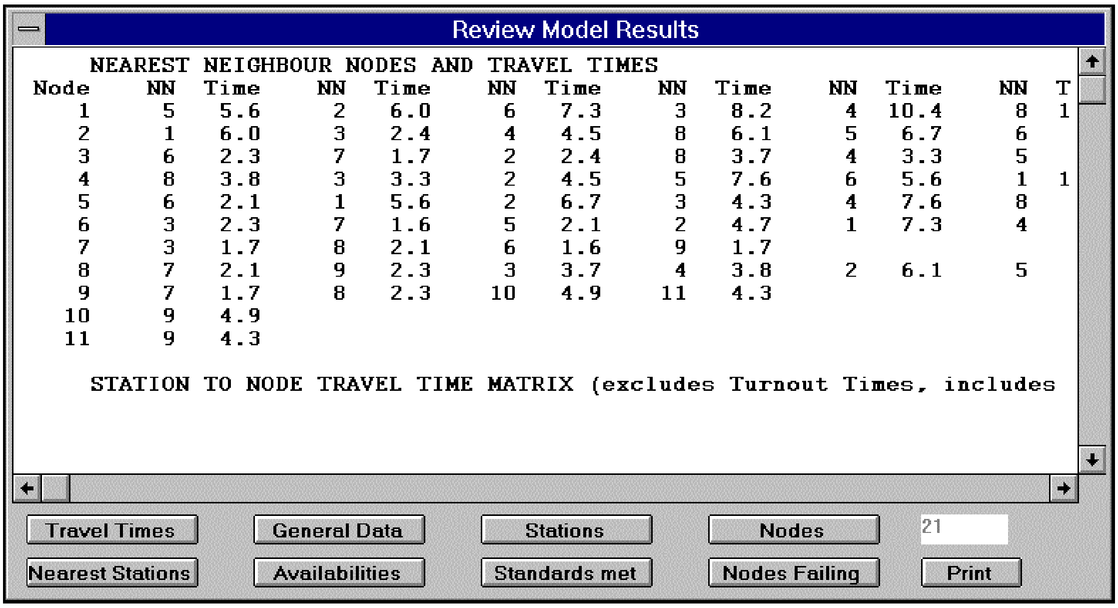

Node	NN	Time	NN	Time	NN	Time	NN	Time	NN	Time	NN	T
1	5	5.6	2	6.0	6	7.3	3	8.2	4	10.4	8	1
2	1	6.0	3	2.4	4	4.5	8	6.1	5	6.7	6	
3	6	2.3	7	1.7	2	2.4	8	3.7	4	3.3	5	
4	8	3.8	3	3.3	2	4.5	5	7.6	6	5.6	1	1
5	6	2.1	1	5.6	2	6.7	3	4.3	4	7.6	8	
6	3	2.3	7	1.6	5	2.1	2	4.7	1	7.3	4	
7	3	1.7	8	2.1	6	1.6	9	1.7				
8	7	2.1	9	2.3	3	3.7	4	3.8	2	6.1	5	
9	7	1.7	8	2.3	10	4.9	11	4.3				
10	9	4.9										
11	9	4.3										

Figure 4-13: 'Review Results' Screen

The screen displayed gives details of travel times between each node and its nearest neighbours and each station and all nodes in the model. Other results which may be printed in tabular form are:

General data	default values for roads speeds, turn out times and incident distributions
Stations	a list of station and the number of incidents which they are predicted to attend
Nodes	a list of nodes and the number of incidents at each
Nearest Stations	a list of all nodes and the first 6 appliances to arrive, with times
Availabilities	predicted availabilities for each appliance in the model
Standards met	the percentage of occasions on which standards were predicted to be met, by risk category, and also the performance indexes
Nodes failing	a list of all nodes which are predicted to fail the fire cover standards and the amount of time by which they fail.

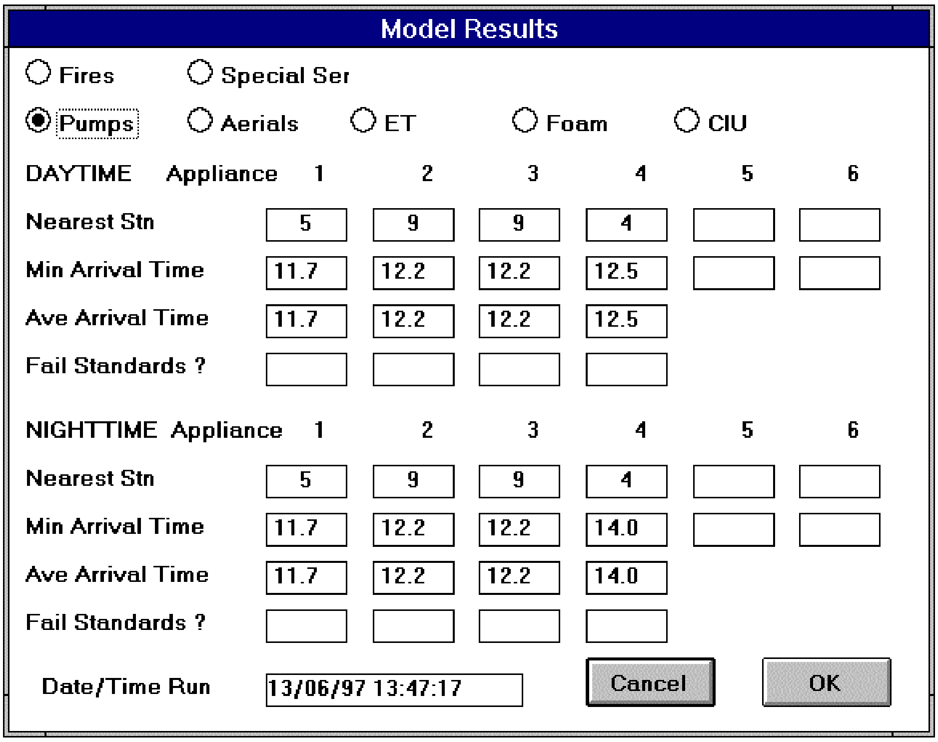

Figure 4-14: Node Results Screen

Results for each node can be queried on the node record - here the arrival times for the first four appliances are shown.

One of the principle advantages of using a Geographical Information System with the model is that the results from the model can be displayed in a format which is readily comprehensible and accessible. Figure 4-15 shows an overlay which queries the results of the model to show which of the zones are predicted to fail the standards defined by the user.

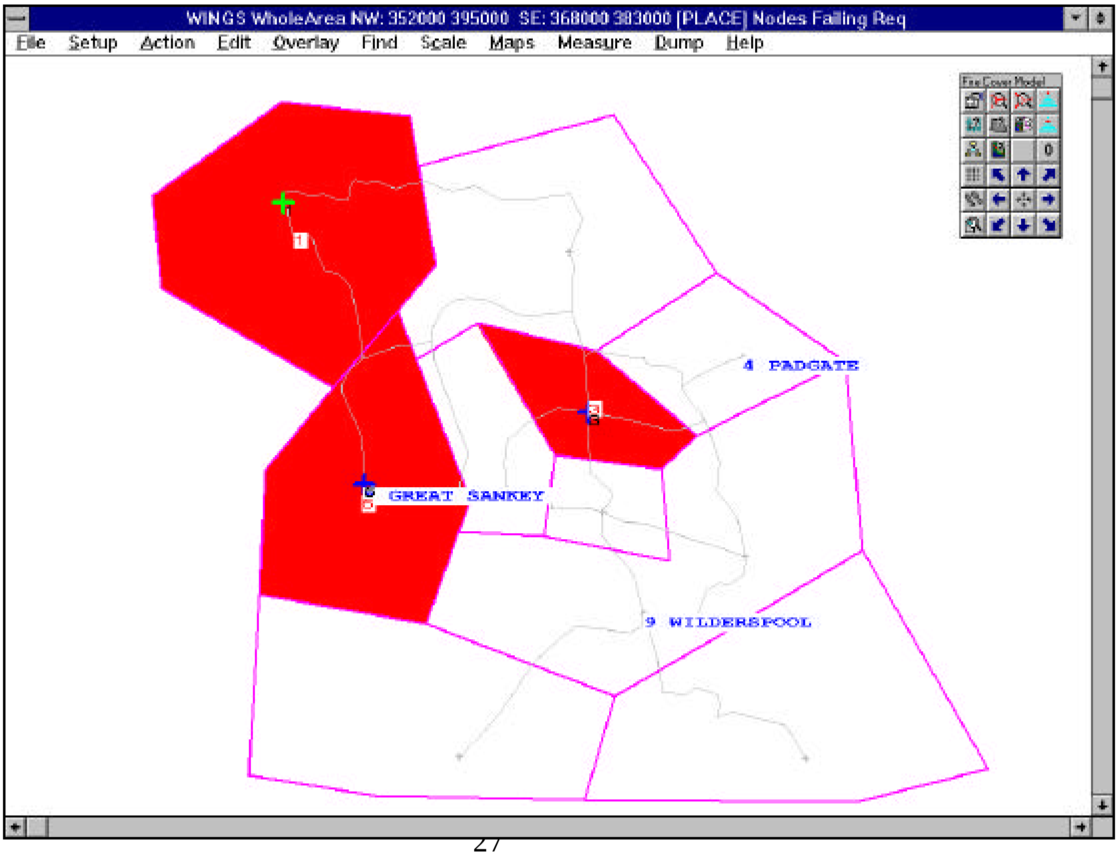

Figure 4-15: Overlay highlighting zones which are predicted to fail standards

4.5 Current Use

Currently there are 25 brigades which have the model and most have taken up the offer of Ordnance Survey 1:50,000 scale raster maps to use with it.

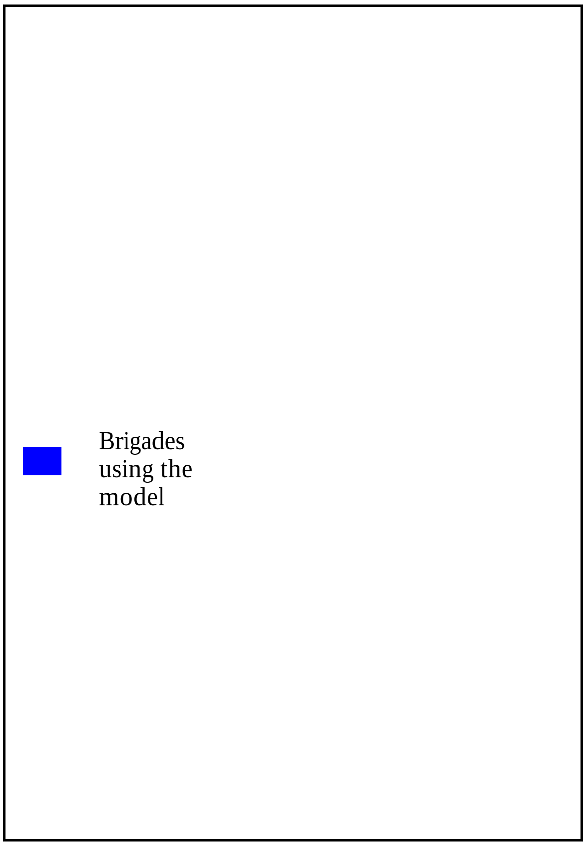

5 THE FUTURE

Following the publication of the 1995 Audit Commission Report "In the Line of Fire", the Home Office has initiated a major programme of work to review the provision of fire cover. The review is being undertaken in three main areas: risk assessment, the provision of operational fire service response, and the effects of fire safety.

The outcome of this review has been a report to the Central Fire Brigades Advisory Committee (CFBAC) which advocates a more flexible approach to the provision of fire cover. The CFBAC has accepted this and asked that development of such an approach should continue.

Although it is clear that the fire cover model will need to be changed to meet the new requirements, it is likely that the fundamentals and the GIS environment will be just as appropriate under the new scheme.

6 REFERENCES

1. Methods of Planning Fire Cover Using Cost Effectiveness Criteria, R Rutstein, Home Office Scientific Advisory Branch, Fire Research Report 7/75.

2. A Distribution Model for an Emergency Service, JM Hogg, Home Office Scientific Advisory Branch, Report 8/70.

3. Fire Cover Modelling in Cleveland, GH Dessent, Home Office SRDB Report 24/83

4. SRDB Fire Cover Model 1982 Review, J Harwood, M Taylor, Home Office SRDB Report 83/83.

5. Assessing the Resource Implications of Alternative Standards of Fire Cover using the SRDB Fire Cover Model, G Dessent, Home Office SRDB Report 55/84.