buildingEXODUS v6.3

THEORY MANUAL

BY

E.R. GALEA, P.J. LAWRENCE, S. GWYNNE,
L. FILIPPIDIS, D. BLACKSHIELDS and D. COONEY

May 2017
Revision 1.0

buildingEXODUS is a product of GUEL, a subsidiary of the University of Greenwich.

Fire Safety Engineering Group
University of Greenwich
London SE10 9LS
UK

CHAPTER 1: INTRODUCTION TO buildingEXODUS

In this chapter the theoretical basis of the EXODUS software is described [1-130]. Section 1.1 provides a general overview of EXODUS, Section 1.2 lists the various option levels available within the buildingEXODUS range, CHAPTER 2:-CHAPTER 7: describe buildingEXODUS in greater detail, while CHAPTER 8: describes the signage model and the visibility catchment area concept. Finally, CHAPTER 9:describes the lift/elevator model. Readers wishing to improve their basic understanding of human behaviour during evacuation are referred to various chapters within the SFPE handbook [34, 171]. A useful overview of published material concerning human behaviour during evacuation has also been prepared by the Fire Safety Engineering Group of the University of Greenwich [35, 56].

1.1 EXODUS Overview

EXODUS [1-15, 37-39, 48-50, 54-58] is a suite of software tools designed to simulate the evacuation and movement of large numbers of individuals within complex structures. The EXODUS family of evacuation models currently consists of airEXODUS, buildingEXODUS and maritimeEXODUS.

airEXODUS is designed for applications in the aviation industry including, aircraft design, compliance with 90 second certification requirements, staff training, development of staff procedures, resolution of operational issues and accident investigation.

buildingEXODUS is designed for applications in the built environment and is suitable for application to supermarkets, hospitals, cinemas, rail stations, airport terminals, high rise buildings, schools etc. buildingEXODUS can be used to demonstrate compliance with building codes, evaluate the evacuation capabilities of all types of structures and investigate population movement efficiencies within structures.

The EXODUS software takes into consideration *people-people*, *people-fire* and *people-structure* interactions. The model tracks the trajectory of each individual as they make their way out of the enclosure, or are overcome by fire hazards such as heat, smoke and toxic gases. The EXODUS software has been written in C++ using Object Orientated techniques and rule-base concepts to control the simulation. Thus, the behaviour and movement of each individual is determined by a set of heuristics or rules. For additional flexibility these rules have been categorised into five interacting sub-models, the OCCUPANT, MOVEMENT, BEHAVIOUR, TOXICITY and HAZARD sub-models (see Figure 1-1). These sub-models operate on a region of space defined by the GEOMETRY of the enclosure. Each of these components will be briefly described in turn.

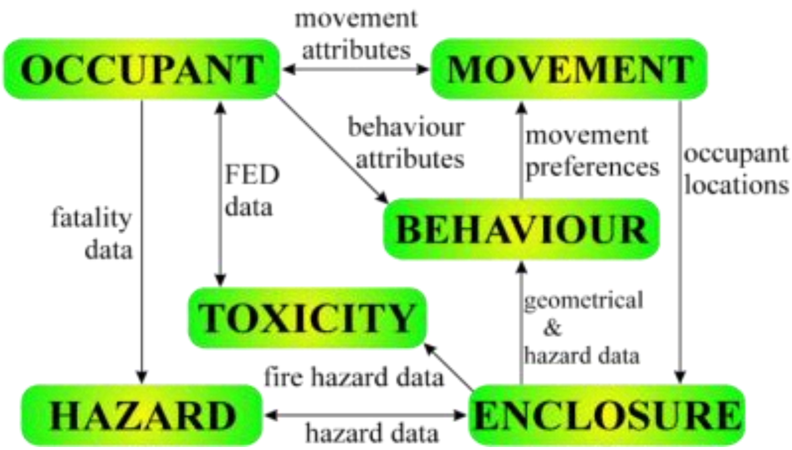

Figure 1-1: EXODUS sub-model interaction

The GEOMETRY (see Section 2.2) of the enclosure can be defined in several ways. It can be (i) constructed interactively using the tools provided, (ii) imported from third party geometry files (i.e. DXF, IFC, FDS or SMF). Internally the entire space of the geometry is covered in a mesh of nodes that are typically spaced at 0.5m intervals. The nodes are then linked by a system of arcs. Each node represents a region of space typically occupied by a single occupant, or (iii) read from a geometry library.

The MOVEMENT SUB-MODEL (see CHAPTER 4:) controls the physical movement of individual occupants from their current position to the most suitable neighbouring location, or supervises the waiting period if one does not exist. The movement may involve such behaviour as overtaking, side-stepping, or other evasive actions.

The BEHAVIOUR SUB-MODEL (see CHAPTER 7:) determines an individual's response to the current prevailing situation on the basis of his/her personal attributes, and passes its decision on to the movement sub-model. The behaviour sub-model functions on two levels:, global and local. The local behaviour determines an individual’s response to his/her local situation while the global behaviour represents the overall strategy employed by the individual. This may include such behaviour as, exit via the nearest serviceable exit or exit via most familiar exit.

The OCCUPANT SUB-MODEL (see CHAPTER 3:) describes an individual as a collection of defining attributes and variables such as gender, age, fast walking speed, walking speed, response time, agility, etc. Some of the attributes are fixed throughout the simulation while others are dynamic, changing as a result of inputs from the other sub-models.

The HAZARD SUB-MODEL (see CHAPTER 5:) controls the atmospheric and physical environment. It distributes pre-determined fire hazards such as heat, smoke and toxic products throughout the atmosphere and controls the opening and closing of exits and the availability of exits.

The TOXICITY SUB-MODEL (see CHAPTER 6:) determines the effects on an individual exposed to toxic products distributed by the hazard sub-model. These effects are communicated to the behaviour sub-model which, in turn, feeds through to the movement of the individual.

EXODUS was originally designed primarily for use with transport systems such as aircraft. However, its modular format makes it ideally suited for adaptation to other types of environment. To achieve realistic predictions in environments other than aircraft, the main component of EXODUS that needed to be adapted was the BEHAVIOUR Sub-model. In addition, EXODUS geometric capabilities required expansion to include multiple floors and staircases, the ability import geometries defined in third party formats such as those used by CAD packages, BIM/IFC and CFD models (i.e. DXF, IFC, FDS and SMF), and the ability to represent obstacles typically found in other types of enclosures such as tables and chairs. The incorporation of these differences has led to the development of the current range of EXODUS products namely, airEXODUS and buildingEXODUS.

To aid in the interpretation of the results produced by buildingEXODUS several data analysis tools have been developed, namely askEXODUS and vrEXODUS. The askEXODUS tools are intended to be used once a simulation has been completed and enable large data output files to be searched and specific data selectively and efficiently extracted. The vrEXODUS tool is a post-processor virtual-reality graphics environment enabling an animated three-dimensional representation of the evacuation simulation to be produced (see Figure 1-2).

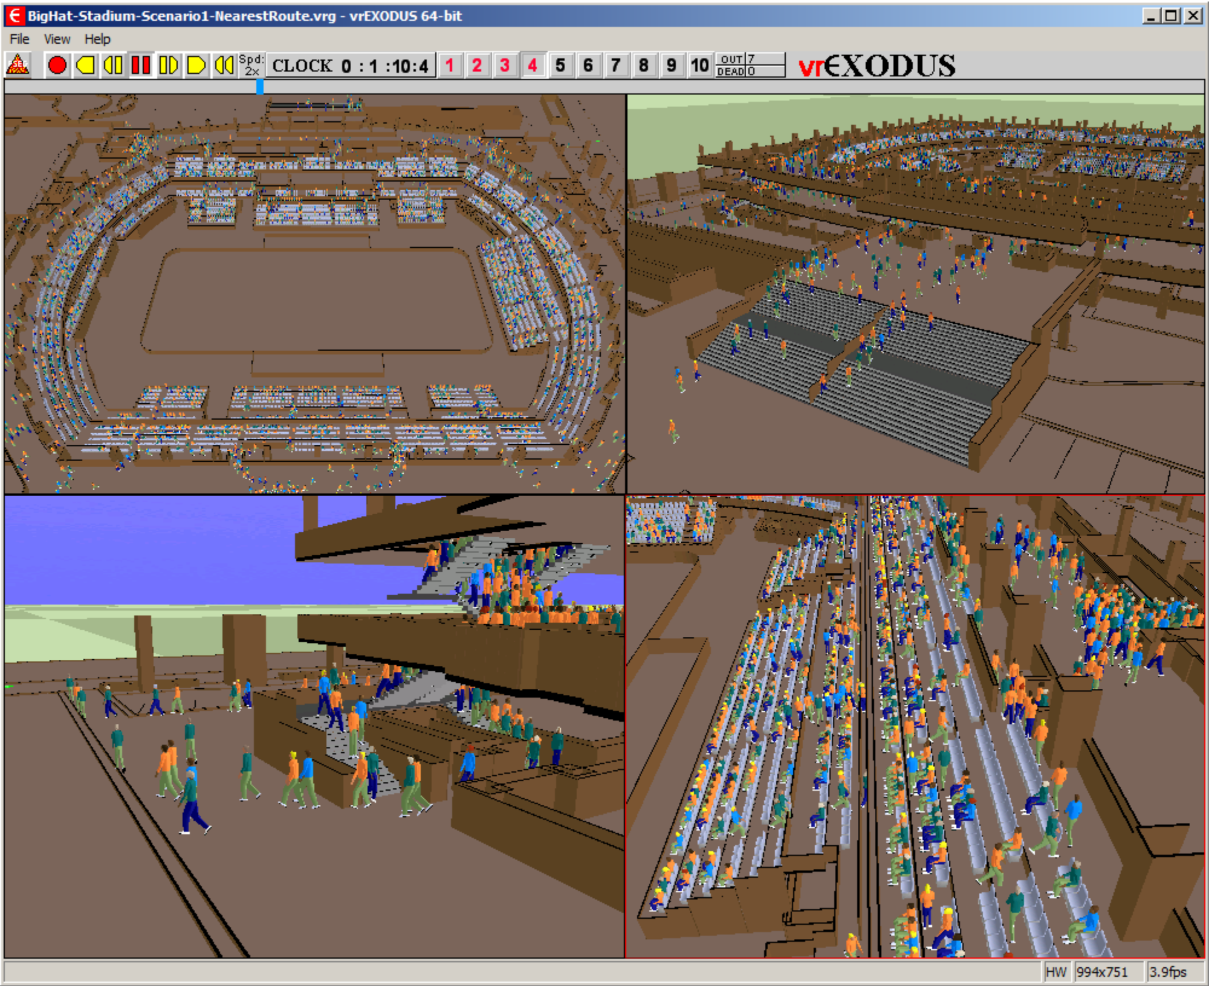

Figure 1-2: Post-processor vrEXODUS representation of a buildingEXODUS simulation.

1.2 buildingEXODUS level options

Version 6.3 of the buildingEXODUS software is available in three distinct capability levels. These are,

Level A:

Level A can handle multiple floors and unlimited population sizes, can simulate emergency and circulation situations, incorporates an abandonment model and includes the movie player facility and the data analysis tool askEXODUS. Limitations are dictated by the capabilities of the host computer. This version does not include a toxicity sub-model and possesses a limited capability hazard sub-model.

Level B:

As Level A but includes the ability to produce virtual reality representations of the structure/simulations. This includes the ability to produce output capable of being read by the post-processor virtual reality software vrEXODUS, as well as the *3D View* window within buildingEXODUS itself.

Level C:

As Level B but includes a toxicity model that allows the inclusion of the fire hazards of smoke, heat, narcotic and irritant gases within the simulation. Level C encompasses the full capability of buildingEXODUS as described in Section 1.1.

CHAPTER 2: GEOMETRY SUB-MODEL

In this section each of the various components and sub-models of buildingEXODUS will be described in turn. The discussion focuses on the full capabilities of the buildingEXODUS software (i.e. level C software). Readers are referred to Section 1.2 for an indication of the limitations of the various buildingEXODUS products. Unless stated otherwise, the capabilities described are common to all three software levels (i.e. A, B and C). The software level of your buildingEXODUS implementation is indicated in the first dialogue box at software start-up. Access to the sub-models and geometry are gained through the various EXODUS modes of operation. These modes are the *GEOMETRY MODE*, *POPULATION MODE*, *SCENARIO MODE and SIMULATION MODE* (see the User Guide, Chapter 2 for details).

NOTE:
Throughout the EXODUS software various parameters and variables are assigned pre-set DEFAULT values. In virtually all these cases the user has the option to alter these default values. While every effort has been made to assign sensible default values, it is unlikely that these will be suitable for all applications. It is therefore essential that the user reviews these default values and ensures that they are appropriate for their intended application.

2.1 Time Description

Within EXODUS time is measured by the *Simulation Clock* (SC). The SC is the master control of the model. Decisions and actions can only occur with each tick of the SC. Each tick of the SC is 1/12 of a second. The accumulation of ticks to exit or expiration for each individual is called the *Personal Elapsed Time* (PET).

NOTE:
It should be noted that the PET generated is not necessary a multiply of 1/12. The use of the 1/12 time period denotes the decision times rather than the consequences of these decisions.

2.2 Space Description

The region of space through which occupants move is termed the geometry. Geometries are defined in the *GEOMETRY MODE* of EXODUS operation. Geometries within EXODUS are represented as two-dimensional grids. The grid can be constructed manually using the interactive tools provided or automatically generated in order to fill a geometry boundary (i.e. circulation space) defined and imported from third party formats such as those used by CAD packages, BIM/IFC and CFD models (i.e. DXF, IFC, FDS or SMF). Once constructed, the enclosure can be saved into a geometry library for later use.

Each location on a grid is called a *Node*, and each node may be linked to its nearest neighbours by a number of *Arcs* (see Figure 2-2). There is no limit to the number of arcs emanating from a node and all nodes need not possess the same number of arcs. Typically, a node will possess four or eight arcs (see Figure 2-1). Occupants travel from node to node along the arcs. During the execution of a scenario (in *SIMULATION* mode), nodes and arcs are usually hidden from the user, leaving simply the outline of the structure (see Figure 2-3). Nodes and arcs possess a range of distinguishing attributes.

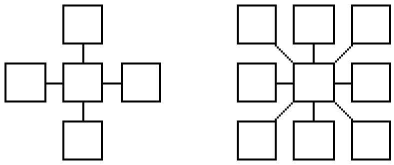

Figure 2-1: (left) Central node is connected with four arcs. (right) Central node is connected with eight arcs.

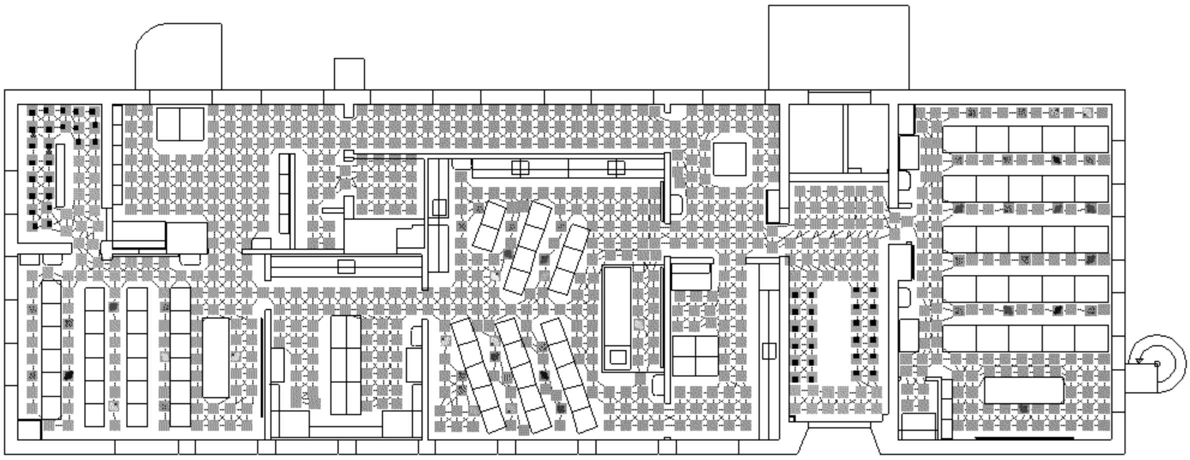

Figure 2-2: Typical outline of a building showing nodes and arcs (the DXF file is also displayed).

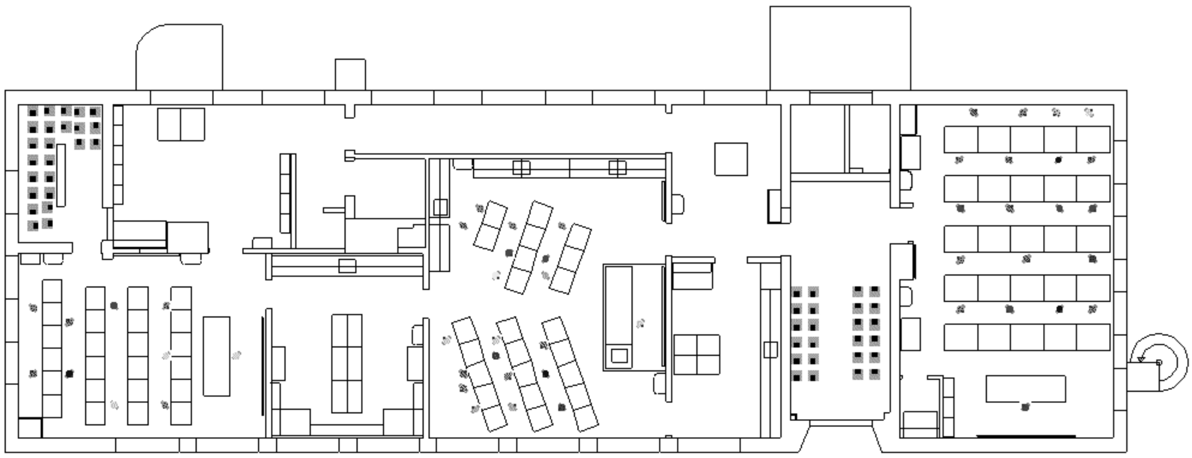

Figure 2-3: Outline of building in Figure 2.2 in Simulation mode, nodes and arcs are not visible (the DXF file is also displayed).

2.3 Node attributes

Associated with each node is a set of attributes that are used to define the nodes terrain type, environmental state and location. The attributes associated with a node are important as they may exert an influence over the person traversing the node. Nodes have a set of core attributes - common to all nodes - and a set of specialist attributes, associated with their special roles Table 2-1 lists the set of core attributes associated with each node. Specialist attributes are described in later sections.

Nodes that have common distinguishing features may be assigned to a node terrain type, for example, nodes which correspond to *Stairs* have different features to those nodes associated with *Free-Space*. Stair nodes are thus in a different terrain type to free-space nodes. There are 14 terrain types in EXODUS. These are STAIRS, LANDING, SEATS, EXTERNAL EXITS, INTERNAL EXITS, FREE-SPACE, CENSUS REGIONS, BOUNDARY, ATTRACTOR, REDIRECTION, SOURCE, DIRECTION, DISCHARGE and TRANSIT. The nature of the terrain type will influence the behaviour and maximum travel speed of the occupant passing over the node (see Table 2-2). Information concerning the terrain type is thus passed onto the BEHAVIOUR sub-model and the OCCUPANT sub-model.

Associated with each node is a set of attributes defining the environmental state of the node. These are, concentration of HCN (ppm), CO (ppm), CO_2 (%), oxygen depletion (%), smoke (l/m), temperature (ºC), HCL(ppm), HBr(ppm), HF(ppm), SO_2(ppm), NO_2(ppm), CH_2CHO (*Acrolein*) (ppm), HCHO (*Formaldehyde*) (ppm) and Radiative Flux (kW/m^2). With the exception of Radiative Flux, for each of these variables, two values are stored, representing the value at head height (e.g. can be arbitrarily set to 1.7m) and near floor level (e.g. can be arbitrarily set to 0.5m).

Table 2-1: List of node attributes used in EXODUS

* indicates attributes that have two values, an *upper* and *lower* value, that indicate the value at two heights

Attribute	Unit	Default	Updated By	Editable Mode	Node Types	Level
Title	-	Node type x x = total nodes	Auto	Geometry	All	A,B,C
Type	-	Free-Space	User	Geometry	All	A,B,C
Potential	-	0	User / Auto	Scenario	Attractor, Int. Exit	A,B,C
Node Dir	°	90	User	Geometry	Seat, Stair, Int. Exit	A,B,C
Collapsed	-	Upright	User / Auto	Geometry	Seat	A,B,C
Min UFR	occ/m/s	999	User	Scenario	Attractor, Int. Exit	A,B,C
Max UFR	occ/m/s	999	User	Scenario	Attractor, Int. Exit	A,B,C
Width	m	0.5	User / Auto	Geometry	Attractor, Int. Exit	A,B,C
Temperature*	ºC	20	User / Auto	Scenario	All	C
O_2*	%	21	User / Auto	Scenario	All	C
Smoke*	1/m	0	User / Auto	Scenario	All	C
CO*	ppm	0	User / Auto	Scenario	All	C
CO_2*	%	0	User / Auto	Scenario	All	C
HCN*	ppm	0	User / Auto	Scenario	All	C
Radiative Flux	kW/m^2	0	User / Auto	Scenario	All	C
HCL*	ppm	0	User / Auto	Scenario	All	C
HBr*	ppm	0	User / Auto	Scenario	All	C
HF*	ppm	0	User / Auto	Scenario	All	C
SO2*	ppm	0	User / Auto	Scenario	All	C
NO2*	ppm	0	User / Auto	Scenario	All	C
CH_2CHO* (*Acrolein*)	ppm	0	User / Auto	Scenario	All	C
HCHO* (*Formaldehyde*)	ppm	0	User / Auto	Scenario	All	C

NOTE:
The environmental attributes are used only in level C.

NOTE:
*The Smoke attribute is measured in units of extinction coefficient (K). In order to convert from extinction coefficient to optical density per metre (OD/m) simply multiply OD by 2.3 [25], i.e. K = OD/m * 2.3.*

The spatial and temporal variation of the environment is defined in the *SCENARIO MODE*. An occupant located on a node will experience the environmental state present at that node for as long as he/she remains at that location and for as long as that state persists. Environmental information is passed onto the TOXICITY SUB-MODEL, which determines the physiological response to the hazards for each individual, and the BEHAVIOUR SUB-MODEL, which modifies his/her physical behaviour. The environmental state of a node is controlled by the HAZARD SUB-MODEL.

As noted in Table 2-1, each node possesses an attribute known as the *Potential*. The *Potential* is a measure of the node's distance from the nearest exit. As with the environmental attributes, the *Potential* attribute may be modified by the HAZARD SUB-MODEL. *Potential*s are *grown* from each exit and increases with each step from the seed exit. Once a geometry has been constructed and the exits defined, EXODUS will automatically create the potential map. This process is performed in *SIMULATION MODE.*

As noted in Table 2-1, each node possesses an attribute known as the *Potential*. The *Potential* is a measure of the node's distance from the external door. As with the environmental attributes, the *Potential* attribute may be modified by the HAZARD SUB-MODEL. *Potential*s are *grown* from each external door and increases with each step from the seed door. Once a geometry has been constructed and the doors defined, EXODUS will automatically create the potential map. This process is performed in *SIMULATION MODE.*

Table 2-2: Node types available within buildingEXODUS.

Node Type	Description of use and influence on behaviour
FREE-SPACE	Allows unhindered movement and represents unobstructed horizontal terrain
BOUNDARY	Occupants attempt to avoid if possible, if not possible to avoid, occupants traverse at reduced speed (WALK SPEED).
ATTRACTOR	Used to locally modify potential map and control occupant flow rate (in concert with *Discharge* nodes).
DISCHARGE	Used to locally modify potential map and control occupant flow rate (in concert with *Attractor* nodes).
SEAT	Represents seating area and either forces occupant to engage in hindered movement or leaping behaviour when forced to traverse.
STAIR	Used to represent staircase and forces the occupant speed to be reduced, according to the direction of travel. Occupant forced to adhere to a range of behaviour rules associated with stairs.
LANDING	Replicates the behavioural impact of the Free-Space nodes, however is specifically designed to connect staircases and to remain visible in Boundary mode. It also reduces the travel speed of the individual to *Walk Rate*.
EXTERNAL EXIT	Represents final exit point for the occupant population. Once an occupant has reached this point, they are assumed to have completed the evacuation. It may be manipulated to control the flow capabilities, attractiveness and availability of the particular exit point. The operator may import typical values for the attributes of the exit nodes from the Appendix located at the end of the User Guide.
INTERNAL EXIT	Used to represent an exit within the building, i.e. that does not immediately lead to the outside of the structure.
CENSUS REGION	Has no impact upon occupant behaviour, but is designed to collect data concerning occupant arrivals and flow rates over the identified region.
SOURCE	Means of automatically generating people during a simulation
REDIRECTION	Redirection nodes are decision points that allow the inclusion of sub-itineraries into circulating individuals during the simulation
DIRECTION	Direction nodes are a means by which to control the direction of movement across a nodal location
TRANSIT NODE	Used to represent stair, escalators, travelators, lift shaft openings, metered gates and corridors within buildingEXODUS.

Figure 2-4 depicts two potential maps for a simple enclosure, the difference between the two figures being the width of the exit point and the subsequent number of connections from each exit point to the enclosure. Note that for the sake of ease of graphical representation and description, only 90° connections and integer distances are used to represent this map. By default, the potential map is used by the BEHAVIOUR SUB-MODEL to define and influence the global escape behaviour of the occupants (see Section 7.1(a)).

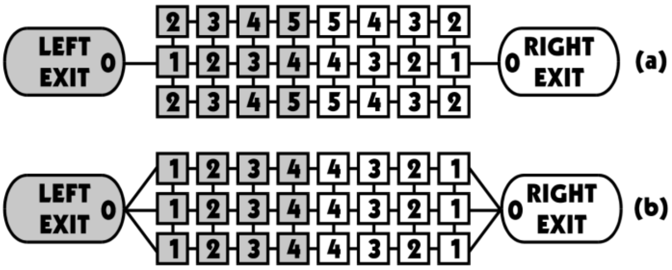

Figure 2-4: Example integer maps for a simple enclosure with (a) one connection from exit to enclosure, and (b) three connections from exit to enclosure

NOTE:
The Potential map is not used by occupants in the navigation process if the occupant has a target exit or if the OEK system is used (see Section 7.1). Under these circumstances the occupants move according to a distance map where nodal attractiveness is calculated purely according to its distance from a specific exit.

2.4 Arc attributes

Associated with each arc are two attributes. The first attribute is the length attribute. This represents the actual physical distance between nodes. In most cases this distance is set at 0.5m.

Table 2-3: Arc attributes in buildingEXODUS

Attribute	Default
Length	0.5 m
Obstacle	
Open space	0
Light debris	1
seat row (within row)	1
seat row (between rows)	7

The second attribute is known as the *Obstacle*. The *Obstacle* attribute is an integer measure of the degree of difficulty in passing over the node. *Obstacle* values can be used to differentiate node types. For example, nodes representing open space are linked with arcs which have an obstacle value of 0, while nodes littered with debris may have a higher obstacle value of 1 or 2. A geometry consisting of densely packed seats such as a cinema, theatre or aircraft could have the nodes within a seat row linked by arcs with an *Obstacle* value of 1. This slightly elevated value reflects the increased effort required to shuffle between seats compared to that required to traverse open space. The arcs linking nodes between seat rows may have an even higher *Obstacle* value of 7. This large value reflects the considerable difficulty in passing over an upright seat.

NOTE:
The obstacle attribute will affect the Travel Speed of an occupant traversing it.

As an occupant passes over a node or contemplates the possibility of doing so, information concerning the obstacle attribute are passed on to the BEHAVIOUR SUB-MODEL (see Section 7.2(3(iii))) and the OCCUPANT SUB-MODEL (see CHAPTER 3:). The obstacle information combined with the individual's personal attributes will influence their behaviour. Impassable obstacles such as walls, bulkheads, internal compartments etc., are formed by simply removing the arcs connecting the nodes across the obstacle.

2.5 Interpreting Nodes and Arcs

Each node within a geometry represents a portion of space whose size is dictated by the length of the arcs emanating from the node, and may only be occupied by one occupant at any one time. Thus each node can be considered as the space available to a single occupant. (This is true of all nodes except *Transit Nodes*, which represent components that occupy user-defined areas of space and which can be occupied by a specified number of people simultaneously). The length of the arcs between nodes reflects the distance between the centres of the connected nodes and by default this distance is 0.5m. Hence, each node represents a portion of space measuring 0.5m x 0.5m. Thus the nodes can be considered to tile the entire physical space of the geometry. However, in order to represent the connectivity of the grid, the size of each node is reduced in the graphical representation displayed on the computer screen. This concept is shown in Figure 2-5, where the node representation of an enclosure measuring 1m x 1m is shown. In this case, two nodes in each direction are used to represent the space, separated by 0.5m arcs. In reality the floor space would be completely covered by the four nodes. However, in the buildingEXODUS graphical representation of the space, the four nodes have the appearance of being reduced in size so that they do not touch each other.

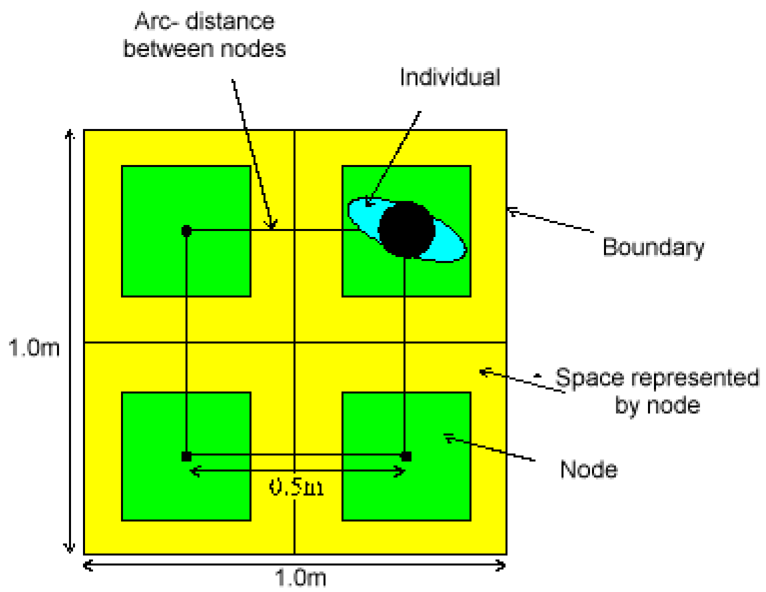

Figure 2-5: Node representation of a 1m x 1m enclosure

Furthermore, in the buildingEXODUS representation of the enclosure, a person placed in the enclosure would occupy a node centre and thus the centre of the person would be a distance of 0.25m away from his nearest two walls. In moving to an adjacent node, the person would travel a distance of 0.5m and again his/her body centre would be a distance of 0.25m away from the nearest walls. While seeming a little odd at first glance, in principle, this is in fact very close to reality as people have a finite thickness.

NOTE:
The space occupied by each occupant in buildingEXODUS is scaled to fit on a node.

NOTE:
If the default arc length of 0.5m is used, this restricts the packing density of people to a maximum of 4 people/m^2. Applications expected to result in greater packing densities cannot be simulated accurately under these conditions.

NOTE:
If a geometry is meshed with an arc length different to the default value, it is vital that key parameters such as travel speeds and conflict times are reassessed for their appropriateness (see Section 7.2 part 1 (ii) for further details).

2.6 buildingEXODUS Software Structure

As described in Section 1.1, buildingEXODUS comprises five core interacting sub-models, the OCCUPANT, MOVEMENT, BEHAVIOUR, TOXICITY and HAZARD sub-models. The software describing these sub-models utilises *rule-base* concepts, with the behaviour of each individual being determined by a set of heuristics or rules. This is very different to the majority of engineering software, which are generally made up of a collection of equations and formulations.

The rules governing the simulation have been categorised into these five sub-models. The software structure allows each of the rules that make up the sub-models to be easily modified. It is this flexible modular structure that enables EXODUS to be adapted for applications as different as a multi-storey school building and a B747.

CHAPTER 3: OCCUPANT SUB-MODEL

The Occupant sub-model defines each individual as a collection of attributes which broadly fall into four categories, physical (such as *Age, Gender, Agility* etc.), psychological (such as *Patience, Drive* etc.), experiential (such as *Distance, PET* etc.) and hazard effects (such as *FIN, FICO2, FIH* etc.). These attributes have the dual purpose of defining all occupants as individuals while allowing their progress through the enclosure to be tracked. Some of the attributes are fixed throughout the simulation while others are dynamic, changing as a result of inputs from the other sub-models. Furthermore, some of the attributes require the user to either manually set defining values or accept the default settings while other attributes are calculated during the simulation. (It should be noted that in this section all of those located within the structure are assumed to be represented by the Occupant sub-model, irrespective of whether they are staff or fee-paying occupants).

A population can be defined using one of several methods. Populations are defined in *POPULATION MODE* and can be created by,

(i) Manually setting the parameters of each individual (see the User Guide, Section 4.2)
(ii) Creating a block of *n* individuals with randomly generated parameters (see the User Guide, Section 4.2)
(iii) Creating a block of *n* individuals whose parameters are randomly distributed between user defined upper and lower bounds (see the User Guide, Section 4.3)
(iv) A combination of techniques (i) to (iii)
(v) Selection from a user defined population library (see the User Guide, Sections 4.3 and 4.4)

Once a population is created, individuals within the population may be edited in order to modify particular attributes.

NOTE:
Where attribute default values are given in this section, the values used refer to the default person (see User Guide, Section 4.2.1 part (a)).

NOTE:
When the attributes of an individual occupant are altered, the alterations do not feed through to other relevant values as in the Population Panel. Therefore, special attention should be paid when adjusting the attributes of individual occupants.

In the following section each of the occupant attributes will be described and each attribute's properties will be summarised. This comprises of the following descriptors:

Range: Describes the allowable range of values for the attribute.
Default: Describes the default value assigned to the attribute.
Influenced by: Describes which other attributes or properties exert a direct influence on the attribute in question.
Influences: Describes which other attributes or properties are directly affected by the attribute in question.
Used in level: Describes which software level the attribute exerts an influence in.

3.1 Physical Attributes

(1) *Gender*, *Age*, *Height* and *Weight* attributes

These attributes are used to assist in distinguishing one individual from another and in providing a rationale for assigning various attributes. This type of distinction is necessary because, on average, values for the defining characteristics are dependent on age, gender and weight.

Attribute	: *Gender*.
Range	: Male or Female.
Default	: Male.
Influenced by	: None.
Influences	: Agility, Drive, Travel Speed, Height and Weight.
Used in level	: A, B and C
Note	: Used to represent the gender of each occupant.

Attribute	: *Age.*
Range	: 1 - 100 years.
Default	: 25 years.
Influenced by	: None.
Influences	: Travel Speed, Height and Weight
Used in level	: A, B and C
Note	: Used to represent the age of each occupant.

Attribute	: *Weight.*
Range	: 1.0 - 200 KGs.
Default	: 80 KGs.
Influenced by	: Age and Gender.
Influences	: None
Used in level	: A, B and C
Note	: Used to represent the weight of each occupant. This variable does not influence any other agent parameter. Users may use this attribute as a guide to setting other variables e.g. users may wish to correlate walking speed with weight.

Attribute	: *Height.*
Range	: 1.0 – 2.0 m.
Default	: 1.8 m.
Influenced by	: Age and Gender.
Influences	: *FIN, FIH, FIC, FLD*
Used in level	: C
Note	: Used to represent the height of each occupant. This attribute can be used when fire hazards are defined, in particular when a hazard description is imported from the CFAST model (see the User Guide, Chapter 5).

Attribute : *Leader*.
Range : Yes or No.
Default : No
Influenced by : None
Influences : Communication
Used in level : A, B and C
Note : Used to represent the seniority of an occupant's role within a social group. This influences the importance afforded to any information shared within the group, with information provided by a *Leader* given priority.

Through the POPULATION PANEL (see the User Guide, Section 4.3), the user has complete control over the distribution used to represent the *Height* and *Weight* attributes. When generating populations using the RANDOM generate option (see the User Guide, Section 4.2), the *Height* and *Weight* attribute of the generated population (see the User Guide, Chapter 4) is dependent upon the *Age* and *Gender* of the occupant. These relationships are extracted from British medical statistics relating to the ideal height and weight of people up to adulthood [59]. These figures are then a conservative estimate of what may be expected within an actual population.

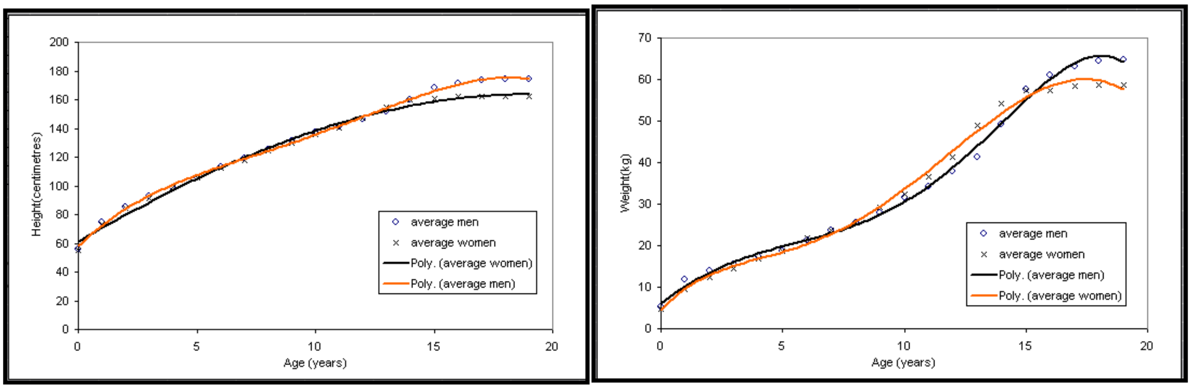

Figure 3-1: Functions relating Weight and Height to Age (only the averages are shown).

The functions (see Figure 3-1) concern data relating to occupants between the ages of 0 - 20 years (adults). The functions include distributions about the average height and weight for each age, so that at a specific age the height and weight of an occupant falls within a range of values, centred around the average shown in Figure 3-1. Occupants older than 20 years are treated as being of adult proportions. This data should be treated as indicative of the expected height and weight of an individual rather than as an accurate representation. The user should provide additional information if this factor is considered important and should not solely rely on the data within the model.

(2) *Mobility* attribute.
The *Mobility* attribute is a multiplicative factor used in conjunction with the *Travel Speed* attribute (see Section 3). It has two functions: firstly it is intended to allow the introduction of physical disability into the occupant description. An occupant not suffering from any disability will have an initial *Mobility* of 1.0, while an occupant with a minor disability, such as an arm in plaster, will have a slightly reduced *Mobility* value of for example 0.9. A major disability, such as blindness or a broken leg, will result in a considerable reduction to say 0.2.

The second function of the *Mobility* attribute is to reduce the occupants' *Travel Speed* in response to their growing exposure to the narcotic agents, irritant agents and smoke concentration (either irritant or non-irritant smoke). These capabilities are only available in Level C of the EXODUS software. The *Mobility* may vary from its initial value (no detrimental effects), to zero (individual has expired). The *Mobility* decreases as:

1) FIN - determined by the toxicity sub-model - increases, and/or
2) FIC increases, and/or
3) Smoke concentration increases (either irritant or non-irritant smoke).

The effects of the irritant gases, narcotics and smoke on an agent are calculated individually. Each hazard (i.e. narcotic agents, irritant agents and smoke concentration) produces its own *Mobility Degradation Factor* defining the effect of the respective hazard on the agent's *Mobility*. In the case of smoke, the method of calculating its effect on the agent is dependent upon whether the *Irritant* model is enabled or disabled. If the *Irritant* model is enabled, then irritant gases are explicitly represented in the fire hazard, and hence are assumed not to be present in the smoke (i.e. the smoke is assumed non irritant). Conversely, if the *Irritant* model is disabled, then irritant gases are NOT explicitly represented in the fire hazard, and hence are assumed to be present in the smoke (i.e. the smoke is assumed irritant). The most severe impact upon the individual's *Mobility* and health resulting from exposure to irritants, narcotics and smoke is then adopted (i.e. the smallest *Mobility Degradation Factor* is then used in the calculation of *Mobility* and hence also *Travel Speed*).

If the *Irritant* model is enabled (i.e. if irritants are explicitly represented) then the *Mobility Degradation Factor* is calculated as follows:

$$\textit{Mobility Degradation Factor} = \mathrm{Min}(\textit{Mobility Degradation Factor}_{FIN},\ \textit{Mobility Degradation Factor}_{FIC},\ \textit{Mobility Degradation Factor}_{NonIrritantSmoke})$$

If however the *Irritant* model is disabled (i.e. if irritants are not explicitly represented) then the *Mobility Degradation Factor* is calculated as follows:

$$\textit{Mobility Degradation Factor} = \mathrm{Min}(\textit{Mobility Degradation Factor}_{FIN},\ \textit{Mobility Degradation Factor}_{IrritantSmoke})$$

This *Mobility Degradation Factor* is then multiplied by the occupant's *Initial Mobility* level to form a dynamically calculated *Mobility*, dependent upon the surrounding conditions, reflected in

$$Mobility\ =\ Initial\ Mobility\ *\ Mobility\ Degradation\ Factor \tag{1}$$

Therefore, an occupant with an *Initial Mobility* level set to less than 1.0 will still have their *Mobility* level reduced in proportion to the surrounding environmental conditions.

The calculation of the *Mobility Degradation Factor* resulting from exposure to narcotic agents (FIN), irritant agents (FIC) and smoke concentration will now be discussed:

i) Relationship between FIN and occupant Mobility Degradation Factor

The link between the degradation in occupant attributes and increasing exposure to narcotic gases (FIN) has not been researched thoroughly. As a result, guidelines to the suggested settings cannot be found in the literature. It is however expected that significant degradation only occurs at extremely high values of FIN. The default values in Table 3-1 reflect this expected link.

Table 3-1: Relationship between FIN and occupant Mobility Degradation Factor$_{FIN}$

FIN	Mobility Degradation Factor$_{FIN}$
0.00 - 0.89	1.00
0.90 - 0.95	0.90
0.96 - 1.00	0.80

NOTE:
As not much is known for certain concerning the linkage between FIN and Mobility, users have the option of activating or deactivating this link via the Feedback option within the Hazard Control dialogue box. If the link is deactivated, individuals exposed to narcotic gases will remain fully mobile until incapacitation is predicted.

ii) Relationship between FIC and occupant Mobility Degradation Factor

In a similar manner to the narcotic gases, the irritant gases also directly affect the *Mobility* of the individual, as well as their well-being. It should be emphasised that only the instantaneous (FIC) impact of the irritant products influence the *Mobility* of an individual. As the combined FIC value increases, so the *Mobility* of the individual decreases, reducing the *Travel Speed* of the individual (it should be borne in mind that within buildingEXODUS the *Mobility* attribute of an individual is a coefficient of their travel speed). Equation 2 below (adapted from the Weibull function provided by Purser [87]) represents the function that approximates the sigmoidal function produced by Purser [87].

$$Mobility\ Degradation\ Factor_{FIC} = \frac{\left(e^{-((FIC*1000)/160)^2} + (-0.2*FIC+0.2)\right)}{1.2} \quad \textbf{(2)}$$

The relationship between the FIC level and the *Mobility Degradation Factor$_{FIC}$* (and hence the individual's *Mobility*) is shown in Figure 3-2.

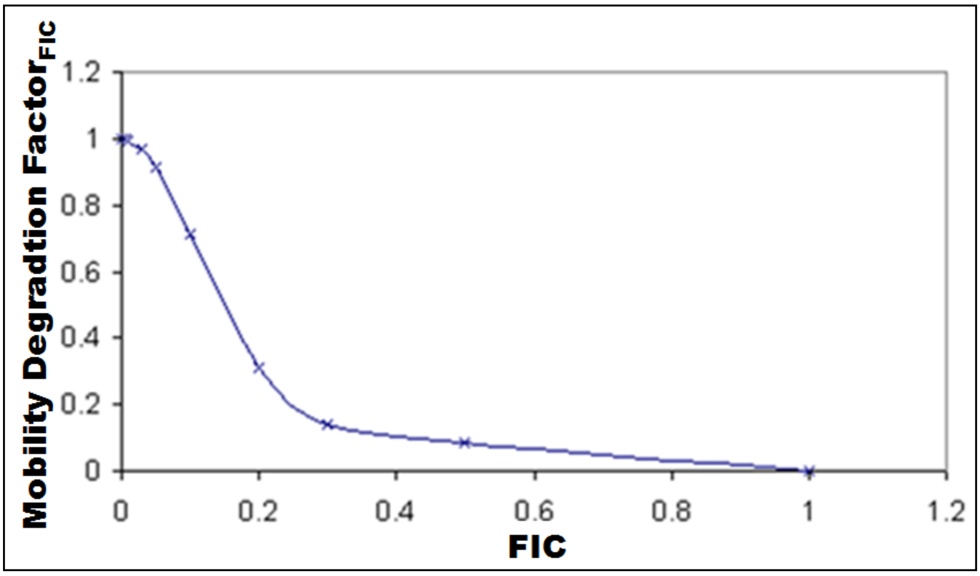

Figure 3-2: Impact of FIC upon the Mobility Degradation Factor$_{FIC}$.

iii) Relationship between Smoke and occupant Mobility Degradation Factor

Smoke has the effect of obscuring vision and irritating the eyes thus impairing the ability of an individual to escape. Several studies [17,18] have suggested that a victim's movement rate decreases as the smoke concentration increases. This effect is thought to be concentration related and does not increase with prolonged exposure. Within EXODUS, the smoke density is linked to the *Mobility* attribute.

NOTE:

In addition to affecting an occupant's Travel Speed, the Smoke density may also exert an influence on the occupant's navigation efficiency (see Section 7.2(2(iii))).

The impact of smoke upon the individual's *Mobility* is related to the representation of irritants within the simulation. If the irritant gases are not explicitly represented in the fire hazard (i.e. if the *Irritant* model is disabled), the Jin 'irritant' data-set is used to describe the complete impact of the smoke and irritant gases on the movement rates of exposed individuals. This does not require the specification of irritant gas concentrations. The applied relationship is intended to approximate the reduction in the individuals travel speed due to the impact of irritant smoke (including the obscuration effect of smoke). However, it does not directly impact the well-being of the exposed individuals.

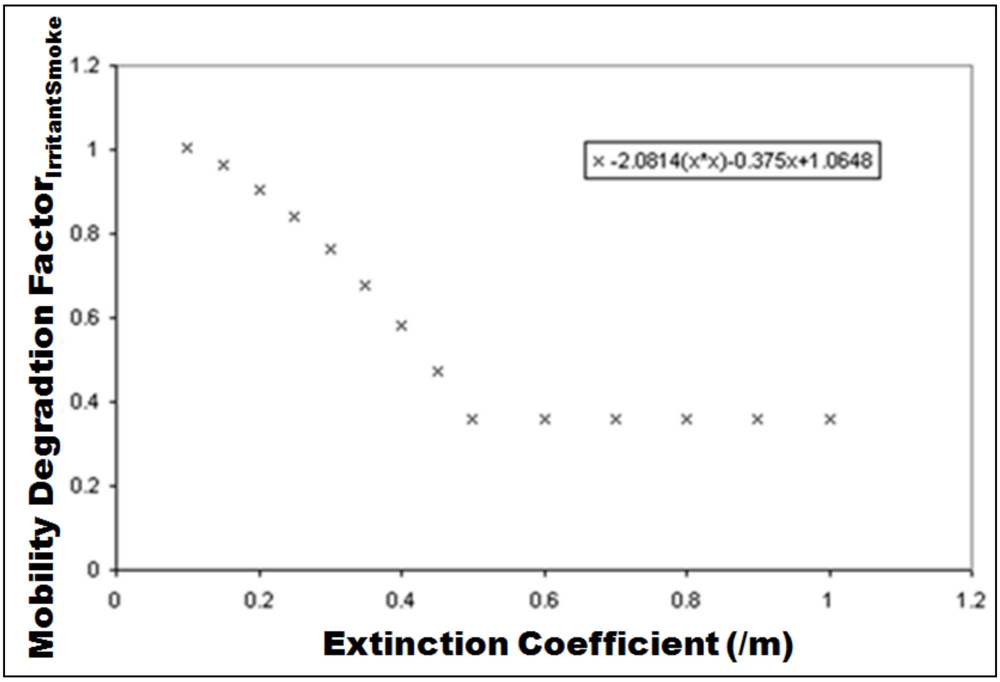

Figure 3-3: The impact of smoke upon the Mobility Degradation Factor$_{\mathbf{IrritantSmoke}}$ derived from the work of Jin [17,18].

The *Mobility Degradation Factor$_{IrritantSmoke}$* (and hence the *Mobility* attribute) is kept constant up to smoke concentrations of 0.1 /m after which point it is calculated according to the following function,

$$Mobility\ Degradation\ Factor_{IrritantSmoke} = -2.0814\ K^2 - 0.375K + 1.0648 \qquad \textbf{(3)}$$

where *K* represents the extinction coefficient (/m) of the smoke. This function has been derived from the work of Jin (see Figure 3-3) [17,18]. For a smoke concentration of 0.45 /m, the above formulation decreases the *Mobility Degradation Factor$_{IrritantSmoke}$* to approximately half of its original value. For smoke concentrations above 0.5 /m (i.e. *Mobility Degradation Factor$_{IrritantSmoke}$* values less than 0.36) occupant escape abilities are severely limited and the model assumes a maximum *Travel Speed* equivalent to the *Crawl Rate* (see Table 3-2) rather than establishing the *Travel Speed* according to the occupant *Mobility*. Users should note that the default *Crawl Rate* is set to 20% of the occupant's *Fast Walk* speed (see Table 3-3). However, if the *Crawl* option is disabled (in the *BEHAVIOUR CONTROL* dialogue box), and the smoke concentration exceeds 0.5 /m, the occupant's *Mobility Degradation Factor$_{IrritantSmoke}$* will be maintained at 0.36, regardless of the smoke concentration.

NOTE:
Occupants moving through smoke must not be allowed to improve their travel speeds when forced to crawl. The default settings for travel speeds produces a crawl speed that is 20% of their original Fast Walk travel speed and so this situation cannot arise with the default settings. Users should exercise care to ensure that this relationship is maintained when the default speed settings are changed.

Table 3-2: Relationship between smoke concentration and Mobility Degradation Factor

Smoke Concentration (1/m)	Mobility Degradation Factor$_{IrritantSmoke}$
0.0 - 0.1	1.0
0.2	0.92
0.3	0.76
0.4	0.57
0.5	0.36
>0.5	0.36*

*if enabled occupant is forced to *Crawl* and the *Travel Speed* will be dependent upon the occupant's *Crawl Rate* rather than upon the *Fast Walk* rate and the *Mobility*

If irritant gas concentrations are specified (i.e. the *Irritant* model is enabled) then a more comprehensive model is utilised in relation to the impact of smoke obscuration upon the *Mobility* of the individual (see Figure 6-5). This examines the concentration of several irritants and determines the impact upon the individual accordingly. Initially the Jin data relating to experiments involving 'non-irritant' gases is used. This is assumed to represent the impact of the visual obscuration of the smoke alone, without representing any of the irritant effects of the smoke present. This produces a slight decrease in travel speed resulting from the obscuration affects of the smoke at sufficiently high levels (see Figure 6-6).

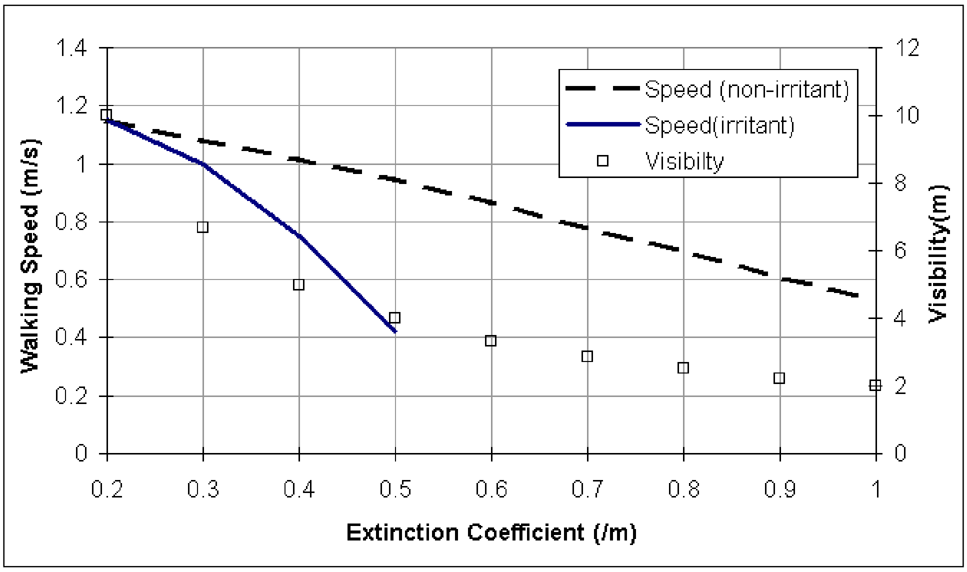

Figure 3-4: The walking speed of volunteers through irritant (solid line) and non-irritant (dashed line) smoke. The squares indicate the approximate visibility levels, purely according to the extinction coefficient of the smoke (redrawn from the original [10]).

The derived curve which represents this phenomenon is

$$Mobility\ Degradation\ Factor_{NonIrritantSmoke} = -0.161K^2 - 0.488K + 1.105 \qquad \textbf{(4)}$$

where K is the extinction coefficient of the *Smoke* and is valid between extinction coefficients of 0.2/m and 1.0/m. This curve represents the impact of the environment, specifically in relation to reduced visibility.

NOTE:
Users have the option of activating or deactivating the linkage between smoke concentration and Mobility. If the link is deactivated, individuals exposed to smoke will remain fully mobile and experience no degradation in Travel Speed.

Attribute	: *Mobility*.
Range	: 0.0 - 1.0.
Default	: 1.0.
Influenced by	: FIN, FIC and Smoke Concentration (Level C only)
Influences	: Travel speed.
Used in level	: A, B and C
Note	: Used to represent movement impairing disability and the impact of fire hazards.

(3) *Travel Speed* attribute
The *Travel Speed* attribute reflects the occupant's current travel speed. This is dependent on the occupant's initial maximum *Travel Speed*, their *Mobility* and the terrain being travelled over.

Within buildingEXODUS the user sets - for each occupant - six levels of travel speed. These travel speeds are identified as *Fast Walk*, *Walk, Leap, Crawl, Stairs-Up* and *Stairs-Down*. These represent the maximum unhindered speed the occupant can attain under a variety of conditions. Initially each occupant is assigned a maximum *Fast Walk* speed and the *Walk, Leap* and *Crawl* rates are automatically determined as an arbitrary percentage of this speed (see the User Guide, Section 4.2.1). The arbitrary percentage breakdown is set as default values that can be altered by the user (see Table 3-3). The *Stairs-Up* and *Stairs-Down* rates are dependent on age and gender and are determined from data generated by Fruin [19] (see Table 3-4). The *Escalators-Up* and *Escalators-Down* data is derived from UK empirical data collected by FSEG staff [124] (see Table 3-5). It should be noted that the escalator travel speeds shown Table 3-5 are average values, representing the range of values employed within the model.

NOTE:
A person's Travel Speed on stairs is also dependent on the geometry of the stair, particularly the depth and height of the riser. buildingEXODUS is not sensitive to this dependency. The default stair travel speeds used in buildingEXODUS are based upon averaged data for the two staircase configurations studied by Fruin [19].

NOTE:
The maximum Travel Speed that can be set for an occupant is dependent on the minimum arc length. The time taken for an occupant to traverse any arc must ≥ 1/12 of a second.

NOTE:
In addition to the factors highlighted, the orientation of the structure will also have an impact upon the Travel Speeds attained, according to the behavioural flags enabled and the data-set used.

NOTE:
The Travel Speed can be modified during the simulation as part of the social grouping process (i.e. in an attempt to maintain the proximity of group members) given that the Social Movement flag has been enabled in the Behavioural Options dialogue box in Simulation Mode.

Table 3-3: The default Walk, Leap and Crawl travel speeds are set as an ARBITRARY percentage of Fast Walk

Travel Speed Type	Default Speed m/s
FAST WALK	1.5
WALK	1.5 * 90%
LEAP	1.5 * 80%
CRAWL	1.5 * 20%

Table 3-4: Default stair travel rate as derived from Fruin [19]

Gender	Age (years)	Down avg (m/s)	Up avg (m/s)
Male	<30	1.01	0.67
Female	<30	0.755	0.635
Male	30 - 50	0.86	0.63
Female	30 - 50	0.665	0.59
Male	>50	0.67	0.51
Female	>50	0.595	0.485

NOTE:
The default stair speed values are an indication of ***MAXIMUM*** *recommended values.*

Table 3-5: Default escalator travel rate. Empirical data collected by FSEG staff at a London Underground station.[124]

Gender	Age (years)	Down avg (m/s)	Up avg (m/s)
Male	<30	0.84	0.72
Female	<30	0.75	0.67
Male	30 - 50	0.84	0.72
Female	30 - 50	0.75	0.67
Male	>50	0.84	0.72
Female	>50	0.75	0.67

The terrain and atmospheric conditions govern which one of the six travel speed rates is appropriate. This is determined according to Table 3-6.

Table 3-6: Terrain-travel speed dependence

Terrain	Travel Speed
Free-Space - Free-Space	*Fast walk*
Aisle – Aisle	*Fast walk*
Free-Space-Direction/Redirection/Source- Free-Space	*Fast walk*
Free-Space – Seat/Seat - Free-Space	*Walk*
Seat – Seat (within row)	*Walk*
Seat – Seat (between rows)	*Leap*
X – Boundary/Boundary – X	*Walk*
X – Discharge-node/ Discharge –node - X	*Walk*
X – Attractor-node/ Attractor –node - X	*Walk*
X- Landing/Landing-Landing/Landing-X	*Walk*
X – External Exit	*Walk*
Stair	*Stairs-up/ Stairs-down*
Escalator	*Escalator-up/ Escalator-down*
Free-Space-Census /Census- Free-Space	*Fast walk*
Free-Space-Internal Exit/Internal Exits- Free-Space	*Fast walk*
Free-Space-Transit Node- Free-Space	*Depends on the terrain*

NOTE:
The Crawl speed is used only when crawling is enabled within the model and the environmental conditions to which the individual is exposed exceeds defined thresholds (see item 2 in this Section, CHAPTER 5: and Section 7.2(2)). Users should also refer to the User Guide, Chapter 6 for a detailed explanation of the environmental conditions necessary to induce crawling behaviour.

NOTE:
The occupant's travel speed is altered when traversing arcs that have an associated obstacle value greater than zero

The final contribution to the travel speed is derived from the occupant's *Mobility* attribute. Reductions in *Mobility* cause a reduction in *Travel Speed*. An occupant's *Travel Speed* at any point in time is determined by the relation:

$$Travel\ Speed\ =\ Initial\ Travel\ Speed\ *\ Mobility \qquad \textbf{(5)}$$

When caught in a crowd or queue, EXODUS regulates an individual's *Travel Speed* through the *Conflict Resolution* procedure in which a time penalty is added to each occupant involved in a conflict (see CHAPTER 7:). Thus, even if an individual is assigned a high *Fast Walk* rate, if the individual is caught in a slowly moving crowd or exit queue, he will automatically move with a reduced speed. The reduced speed is intended to represent the speed of the crowd or queue resulting from the congestion.

NOTE:
EXODUS DOES NOT make use of well known congestion related travel speed correlations from for example the observations of Predtechenskii and Milinskii [51] or the Level of Service approach of Fruin [19].

Attribute	: *Travel speed.*
Range	: 0.0 - 10.0 m/s.
Default	: Fast walk 1.5 m/s.
	Walk 90% of Fast walk
	Leap 80% of Fast walk
	Crawl 20% of Fast walk
	Stairs-up See Table 3-4
	Stairs-down See Table 3-4

Influenced by	: *Mobility*, Terrain, Atmospheric conditions.
Influences	: Performance and behaviour.
Used in level	: A, B and C
Note	: Used to represent the maximum travel speed of an occupant.

(4) *Agility* Attribute.
Agility is intended to represent the physical prowess of the individual in tackling obstacles such as movement over seat backs.

The *Agility* of an agent is (like their *Mobility*, see Section 2) reduced in response to their exposure to the narcotic gases, irritant gases and irritant smoke concentration. These capabilities are only available in Level C of the EXODUS software. The *Agility* may vary from its initial value (no detrimental effects), to zero (individual has expired). The *Agility* decreases as:

1) FIN - determined by the toxicity sub-model - increases, and/or
2) FIC increases, and/or
3) Irritant smoke concentration increases.

The effects of the irritants gases, narcotics gases and irritant smoke on an agent are calculated individually. Each hazard (i.e. narcotic gases, irritant gases and irritant smoke concentration) produces its own *Mobility Degradation Factor* defining the effect of the respective hazard on the agent's *Mobility*. The calculation of the *Mobility Degradation Factor* for each hazard is outlined in the *Mobility* section above (i.e. Section 2). The most severe impact upon the individual's *Mobility* and health resulting from exposure to irritant gases, narcotic gases and irritant smoke is then adopted (i.e. the smallest *Mobility Degradation Factor* is then used in the calculation of *Agility*).

If the *Irritant* model is enabled (i.e. if irritants are explicitly represented) then the *Agility Degradation Factor* is calculated as follows:

$$\textit{Agility Degradation Factor} = \mathrm{Min}(\textit{Mobility Degradation Factor}_{FIN},\ \textit{Mobility Degradation Factor}_{FIC})$$

If however the *Irritant* model is disabled (i.e. if irritants are not explicitly represented) then the *Agility Degradation Factor* is calculated as follows:

$$\textit{Agility Degradation Factor} = \mathrm{Min}(\textit{Mobility Degradation Factor}_{FIN},\ \textit{Mobility Degradation Factor}_{IrritantSmoke})$$

This *Agility Degradation Factor* is then multiplied by the occupant's *Initial Agility* level to form a dynamically calculated *Agility*, dependent upon the surrounding conditions, reflected in

$$Agility\ =\ Initial\ Agility\ *\ Agility\ Degradation\ Factor \qquad \textbf{(6)}$$

Therefore, an occupant with an *Initial Agility* level set to 6.0 will have their *Agility* level reduced in proportion to the surrounding environmental conditions.

It is important to note that the effect of irritant gases, narcotic gases and smoke on an agent's *Agility* is similar to their effect on agent *Mobility* (see Section 2). However, while non irritant smoke affects an agent's *Mobility* it has no corresponding effect on their *Agility*. This is because non irritant smoke only causes visual obscuration. While this obscuration may affect an agents ability to move within a structure (i.e. their *Mobility* and *Travel Speed*) it is felt unlikely to affect their ability to negotiate obstacles (i.e. their *Agility*).

NOTE:
In versions of buildingEXODUS prior to v6.0 an agent's exposure to non-irritant smoke would also have reduced their corresponding Agility value, meaning that their ability to negotiate obstacles would previously have been reduced.

The *Initial Agility* distribution should be set in conjunction with the arc obstacle values. If the default arc obstacle values are not used, the occupant *Initial Agility* distribution should be reconsidered. Failure to do this may result in agents being unable to move between regions within the structure (i.e. traverse arcs) and hence becoming trapped within the geometry.

Attribute	: *Agility*.
Range	: 0.0 - 7.0
Default	: 5.
Influenced by	: FIN, FIC and Smoke Concentration (Level C only)
Influences	: Performance and behaviour.
Used in level	: A, B and C
Note	: Used to represent movement impairing disability and impact of fire hazards on exposed occupants. Factors such as an occupant's age, gender, weight and degree of mobility can be taken into consideration when assigning the *Agility* attribute.

(5) *Respiratory Minute Volume (RMV)* Attribute.
The volume of air breathed per minute (or minute volume) is a measure of the volume of air taken into the lungs (litres/min). It is used by the TOXICITY SUB-MODEL (see CHAPTER 6:) to calculate the *FICO* (Carbon Monoxide dose). The *RMV* is typically dependent on *Gender*, *Weight*, *Age* and type of activity the individual is involved in. For example a 70kg male involved in light work has an *RMV* of about 25 l/min, while at rest, it falls to 8.5 l/min and while involved in heavy work it increases to 50 l/min [20]. The default values used by EXODUS are depicted in Table 3-7. In the current implementation of EXODUS the *RMV* shows only a dependence on *Gender* and activity.

NOTE:
The values for female RMV are arbitrarily set at 90% of the values for males. When more definitive data becomes available these will be modified.

Table 3-7: Default occupant RMV values used by EXODUS

Activity	Males (l/min)	Females (l/min)	Work Load
Rest	8.5	7.65	Occupant is static
Light	25.0	22.5	Occupant is using WALK SPEED, or STAIR DOWN SPEED or is moving laterally on staircases
Heavy	50.0	45.0	Occupant is using FAST WALK SPEED, or LEAP SPEED, STAIR UP SPEED, is travelling across an arc with an obstacle greater than 1.0 or is CRAWLING.

NOTE:
It is assumed that an occupant will be crawling only if they have encountered severe conditions.

Attribute : *RMV*.
Range : 0.0 - 50 l/min.
Default : See Table 3-7
Influenced by : Activity, *Gender*.
Influences : FICO, FIN, *Mobility*, *Agility*, *Travel Speed*, Performance and behaviour.
Used in level : C
Note : Used in calculation of carbon monoxide up-take and in prediction of incapacitation. Influences FIN calculation, which in turn affects the calculation of the *Mobility Degradation Factor*, and hence the calculation of *Mobility*. Similarly, its influence on the FIN calculation in turn affects the calculation of the *Agility Degradation Factor*, and hence also the calculation of *Agility*.

3.2 Psychological Attributes

(1) *Drive* Attribute.
The *Drive* attribute is a measure of the assertiveness of an occupant. It is used as a basis for conflict resolution (see 7.2 Part 1(ii)). In situations where occupants compete to occupy a node, the occupant's *Drive* will resolve possible conflicts. The *Drive* attribute is assigned values from 1 (low drive) to 15 (high drive). There is some evidence to support the belief that young males generally have the highest drive while older females tend to have the lowest drive. This is suggested from competitive evacuation trials involving financial payments to evacuees who are amongst the first few to successfully evacuate [21].

Attribute : *Drive.*
Range : 1 - 15.
Default : 10.0
Influenced by : None
Influences : Conflict resolution, Performance and behaviour.
Used in level : A, B and C
Note : Used in conflict resolution by BEHAVIOUR Sub-model. Factors such as an occupant's *Age*, *Gender* and social affiliation can be taken into consideration when assigning the *Drive* attribute.

(2) *Patience* Attribute.
During a simulation, various occupants will be forced to remain stationary while in a queue or while attempting to join a flow of moving people. The *Patience* attribute is a measure of the amount of time (in seconds) that an occupant is prepared to wait before considering or attempting an alternative action, such as moving away from the queue, attempting to jump over a row of seats, etc. (see Section 7.2). Compliant occupants will have a long *Patience* while less compliant occupants will have a relatively short *Patience* attribute. In simulations where occupants are expected to wait their turn and not display any "extreme behaviour" (see CHAPTER 7:) the *Patience* attribute should be set to extremely large values such as 1000 seconds.

Attribute : *Patience.*
Range : 1 - 1000 seconds.
Default : 1000 seconds.
Influenced by : None.
Influences : Performance and behaviour.
Used in level : A, B and C
Note : Used by Behaviour Sub-model.

NOTE:
The occupant Patience level can be overridden by enabling the Impatient flag, located on the BEHAVIOUR OPTIONS dialogue box. (see the User Guide, Chapter 6) . This guarantees that the occupant will exhibit Extreme Behaviour, if the occupant waits for any amount of time.

(3) *Response Time* Attribute.
The *Response Time* is intended to be a measure of the pre-evacuation movement time incurred by the occupant. It represents the difference in time between the time the occupant begins to actively evacuate and the time at which the call to evacuate was issued. Occupants displaying behavioural in-action are simulated by using an extremely long *Response Time*, for example 10000 seconds.

The appropriate *Response Time* distribution to impose on a simulation is dependent on a number of scenario related factors including: the nature of the activity the occupants are involved in, the nature of the occupied structure space, the type of alarm system employed, the quality of the management structure employed in the structure, the level of training the target population are exposed to etc. Typical response times can vary from seconds (occupants are awake, trained, familiar with the structure, alarm systems and procedures) to many minutes (situations where occupants may require assistance such as in the medical office.) [53]. *Response Time*s can be individually specified or specified according to a range or a pre-determined distribution.

Attribute : *Response Time.*
Range : 0 - 10000 seconds.
Default : 0 seconds.
Influenced by : None.
Influences : Performance and behaviour.
Used in level : A, B and C
Note : Used by Behaviour Sub-model.

NOTE:
The Response Time can be overridden by enabling the Specified Response option in the BEHAVIOUR OPTIONS dialogue box (see the User Guide, Chapter 6). This guarantees that the entire population will only start to evacuate after a user defined time within the simulation. This does not replace the Response Time attribute of the individuals. Instead these are not considered and the population will only start to evacuate when the specified time is reached within the simulation. If then disabled, the population will then respond according to their individual Response Times. In addition, the Response times of individuals may be overridden locally by the use of Response Zones (see the User Guide, Chapter 5) or according to the environmental conditions around them (see the User Guide, Chapter 6). Response Times may also be shared through the communication process, if the Social Response or Social Movement flags have been enabled in the Behavioural Options dialogue box in Simulation Mode.

(4) *Gene* Attribute
The *Gene* attribute is used when determining the 'identity' of the individual; specifically in relation the ability of occupant's to communicate information.

Attribute	: *Gene.*
Range	: Greater than 0.
Default	: 0.
Influenced by	: None.
Influences	: Performance and behaviour.
Used in level	: A, B and C
Note	: Used by Behaviour Sub-model.

NOTE:
A Gene value greater than zero is used to represent related groups. A gene value of zero indicates that this person is not related to any other member within the structure and therefore will not communicate any information.

3.3 Experiential Attributes

(1) *Personal Elapsed Time (PET)* Attribute.
The *PET (Personal Elapsed Time)* is a dynamic attribute that is calculated continuously by EXODUS for each occupant. It is a measure of the time spent by the occupant in the evacuation. If the occupant is interrogated at any time, the *PET* indicates the amount of time the occupant has spent in getting to his/her current location. At the end of the simulation, it measures the time to exit or incapacitation. *PET* is initially set to zero.

Attribute	: *PET.*
Range	: 0 - indefinite seconds.
Default	: 0 seconds.
Influenced by	: None.
Influences	: None.
Used in level	: A, B and C
Note	: The PET is a measure of time. At the end of the simulation it indicates the time to evacuate or to incapacitation.

(2) *Distance Travelled Attribute (Dist. Trav.)*
Distance Travelled is a dynamic attribute that is calculated continuously by EXODUS for each occupant. The *Distance Travelled* attribute is a measure of the total distance (measured in metres) travelled by the occupant at any point in time. At the end of the simulation, it measures the distance travelled to an exit point, or the point of incapacitation. The *Distance Travelled* attribute is initially set to zero.

Attribute	: *Distance Travelled (Dist Trav.)*
Range	: 0 - indefinite metres.
Default	: 0 metres.
Influenced by	: None.
Influences	: None.
Used in level	: A, B and C
Note	: The *Distance Travelled* attribute is a measure of the distance travelled.

(3) *Distance Remaining* Attribute (*Dist. Rem.*)
The *Distance Remaining* is a dynamic attribute that is calculated continuously by EXODUS for each occupant. If the occupant has no tasks to perform prior to exiting then the figure shown represents the remaining distance (measured in metres) that the occupant must travel in order to reach their exit point. Otherwise, the figure shown will be the distance to the location of their next task. Once the occupant's tasks are complete the figure will then represent the distance to the occupant's exit point. If the occupant successfully evacuates from the enclosure, at the end of the simulation, it is zero. If the occupant is incapacitated the *Distance Remaining* attribute specifies how close to an exit point he/she was at the time of incapacitation. The *Distance Remaining* attribute is initially set to zero.

Attribute	: *Distance Remaining. (Dist. Rem.)*
Range	: 0 - indefinite metres.
Default	: 0 metres.
Influenced by	: None.
Influences	: None.
Used in level	: A, B and C
Note	: The *Distance Remaining* attribute is a measure of the remaining distance (measured in metres) that the occupant must travel in order to reach an exit.

(4) *Wait* Attribute.
The *Wait* attribute is a dynamic attribute that is calculated by EXODUS for each stationary occupant. The *Wait* attribute is a measure of the time (in seconds) that an occupant remains stationary after he/she has started to evacuate. Once the waiting occupant begins to move the *Wait* counter is reset to zero. The *Wait* counter is continuously compared with the *Patience* attribute by the BEHAVIOUR SUB-MODEL (see CHAPTER 7:). The *Wait* attribute is initially set to zero.

Attribute : *Wait.*
Range : 0 - indefinite seconds.
Default : 0 seconds.
Influenced by : None.
Influences : Behaviour and performance.
Used in level : A, B and C
Note : The *Wait* attribute measures the time an occupant remains stationary. This is reset to zero every time the occupant moves.

(5) *Cumulative Wait Time (CWT)* Attribute.
The *CWT* attribute is a dynamic attribute that is calculated by EXODUS for each occupant. The *CWT* is a measure of the total time (in seconds) that an occupant remains stationary after he/she has started to evacuate. It represents a summation of all the *Wait* attributes incurred by an occupant. The *CWT* attribute is initially set to zero.

Attribute : *CWT.*
Range : 0 - indefinite seconds.
Default : 0 seconds.
Influenced by : None.
Influences : None.
Used in level : A, B and C
Note : The CWT attribute measures the *total* time an occupant remains stationary

(6) *Target Door* Attribute.
The *Target Door* attribute allows an occupant to be directed to a specific external door, regardless of the potential map. The values this attribute may take are the names of the exits in the geometry. The default value is "Nearest Door", i.e. follow the potential map. It should be noted that should an occupant exhibit *Extreme* behaviour whilst en-route to a target exit, the occupant ignores the *Target* and follows the Potential map. However, this behaviour may be overridden by enabling the *Maintain Target Exit* option on the *BEHAVIOUR CONTROL* dialogue box (see the User Guide, Section 6.2.1).

Attribute : *Target Door.*
Range : Nearest Door or List of Doors.
Default : Nearest Door.
Influenced by : Local Behaviour.
Influences : Behaviour and Performance.
Used in level : A, B and C
Note : The *Target Door* attribute allows an occupant to be directed to a specific exit.

NOTE:
Target Doors can be shared as part of the communication process if the Social Response or Social Movement flags have been enabled in the Behavioural Options dialogue box in Simulation Mode.

(7) *Occupant Exit Knowledge (OEK)*
The *Occupant Exit Knowledge (OEK)* allows an occupant to possess a localised understanding of the structure, in the form a list of external exits. The *OEK* is activated if the user selects the *Local Familiarity* option on the *BEHAVIOUR OPTION* dialogue box accessible in Simulation mode. The *OEK* is a user-definable attribute or the user can elect to have a list automatically determined. An occupant cannot have zero entries in the *OEK.* If no exits are specified by the user, the occupant is given knowledge of the nearest exit by default.

Range	: Nearest Exit to All Available Exits.
Default	: Nearest exit
Influenced by	: None.
Influences	: Behaviour and Performance.
Used in level	: A, B and C
Note	: The *Occupant Exit Knowledge* is a measure of the occupant familiarity with the exit locations.

NOTE:
Occupant Exit Knowledge can be shared as part of the communication process if the Social Response or Social Movement flags have been enabled in the Behavioural Options dialogue box in Simulation Mode.

(8) *Occupant Itinerary List (OIL)*
The *Occupant Itinerary List* (OIL) allows an occupant to possess a finite number of tasks that are performed prior to evacuation. These tasks take the form of a nodal location (where the task is assumed to be performed) and an associated time of delay (the length of time the task requires). By default the occupant is initially assumed to have no tasks to perform.

Range	: No tasks - unlimited number of tasks (constrained by host computer capabilities)
Default	: Null
Influenced by	: None.
Influences	: Behaviour and Performance
Used in level	: A, B and C
Note	: The *Occupant Itinerary List* is a representation of the tasks that an occupant might be expected to conduct either en route or prior to an evacuation.

NOTE:
Occupant Itinerary List can be shared as part of the communication process given that the Social Movement flag has been enabled in the Behavioural Options dialogue box in Simulation Mode.

3.4 Hazard Effect Attributes

3.4.1 Heat

(1) *FIH* attribute.
The *FIH* attribute measures the occupant's cumulative exposure to radiative and convective heat. It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). When *FIH* is equal to 1.0, the occupant is incapacitated due to heat exposure. The default value for *FIH* is zero.

Attribute : *FIH*.
Range : 0 - 1.
Default : 0
Influenced by : *FIHc*, *FIHr*.
Influences : incapacity status.
Used in level : C
Note : The *FIH* attribute measures the occupant's combined cumulative exposure to convective and radiative heat.

(2) FIHc
The *FIHc* measures the occupant's cumulative exposure to convective heat. It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). When *FIHc* is equal to 1.0, the occupant is incapacitated due to convective heat exposure. The default value for *FIHc* is 0.0. The *FIHc* attribute is one of the components that affect the *FIH* attribute.

Attribute : *FIHc*.
Range : 0 - 1.
Default : 0
Influenced by : Temperature.
Influences : FIH.
Used in level : C
Note : The FIH_C attribute measures the occupant's cumulative exposure to convective heat.

(3) FIHr
The *FIHr* measures the occupant's cumulative exposure to radiative heat. It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). When *FIHr* is equal to 1.0, the end effect of the radiation exposure is determined to occur. Essentially the occupant is considered to be incapacitated, but the nature of the endpoint is determined by the D_r value used in the FIH_r expression. The endpoint can be based on the pain threshold, incapacitation threshold or can be user defined.

Attribute : *FIHr*.
Range : 0 - 1.
Default : 0
Influenced by : Radiative Flux, Radiative Denominator.
Influences : *FIH*.
Used in level : C
Note : The *FIHr* attribute measures the occupant's cumulative exposure to radiative heat.

(4) D_r

D_r (the Radiative Denominator) is the dose of radiation required to cause the desired effect and has units of $[s(kW/m^2)^{4/3}]$. It is a user defined attribute used in the TOXICITY SUB-MODEL (see CHAPTER 6:). Within EXODUS two values for D_r are provided, these represent the critical value for "pain threshold" $D_r = 80$ and the critical value for "incapacitation", $D_r = 1000$. It is also possible for the user to specify any desired value. The default value for D_r is 80.

Attribute : D_r.
Range : unlimited.
Default : 80
Influenced by : None.
Influences : *FIHr*.
Used in level : C
Note : The D_r value is a measure of the dose of radiation necessary to cause the desired effect.

(5) Radiant Heat Threshold (*RHT)*

The Radiant Heat Threshold (RHT) is the critical value of the radiative flux, measured in kW/m^2, below which the instantaneous radiative flux has no effect on the FIH_r. It is a user defined attribute used in the TOXICITY SUB-MODEL (see CHAPTER 6:). The default value for RHT is 1.7 kW/m^2.

Attribute : *RHT*.
Range : unlimited.
Default : 1.7
Influenced by : None.
Influences : *FIHr*.
Used in level : C
Note : The *RHT* value will impact the calculation of the FIH_r, increasing the value will effectively make agents more tolerate to radiative heat, while decreasing the value will make agents less tolerant.

3.4.2 Narcotic Gases

(1) *Personal Incapacitation Dose (PID*) attribute.

The *Personal Incapacitation Dose* (*PID*) is a measure of the carboxyhaemoglobin (COHb) concentration necessary to cause incapacitation. It is used by the TOXICITY SUB-MODEL (see CHAPTER 6:) to calculate the FICO (see part 3 this section). The incapacitation dose is known to be dependent on age, gender, body size, state of health and level of activity [20,23-27,52,173]. In the present implementation of EXODUS, each agent is randomly assigned a PID% from the default SFPE distribution [173] (see Table 6-1). In addition, the user has two other options, they can assign each agent a fixed PID value of 30%, or allocate the PID of agents according to a user defined distribution.

Attribute : *PID*
Range : 0 – 100 %.
Default : SFPE Distribution [173]
Influenced by : None.
Influences : FICO.
Used in level : C
Note : The *Personal Incapacitation Dose* is a measure of the carboxyhaemoglobin (COHb) concentration necessary to cause incapacitation.

NOTE:
Levels of blood COHb in non-fire CO related fatalities can vary from 20% to over 90% with the majority of fatalities occurring in the range 50 to 90% [52]. However, in fire related incidents, incapacitation may occur at quite low levels of COHb with a small proportion of the population likely to be at risk of collapse (incapacitation) at low blood COHb, in the range of 5-20% COHb and very few are likely to withstand greater than 40% COHb, with few surviving greater than 50%COHb [173].

NOTE:
The PID of agents in versions of buildingEXODUS prior to v6.3 were by default randomly assigned between 25% and 35% according to a random uniform distribution.

(2) *FICO* attribute.
The *FICO* attribute measures the occupant's cumulative exposure to carbon monoxide (*CO*). It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). When *FICO* is equal to 1.0, the occupant is incapacitated due to *CO* poisoning. The default value for *FICO* is 0.0.

Attribute : *FICO.*
Range : 0 - 1.
Default : 0
Influenced by : *CO* concentration, *RMV* and *Personal Incapacitation Dose (PID*).
Influences : FIN, incapacity status.
Used in level : C
Note : The *FICO* attribute measures the occupant's cumulative exposure to *CO*.

(3) *FICN* attribute.
The *FICN* attribute measures the occupant's cumulative exposure to hydrogen cyanide (*HCN*). It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). When *FICN* is equal to 1.0, the occupant is incapacitated due to *HCN* poisoning. The default value for *FICN* is 0.0.

Attribute : *FICN*.
Range : 0 - 1.
Default : 0
Influenced by : *HCN* concentration and *RMV*.
Influences : *FIN*, incapacity status.
Used in level : C
Note : The *FICN* attribute measures the occupant's cumulative exposure to *HCN*.

(4) *FIO* attribute.
The *FIO* attribute measures the occupant's cumulative exposure to low oxygen (O_2). It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). When *FIO* is equal to 1.0, the occupant is incapacitated due to lack of oxygen. The default value for *FIO* is 0.0.

Attribute : *FIO*.
Range : 0 - 1.
Default : 0
Influenced by : O_2 concentration.
Influences : *FIN*, incapacity status.
Used in level : C
Note : The *FIO* attribute measures the occupant's cumulative exposure to low oxygen.

(5) VCO_2 attribute.
The VCO_2 attribute is an estimate of the hyperventilation effect caused by the occupant's exposure to carbon dioxide gas (CO_2). It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). As VCO_2 increases the ventilation rate increases and so the intake of the toxic gases increases. The default value for VCO_2 is 1.0.

Attribute : VCO_2.
Range : 0 - 20.
Default : 1
Influenced by : CO_2 concentration.
Influences : *FIN*.
Used in level : C
Note : VCO_2 estimates the hyperventilation effect caused by exposure to CO_2.

(6) $FICO_2$ attribute.
The $FICO_2$ attribute measures the occupant's cumulative exposure to carbon dioxide gas (CO_2). It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). When $FICO_2$ is equal to 1.0, the occupant is incapacitated due to carbon dioxide. The default value for $FICO_2$ is 0.0.

Attribute : *$FICO_2$*.
Range : 0 - 1.
Default : 0
Influenced by : *CO_2* concentration.
Influences : Incapacity status.
Used in level : C
Note : The *$FICO_2$* attribute measures the occupant's cumulative exposure to carbon dioxide.

(7) *FIN* attribute.
The *FIN* attribute measures the occupant's combined cumulative exposure to low *O_2*, *HCN*, *CO* and *CO_2*. It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). When *FIN* is equal to 1.0, the occupant is incapacitated due to the combined effect of these gases. As the FIN increases the *Mobility Degradation Factor* (and hence *Mobility* and *Travel Speed*) decrease (see Section 3.1 part 2). Similarly, as FIN increases the *Agility Degradation Factor* (and hence *Agility*) also decrease (see Section 3.1 part 4). The default value for *FIN* is 0.0.

Attribute : *FIN*.
Range : 0 - 1.
Default : 0
Influenced by : *FICO, FICN, FIO, FLD, VCO_2*.
Influences : *Mobility*, *Agility*, *Travel Speed* and incapacity status.
Used in level : C
Note : FIN measures the combined cumulative exposure to low O_2, HCN, CO and CO_2

3.4.3 Irritant Gases
The instantaneous impact of the irritant gases are described by the attributes with the prefix FIC.

(1a) *FIC* attribute
The *FIC* attribute measures the occupant's instantaneous exposure to all of the irritant gases. It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). When *FIC* is equal to 1.0, the end effect is determined to occur. The nature of the end effect is dependent on the Tolerance Factor (*TF*) selected and can represent severe impediment to evacuation or incapacitation. The default value for *FIC* is 0.0.

Attribute : *FIC*
Range : 0 - 1.
Default : 0
Influenced by : the concentration, in ppm of *HCL, HBr, HF, SO_2, NO_2, CH_2CHO (Acrolein), HCHO (Formaldehyde)* and *TF* associated with each irritant.
Influences : -
Used in level : C
Note : The *FIC* attribute measures the occupant's combined instantaneous exposure to the irritant gases.

(1b) *TF* attribute

The *TF* attribute is a measure of the concentration of irritant gas (ppm) required to cause a given endpoint. It is a user defined attribute used in the TOXICITY SUB-MODEL (see CHAPTER 6:). The endpoint can be severe mobility impairment or incapacitation depending on the *TF* selected. Furthermore, *TF* varies between individuals within a population, with some individuals being more or less tolerant to the irritant gases. *TF* values for the various irritant gases are presented in Table 6-4 and represent those values to cause the end effect for 50% of the population.

Attribute : *TF*
Range : Distribution based on values in Table 6-4.
Default : SFPE Escape Impaired values (Distribution)
Influenced by : None.
Influences : FIC.
Used in level : C
Note : The *Tolerance Factor* can be set as equal for all occupants or can be distributed with each occupant being assigned a unique value.

(2a) FIC_{HCL} attribute

The FIC_{HCL} attribute measures the occupant's instantaneous exposure to HCL (Hydrogen Chloride). It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). When FIC_{HCL} is equal to 1.0, the end effect due to HCL occurs. The default value for FIC_{HCL} is 0.0.

Attribute : FIC_{HCL}
Range : 0 - 1.
Default : 0
Influenced by : *HCL* concentration and *TF* associated with *HCL*.
Influences : *FIC*.
Used in level : C
Note : The FIC_{HCL} attribute measures the occupant's exposure to HCL.

(2b) TF_{HCL} attribute

The TF_{HCL} attribute measures the occupant's tolerance to an instantaneous exposure to HCL (Hydrogen Chloride). It is a user defined attribute used in the TOXICITY SUB-MODEL (see CHAPTER 6:).

Attribute : TF_{HCL}
Range : Distribution based on values in Table 6-4.
Default : SFPE Escape Impaired values (Distribution)
Influenced by : None.
Influences : *FIC* and FIC_{HCL}.
Used in level : C
Note : The TF_{HCL} attribute measures the occupant's tolerance to HCL.

(3a) FIC_{HBr} attribute

The FIC_{HBr} attribute measures the occupant's instantaneous exposure to HBr (Hydrogen Bromide). It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). When FIC_{HBr} is equal to 1.0, the end effect due to HBr occurs. The default value for FIC_{HBr} is 0.0.

Attribute : FIC_{HBr}
Range : 0 - 1.
Default : 0
Influenced by : *HBr* concentration and *TF* associated with *HBr*.
Influences : *FIC*.
Used in level : C
Note : The FIC_{HBr} attribute measures the occupant's exposure to HBr.

(3b) TF_{HBr} attribute

The TF_{HBr} attribute measures the occupant's tolerance to an instantaneous exposure to HBr (Hydrogen Bromide). It is a user defined attribute used in the TOXICITY SUB-MODEL (see CHAPTER 6:).

Attribute : TF_{HBr}
Range : Distribution based on values in Table 6-4.
Default : SFPE Escape Impaired values (Distribution)
Influenced by : None.
Influences : *FIC* and FIC_{HBr}.
Used in level : C
Note : The TF_{HBr} attribute measures the occupant's tolerance to HBr.

(4a) FIC_{HF} attribute

The FIC_{HF} attribute measures the occupant's instantaneous exposure to HF (Hydrogen Fluoride). It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). When FIC_{HF} is equal to 1.0, the end effect due to HF occurs. The default value for FIC_{HF} is 0.0.

Attribute : FIC_{HF}
Range : 0 - 1.
Default : 0
Influenced by : *HF* concentration and *TF* associated with *HF*.
Influences : *FIC*.
Used in level : C
Note : The FIC_{HF} attribute measures the occupant's exposure to HF.

(4b) TF_{HF} attribute

The TF_{HF} attribute measures the occupant's tolerance to an instantaneous exposure to HF (Hydrogen Fluoride). It is a user defined attribute used in the TOXICITY SUB-MODEL (see CHAPTER 6:).

Attribute : TF_{HF}
Range : Distribution based on values in Table 6-4.
Default : SFPE Escape Impaired values (Distribution)
Influenced by : None.
Influences : *FIC* and FIC_{HF}.
Used in level : C
Note : The TF_{HF} attribute measures the occupant's tolerance to HF.

(5a) FIC_{SO2} attribute
The FIC_{SO2} attribute measures the occupant's instantaneous exposure to SO_2 (Sulphur Dioxide). It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). When FIC_{SO2} is equal to 1.0, the end effect due to SO_2 occurs. The default value for FIC_{SO2} is 0.0.

Attribute : FIC_{SO2}
Range : 0 - 1.
Default : 0
Influenced by : SO_2 concentration and *TF* associated with SO_2.
Influences : *FIC*.
Used in level : C
Note : The FIC_{SO2} attribute measures the occupant's exposure to SO_2.

(5b) TF_{SO2} attribute
The TF_{SO2} attribute measures the occupant's tolerance to an instantaneous exposure to SO_2 (Sulphur Dioxide). It is a user defined attribute used in the TOXICITY SUB-MODEL (see CHAPTER 6:).

Attribute : TF_{SO2}
Range : Distribution based on values in Table 6-4.
Default : SFPE Escape Impaired values (Distribution)
Influenced by : None.
Influences : *FIC* and FIC_{SO2}.
Used in level : C
Note : The TF_{SO2} attribute measures the occupant's tolerance to SO_2.

(6a) FIC_{NO2} attribute
The FIC_{NO2} attribute measures the occupant's instantaneous exposure to NO_2 (Nitrogen Dioxide). It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). When FIC_{NO2} is equal to 1.0, the end effect due to NO_2 occurs. The default value for FIC_{NO2} is 0.0.

Attribute : FIC_{NO2}
Range : 0 - 1.
Default : 0
Influenced by : NO_2 concentration and *TF* associated with NO_2.
Influences : *FIC*.
Used in level : C
Note : The FIC_{NO2} attribute measures the occupant's exposure to NO_2.

(6b) TF_{NO2} attribute

The TF_{NO2} attribute measures the occupant's tolerance to an instantaneous exposure to NO_2 (Nitrogen Dioxide). It is a user defined attribute used in the TOXICITY SUB-MODEL (see CHAPTER 6:).

Attribute	: TF_{NO2}
Range	: Distribution based on values in Table 6-4.
Default	: SFPE Escape Impaired values (Distribution)
Influenced by	: None.
Influences	: *FIC* and FIC_{NO2}.
Used in level	: C
Note	: The TF_{NO2} attribute measures the occupant's tolerance to NO_2.

(7a) FIC_{CH2CHO} attribute

The FIC_{CH2CHO} attribute measures the occupant's instantaneous exposure to CH_2CHO (*Acrolein*). It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). When FIC_{CH2CHO} is equal to 1.0, the end effect due to CH_2CHO (*Acrolein*) occurs. The default value for FIC_{CH2CHO} is 0.0.

Attribute	: FIC_{CH2CHO}
Range	: 0 - 1.
Default	: 0
Influenced by	: CH_2CHO (*Acrolein*) concentration and *TF* associated with CH_2CHO.
Influences	: *FIC*.
Used in level	: C
Note	: The FIC_{CH2CHO} attribute measures the occupant's exposure to CH2CHO (*Acrolein*).

(7b) TF_{CH2CHO} attribute

The TF_{CH2CHO} attribute measures the occupant's tolerance to an instantaneous exposure to CH_2CHO (*Acrolein*). It is a user defined attribute used in the TOXICITY SUB-MODEL (see CHAPTER 6:).

Attribute	: TF_{CH2CHO}
Range	: Distribution based on values in Table 6-4.
Default	: SFPE Escape Impaired values (Distribution)
Influenced by	: None.
Influences	: *FIC* and FIC_{CH2CHO}.
Used in level	: C
Note	: The TF_{CH2CHO} attribute measures the occupant's tolerance to CH2CHO.

(8a) FIC_{HCHO} attribute

The FIC_{HCHO} attribute measures the occupant's instantaneous exposure to HCHO (*Formaldehyde)*. It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). When FIC_{HCHO} is equal to 1.0, the end effect due to HCHO (*Formaldehyde)* occurs. The default value for FIC_{HCHO} is 0.0.

Attribute : FIC_{HCHO}
Range : 0 - 1.
Default : 0
Influenced by : HCHO (*Formaldehyde)* concentration and *TF* associated with HCHO.
Influences : *FIC*.
Used in level : C
Note : The FIC_{HCHO} attribute measures the occupant's exposure to HCHO (*Formaldehyde).*

(8b) TF_{HCHO} attribute
The TF_{HCHO} attribute measures the occupant's tolerance to an instantaneous exposure to HCHO (*Formaldehyde).* It is a user defined attribute used in the TOXICITY SUB-MODEL (see CHAPTER 6:).

Attribute : TF_{HCHO}
Range : Distribution based on values in Table 6-4.
Default : SFPE Escape Impaired values (Distribution)
Influenced by : None.
Influences : *FIC* and FIC_{HCHO}.
Used in level : C
Note : The TF_{HCHO} attribute measures the occupant's tolerance to HCHO.

The impact of an exposure dose to irritant gases over a period of time is described by the attributes with the prefix FLD which stands for Fractional Lethal Dose of irritants.

(9a) *FLD* attribute
The *FLD* attribute measures the occupant's cumulative exposure to the irritant gases. It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). When FLD = 1.0 fatality is expected to occur however, death is not likely to occur during the evacuation but may occur several hours or several days after receiving the lethal dose. The FLD is thus not intended to be used as a critical factor describing evacuation or as a design limit but for estimations of the extent to which post-exposure deaths from lung oedema and inflammation are likely to occur [173]. As a result, agents are not incapacitated due to excessive cumulative exposure to irritant gases (i.e. FLD ≥ 1.0). The default value for *FLD* is 0.0.

NOTE:
In versions of buildingEXODUS prior to v6.3 agents were incapacitated and hence removed from the simulation and added to the mortuary when their corresponding FLD ≥ 1.0.

Attribute : *FLD*
Range : 0 or greater.
Default : 0
Influenced by : The dose *(ppm.min) of HCL, HBr, HF, SO2, NO2, CH2CHO(Acrolein), HCHO (Formaldehyde) and the CD associated with each irritant*.
Influences : -
Used in level : C
Note : The *FLD* attribute measures the occupant's cumulative exposure to the irritant gases.

(9b) *CD* attribute
The Critical Dose (*CD)* attribute is a measure of the critical dose of irritant gas (ppm.min) required to cause a given endpoint. It is a user defined attribute used in the TOXICITY SUB-MODEL (see CHAPTER 6:). The endpoint is usually death. Furthermore, while *CD* varies between individuals within a population, with some individuals being more or less tolerant to the irritant gases the nature of the distribution is not known with certainty. *CD* values for the various irritant gases are presented in Table 6-4.

Attribute : *CD*
Range : Fixed value based on values in Table 6-4.
Default : SFPE Critical Dose (Fixed)
Influenced by : None.
Influences : FLD.
Used in level : C
Note : There are three choices for specifying the *Critical Dose*.

(10a) FLD_{HCL} attribute
The FLD_{HCL} attribute measures the occupant's cumulative exposure to HCL (Hydrogen Chloride). It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). The default value for FLD_{HCL} is 0.0.

Attribute : FLD_{HCL}
Range : 0 or greater.
Default : 0
Influenced by : *HCL* dose and CD for *HCL*.
Influences : *FLD*.
Used in level : C
Note : The FLD_{HCL} attribute measures the occupant's cumulative exposure to HCL.

(10b) CD_{HCL} attribute
The CD_{HCL} attribute measures the occupant's tolerance to a dose of HCL (Hydrogen Chloride). It is a user defined attribute used in the TOXICITY SUB-MODEL (see CHAPTER 6:).

Attribute : CD_{HCL}
Range : Fixed value based on values in Table 6-4.
Default : SFPE Critical Dose (Fixed)
Influenced by : None.
Influences : *FLD* and FLD_{HCL}.
Used in level : C
Note : The CD_{HCL} attribute measures the occupant's tolerance to a dose of HCL

(11a) FLD_{HBr} attribute
The FLD_{HBr} attribute measures the occupant's cumulative exposure to HBr (Hydrogen Bromide). It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). The default value for FLD_{HBr} is 0.0.

Attribute : FLD_{HBr}
Range : 0 or greater.
Default : 0
Influenced by : *HBr* dose and CD for *HBr*.
Influences : *FLD*.
Used in level : C
Note : The FLD_{HBr} attribute measures the occupant's cumulative exposure to HBr.

(11b) CD_{HBr} attribute
The CD_{HBr} attribute measures the occupant's tolerance to a dose of HBr (Hydrogen Bromide). It is a user defined attribute used in the TOXICITY SUB-MODEL (see CHAPTER 6:).

Attribute : CD_{HBr}
Range : Fixed value based on values in Table 6-4.
Default : SFPE Critical Dose (Fixed)
Influenced by : None.
Influences : *FLD* and FLD_{HBr}.
Used in level : C
Note : The CD_{HBr} attribute measures the occupant's tolerance to a dose of HBr

(12a) FLD_{HF} attribute
The FLD_{HF} attribute measures the occupant's cumulative exposure to HF (Hydrogen Fluoride). It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). The default value for FLD_{HF} is 0.0.

Attribute : FLD_{HF}
Range : 0 or greater.
Default : 0
Influenced by : *HF* dose and CD for *HF*.
Influences : *FLD*.
Used in level : C
Note : The FLD_{HF} attribute measures the occupant's cumulative exposure to HF.

(12b) CD_{HF} attribute
The CD_{HF} attribute measures the occupant's tolerance to a dose of HF (Hydrogen Fluoride). It is a user defined attribute used in the TOXICITY SUB-MODEL (see CHAPTER 6:).

Attribute : CD_{HF}
Range : Fixed value based on values in Table 6-4.
Default : SFPE Critical Dose (Fixed)
Influenced by : None.
Influences : *FLD* and FLD_{HF}.
Used in level : C
Note : The CD_{HF} attribute measures the occupant's tolerance to a dose of HF

(13a) FLD_{SO2} attribute
The FLD_{SO2} attribute measures the occupant's cumulative exposure to SO_2 (Sulphur Dioxide). It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). The default value for FLD_{SO2} is 0.0.

Attribute : FLD_{SO2}
Range : 0 or greater.
Default : 0
Influenced by : SO_2 dose and CD for SO_2.
Influences : *FLD*.
Used in level : C
Note : The FLD_{SO2} attribute measures the occupant's cumulative exposure to SO_2.

(13b) CD_{SO2} attribute
The CD_{SO2} attribute measures the occupant's tolerance to a dose of SO_2 (Sulphur Dioxide). It is a user defined attribute used in the TOXICITY SUB-MODEL (see CHAPTER 6:).

Attribute : CD_{SO2}
Range : Fixed value based on values in Table 6-4.
Default : SFPE Critical Dose (Fixed)
Influenced by : None.
Influences : *FLD* and FLD_{SO2}.
Used in level : C
Note : The CD_{SO2} attribute measures the occupant's tolerance to a dose of SO_2

(14a) FLD_{NO2} attribute
The FLD_{NO2} attribute measures the occupant's cumulative exposure to NO_2 (Nitrogen Dioxide). It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). The default value for FLD_{NO2} is 0.0.

Attribute : FLD_{NO2}
Range : 0 or greater.
Default : 0
Influenced by : NO_2 dose and CD for NO_2.
Influences : *FLD*.
Used in level : C
Note : The FLD_{NO2} attribute measures the occupant's cumulative exposure to NO_2.

(14b) CD_{NO2} attribute
The CD_{NO2} attribute measures the occupant's tolerance to a dose of NO_2 (Nitrogen Dioxide). It is a user defined attribute used in the TOXICITY SUB-MODEL (see CHAPTER 6:).

Attribute : CD_{NO2}
Range : Fixed value based on values in Table 6-4.
Default : SFPE Critical Dose (Fixed)
Influenced by : None.
Influences : *FLD* and FLD_{NO2}.
Used in level : C
Note : The CD_{NO2} attribute measures the occupant's tolerance to a dose of NO_2

(15a) FLD_{CH2CHO} attribute
The FLD_{CH2CHO} attribute measures the occupant's cumulative exposure to CH_2CHO (*Acrolein*). It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). The default value for FLD_{CH2CHO} is 0.0.

Attribute : FLD_{CH2CHO}
Range : 0 or greater.
Default : 0
Influenced by : CH_2CHO (*Acrolein*) dose and CD for CH_2CHO.
Influences : *FLD*.
Used in level : C
Note : The FLD_{CH2CHO} attribute measures the occupant's cumulative exposure to CH_2CHO (*Acrolein*).

(15b) CD_{CH2CHO} attribute
The CD_{CH2CHO} attribute measures the occupant's tolerance to a dose of CH_2CHO (*Acrolein*). It is a user defined attribute used in the TOXICITY SUB-MODEL (see CHAPTER 6:).

Attribute : CD_{CH2CHO}
Range : Fixed value based on values in Table 6-4.
Default : SFPE Critical Dose (Fixed)
Influenced by : None.
Influences : *FLD* and FLD_{CH2CHO}.
Used in level : C
Note : The CD_{CH2CHO} attribute measures the occupant's tolerance to a dose of CH_2CHO

(16a) FLD_{HCHO} attribute
The FLD_{HCHO} attribute measures the occupant's cumulative exposure to HCHO (*Formaldehyde*). It is a dynamic attribute calculated by the TOXICITY SUB-MODEL (see CHAPTER 6:). The default value for FLD_{HCHO} is 0.0.

Attribute : FLD_{HCHO}
Range : 0 or greater.
Default : 0
Influenced by : HCHO (*Formaldehyde)* dose and CD for HCHO.
Influences : *FLD*.
Used in level : C
Note : The FLD_{HCHO} attribute measures the occupant's cumulative exposure to HCHO (*Formaldehyde*).

(16b) CD_{HCHO} attribute
The CD_{HCHO} attribute measures the occupant's tolerance to a dose of HCHO (*Formaldehyde*). It is a user defined attribute used in the TOXICITY SUB-MODEL (see CHAPTER 6:).

Attribute : CD_{HCHO}
Range : Fixed value based on values in Table 6-4.
Default : SFPE Critical Dose (Fixed)
Influenced by : None.
Influences : *FLD* and FLD_{HCHO}.
Used in level : C
Note : The CD_{HCHO} attribute measures the occupant's tolerance to a dose of HCHO

(17) *Critical Irritant Incapacitation Factor (CIIF)* attribute
The *CIIF* attribute defines the FIC value required to cause incapacitation, resulting in the exposed agent being removed from the simulation and added to the mortuary. By default the value of this attribute is 4.5 as the default Tolerance Factors (TF) defining agent tolerance to instantaneous irritant exposure are derived from the SFPE Escape Impaired values. As the name suggest, these values define the irritant concentrations required to incapacitate an agent. The value of 4.5 is used as the concentrations of irritant gas required to cause incapacitation are approximately 4-5 times greater than those required to cause Escape Impairment (see Table 6-4).

Attribute	: *Critical Irritant Incapacitation Factor*
Range	: Greater than or equal to zero.
Default	: 4.5
Influenced by	: Which TF option is selected.
Influences	: Number of incapacitated (mortuary count).
Used in level	: C
Note	: The *CIIF* attribute defines the FIC value required to cause incapacitation, and hence result in the removal of the agent from the simulation and added to the mortuary.
Note	: If TF for Incapacitation is used, CIIF = 1.0 by default.
Note	: If user defined TF is used, CIIF = 0.0 by default, requiring the user to select an appropriate value.

CHAPTER 4: MOVEMENT SUB-MODEL

The MOVEMENT sub-model is only active during *SIMULATION* mode. It is primarily concerned with the physical movement of the occupants through the different terrain types. It consists of a number of rules, the main function of which is to determine the appropriate travel speed for the current terrain type. For example, *Leap Speed* is selected for occupants who have decided to climb over a row of seats, while *Fast Walk Speed* is selected for an occupant who is travelling through open space. In addition, the MOVEMENT sub-model ensures that the occupant has the capability of performing the requested action, for example it checks if the occupant *Agility* is sufficient to allow travel over nodes with particular *Obstacle* values.

While the movement sub-model is responsible for moving the occupant, it is the BEHAVIOUR sub-model (see CHAPTER 7:) that selects the direction of travel. If a suitable move is not available to the occupant, the MOVEMENT sub-model will supervise a *Wait* period (see CHAPTER 7:). During the *Wait* period the occupant remains stationary until a suitable move becomes available.

Movement rules are fired on even ticks of the SC, while selection rules (from the BEHAVIOUR sub-model) are fired on odd ticks. The wait rules are fired continuously, as the occupant always has the option to wait. Movement decisions and actions will only take place if the SC shows a time that is at least as large as the *PET* (see CHAPTER 3: and CHAPTER 7:).

When a move decision has been made, EXODUS waits until the next tick of the SC, and then moves the occupant to the location. Using the occupant's *Travel Speed* (see CHAPTER 3:) and distance travelled (see CHAPTER 3:), EXODUS calculates the travel time and advances his/her *PET* by the appropriate amount. The occupant now sits on the selected node until the SC catches up with the *PET*, at which time another movement decision may be taken. If the occupant is forced to wait at the current location the *PET* is updated with each tick of the SC.

In the current version of EXODUS there is no provision for occupant’s to push past each other while in queues. In these circumstances, the person can wait for his/her turn to move, or if their *Patience* has expired, take a detour round the obstruction.

CHAPTER 5: HAZARD SUB-MODEL

The HAZARD SUB-MODEL is accessed by the user during *SCENARIO* mode and it is utilised by EXODUS during *SIMULATION* mode. The HAZARD sub-model controls the development of the atmospheric and physical environment. The atmospheric aspects comprise the distribution of fire hazards *CO_2, CO, HCN, O_2* depletion, *Heat* (radiative and conductive) and *Smoke*, as well as the irritant gases HCl, HBr, HF, SO_2, NO_2, CH_2CHO (*Acrolein*) and HCHO (*Formaldehyde*). The physical aspects include setting of opening and closing times for exits.

The primary function of the hazard sub-model is to distribute the fire hazards. EXODUS does not possess a component such as a zone or field fire model [22] to predict the generation and spread of fire hazards. There are however several means by which fire hazard data may be included. These are manual data entry, arbitrary calculation, library data, direct import of history files (.HI) from the CFAST (version 4.0.1 to 5.1.1) zone model, direct import of data files (.DAT) from the SMARTFIRE V4.0 fire field model and the importation of data files converted to the SMARTFIRE data file format from CFAST (version 6.0.8 – 7.0.1) data output files (.OUT).

Within buildingEXODUS fire hazards operate at two heights: the upper and lower height. The definition of the upper and lower height is dependent on the approach used to specify the hazards. When using the user-defined option, evacuees are continually exposed to hazards at the upper height while they assume the standing position. The hazards defining the upper height conditions should represent those hazards existing at a nominal head height e.g. 1.7m. When evacuees are forced to crawl, they are then exposed to the hazard values at the lower height and so lower height hazard values should represent those hazards existing near the floor e.g. about 1.0 m above the floor. However, when CFAST generated hazards are used, the height of the upper hazard layer is also determined, thus the height of the upper layer changes with time as do the hazard values. In this case evacuees do not typically come into contact with the upper layer until it has descended to a distance equal to their height (see the User Guide, Section 5.3.6). When using the SMARTFIRE defined hazards, the rationale used is similar to that used when making use of user-defined data. When people are standing they are continually exposed to the environmental conditions that exist at the upper height. When evacuees are forced to crawl, they are then exposed to the hazard values at the lower height and so lower height hazard values should represent those hazards existing near the floor. The position and depth of both the upper and lower regions are defined, by the user, within SMARTFIRE prior to running the simulation. Hence it is up to the user to ensure that the upper and lower regions are representative of the regions to which standing and crawling people would be exposed. The exception to this is the exposure to radiative flux, which is provided as constant in both the upper and lower regions (see below). It is suggested that the hazards defining the upper height conditions should represent those hazards existing at a nominal head height e.g. 1.7m while those for crawling individuals should represent a height of about 1.0 m above the floor.

The methods of data entry are now described.

(1) Manual data entry.
For a particular scenario, information relating to fire hazard spread may be known in detail. This could be generated from for example, experimental data or output from zone or field fire models. In these circumstances, the data has to transferred manually and the regions manually defined.

(2) Arbitrary calculation.
Within the HAZARD sub-model it is possible to define polynomial functions which determine the increase in fire hazard concentration over specified regions for a specified period of time. This requires the user to,

(i) define the zones, for example:
zone 1 is room 1
zone 2 rear quarter of room 2
zone 3 remainder of room 3
zone 4 remainder of enclosure,

(ii) select a function that defines the growth rate for each hazard , for example:

temperature = ambient temperature + 0.01 * time^2	for $0 < time < 10$
temperature = ambient temperature + 0.1 * time$^{2.5}$	for $10 < time < 100$
O_2 concentration = ambient concentration	for $0 < time < 50$
O_2 concentration = ambient concentration - 0.01*time	for $50 < time < 100$

etc.

(3) Library data.
A fire atmosphere can be saved with a geometry and recalled for later use from the library and either modified or used as is.

(4) Direct import of History files from the CFAST (version 4.0.1 to 5.1.1) zone model.
History files (.HI) generated by the CFAST zone model (versions 4.0.1-5.1.1) [60,61] can be imported directly into buildingEXODUS. Once CFAST has been run, the CFAST generated .HI file (i.e. a history file) contains a record of the simulation results. This file can then be imported into buildingEXODUS as a fire hazard, which then requires association with an area within the geometry (see the User Guide, Chapter 5 for details). The imported data can consist of the fire products that buildingEXODUS makes use of in its hazard calculations (i.e. *Smoke Concentration, Temperature, Radiative Flux, HCN, HCL, CO,* CO_2 *and* O_2 concentrations). The units that CFAST uses for these variables are similar to those used by buildingEXODUS except for the *Smoke Concentration* - where CFAST makes use of the optical density rather than extinction coefficient - *Temperature* - where CFAST uses Kelvin and buildingEXODUS uses Centigrade and radiative heat flux where buildingEXODUS uses kWatts/m^2 and CFAST uses Watts/m^2. Once buildingEXODUS reads the CFAST data, these units are automatically converted to the appropriate units and therefore require no further user attention. The structure geometry can consist of any multi-compartment structure that both CFAST and buildingEXODUS can accommodate.

CFAST produces predictions for the radiated heat present within each compartment, unlike the other environment variables, radiant heat is not layer based. There are several different radiative components predicted by CFAST. The radiated heat values are calculated as the net flux to a point in the centre of the compartment floor. Within CFAST, this value can be determined taking into account the actual temperature of the floor element (the so-called "*Flux to Target*"), or assuming that the floor element is at ambient temperature (the so-called "*Ambient Target*" flux). It is this later alternative that is used within buildingEXODUS.

In using this value it is assumed that up to the point of physical pain, injury and incapacitation, the core body temperature remains near its ambient value. Thus it can be assumed that the bulk average temperature of the person target is closer to ambient than the typically higher surface temperature of the compartment floor.

NOTE:
The Radiative Flux generated by CFAST and used by buildingEXODUS is calculated as the net flux to a point in the centre of the compartment floor, where the floor is assumed to be at ambient temperature.

NOTE:
Any History files (.HI) generated using CFAST versions greater than 5.1.1 can NOT currently be used within maritimeEXODUS (see Section 6).

(5) Direct import of data files from the SMARTFIRE fire field model.
buildingEXODUS will accept data from the CFD based fire simulation model SMARTFIRE. This enables evacuation analysis to benefit from the greater modelling accuracy that CFD fire field modelling offers fire simulation. In order to simplify the data importing process, from the user's point of view, the EXODUS-SMARTFIRE link has been implemented so as to be as consistent as possible with the existing CFAST data importing mechanism, described in the previous section. The data link is achieved using a zone-filter that processes some of the data produced by SMARTFIRE and allows the required "zoned" data to be loaded into EXODUS and used as EXODUS evacuation "hazards".

buildingEXODUS will only read data from SMARTFIRE V4.0 or later. SMARTFIRE V4.0 [88-102] is an open architecture CFD environment, written in C++, comprising four major components: the CFD numerical engine, various Graphical User Interfaces, an automated meshing tool and the Intelligent Control System. The SMARTFIRE system has been described in previous publications [88-102], and so only a brief outline is presented here. SMARTFIRE includes a six-flux radiation model, a multiple ray radiation model, provision for heat transfer through walls, a volumetric heat release model or gaseous combustion model (using the eddy dissipation model) to represent fires, smoke modelling and turbulence (using a two equation K-Epsilon closure with buoyancy modifications). SMARTFIRE uses three-dimensional unstructured meshes, enabling complex irregular geometries to be meshed. The code uses the SIMPLE algorithm and can solve turbulent or laminar flow problems under transient or steady state conditions.

SMARTFIRE V4.0 produces a data file (named "casename".dat), which is formatted in such a way that it can be directly imported into buildingEXODUS. The format is consistent with the CFAST output format. In the present implementation of the EXODUS-SMARTFIRE interface, the data imported from SMARTFIRE consists of the following fire hazards: *Smoke Concentration*, *Temperature, Radiative Flux* and *toxic gases (i.e.* O_2 *depletion, CO,* CO_2 *and HCL).*

NOTE:
The native units that SMARTFIRE uses for the fire hazards Smoke Concentration, Temperature and Radiative Flux are different to those required by EXODUS. SMARTFIRE makes use of optical density for Smoke Concentration while EXODUS requires extinction coefficient, for the temperature, SMARTFIRE uses Kelvin while EXODUS uses Centigrade and for radiative heat flux, SMARTFIRE uses Watts/m2, whereas EXODUS makes use of KWatts/m2. The filter within the EXODUS software automatically converts the SMARTFIRE data into the appropriate units and so no user intervention is required to convert these units.

While EXODUS works on a nodal system similar to SMARTFIRE, for evacuation analysis it is not generally necessary to have unique and individual hazard identification at each nodal location. The EXODUS simulation therefore uses a zonal system for the specification of hazard information. Furthermore, within EXODUS, hazard information is only required at two characteristic heights known as “upper height” and “lower height”. Therefore, before the hazard data from SMARTFIRE can be used by EXODUS it must be averaged over the same spatial zones as defined within the EXODUS simulation and at the required vertical locations.

NOTE:
Conversion of the SMARTFIRE nodal data into zonal data is performed within the SMARTFIRE software. Within SMARTFIRE, the user may use a default averaging approach or may specify their preferred averaging algorithm. It is recommended that the user selects the default averaging algorithms as there is a distinct possibility that manual configuration will create averaged values with incompatible units to those required by EXODUS.

Within SMARTFIRE the radiative flux is determined for a standing individual – which is represented as an elongated cuboid – and is summed over all the components of radiation intersecting the surface area of the standing individual. Therefore irrespective of the posture adopted by the individual they will always be exposed to the standing value for the thermal radiation.

NOTE:
SMARTFIRE provides upper and lower region values for the Temperature and the Smoke concentration levels. However, SMARTFIRE provides EXODUS with a single radiative flux value, which is assumed to be constant for both the upper and lower layers.

A volume averaging technique is used to harmonize the three-dimensional control-volume discretisation used within SMARTFIRE with the meshing and zoning system used within EXODUS. This technique effectively groups together potentially large numbers of cells and averages the data within them to produce representative values for the hazards over the specified zone.

The zoned data is generated at each time-step within the SMARTFIRE simulation. An ambient set of initial data is produced at a time of 0 seconds – before the simulation commences. The size of the time-step within SMARTFIRE is user defined and will be dependent – to a large extent – upon the nature of the simulation, the geometry, the meshing and the numerical stability of the scenario being examined. The amount of data generated by SMARTFIRE and transferred to EXODUS can be further reduced by increasing the size of the SMARTFIRE time-steps (only possible if the simulation has sufficient stability) or by requesting that SMARTFIRE perform output data saves less frequently.

In addition, the SMARTFIRE time step size can be adjusted automatically by SMARTFIRE due to predicted or detected difficulties at a particular stage of the simulation. In order for EXODUS to be able to interpret the possibly inconsistent temporal data, the EXODUS zoned data reader is able to linearly interpolate the imported data, between available data times, if a required matching simulation time is unavailable.

Once these parameters have been specified and the SMARTFIRE simulation has been completed, SMARTFIRE will have generated an output file containing the appropriate hazard information specified at the appropriate locations that can be read by EXODUS. If the zonal definition has been correctly specified within SMARTFIRE, this file will be compatible with the EXODUS hazard zone specification. ***It is the user's responsibility to ensure that the zones defined in the SMARTFIRE model and those used in the buildingEXODUS simulation coincide.*** If this is not the case then anomalous results may occur. Care should be taken to account for any extended regions that might be created when the CFD scenario is meshed as these may alter the coordinates of the geometry and hence can change the user's geometric understanding of the scenario.

NOTE:
The user must generate a fire scenario that is compatible with both buildingEXODUS and SMARTFIRE. This means that the data created for the zones in SMARTFIRE should be appropriate (i.e. should represent the same geometrical space) for the zones that it is intended to represent within buildingEXODUS. This alignment is not automated and will therefore be the user's responsibility to ensure that the fire being modelled within SMARTFIRE is appropriate for the structure in which it will be located within buildingEXODUS.

(6) Direct import of data files converted to the SMARTFIRE data file format from CFAST (versions 6.0.8 – 7.0.1) data output files.
At present buildingEXODUS only has the ability to directly load hazard data contained within binary History files (.HI) generated using CFAST version 4.0.1 to 5.1.1. **Consequently, any History files generated using CFAST versions greater than 5.1.1 can therefore NOT be used within buildingEXODUS.** For CFAST versions greater than 5.1.1, buildingEXODUS instead utilises the data contained within the ASCII output files (.OUT) produced automatically by CFAST upon completion of any given simulation. These output files typically contain information about all aspects of any given CFAST simulation (including those aspects not relevant to buildingEXODUS) in a text based format designed to be read by people, as opposed to computers. To overcome the un-optimised format and excess simulation information that these output files (.OUT) contain, a conversion utility called CFAST2SMF was produced to convert their relevant hazard data into corresponding optimised data files in the existing SMARTFIRE file format. As with conventional binary History files, converted data files will contain each of the CFAST fire products relevant to buildingEXODUS (i.e. *Smoke Concentration, Temperature, Radiative Flux, HCN, HCL, CO,* CO_2 *and* O_2 concentrations). The units of those fire products will also be automatically converted (where necessary) from the units used within CFAST to the units appropriate to SMARTFIRE data files (see Section 4). Once converted into a SMARTFIRE data file, the hazard data originally generated using CFAST can then be loaded into buildingEXODUS in the same manner as conventional SMARTFIRE generated data (see Section 4).

Thus, using one of the six techniques defined above, each node in the geometry is assigned hazard data. Hazard data is defined for two heights, *Head Height* and near *Floor Height* (see Section 2.3).

Physical aspects that are controlled by the HAZARD sub-model include the setting of exit opening/closing times. During the course of an evacuation the status of exits may alter as a result of the progress of the fire or equipment malfunction. Some exits may not be open throughout the evacuation due to difficulties with the equipment, or inability of the occupants to reach the exit. Exits also become unserviceable due to the spread of fire inside or outside the enclosure. Likewise, certain areas in the enclosure may become no-go areas due to the presence of the fire or structural damage. These sequences of events may be specified in the HAZARD sub-model that notifies the MOVEMENT sub-model when/which door is opened/closed and where/when no-go areas occur.

CHAPTER 6: TOXICITY SUB-MODEL

The TOXICITY sub-model is accessed by the user during *SCENARIO* mode and it is utilised by EXODUS during *SIMULATION* mode. It is only available with software level C. To determine the effect of the fire hazards on occupants, EXODUS uses a Fractional Effective Dose (FED) toxicity model [2, 20, 23-27, 34, 52, 62, 173]. FED models assume that the effects of certain fire hazards are related to the dose received rather than the exposure concentration. The model calculates, for these agents, the ratio of the dose received over time to the effective dose that causes incapacitation or death, and sums these ratios during the exposure. When the total reaches unity, the toxic effect is predicted to occur. These effects are communicated to the BEHAVIOUR SUB-MODEL which, in turn, feeds through to the movement of the individual. For example, as the FED model for narcotic agents (i.e. FIN) approaches unity the occupant's *Mobility*, *Agility*, and travel speeds can be reduced making it more difficult for the affected occupant to escape (see Section 3.1).

The core toxicity model implemented within EXODUS is the FED model of Purser [20, 23-25,62]. This model considers the toxic and physical hazards associated with elevated temperature, *HCN*, *CO*, CO_2 and low O_2 and estimates the time to incapacitation. The toxicity model was updated in 2017 to ensure compatibility with the most recent formulation presented in the SFPE Handbook, 5th Edition [173].

In each of the following expressions, *t* is the exposure time (minutes). The fractional incapacitating dose (FID) for each of the agents is calculated as follows:

(i) *CO* (measured in ppm):

$$FICO = 3.317 \times 10^{-5} \times CO^{1.036} \times RMV \times \frac{t}{PID} \qquad \textbf{(7)}$$

where *RMV* is the minute volume (litres/minute) and PID is the *Personal Incapacitation Dose* (%) (see CHAPTER 3:).

The PID is not the same value for each member of the population. While the range 30-40% COHb is typically used to represent the distribution within the population, this only applies to average healthy adults. Some members of the population may be more susceptible while others may be more tolerant. Unfortunately, there are no definitive studies that define how PID is distributed in the general population however, the SFPE have suggested a distribution [173] that reflects that a small fraction of the population are at risk of collapse with small concentrations of COHb (5% - 20%) while most of the population are susceptible in the range 25% - 35% COHb while none can withstand greater than 45% COHb. The SFPE recommended distribution is shown in Figure 6-1 and Table 6-1.

NOTE:
The value of PID used within EXODUS represents the endpoint of collapse (incapacitation) and not death.

NOTE:
Levels of COHb in excess of 10% indicates that there has been some smoke inhalation, with 20% suggesting severe smoke inhalation. Survival rates are high for inhalation of smoke resulting in COHb up to 40%. A level of COHb in excess of 40% presents a high risk of death, and a level in excess of 50% represents a very high risk of death. However, some people have been known to survive with greater than 50% COHb [174].

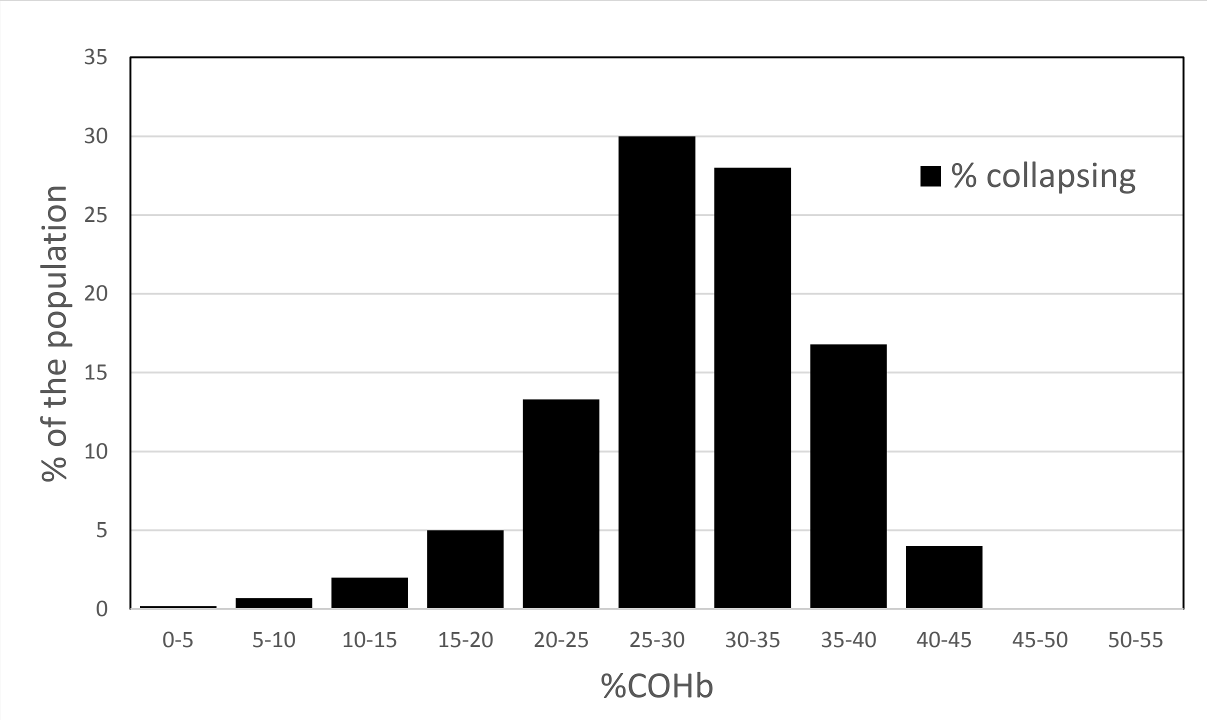

Figure 6-1: Distribution of % COHb to cause incapacitation within the population based on data from the SFPE Handbook [173].

Table 6-1: Approximate proportion of population with various %COHb required to cause incapacitation extracted from Figure 6-1 (based on [173])

% COHb	% of Population	% COHb	% of Population
0 - 5	0	30-35	0.28
5-10	0.01	35-40	0.17
10-15	0.02	40-45	0.04
15-20	0.05	45-50	0
20-25	0.13	50-55	0
25-30	0.30		

Within EXODUS there are three options available to specify the PID within the simulation:

- SFPE PID:
  Each member of the population is assigned a *PID* from the SFPE distribution presented in Table 6-4. Thus using this approach, each member of the population will have their own unique *PID* based on the SFPE values. This is the default setting for the *PID*.

- User Defined PID:

  Each member of the population is assigned a *PID* determined by the user. This requires the user to populate a distribution table similar to Table 6-1.

- Fixed Value PID:

  Each member of the population is assigned the same single value for PID. Thus each member of the population has an identical PID. The default value is 30% COHb,

NOTE:
Versions of buildingEXODUS prior to v6.3 had two options for setting the PID values. They could provide every member of the population with the same PID (set to 30% COHb, as in option 3) or they could assign the population with a uniform random distribution of PID set between an upper and lower limit (which was by default set to 25-35% COHb).

NOTE:
The FICO expression (equation 7) is unreliable for small adults or children.

NOTE:
The FICO model assumes that inhaled CO is immediately converted to COHb. In reality there may be a delay.

NOTE:
The FICO model cannot be used reliably in situations where the CO concentration is decreasing.

NOTE:
This expression is more complex than that suggested for use in ISO13571 [172] which is given simply by:

$$FICO = CO \times \frac{t}{35000} = 2.86 \times 10^{-5} \times CO \times t \qquad \textbf{(7b)}$$

In this expression, PID is equivalent to 30% blood carboxyhaemoglobin, with a respiratory minute volume of 20 l/min corresponding to a compromising tenability dose (C × t) of 35,000 ppm.min. If these values for PID and RMV are substituted within equation 7 we have:

$$FICO = 2.21 \times 10^{-5} \times CO^{1.036} \times t \qquad \textbf{(7c)}$$

As can be seen, equation 7c is similar to the ISO equation 7b, with the 1.036 power having a modest effect for small values of CO concentration. Presented in the Figure 6-2 *is a plot of FICO as a function of CO assuming a 1 minute exposure using the ISO recommended curve (equation 7b) and the SFPE recommended curve (equation 7c) with equivalent PID and RMV values to those assumed in equation 7b. As can be seen, the 1.036 power becomes more significant only when CO is large (in the tens of thousands ppm). Thus the use of the SFPE equation (equation 7), using the same values for PID and RMV as assumed within the ISO standard, produces a more conservative prediction for FICO i.e. incapacitations will occur at lower accumulated dose values than predicted using the ISO equation.*

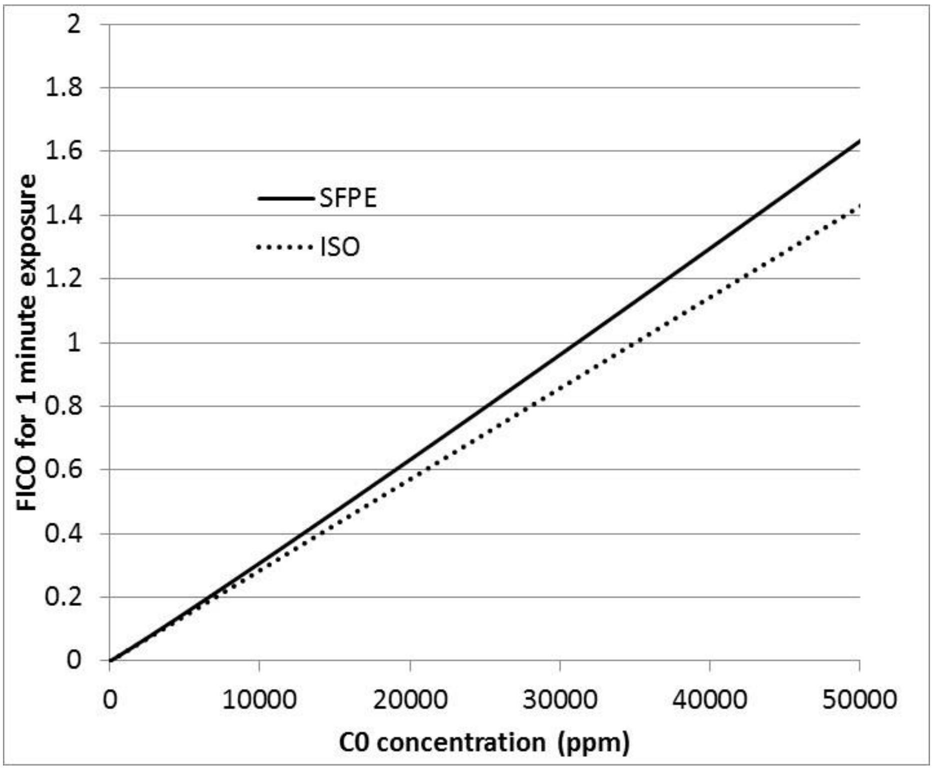

Figure 6-2: Comparison of FICO determined using the ISO (equation 7b) and SFPE (equation 7c) formulations

(ii) *HCN* (measured in ppm):

$$FICN = HCN^{2.36} \times RMV \times \frac{t}{2.43 \times 10^{7}} \quad \textbf{(8)}$$

NOTE:
The FICN expression (equation 8) is unreliable outside the range 80 - 180 ppm HCN. At concentrations less than 80 ppm, HCN will only have a minor effect for exposures of up to an hour, significantly longer than normally required for a building evacuation, and at concentrations above approximately 180 ppm, incapacitation will be very rapid (within 2 minutes). (see CHAPTER 3:[173]).

NOTE:
The FICN expression (equation 8) was updated in buildingEXODUS v6.3, prior versions used an exponential representation.

NOTE:
This expression is more complex than that suggested for use in ISO13571 [172] which is given simply by:

$$FICN = HCN^{2.36} \times \frac{t}{1.2 \times 10^{6}} \quad \textbf{(8b)}$$

Equation 8b, used in ISO13571 [172] is based on data generated for resting primates and is considered to be a good approximation for humans involved in moderate exertion. A more general expression for humans engaged in different levels of physical activity taking into account the 2.83 metabolism factor and the level of activity is given by equation 8. In this equation the change in trend line constant is given by:

$$2.43 \times 10^7 = 1.21 \times 10^6 \times 2.83 \times 7.1$$

Where:
2.83 = ratio of basal metabolic rate kJ/h/kg bodyweight between monkey and human
7.1 = resting human RMV (L/min).

NOTE:
The FICN expression used in EXODUS (equation 8) replaces the previous FICN equation used within EXODUS which was given by the following expression:

$$\mathrm{FICN} = \left(\frac{e^{\frac{[\mathrm{HCN}]}{43}}}{220} - 0.0045 \right) \times t \qquad \textbf{(8c)}$$

Typical RMV values vary in a range from 8 L/min for resting people and can be as large as 50 L/min for people involved in heavy work (see CHAPTER 3:). Times to incapacitation calculated from the two FICN models are presented in Figure 6-3. *As seen in* Figure 6-3, *given the HCN concentrations, the required time to incapacitation predicted by the old HCN model (Model II) is similar to the times predicted by the new model (Model I) for people involved in light work with RMV values 25 L/min.*

Thus for people resting, the old model (Model II) is more conservative than the new model (Model I), predicting incapacitation will occur sooner, but for people involved in heavy work, the new model (Model I) is more conservative. Both models produce approximately the same time to incapacitation for those involved in light work.

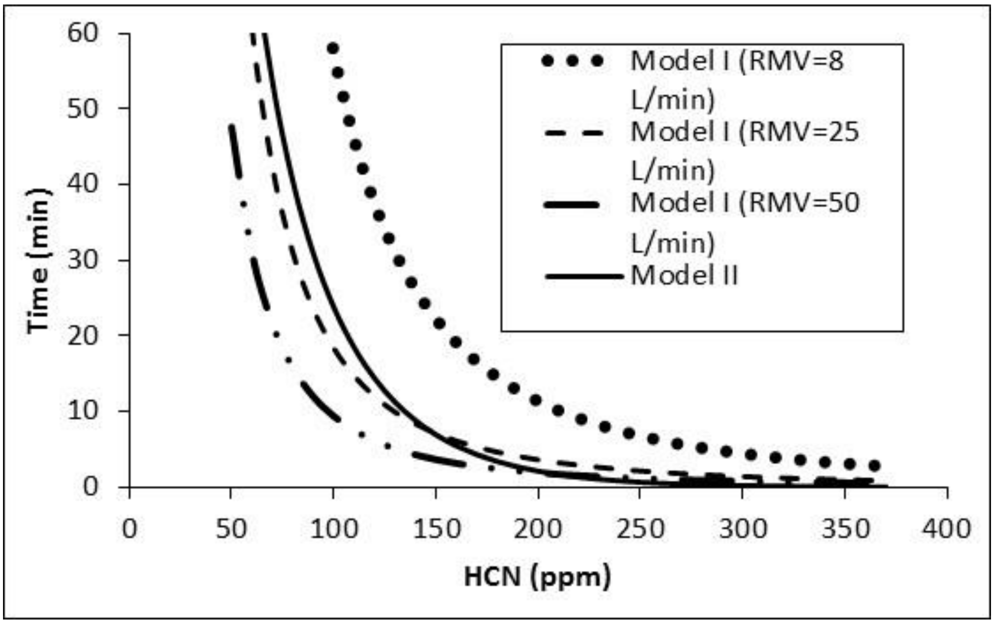

Figure 6-3: Time to incapacitation from exposure to HCN using the current (Model I) and old (Model II) model within EXODUS

(iii) Low O_2 (measured in %):

$$FIO = \frac{t}{e^{(8.13-0.54\times(20.9-o_2))}} \tag{9}$$

(iv) CO_2 (measured in %):

$$FICO_2 = \frac{t}{e^{(6.1623-0.5189\times CO_2)}} \tag{10}$$

Another effect that CO_2 has is to increase an exposed person's *RMV* and thus increase their rate of uptake of other toxic gases.

The FED model considers the combined effect of these agents in the following way,

$$FIN = (FICO + FICN + FLD) \times VCO_2 + FIO \tag{11}$$

it should be noted that the *Fractional Lethal Dose* of irritant gasses (FLD) is described in detail in CHAPTER 6:(vii)

where,

$$VCO_2 = e^{(CO_2\,/5.0)} \tag{12}$$

is a multiplicative factor which measures the increased uptake of *CO*, *HCN* and the effect on FLD due to CO_2-induced hyperventilation (see CHAPTER 3:).

NOTE:
*For a typical adult their Respiratory Minute Volume would not be expected to exceed around 70 lit/min. Hence, within EXODUS, when VCO2*RMV = 70 lit/min, VCO2 is considered to have reached its maximum value and does not increase any further. The software continually checks the VCO2*RMV value, so that if RMV changes or if VCO2 changes such that VCO2*RMV is less than 70 lit/min then the VCO2 value is allowed to vary as determined by equation 12.*

The final hazard considered is due to heat. There are two contributions to this relationship, convective heat (i.e. elevated temperature) and radiative flux,

(vi) Convected Heat:

$$FIH_c = t \times 2.0 \times 10^{-8} \times T^{3.4} \quad \textbf{(13)}$$

where T is the temperature (°C) (see Section 3.4).

NOTE:
From equation 13, the impact of convective heat would be apparent for any value of the air temperature. This means that for relatively long exposures to low air temperatures e.g. 35^oC, could result in incapacitation. To avoid this within EXODUS, a lower limit of temperature is included below which the FIH_c is not calculated. This lower limit value is set to 50^oC.

This model predicts incapacitation when convective heat temperature is at 50^{o}C at approximately 84 minutes.

(vii)Radiative Heat:

$$FIHr = \frac{q^{1.33}}{D_r} \times t \times 60.0 \quad \textbf{(14)}$$

where q is the radiative flux (kW/m^2) (see Section 3.4) and D_r is the radiative denominator. D_r is the Dose of radiation required to cause the desired effect and has units of [s(kW/m^2)$^{4/3}$]. Within EXODUS two values for D_r are provided, these represent the critical value for “pain threshold” D_r = 80 and the critical value for “incapacitation”, D_r = 1000. Both values are subjective and depend on many variables such as age of the occupant, state of health, amount and type of clothing worn, amount of skin exposed etc. As a result these values are only intended to be indicative.

NOTE:
For an average adult, if the head, neck and hands are exposed, this is the equivalent of about 14% of the body surface area. If arms and legs are also exposed, this increases the exposed surface area to approximately 61%

NOTE:
The old are more susceptible to burn injuries than the young. The predicted mortality rate from burns as a function of % area burned and age suggests that a 20 year old with 20% area burned has a 1% chance of death which increases to 5% at age 40 and 31% at age 70. For 65% area burned, the chance of death is 67%, 90% and 99% for 20, 40 and 70 year olds respectively [82].

NOTE:
Within equation 14, t is measured in minutes. To convert time to seconds within the equation we use the factor 60. Thus D_r *in equation 14 is expressed in units of* $s(kW/m^2)^{4/3}$*. However, in the literature, for example the SFPE handbook [173] and the ISO 13571 [172] it is often expressed in units of* $min(kW/m^2)^{4/3}$*. In order to covert from the* D_r *quoted in the EXODUS manual to that in other sources, divide by 60 (see Table 6-2).*

Table 6-2: Thermal Dose Required to Cause Effect values in various units

D_r associated with various end points	**EXODUS $s(kW/m^2)^{4/3}$**	**EXODUS $min(kW/m^2)^{4/3}$**	**SFPE $s(kW/m^2)^{4/3}$**	**SFPE $min(kW/m^2)^{4/3}$**
Pain	*80*	*1.33*	*80*	*1.33*
Incapacitation	*1000*	*16.67*	*600*	*10*
Death	-	-	*1000*	*16.67*

NOTE:
Within EXODUS, the D_r *value used to define Incapacitation is the same value as SFPE uses to define Death.*

NOTE:
Within EXODUS, when the Incapacitation threshold is specified for the thermal radiation fractional effective dose, individuals will be MORE tolerant to thermal radiation than if the SFPE threshold is used. Thus the SFPE Incapacitation threshold is more conservative than the one used in EXODUS i.e. will predict incapacitation at a LOWER thermal radiation dose.

NOTE:
When using D_r = 1000 within *EXODUS, a value of* $FIH_r = 0.7$ *is the equivalent of* $FIH_r = 1.0$ *when using a* $D_r = 600$ *(i.e. the SFPE 'incapacitation' endpoint). Thus, within EXODUS when using* $D_r = 1000$ *the value of* $FIH_r = 0.7$ *could thus be used as an indication of:*

- *Likely incapacitation if the SFPE incapacitation threshold were used or,*
- *The onset of severe injury amongst 'survivors'.*

In reality, the degree of thermal radiation a person is subjected to is likely to influence the person's behaviour as well as their physical ability to evacuate. It is thus likely that several values of D_r may be required to accurately simulate the evacuation process. For instance, during a fire evacuation scenario, a person may be less likely to elect to enter an environment in which they will receive a moderately high dose of thermal radiation (e.g. that required to cause the onset of pain). Faced with such a situation, the person may be more likely to select an alternative exit route if one exists or remain in a place of relative safety. However, if the person is already exposed to the fire environment and is in the process of evacuating, they are likely to tolerate higher exposures to thermal radiation than that required to cause pain before they are rendered incapable of evacuating. In this case it may be appropriate to use higher values of D_r (e.g. the incapacitation value).

NOTE:
The influence of thermal radiation and elevated temperatures on human behaviour during evacuation is currently poorly understood and as a result no attempt has been made to model the influence of thermal radiation on the decision making process.

To select the appropriate value of D_r used in a simulation it is necessary to consider the purpose of the *FIHr* equation within EXODUS. The primary intention of the *FIHr* equation is to indicate when the occupant is likely to be unable to continue to evacuate efficiently due to exposure to thermal radiation.

The so-called "pain threshold" value of D_r (D_r = 80), is the equivalent to a thermal radiation exposure of 2.5 kW/m^2 for 24 seconds. This is a conservative value for D_r and is the value recommended by Purser [62]. The rationale for using this value is that the occupant has received a cumulative dose of thermal radiation equivalent to that required to cause the onset of pain (but not burns) and as a result is rendered incapable of continuing to evacuate effectively.

An arguably more representative value for D_r is the so-called "incapacitation threshold", which in EXODUS is taken as D_r = 1000. This is equivalent to an exposure of 8.25 kW/m^2 for one minute. A cumulative dose of radiation equivalent to 8.6 kW/m^2 for one minute is expected to cause mortality in 1% of the population [81]. The rationale for using this value is that the occupant has received a severe cumulative dose of thermal radiation and is thus likely to be suffering from second degree burns to exposed skin resulting in severe pain and thus will find it difficult to continue to evacuate.

To compare the impact of the two formulations, consider an occupant exposed to a constant radiative heat flux of approximately 2.5 kW/m^2. An exposure to 2.5 kW/m^2 for 24 seconds is the recommended tolerance time according to Purser [62]. According to Hymes et al [81] an exposure to:

- 2.6 kW/m^2 for 1 minute is the lower limit threshold for second degree blistering of exposed skin,
- 2.6 kW/m^2 for 5 minutes can result in a 1% mortality to average dressed subjects,
- 2.8 kW/m^2 for 5 minutes is the lower limit for second degree plus burns (>0.1mm deep) to exposed skin,
- 3.3 kW/m^2 for 5 minutes is the lower limit for the melting of nylon and polyester fabrics and is the lower limit of potentially severe burns from molten fabrics,
- 5 kW/m^2 for 5 minutes results in third degree burns to exposed skin and can result in 50% mortality to average dressed subjects,
- 15.5 kW/m^2 for 1 minute can result in an expected 50% mortality,
- 21 kW/m^2 for 1 minute is the lower limit for the auto-ignition of everyday clothing.

Within EXODUS, using D_r = 80, an occupant exposed to 2.5 kW/m^2 for 24 seconds is considered incapable of evacuating. Using D_r = 1000, an occupant exposed to 2.5 kW/m^2 for 5 minutes is considered incapable of evacuating. For comparison purposes, using D_r = 600 (SFPE definition of the incapacitation threshold), an occupant exposed to 2.5 kW/m^2 for 3.4 minutes is considered incapable of evacuating.

NOTE:
While EXODUS does not use a D_r to represent the death endpoint, using the expected 50% mortality figure from Hymes et al [81] i.e. 15.5 kW/m^2 for 1 minute can result in an expected 50% mortality, this is the equivalent to D_r = 2300 min(kW/m^2)$^{4/3}$ or D_r = 38 sec(kW/m^2)$^{4/3}$.

NOTE:
Using D_r = 1000, FIH_r = 2.3 is equivalent to 50% mortality (15.5 kW/m^2 for 1 minute).

Below a critical value of the radiant heat individuals can withstand indefinite radiant heat exposures. Thus the determination of the FIH_r is only activated when the instantaneous thermal radiation (kW/m^2) an agent is exposed to exceeds the *Radiant Heat Threshold (RHT)* value. The default value of the RHT is 1.7 kW/m^2 [122,123]. However, the RHT parameter can be set by the user.

NOTE:
Within earlier versions of buildingEXODUS the RHT was an internal parameter set to the current default value of 1.7 kW/m^2.

The FED model considers the combined effect of these agents in the following way,

$$FIH = FIH_r + FIH_c \tag{15}$$

When *FIN* or *FICO2* or *FIH* equal or exceed 1.0, the affected occupant is assumed to be incapacitated. The EXODUS model considers fire hazard data located at two heights, head and near floor height (see Section 2.3 and CHAPTER 5:).

While the Purser model is typical of FED models, other formulations have been suggested, Speitel [26, 27] for example has developed an alternative model. In addition to the quantities specified in the Purser model, Speitel considers the gases, HCL, HF, HBr, CH_2CHO (*Acrolein*) and NO_2. Furthermore, expressions for the *CO* and Heat contribution to the FED calculation are significantly different from that specified in the Purser model.

In Purser's model the FIH_C acquired each minute (equation 13) is based on data using subjects with exposed skin, whereas in the Speitel model the *FIH* calculation, is based on data using clothed subjects.

$$FIH_c = t \times 2.4 \times 10^{-09} \times (T°C)^{3.61} \tag{16}$$

The Purser model predicts incapacitation at significantly lower temperatures than the Speitel model. For example, according to equation 13, a one-minute exposure to 185^{o}C results in incapacitation, whereas using equation 16 temperatures in excess of 240^{o}C are required to produce the same result.

A possible deficiency in both models concerns the exclusion of the thermal effects due to humid rather than dry air. The incapacitating effects of air with a high water vapour content are more severe than dry air as it reduces heat loss through sweat and delivers more heat to exposed skin. Furthermore, due to its higher heat capacity, inhaled hot air with a high water vapour content can cause more severe damage to the respiratory tract than dry air at the same temperature [23].

Both the Purser and Speitel models incorporate a factor that takes into account the increased respiration rate that results from the presence of CO_2. The hyperventilation factor, VCO_2 (see equation 12), is used in the Purser model to represent the increase in uptake of *CO* and *HCN* and in the Speitel model it serves a similar function for *CO*, *HCN*, HCL, HF, HBr, NO_2 and CH_2CHO (*Acrolein*). Using equation 12, a 5% atmosphere of CO_2 will increase the *RMV* by 2.72. This will have a significant effect on the *FIN* calculation in both models.

While CO_2 induced hyperventilation has a measurable effect under laboratory conditions and possibly in circumstances where the subject is unaware of the fire, such as a sleeping victim of a domestic fire, it is unclear if it is appropriate that it should be factored into evacuation models in its present form. This is particularly true if the *RMV* used in the *CO* equation is already set to a large value appropriate to heavy work. Clearly the VCO_2 factor has a dramatic impact on the outcome of the simulation and may lead to a severe over estimation of the number of fatalities. To address this problem, within buildingEXODUS the product VCO2*RMV cannot exceed 70 lit/min (see earlier note below equation 12).

Table 6-3: CO_2 induced hyperventilation factors generated by equation 12

CO_2 (%)	VCO_2 using Equation 12
1	1.22
3	1.82
5	2.72

In addition, the user has the option of setting VCO_2 to 1, thereby removing it from consideration. If the *Purser* option is selected the formulation specified in (11) is used, so that the VCO_2 calculation directly influences the *FICO*, *FICN* and *FLD* values. If the alternative model is selected (the "1" option) then the formulation is simplified to:

$$FIN = FICO + FICN + FLD + FIO \tag{17}$$

Given the above choice in toxicity model composition, the TOXICITY sub-model within EXODUS provides users with several options. These are,

(1) The standard Purser FED model (i.e. equations 7-11), with the VCO_2 hyperventilation model enabled, influencing the intake of CO, HCN and irritant gases

(2) As (1) with the FIH_c equation replaced with the Speitel formulation (i.e. equation 16 replacing equation 13),

(3) As (1) with $VCO_2 = 1.0$, so equation 17 replaces equation 11.

(4) Any combination of the above.

The default setting in the FED model is the Purser equation with the product VCO2*RMV capped to 70 lit/min and the VCO2 capped to an appropriate value based on the current RMV.

(vii) Irritant Gases

A significant component of the fire effluent produced by the fires is the irritant gases. The irritant model implemented within buildingEXODUS is based on the model originally developed by Purser from a variety of experimental data-sets [23, 83-87, 173] to a particular irritant gas as well as the accumulated dose that is acquired during the evacuation process. The Purser data has been pooled with the data produced by Jin [17] to form an approximation of the impact of an exposure to irritant substance during an evacuation.

The Purser irritant model represents the impact of the following irritant gases: HCl, HBr, HF, SO_2, NO_2, CH_2CHO (*Acrolein*) and HCHO (*Formaldehyde*). The evolution and propagation of these irritant gases is represented within buildingEXODUS in an identical manner to the other fire hazards (i.e. their development can be manually described, or may be automatically imported from the CFAST zone model).

NOTE:
CFAST is only able to produce data relating to HCL of the irritant products that may be represented within EXODUS.

In the irritant model, each occupant has two additional indices describing their interaction with a particular irritant gas. For each of the irritant gases, the impact of the concentration and the dose received is stored. These are then summed to form the *Fractional Irritant Concentration* (*FIC*) and *Fractional Lethal Dose* (*FLD*) index for each individual. The FIC reflects the instantaneous impact of the concentration of the irritant gases to which the individual is exposed. The FLD reflects the dose (i.e. the cumulative impact) of the irritant gases upon an individual (this is analogous the FED model in relation to the representation of the narcotic gases). FIC represents the impact of the irritant exposure on the individual, which directly influences their travel speed (and long term health), whereas the FLD is a dose that determines whether the individual has survived or succumbed to the exposure.

Each individual within the population will be susceptible to differing levels of instantaneous exposures of the various irritant gases. The value depends on natural variations in the population such as age, gender, state of health, strength, etc. The range of values for the concentration to cause effect, or tolerance values typically found in the population are presented in Table 6-4 based on data presented by Purser [87, 173] and ISO13571 [172]**.** It is noted that the ISO13571 incapacitation concentrations are generally less conservative than those recommended by the SFPE. For evacuation, two types of irritant endpoints are considered, escape impairment and incapacitation. It is important to clearly define what is meant by these two irritant endpoints:

- **Escape impairment due to irritants**: Concentration of fire irritants severely affects escape capability of people exposed. This will result in a severe reduction in the travel speed of those exposed, but they can continue with their evacuation unless the travel speed is reduced to zero. However, the person can recover from this situation if the irritant concentration is reduced and hence their FIC becomes less than 1.0.
- **Incapacitation due to irritants**: Concentration of fire irritants is so painful and results in such breathing problems so as to effectively permanently incapacitate those exposed. This will result in the travel speed of those exposed effectively being reduced to zero.

The effects of the irritant fire gases represent an immediate effect due to exposure, their impact on the exposed individual is expected to decrease if the exposure concentrations are decreased, by for example, if the exposed person moved out of the contaminated area or if the spatial atmospheric concentration in the immediate area decreases. However, if the person is considered incapacitated due to exposure of irritants (FIC = 4.5 if using the Escape Impairment

TF or FIC = 1.0 if using the Incapacitation TF), the person will not be able to recover even if the irritant concentration decreases.

NOTE:
For ethical reasons, most experimental work has been carried out on animals, mostly rodents. As a result there are considerable uncertainties in extrapolating from the existing animal data to the concentrations likely to cause an effect in the average healthy human subjects. Furthermore, there is likely to be considerable variation in susceptibility within the human population.

The FIC values presented in Table 6-4 are intended to represent those values to cause the end effect for 50% of the population. Thus, for HCL, 50% of the population is likely to experience severe mobility impairments (escape impaired) when experiencing a concentration of 200 ppm. However, within a population there may be individuals that may be susceptible to concentrations as low as 60 ppm (0.3 times the population mean) while others may be susceptible to concentrations as high as 400-600 ppm (2 to 3 times the population mean). Similarly, considering the incapacitation end point for HCL, 50% of the population is likely to be incapacitated when experiencing a concentration of 900 ppm. However, within a population, there may be individuals that may be incapacitated when experiencing 270 ppm while others may be incapacitated by concentrations as high as 1800-2700 ppm.

For the FIC calculations used in buildingEXODUS there are five options available to distribute the tolerance factor (TF), the default option utilises an assumed distribution related to the SFPE Escape Impaired effect (see below). Alternatively, the TF values can be fixed (i.e. the same for the entire population), with the values either being set to the SFPE Escape Impaired values, the SPFE Incapacitation values or the ISO 13571 values (see Table 6-4). Finally, the fifth option enables a user defined distribution to be specified for each irritant, which can be either a fixed value or a distribution.

- **SFPE Escape Impaired (Distribution):**
  Each member of the population is randomly assigned a *TF* for each irritant gas component from an assumed distribution. This forms the denominator for the FIC, with the current irritant exposure concentration for a particular gas that the agent is exposed to forming the numerator in the FIC calculation. A TF distribution model is provided to reflect that individuals within a population may be more or less tolerant to the irritant gases. As the nature of the TF distribution is not known, an arbitrary distribution model is used. Within the distribution model, the lower and upper values of the *TF* follow SFPE guidance [173], with the lower value being 0.3 times the mean value and the upper value being 2 times the mean value. The mean value is derived from Table 6-4. The population is arbitrarily divided into one of three TF bands, a lower-band containing 15% of the population, representing the most susceptible proportion of the population, a mid-band containing 70% of the population and an upper-band containing 15% of the population, representing the most resilient. The assigned band is assumed to apply for all the irritant products. Within each band, the tolerance factor is distributed according to a random uniform distribution between an upper and lower limit.
    - Within the mid-band, these limits are determined by taking +/-25% from a mean value. By default the mean value is taken as the SFPE Escape Impaired value in Table 6-4.

- Within the lower-band, the lower limit is assumed to be 0.3 times the mean value of the mid-band. The upper limit is assumed to be equal to the lower limit of the mid-band.
- Within the upper-band, the upper limit is assumed to be 2 times the mean value of the mid-band. The lower limit is assumed to be equal to the upper limit of the mid-band.

- **SFPE Escape Impaired (Fixed):**

Each member of the population is assigned a fixed *TF* for each irritant gas component corresponding to the *SFPE Escape Impaired* mean values (see Table 6-4). Thus using this approach, each member of the population has an identical *TF* for each irritant (i.e. all agents will be assumed to have a TF for HCL of 200ppm, 24ppm for SO_2 etc.).

- **SFPE Incapacitation (Fixed):**

Each member of the population is assigned a fixed *TF* for each irritant gas component corresponding to the *SFPE Incapacitation* mean values (see Table 6-4). Thus using this approach, each member of the population has an identical *TF* for each irritant (i.e. all agents will be assumed to have a TF for HCL of 900ppm, 120ppm for SO_2 etc.).

- **ISO 13571 Incapacitation (Fixed):**

Each member of the population is assigned a fixed *TF* for each irritant gas component corresponding to the *ISO 13571 Incapacitation* mean values (see Table 6-4). Thus using this approach, each member of the population has an identical *TF* for each irritant (i.e. all agents will be assumed to have a TF for HCL of 1000ppm, 150ppm for SO_2 etc.)

- **User Defined:**

The user can specify the nature of the TF distribution to be applied for each irritant. This can either be a fixed value, a range or a full distribution.

An FIC value is calculated for each of the irritant gases that the individual is currently exposed to. If any of these individual FIC values exceeds 1.0 or if the combined impact of all of these FIC ratios is greater than 1.0, then the individual is assumed to succumb to the impact of the particular irritant component or the combined effect of the irritant gases. The FIC is given by equation 18:

$$FIC = \sum_{X=1}^{n} \frac{C_X}{TF_X} \quad \textbf{(18)}$$

Where x represents each of the irritant gases present. The numerator in the summation is the instantaneous concentration (C) of irritant the person is exposed to, measured in ppm, and the denominator is the Tolerance Factor (TF) derived from Table 6-2.

The default *TF* model represents an endpoint of *Escape Impairment*; this means that when an agent's FIC is equal to one, they are considered unable to continue. It is noted that the SFPE *TF* representing *Incapacitation* is between 4.5 and 5.0 times greater than that for escape impairment. Thus, if the FIC of the impaired individual reaches a default value of 4.5 the person is considered to be incapacitated due to the instantaneous effects of the irritant gases.

Thus, for each of the five TF model options, in addition to the tolerance factors for each irritant a corresponding *Critical Irritant Incapacitation Factor (CIIF)* value is also defined. The *CIIF* value defines the FIC value at which agents are assumed to become incapacitated, and hence taken out of the simulation. The *CIIF* value is dependent on the TF defining the endpoint that is used in the simulation (i.e. escape impairment or incapacitation). By default, when the escape impairment option is used, the *CIIF* value is 4.5 as the TF for incapacitation is approximately 4.5 times that for escape impairment. Thus when an agent's FIC = 4.5 the individual is considered incapacitated even though the escape impairment model is being used. If the Incapacitation endpoint is specified, *CIIF* is automatically set to 1.0 (i.e. agents are incapacitated when FIC = 1.0). Hence the *CIIF* value can be viewed as the ratio between the current effect being modelled (i.e. by default *Escape Impaired*) and *Incapacitation*.

NOTE:
The CIIF value is automatically set depending upon the TF model selected:

- *If using the SFPE Escape Impairment endpoint for irritant exposure, CIIF = 4.5*
- *When using Incapacitation as the endpoint for irritant exposure, CIIF = 1.0*
- *When a User Defined TF model is defined the CIIF is automatically set to 0.0. In this instance it is the users responsibility to set the CIIF based upon the effect that they are modelling (i.e. Escape Impairment etc.), and hence the relative difference between the TF values defined and those defining Incapacitation.*

NOTE:
Versions of buildingEXODUS prior to v6.3 only had the option of a random uniform distribution for the FIC Tolerance Factor Model, with the endpoint assumed to be Incapacitation. The specified range for each irritant varied considerably e.g. for HCL it ranged from 200 ppm to 5000 ppm based on values specified in early editions of the SFPE handbook [25]. Thus the endpoint due to exposure to irritant gases in the earlier versions of EXODUS was incapacitation and occurred at a higher concentration of irritant.

In the current release there are five methods for assigning the TF, with the endpoints being either Escape Impairment or Incapacitation. By default the endpoint is assumed to be Escape Impairment, this means that while the exposed individual is not incapacitated by the irritant gases, they are considered to be unable to continue (when FIC >= 1.0). In this condition the agent may become incapacitated by heat or other toxic gases. If the irritant concentration at their current location decreases so that FIC < 1.0, they will be able to resume the evacuation process. However, if the irritant concentration at their current location increases, they may become incapacitated by irritant gases if FIC >= 4.5.

Table 6-4: Concentration (TF) and Dose (CD) values required to cause effect for irritant fire products [173].

	Tolerance Factor (TF)			Critical Dose (CD)	
Component	**SFPE Escape Impaired (ppm)**	**SFPE Incapacitation (ppm)**	**ISO13571 Incapacitation (ppm)**	**SFPE (ppm.min)**	**AEGL-3 (ppm.min)**
HCL	200	900.00	1000	114,000	6,300
HBr	200	900.00	1000	114,000	7,500
HF	200	900.00	500	87,000	1,860
SO_2	24	120.00	150	12,000	900
NO_2	70	350.00	250	1,900	750
CH_2CHO (*Acrolein*)	4	20.00	30	4,500	2,100
HCHO (*Formaldehyde*)	6	30.00	250	22,500	--

Furthermore, as the combined FIC value increases, the *Agility* and *Mobility* of the individual decreases. The reduction in *Mobility* in turn reduces the travel speed of the individual (it should be borne in mind that within buildingEXODUS the *Mobility* attribute of an individual is a coefficient of their travel speed). The following equation (adapted from the Weibull function provided by Purser [83]) represents the function that approximates the sigmoidal function produced by Purser [83].

$$Mobility\ Degradation\ Factor_{FIC} = \frac{\left(e^{-((FIC*1000)/160)^2} + (-0.2*FIC+0.2)\right)}{1.2} \quad \textbf{(19)}$$

The relationship between the FIC level and the *Mobility Degradation Factor*$_{FIC}$ (i.e. a multiple of the individual's original *Mobility* level, see Section 3.1 part 2) is described in Figure 6-4.

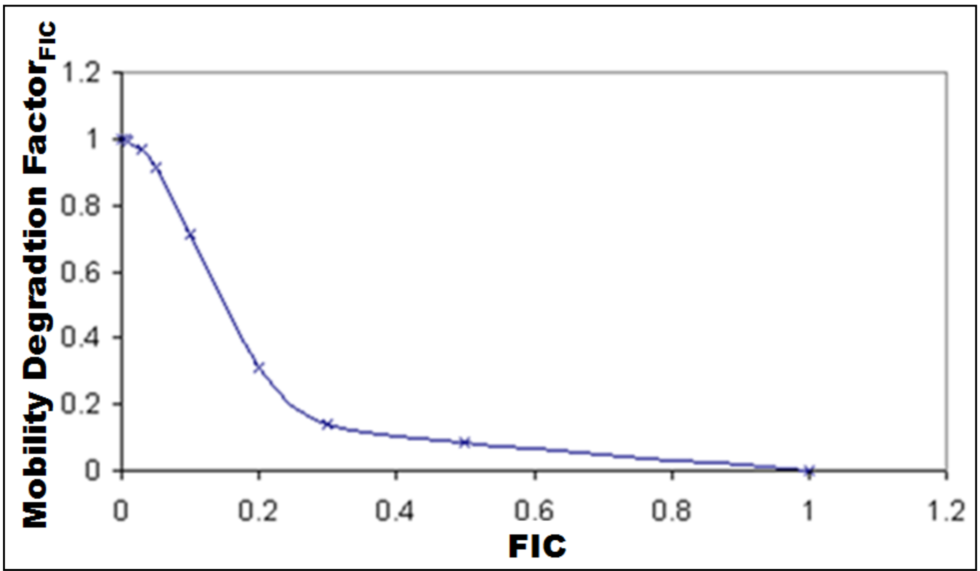

Figure 6-4: Impact of FIC upon the Mobility Degradation Factor.

For example, consider an individual exposed to an atmosphere of HCL, and for simplicity further assume that the fixed tolerance factor Escape Impairment model is used. Thus when the individual is exposed to a concentration of 200 ppm of HCL, their FIC = 1.0, they are assumed to be escape impaired i.e. immobilised and their mobility degradation factor (from equation 19) will be 0.0 i.e. they cannot move. Furthermore, they will be able to recover should the atmospheric concentration of HCL in their location decrease below 200 ppm.

Now consider an individual exposed to an atmospheric concentration of 60 ppm of HCL (200 * 0.3). Their FIC = 0.3 (60/200) and using the Escape Impairment TF this person will experience severe escape impairment. This corresponds to a mobility degradation factor (from equation 19) of 0.14 i.e. their walking speed has been greatly reduced, which is consistent with the 'escape impairment' description of their situation. However, should the local HCL concentration decrease or should they move to a region with lower HCL concentration, they may recover resulting in a smaller value of FIC and a larger mobility degradation factor and hence a larger travel speed.

When the FIC is equal to 4.5 i.e. the default Critical Irritant Incapacitation Factor (CIIF) the agent is considered incapacitated by HCL (i.e. the individual is exposed to a HCL concentration equivalent to the incapacitation value so that FIC = 900/200) and cannot resume their evacuation, even if their FIC falls below 4.5.

If the above example is repeated but with the endpoint of incapacitation due to irritants rather than escape impairment, then the user must specify a TF for HCL of 900 ppm and a *CIIF* value of 1.0. Thus when the individual is exposed to a concentration of 900 ppm of HCL, their FIC = 1.0 (i.e. the individual is exposed to a HCL concentration equivalent to the incapacitation value, thus FIC = 900/900). Furthermore, as the FIC is equal to the *CIIF* (i.e. CIIF = 1.0) the agent is considered incapacitated by HCL and removed from the simulation therefore the agent cannot resume their evacuation, even if their FIC falls below 1.0.

Now consider an individual exposed to an atmospheric concentration of 270 ppm of HCL. Their FIC = 0.3 (270/900) and using the Incapacitation TF this person will experience severe mobility reduction. This corresponds to a mobility degradation factor (from equation 19) of 0.14 i.e. their walking speed has been greatly reduced. However, should the local HCL concentration decrease or should they move to a region with lower HCL concentration, they may recover resulting in a smaller value of FIC and a larger mobility degradation factor and hence a larger travel speed.

The fractional lethal dose (FLD) is calculated in much the same way as the FED calculations for the narcotic gases. The cumulative exposure to irritant gas x is calculated as the product of the exposure concentration of irritant gas x (C_x) and the exposure time (t) – $C_x t$ - and then divided by the critical dose for irritant gas x (CD_x), to give (C_x x t)/CD_x. This is then summed over all the irritant species at a given time period. As the concentration that the person is exposed to is changing over time, the product is summed over the duration of the entire exposure to provide the FLD as shown in equation 20:

$$FLD = \sum_{t=0}^{n} \sum_{x=1}^{n} \frac{(C_{x,t} \text{ x } t)}{CD_x} \quad \textbf{(20)}$$

Where $C_{x,t}$ is the concentration of species x over time period t.

Thus when FLD = 1.0 fatality is expected to occur. However, death is not likely to occur during evacuation but may occur several hours or several days after exposure. The FLD is thus not intended to be used as a critical factor describing evacuation or as a design limit but for estimations of the extent to which post-exposure deaths from lung oedema and inflammation are likely to occur [173].

As with the FIC TF, there may be individuals within a population that are more or less susceptible to the irritant gases. Thus the CD values are likely to be distributed within the population. However, population distributions for CD are not currently available. The CD for each of the irritant gases is provided in Table 6-4. The SFPE values represent the dose over 30 minutes to cause fatality in 50% of the population. Thus, using the SFPE CD for HCL, 50% of the population is likely to perish when they receive a critical dose of 114,000 ppm.min. Also provided in Table 6-2 are the CD based on the AEGL-3 – which in contrast to the SFPE value suggests a critical dose of 6,300 ppm.min is required to cause death.

The AEGL is the Acute Exposure Guideline Levels produced by the US Environmental Protection Agency. They provide guidance for the assessment of likely acute effects of human exposure to a range of substances during accidental industrial releases. The AEGL-3 level represents an airborne concentration which is predicted that the general population, including susceptible individuals could experience life threatening adverse health effects or death. The AEGL values in Table 6-2 are also based on a 30 minute exposure.

Clearly, the AEGL-3 values are much more conservative than the SFPE values as these take into consideration susceptible individuals. The choice of values to be used is dependent on the nature of the simulation to be conducted.

NOTE:
The FLD is not likely to impact the outcome of an evacuation simulation but may be used to determine post-incident mortality amongst those who have survived the evacuation.

As population distributions for the CD are not currently available, within the FLD calculations used in buildingEXODUS a fixed CD model is used. There are three options for this:

- **SFPE FLD Critical Dose:**
  Each member of the population is assigned a *CD* for each irritant gas component from the SFPE distributions presented in Table 6-4. This forms the denominator for the FLD. Thus using this approach, each member of the population has an identical *CD* based on the SFPE values. This is based on values that are likely to cause fatality in 50% of the population. By default, the *CD* is selected from *SFPE Critical Dose* column.

- **AEGL-3 FLD Critical Dose:**
  Each member of the population is assigned a *CD* for each irritant gas component from the AEGL-3 distributions presented in Table 6-4. This forms the denominator for the FLD. Thus using this approach, each member of the population has an identical *CD* based on the AEGL-3 values and so is very conservative, taking into consideration particularly vulnerable members of the public.

- **User Defined FLD Critical Dose:**
  Each member of the population is assigned a *CD* for each irritant gas component based on the user specification.

NOTE:
*Versions of buildingEXODUS prior to v6.3 only had the SFPE FLD Critical Dose (CD) option. These values were arbitrarily distributed with the actual value for each individual being selected from a range of 50%*median to 200%*median.*

NOTE:
In versions of buildingEXODUS prior to v6.3, if FLD = 1.0 the agent was assumed to be incapacitated and not able to continue the evacuation and so were added to the mortuary. This is no longer the case, hence FLD has no impact on an agent's ability to evacuate but can be used to assess the likelihood of them surviving post evacuation.

The impact of narcotic gases is currently modelled within the buildingEXODUS model as is the impact of obscuration due to the presence of smoke upon the ability of individuals to move through the environment. The impact of this and the irritant gases are determined independently, with the larger of the two impacts assumed to be the effect upon the individual.

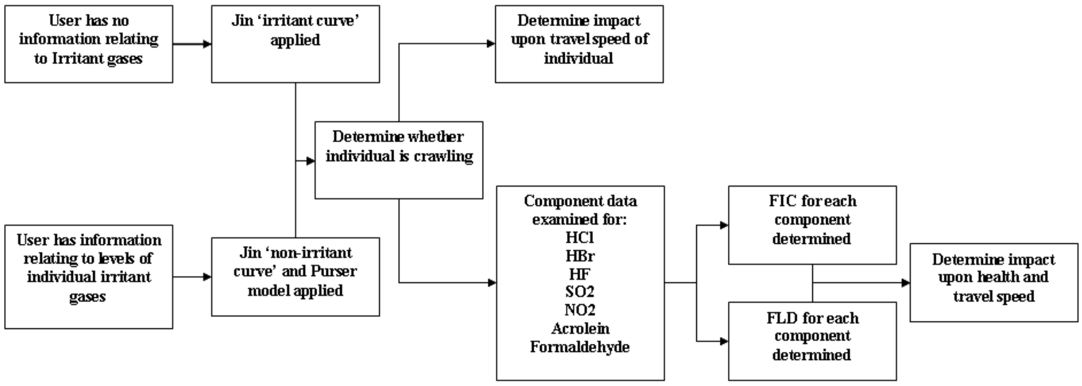

Figure 6-5: Use of the irritant model implemented within buildingEXODUS.

NOTE:
To summarise the impact of FIC on the state of an agent, assuming the default SFPE Escape Impaired values are used as the basis for the TF values):

- *As FIC increases and approaches 1, the travel speed of the agent is decreased according to equation 19.*
- *When the FIC is equal to 1, the agent is severely impaired and cannot continue. Their travel speed is reduced to effectively zero according to equation 19.*
    - *However, if the agents FIC should fall below 1, due to a reduction in the concentration of irritant product at their location, the agent can again resume their evacuation with reduced travel speed determined by equation 19.*
- *When the FIC is equal to 4.5 the agent is considered incapacitated by irritant product and cannot resume their evacuation, even if their FIC falls below 4.5.*

To summarise the impact of FIC on the state of an agent, assuming the SFPE Incapacitation values are used as the basis for the TF values):

- *As FIC increases and approaches 1, the travel speed of the agent is decreased according to equation 19.*
- *When the FIC is equal to 1, the agent is incapacitated and they are placed in the mortuary.*

In order to determine the impact on occupant travel speeds, the model is separated into two separate components, the use of which is dependent upon whether or not the presence of irritant gases are to be considered in the simulation. If irritant gases are not to be explicitly considered in the simulation a *simplified model* is used, while a more *comprehensive model* is used if irritant gases are to be considered in the simulation (see Figure 6-5).

6.1 Simplified model

If the irritant gases are not explicitly represented in the fire hazard (i.e. if the *Irritant* model is disabled), the Jin 'irritant' data-set is used to describe the complete impact of the smoke and irritant gases on the movement rates of exposed individuals. This does not require the specification of irritant gas concentrations. The applied relationship is intended to approximate the reduction in the individuals travel speed due to the impact of irritant smoke (including the obscuration effect of smoke). However, it does not directly impact upon the well-being of the exposed individuals. This effect is applied purely as a reduction factor which is used as a coefficient of their maximum attainable travel speed. This coefficient is maintained at a constant level until smoke concentrations of 0.1 /m, after which point the calculation is made according to the following function,

$$Mobility\ Degradation\ Factor_{IrritantSmoke} = -2.0814K^2 - 0.375K + 1.0648 \qquad \textbf{(21)}$$

where *K* represents the extinction coefficient (/m) of the smoke. This equation is valid between extinction coefficients of 0.1 and 0.5/m. This function was is derived from the work of Jin [17,18]. For a smoke concentration of 0.45 /m, the above formulation decreases the *Mobility Degradation Factor*$_{IrritantSmoke}$ to approximately half of its original value. For smoke concentrations above 0.5/m (i.e. *Mobility Degradation Factor*$_{IrritantSmoke}$ values less than 0.36) individual escape abilities are severely limited and the model assumes that the individual is forced to crawl or is travelling at a speed which is comparable with crawling. Jin did not observe people crawling during his experimental trials, which extended up to conditions involving an extinction coefficient of 0.5/m. Crawling is therefore only assumed to be likely outside of this experimental envelope, which corresponds to a reduction in travel speed of approximately 70% (based on the Jin data).

6.2 Comprehensive model

If irritant gas concentrations are specified (i.e. the *Irritant* model is enabled) a more comprehensive model can be utilised in relation to the impact of smoke obscuration upon the *Mobility* of the individual (see Figure 6-5). This examines the concentration of several irritants and determines the impact upon the individual accordingly. Initially the Jin data relating to experiments involving 'non-irritant' gases is used. This is assumed to represent the impact of the visual obscuration of the smoke alone, without representing any of the irritant effects of the smoke present. This produces a slight decrease in travel speed resulting from the obscuration affects of the smoke at sufficiently high levels (see Figure 6-6).

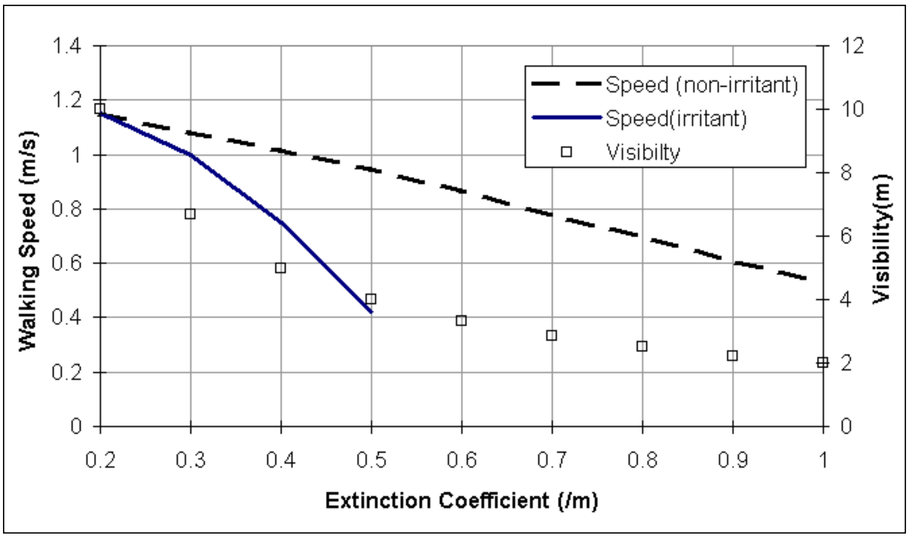

Figure 6-6: The walking speed of volunteers through irritant (solid line) and non-irritant (dashed line) smoke. The squares indicate the approximate visibility levels, purely according to the extinction coefficient of the smoke (redrawn from the original [10]).

The derived curve which represents this phenomenon is

$$Mobility\ Degradation\ Factor_{NonIrritantSmoke} = -0.161K^2 - 0.488K + 1.105 \qquad \textbf{(22)}$$

where K is the extinction coefficient of the smoke and is valid between extinction coefficients of 0.2/m and 1.0/m. This curve represents the impact of the environment, specifically in relation to reduced visibility. The effects of the irritants and narcotics (in terms of the reduction in the occupants travel speed) are then combined with this effect. The most severe impact upon the individual's *Mobility* and health is then adopted (i.e. the smallest *Mobility Degradation Factor* is then used in the calculation of *Mobility* and hence also *Travel Speed*).

buildingEXODUS is also capable of representing the impact of the evacuees being forced to crawl. To implement the crawling behaviour in relation to the non-irritant curve, a level of consistency had to be achieved in relation to the assumptions based on the use of the 'irritant' Jin data (where we assumed that the evacuees crawl only after the end of the data-set, i.e. beyond the conditions of the original experiment) and therefore the evacuees should crawl at 1.0/m, as the 'non-irritant' curve extends further than the 'irritant' curve. However, for an evacuee to crawl whilst using the non-irritant curve (i.e. the effect of the visual obscuration of the smoke alone), the FIC value (i.e. the combined effect of the irritant gases) should also be examined to determine whether it has reached the critical value of 0.3 (this is chosen as at this stage the effect on *Mobility* approximates that of crawling anyway). Therefore either the extinction coefficient or the FIC need to be checked to determine whether an occupant was crawling. This then replicates the behaviour in relation to the original Jin 'irritant' curve which represents all of the possible effects of the irritant gases and the smoke. Once the occupant starts to crawl the maximum that their speed can attain is their crawling speed (irrespective of whether the individual is actually crawling). If the combined effect of the irritant/smoke and narcotic gases on their *Mobility* has reduced their speed still further then this speed is adopted.

The final parameters, which can be accessed through the toxicity sub-model, are called the *Triggering Temperature* and *Triggering Smoke Concentration*. These parameters apply to the entire population. They represent the critical temperature and smoke concentrations at which an occupant's response time attribute is overridden. When the temperature or smoke concentration at the location equals or exceeds these values the person at that location will begin to evacuate regardless of his response time.

Attribute	: *Triggering Temperature.*
Range	: 0 - 1000^{o}C.
Default	: 50^{o}C.
Influenced by	: None.
Influences	: Behaviour.
Used in level	: C
Note	: Temperature at which an individual will begin evacuation, regardless of response time attribute.

Attribute	: *Triggering Smoke Concentration.*
Range	: 0.0 - 1.0 /m.
Default	: 0.1 /m.
Influenced by	: None.
Influences	: Behaviour.
Used in level	: C
Note	: Smoke concentration at which an individual will begin evacuation regardless of response time attribute.

CHAPTER 7: BEHAVIOUR SUB-MODEL

The BEHAVIOUR sub-model is the most complex of all the sub-models and is utilised by EXODUS during *SIMULATION* mode. The behaviour sub-model - which operates on two levels, *global* and *local* - determines an occupant's response to the evacuation scenario. The *global* behaviour provides an overall escape strategy for the occupants while the *local* behaviour governs the occupants' responses to their current situation. While attempting to implement the global strategy, an individual's behaviour can be significantly modified by the dictates of their local behaviour.

7.1 Global Behaviour

In the current release of EXODUS there are two aspects to *global* behaviour. The first involves occupants implementing an escape strategy that leads to their direct escape. The second aspect involves occupants completing a set of tasks prior to their evacuation.

Global behaviour resulting in direct evacuation can result in an escape strategy that leads occupants to exit via their nearest serviceable exit, an assigned exit, or an exit based on their knowledge of the structure.

Several methods are provided by which these strategies can be implemented. The default method can be used when directing occupants to their nearest serviceable exit or most familiar exit. This strategy is achieved through the use of the node potentials and biased exit-potentials (see Section a). The second method is used when directing occupants to an assigned exit (which may be the exit that the occupant is most familiar with). This involves occupants travelling to a particular target exit that has been specified by the user, regardless of the underlying potential map (see Section b). The third method is used when occupants are modelled as relying on their personal knowledge of the structure. Using this approach, the occupant's do not follow the potential map but move towards the nearest exit, selected from their own list of familiar exits (see Section c).

The exits that are available to an individual can also be affected by communication between occupants. This is described in Section d.

The second aspect of *global* behaviour enables the occupants to be attributed with short-term tasks that are completed prior to their exit. This itinerary is user-specified and takes priority over the occupant's immediate evacuation (see Section e). Methods have been introduced so simplify this process and to automate it as much as possible. These are described in Section f.

(a) Occupants following the Potential Map.

In the first instance, occupants move under the influence of the nodal potential, which is a measure of the distance from the node to the most attractive exit point (see Section 2.3). Within EXODUS, occupants will attempt to move in a manner that lowers their potential. However, situations arise where this may not always be possible. In these circumstances the best move will be selected on the basis of the following priority list:

Lower Potential
Equal Potential
Wait

(i) The top priority is to move to a node with a lower potential, thus reducing the distance to the nearest exit point. If more than one node is available that reduces the potential, the node reducing it by the largest amount will be selected.

(ii) If a node of lower potential is unavailable, nodes of equal potential may be considered. This allows occupants to "side step" blockages.

(iii) Where no nodes of lower or equal potential are available, an occupant will wait. Nodes with higher potentials are NOT considered.

Using this approach, occupants will generally move to their nearest serviceable exit (see Section 7.2 part 3). However, in building evacuations, occupants have a tendency to prefer more familiar routes. In this way the closest exit point may be ignored for one which is further away but more familiar to the occupant. This type of behaviour can be incorporated within the global behaviour through the use of biased exit-potentials (see CHAPTER 7: part 3). The exit-potential is effectively the seed for the potential map.

By biasing the exit-potential, exit points can be made more or less attractive. In Figure 2-4 and Figure 7-1 the node shading represents the zone of influence of a particular exit point. In Figure 2-4 both exits are equally biased and so each exit point has an equal zone of influence. In Figure 7-1 the right exit has been biased to become less attractive and so its zone of influence is decreased. The shading in both Figure 2-4 and in Figure 7-1 also indicate to which exit point an occupant at each node would preferentially move towards. Note that in Figure 7-1, three nodes are subject to cross shading, indicating that occupants on these nodes are equally likely to move towards either exit point. In this case, occupants make a random choice as to which exit point to use, subject to local occupant-occupant interactions.

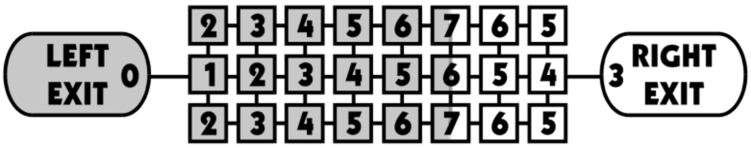

Figure 7-1: Example potential map as shown in Figure 2-4(a) but with LEFT exit more desirable

(b) Occupants assigned Target Exits.

As an alternative to simply following the potential map, the *Behaviour* sub-model allows occupants to be directed to specific exits. Individual occupants or groups of occupants may be assigned a *Target Exit* that may include any of the external exits defined in the geometry. Occupants who are given a *Target Exit* use a similar method to reach their target, but in this case, the actual distance to the exit is followed rather than the potential map, i.e. occupants with a target exit travel to their exit by minimising their distance to that exit. The same priority list applies to this method of travel, the difference being that the distance to the exit is used rather than the potential.

It should be noted that should an occupant exhibit *Extreme* behaviour whilst en-route to a target exit, the occupant ignores the target exit and follows the potential map. However, this behaviour may be overridden by enabling the *Maintain Target Exit* switch (see the User Guide, Section 6.5.3). In this case, the occupant may exhibit *Extreme* behaviour AND maintain their target exit.

(c) Occupant is attributed with a localised exit familiarity.

The behavioural sub-model allows occupants to be attributed with a localised understanding of the exits available, in the form of an exit list or *Occupant Exit Knowledge (OEK)*. In this case the complex structural usage evident in multi-purpose enclosures can be captured through representing the knowledge of individual occupants. Individual occupants or groups of occupants can be attributed with an *OEK* that may consist of any of the exits defined within the geometry (see Figure 7-2). This may be achieved manually, through the selection of individual exits, or automatically determined, where a probability distribution supplied by the user (to the *Attractiveness* check box in the DOOR dialogue box) is interrogated to determine the exact contents of each occupant's exit awareness. Using this automated mechanism the occupant's familiarity may fluctuate between simulations. In no instance will the occupant be without an exit, as if no exit is selected during either the automatic or the manual assignment process, then the nearest exit (as determined at the start of the simulation) is automatically added according to the behavioural regime enabled.

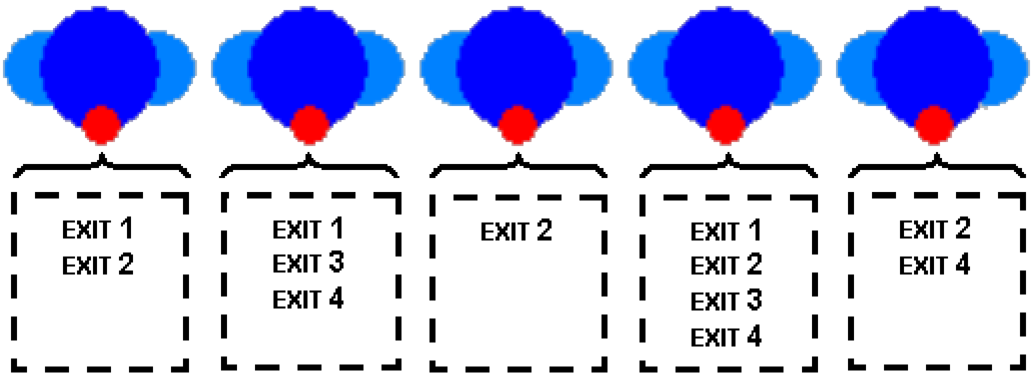

Figure 7-2: Occupant exit familiarity is localised and specific to individual occupants or groups of occupants

Once the *OEK* has been determined, the exit list must be ordered in terms of exit priority. If the *Normal* behavioural regime is selected, the *OEK* is prioritised in terms of *Exit Type* and distance. In this manner, exits that are in constant use (i.e. *Main Exits*) take priority over *Emergency* exits and the list is further ordered in terms of their initial distance from the occupant. Alternatively, if the *Extreme* behavioural regime is enabled (and the occupant is assumed to be involved in egress movement), prioritisation according to *Exit Type* is ignored and the exit will be chosen purely according to distance.

Once this process has been completed, the occupant will follow their initial target in a similar manner as to the *Target Door* highlighted in Section (b) above; the same nodal priority list applies here as did previously with the distance to the exit being used rather than the potential map (as with the assignment of target doors).

The representation of the occupant's complete familiarity with a structure rather than just their present target, has important consequences should the occupant be forced to redirect during the simulation. Instead of interrogating the global list of exits (an action that would be based on the assumption that the occupant had a complete understanding of the geometry), the occupant's local list is used, therefore reflecting their knowledge and their subsequent decision based on this level of information.

NOTE:
Unlike the use of Target Exits described in (b), the Occupant Exit Knowledge is not 'lost' when an occupant becomes Extreme. Instead, due to the increased detail in representation of the occupant awareness, occupants have the opportunity to switch between exits according to the surrounding conditions (see CHAPTER 7:*).*

(d) ***Occupant Itinerary List (OIL)***
The user has the ability to identify a list of tasks (*Occupant Itinerary List*) that an occupant must perform (or has a probability of performing) prior to exiting the structure. This is based on the fulfilment of procedural tasks that are often required of occupants during an evacuation, or events that may occur prior to the commencement of an evacuation that may influence the outcome of an evacuation.

An itinerary can be formed from a number of tasks. These tasks may include:

- Delay / Delay Zone
- Wait / Wait Zone
- Way Point
- Queue
- Remove
- Discount / Discount Mill
- Lift Bank
- Lift Wait
- Form Group / Form Group Zone
- Leave Group / All Leave Group
- Wait For
- Find Via Signage
- Collect Person
- Drop Off Person
- Coordinated Delay

Each task comprises a task type that represents the type of action required to be conducted (i.e. *Delay*, *Wait*, or *Queue* etc.). These tasks define the location where the task is required to be performed (i.e. a node, zone, component or queue) and an associated duration of time required to complete the task, group activities and/or sharing of information. The time period may be controlled by the user or may be randomly distributed between an upper and lower time specified by the user, or may be determined through the interaction with other people (e.g. waiting for their arrival), or with a component (e.g. using a transit node). The user has control over the performance of these tasks using the *Has Tasks* check box on the *PERSON* dialogue box.

Each task is deemed completed when the occupant has remained at the pre-defined location for the duration specified or until the conditions necessary for the completion of the task have been fulfilled. Once a task is completed, the occupant then either moves on to the next task, or if no further tasks exist, evacuates according to the method of familiarity implemented (see Figure 7-3). The list of tasks will be followed while the occupant is patient and while redirective behaviour in response to smoke has not taken place. If the occupant has waited at a location for a period of time exceeding his *Patience*, they will ignore the remainder of their itinerary list and proceed to evacuate directly.

NOTE:
Itinerary lists will be followed until the passenger's Patience has been exceeded – for whatever reason - after which the passenger will no longer attempt to complete the tasks on the itinerary list and will evacuate. This can be prevented from occurring by selecting the Maintain Itinerary option in the Behaviour Options dialogue box, whilst in Simulation mode. This will then force the passenger to attain the location specified in their task, irrespective of whether the passenger's Patience has expired. By default the Maintain Itinerary option is enabled.

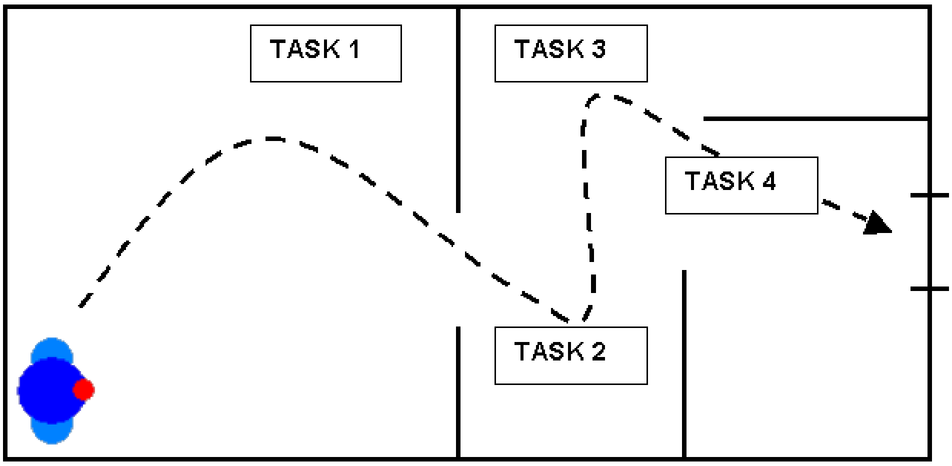

Figure 7-3: Potential path adopted while using task nodes.

Regardless of the global behaviour method employed, local behaviour usually modifies the escape strategy in some way. The degree to which this occurs is dependent upon the number of interactions that each occupant experiences while attempting to implement their global strategy.

(e) The Impact of Communication on Movement

The user has the ability to identify groups of individuals using the *Gene* attribute. When assigned, individuals who share a common gene value will have the ability to communicate some basic information between themselves according to the flags that are enabled. Agents within 2m of each other will be examined to determine whether they share a gene. The nature of the information shared is dependent on the options enabled in *Simulation Mode* (see the User Guide, Chapter 6).

If *Social Response* is enabled, the information that is shared includes an agent's awareness of the exits (their *OEK*) and the response time attribute. An agent's awareness of the structure will be augmented if omissions are evident through comparison with those that share a gene, who fall within a 2m radius of them.

The *Response Times* of those 'communicating' will revert to the lowest *Response Time* detected. Also, if a group member (i.e. someone sharing the same gene) passes within 2 metres of a non-responding group member the non-responding relative becomes active (i.e. they adopt the response time of the agent that has already moved off, effectively causing them to start evacuating). If *Social Movement* is enabled, all of the information highlighted above is shared along with the itineraries of those within the group (see previous section).

The itineraries are shared according to seniority in the group: the group will adopt the itineraries of the most senior pre-determined group *Leader* (according to the Leadership attribute); if no leaders are specified, then the most senior person in the group determined within the model and their itinerary will be followed. In addition, the group will attempt to remain in close proximity through the dynamic modification of travel speeds; e.g. with faster moving agents slowing down and waiting for slower moving agents should they begin to separate. Effectively, when group members are separated by more than 1m, the group member will reduce their speed to the slower moving group member's speed to reduce the gap between them. Caveats are included to ensure that the group members do not iteratively reduce their speed causing their movement to stop altogether. The impact of selecting either the *Social Response* or the *Social Movement* flags is shown in Table 7-1.

Table 7-1: Impact of behavioural flag

<table>
<tr><th rowspan="2">Behavioural Impact</th><th colspan="2">Behavioural Flag Selection</th></tr>
<tr><th>Social Response</th><th>Social Movement</th></tr>
<tr><td>Response Time</td><td colspan="2">Agents can share response times allowing them to the shortest response time encountered</td></tr>
<tr><td>Target Exits</td><td colspan="2">No impact</td></tr>
<tr><td>Occupant Itinerary List</td><td>No impact</td><td>Agents may share itineraries</td></tr>
<tr><td>Adaptive Speed Calculations</td><td>No impact</td><td>Agents adapt their travel speeds to compensate for slower moving agents who have the same Gene attribute.</td></tr>
<tr><td>Exit Awareness</td><td colspan="2">Agents can share exit knowledge to ensure that each agent's eventual knowledge is a superset of all of the exit knowledge of those agents previously encountered who share a Gene attribute.</td></tr>
</table>

The sharing of information follows a precedent according to the identity of the agents involved. This represents the likelihood of information to be shared given the position of the individual within the social structure in question. The user is able to identify whether an individual is senior by using the *Leader* attribute. Information provided by the agent with this attribute enabled will get precedence. For instance, several members of a group may have conflicting itineraries. If one of the group has the *Leader* attribute enabled then their itinerary will be given precedence and will be adopted by the other group members affected by the communication process. Where several individuals are deemed to be a *Leader* or where no group member is deemed to be a *Leader*, then the model will calculate which information has precedence according to the following (arbitrary) formulation. This assumes the following (in order of priority):

- An individual with the *Leader*ship attribute set, is more senior than one that has not
- An adult (between 18 and 50 years old) is more senior than everyone else
- A male is more senior than a female

Therefore, the impact of having being assigned as a *Leader* has a significantly larger impact than being an adult, which in turn has a more significant impact than being male, and so on. This simplistic approach is only adopted where no group leader has been specified, in order to allow the range of group functionality to operate successfully.

There is a subtle, but important, difference between the representation of social interaction using the *Social* flags and that represented using the *Group* functionality employed in the itinerary tasks (see previous section), which does not require the use of the *Social* flags. The *Social* flags represent the interaction between agents who are able/willing to communicate with those sharing a common *Gene*, and the subsequent impact that this communication might have. Although this communication might be hierarchical (i.e. some agents may have more influence than others), it is not pre-planned, occurring as the agents encounter each other during the simulation. The *Group* functionality employed during the itinerary tasks represents activities that are premeditated or the result of direct instruction (e.g. as part of a procedure). Social 'groups' may form and disband as part of these instructions; indeed, these groups may have previously been involved in the more spontaneous communication interactions highlighted above. However, this grouping represents specific actions rather than more general social activities.

NOTE:
This behaviour must first be enabled in Simulation mode by enabling the 'Social Response' Flag or the 'Social Movement' Flag.

NOTE:
A Gene value greater than zero is used to represent related groups. A gene value of zero indicates that this person is not related to any other member within the structure and therefore will not communicate any information.

NOTE:
It should be noted that the model implicitly represents groups; i.e. that the group is represented according to the aligned actions of the group members. As such, the group is not represented as an entity within the model. This allows a greater degree of flexibility in allowing groups to form and disband, but also in allowing a group to cope with complex structures, components and conditions.

7.2 Local Behaviour

The second level of *BEHAVIOUR* sub-model function concerns the occupants' response to local situations. This includes:

(1) *people-people* interactions e.g. overtaking
(2) *people-fire* (environmental) interactions e.g. movement through smoke and
(3) *people-structure* (terrain) interactions e.g. movement on stairs.

Local behaviour is strongly influenced by the occupants' attributes (see CHAPTER 3:) and as certain behaviour rules (e.g. conflict resolution) are probabilistic in nature, EXODUS is unlikely to produce identical results if a simulation is repeated.

There are two operational regimes under which local behaviour rules function, these are known as, *Extreme* and *Normal* behaviour. These are consistent whether the occupant is following the potential map or the distance map (using the *Target Exit* system or the *Local Familiarity* system). Under *Normal* behaviour conditions, the occupants will follow the potential map as described in the previous section i.e. each occupant will attempt to lower their potential/distance and NOT increase it. Under *Extreme* behaviour occupants are given the option to select locations that increase their potential/distance. When in *Extreme* behaviour mode the order of priorities becomes:

Lower Potential (distance)
Equal Potential (distance)
Higher Potential (distance)
Wait

The user selects the *Extreme/Normal* behaviour options for a particular scenario while in *SIMULATION* mode (see the User Guide, Chapter 6). If the *Extreme* behaviour option is selected, occupants will display *EXTREME* behaviour only when their *Wait Counter* attribute (see CHAPTER 3:) exceeds the corresponding *Patience* attribute (see CHAPTER 3:). An occupant will revert to *Normal* behaviour once a move has occurred that lowers their potential.

Extreme behaviour may also impact occupants who have been assigned a *Target Exit*. Such occupants will revert to using the potential map should they experience *Extreme* behaviour. This will result in the occupant ignoring their target exit in preference for their nearest exit or biased exit. However, this behavioural option may be prevented from occurring by using the *Behaviour Option* dialogue box switch *Maintain Target*. If this option is selected, an occupant with a target exit will not lose the target exit if *Extreme* behaviour has been activated (see the User Guide, Section 6.2.1). *Extreme* behaviour has a similar impact upon the *Occupant Itinerary List* (see Section 7.2). This can be overridden by flagging the *Maintain Itinerary* in the *BEHAVIOUR OPTION* dialogue box.

The type of behavioural regime selected will also affect the occupant exit selection if the *Local Familiarity* option is flagged. In the *Normal* behaviour, occupants will prioritise their *OEK* list according to the nature of the exit use. Therefore exits in constant use (a *Main Exit*) will always be adopted prior to emergency exits. Only if the occupant is completely unaware of any general exits will they adopt an emergency exit. This is to represent the propensity of occupants during normal building circulation to use the exits provided, rather than utilise non-reversible emergency exits that may be alarmed and will certainly be used less frequently. If the *Extreme* behavioural regime is selected, then no prioritisation occurs and an exit is selected purely on the basis of distance.

TIP:
It is recommended to use Extreme behaviour when attempting to simulate emergency conditions as occupants are given the option of taking actions that may move them further away from an exit (i.e. jockey) rather than wait patiently for their turn (see the User Guide, Chapter 6).

The local behaviour capabilities of buildingEXODUS will now be discussed. Unless otherwise stated, the following applies to both the *Extreme* and *Normal* behaviour regimes. This discussion will follow the three areas identified at the start of this section.

7.2.1 People-People interaction:

7.2.1.1 (i) Response time.

This is a measure of the time an occupant requires before they positively react to the call for evacuation. An individual’s *Response Time* is part of the occupant attribute parameter set defined for a particular simulation scenario (see CHAPTER 3:). Depending on the behaviour options selected when in *Simulation* mode, the individuals may be static or milling prior to their evacuation movement. Therefore, if the exact pre-evacuation location of an occupant is known then it might be useful to enforce a location upon them. However, if the user is unsure as to the exact location of an individual, has a vague understanding of their pre-evacuation location, or believes that they would not be static prior to the evacuation, Then the occupants should be allowed to mill around prior to the commencement of the evacuation. This movement will not allow them to accidentally evacuate and will not be recorded as part of the evacuation process. The *Milling* option is available via the *Behavioural Options* dialogue box from the *Rulebase* menu in *Simulation* mode and is described in the User Guide, Chapter 6.

NOTE:
The Milling option only influences occupants behaviour before they respond (i.e. before the simulation clock reaches the occupant's Response Time). It does not influence occupant milling when produced by itineraries.

7.2.1.2 (ii) Conflict resolution.

When two or more occupants vie for space (usually in crowds) conflicts arise that must be resolved. Conflict resolution is the procedure by which this occurs within EXODUS.

EXODUS utilises a fine network of nodes to describe an enclosure (see Section 2.2). Each node is intended to represent the smallest amount of free space available for occupancy, essentially it is the space that a single individual can occupy. Thus only one occupant can occupy a node at a time. However, the situation often arises where two or more occupants may wish to occupy a particular node. An illustrated example of this is shown in Figure 7-4, where three occupants wish to occupy the same node. Two occupants (labelled 1 and 2) are attempting to enter an aisle from their seats, while a third occupant (labelled 3), already in the aisle, is attempting to proceed. The three occupants are attempting to occupy the indicated node and thus a three-way conflict arises.

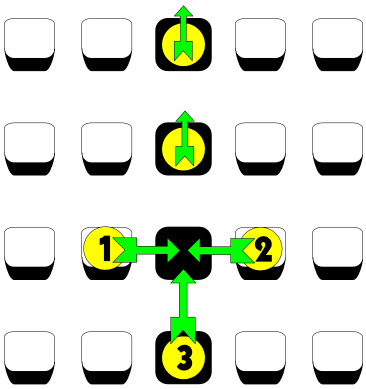

Figure 7-4: Example conflict between three occupants labelled 1, 2 and 3.

At the first level of conflict resolution the travel time for each conflicting occupant to arrive at the node in question is examined. If the occupants are determined to arrive at the contested node during the same tick of the simulation clock they are deemed to be *in conflict*. The resolution of such a conflict is then attempted through an evaluation of the *Drive* attribute for each of the occupants (see CHAPTER 3:). The *Drive* for each occupant involved in the conflict is compared. If one of the occupants has a *Drive* significantly higher than the others, this occupant becomes the winner. However, if the *Drive*s are sufficiently close, the winner is randomly selected. Sufficiently close is here defined as the absolute normalised difference between the various drives being less than 10% (i.e. abs [(*Drive*1 – *Drive*2)/Max(*Drive*1,*Drive*2)]).

All occupants involved in conflicts attract a time penalty that is randomly selected between pre-determined limits (see Table 7-2). The time penalty represents the time lost in the interaction. There are two levels of time penalty. The first level is associated with conflicts that are resolved on clear differences in *Drive*. If the conflict is resolved in a random manner, the second level time penalty (which is longer than the first level) is used.

Conflict losers may continue to wait in the same location until another opportunity to occupy the node arises, or perform another action such as change direction. Within large crowds, many conflicts occur which effectively results in decreased average travel speed. The time penalty associated with conflict resolution effectively reduces the travel speed of those involved in the conflict. Thus occupants caught in high congestion regions will appear to slow down as they will be involved in many conflicts. Furthermore, as the conflict resolution process involves a number of random processes, simulations will not necessarily produce identical results if a repeat of the simulation is performed. Users can choose to accept the default values shown in Table 7-2 or manually redefine the time penalty in the *SIMULATION* mode.

NOTE:
It is extremely difficult to determine a set of conflict time penalties that are appropriate for all circumstances. The default values used in EXODUS are thought to be reasonable and have produced acceptable results during testing and validation. Users should note that in some circumstances, changes to these values may produce significant differences in model predictions. If, as part of a simulation or scenario, the conflict time penalties are altered, it is vital that these changes are noted in the documentation supporting the results.

NOTE:
The default conflict times are appropriate for the default mesh spacing. If a geometry is meshed with a spacing different to the default value, it may prove necessary to reassess and redefine the conflict times and travel speed distributions. This may prove necessary as changes to the mesh may impact the range of packing densities in heavily congested areas allowed by the model (see the Application Manual, Chapter 2 for further details).

Table 7-2: Time penalties associated with conflict resolution

Drive Attribute Difference	Penalty Range
> 10%	0.5s - 0.7s
≤ 10%	0.8s - 1.5s

Another form of conflict may arise under contra-flow conditions. Such situations may occur when *Target Exits* have been specified for some or all the occupants or the *OEK* navigational system is use. To prevent gridlock from occurring in such situations, additional conflict resolution rules are required. Where two occupants face a competition for each other’s node, and no other nodes are available, i.e. a *head-on collision*, a contra-flow situation has developed. Under these circumstances, the occupants are allowed to squeeze past each other, effectively swapping nodes. This process occurs at a reduced travel speed, (***arbitrarily set at 50% of the slowest walk-speed of those involved***), and a time penalty is also added to those involved, equivalent to the first level of conflict resolution penalty (see Section 7.2 part 1 (ii)). Occupants involved in such a conflict who are crawling do not travel at a reduced speed, but still incur the conflict penalty time.

Additional contra-flow rules cater for situations where a multiple (or cyclic) contra-flow condition arises. For example, consider a three way contra-flow in which occupant 1 wishes to move to the node occupied by occupant 2; occupant 2 wishes to move to the node occupied by occupant 3 and occupant 3 wishes to move to the node occupied by occupant 1. An example of this type of conflict is shown in Figure 7-5, where occupant 1 is heading for exit 1, occupant 2 for exit 2, and occupant 3 for exit 3.

Essentially, when buildingEXODUS detects this condition, the occupants involved may swap nodes with each other if a mutual benefit may be obtained, i.e. if both occupants move closer to their target exits. If several options are available, buildingEXODUS will select the option offering the maximum mutual benefit. If however, several options offer the same maximum mutual benefit, then the outcome is randomly selected.

In the example shown in Figure 7-5, there are three possible *swaps*, occupant 1 with occupant 2, occupant 2 with occupant 3 and occupant 1 with occupant 3. Two of these swaps offer maximum mutual benefit, i.e. occupant 1 with occupant 2 and occupant 2 with occupant 3, as both these swaps result in two occupants moving closer to their target exits. The other possible swap, i.e. occupant 1 with occupant 3, results in a benefit only for occupant 1. As there are two equal choices, buildingEXODUS randomly selects the outcome, as shown in Figure 7-5.

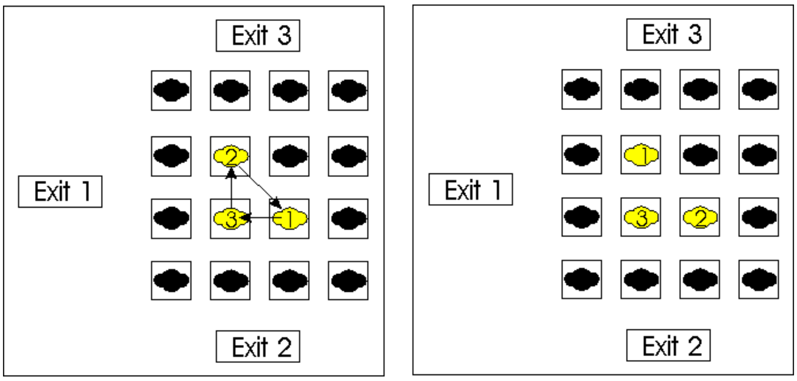

Figure 7-5: Example of a cyclic contra-flow and a possible resolution

7.2.1.3 (iii) Direction changes.

Direction changes may occur as a result of several factors: exits becoming available/unavailable or queuing/crowding. Generally, occupants will move wherever possible, in a straight line. When the path is obstructed, arcs in other directions are considered, but these must still conform to the priority list (see CHAPTER 7:). Major direction changes may occur when an occupant exhibits *Extreme Behaviour*, as this allows occupants to increase their potential.

NOTE:
Occupants with specified target exits will revert to the potential map if their Patience expires, i.e. Wait Counter > Patience.

7.2.1.4 (iv) Overtaking.

Overtaking occurs as a natural consequence of the movement rules (see CHAPTER 3:), specific overtaking algorithms are not required. An occupant blocked by a slower moving occupant will attempt to find an alternative neighbouring empty nodal position in the direction of travel.

The paths of three persons involved in an evacuation from a restaurant area are depicted in Figure 7-6. The portion of the geometry depicted represents the restaurant (with toilets) and a part of the supermarket including shelf region, till region and one of the exits. The nature of the paths depicted demonstrate overtaking, redirection and obstacle avoidance. In this case, a distribution of response times was used so the three occupants shown did not necessarily interact with each other. Occupant 1 started from within the toilet area and hence had a long response time. For this reason, little interaction with other occupants occurred, however three cases of path diversion due to overtaking, can clearly be seen.

Occupant 2 started in the restaurant area and had to make his/her way through the maze of tables and around several seated diners en route to the exit point. When this occupant reached the restaurant internal exit, he/she encountered a large amount of congestion. The occupant attempted to get around this blockage, taking many side steps before making his/her way out.

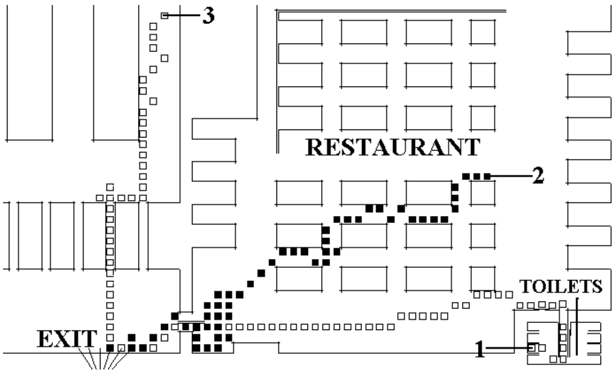

Figure 7-6: Example paths taken by three occupants to evacuate the supermarket/restaurant geometry

Occupant 3 started in a less congested area. While still here, he/she made his/her way around five slower or stationary occupants. Slower moving occupants blocked this occupant's most direct path to the exit point. The occupant side-stepped the blockage and attempted to exit via the second row. Unfortunately this path was also temporarily blocked. After taking several side steps a path cleared and the occupant eventually took this path out.

The behaviour exhibited by these occupants clearly demonstrates the overtaking and route diverting capabilities of buildingEXODUS.

7.2.1.5 (v) Natural Movement (formerly Population Densities and Angle of Movement)

The overtaking procedure can also take into consideration two additional aspects of the surrounding conditions, as perceived by the individual. These are the population densities at potential target locations and the angle of movement of the individual. The impact of these factors will be dependent upon the options enabled in *Simulation* mode (see the User Guide, Chapter 6). Once the appropriate option is enabled then these two factors are taken into consideration in the movement of an individual. The individual will judge their future nodal choices not simply on their distance to a prospective target but also on their population densities; occupants prefer locations with minimum levels of population density and will prioritise nodal locations accordingly. This will only happen to those nodes that are closer to their desired target. Therefore nodes ahead of them will still always be preferred to those further away, but will be ordered according to the population densities as well as their relative distances to the desired target. In addition, the trajectory of the individual can also be included as a factor. Here the nodes, which produce the smallest change in the angle of trajectory of the individual will be preferred to those that incur wild swings in occupant movement. Again this only applies to those nodes closer to the target. This behavioural option can produce far more naturalistic occupant movement than might otherwise be the case.

7.2.1.6 (vi) Wall Proximity

This behaviour allows the individual to prioritise the choice of direction based on the available movement options that each new location provides. Therefore, the likelihood of occupying a location is biased according to its connectivity i.e. the choice of a node will be biased against nodes of low connectivity compared to their adjacent nodes. The avoidance of such nodes, where possible, results in the person avoiding locations that are in close proximity to the enclosure boundary e.g. walls. This behaviour is particularly useful when the user is having difficulties in generating *Boundary Nodes* and when used in conjunction with the behaviours described in (1)(v).

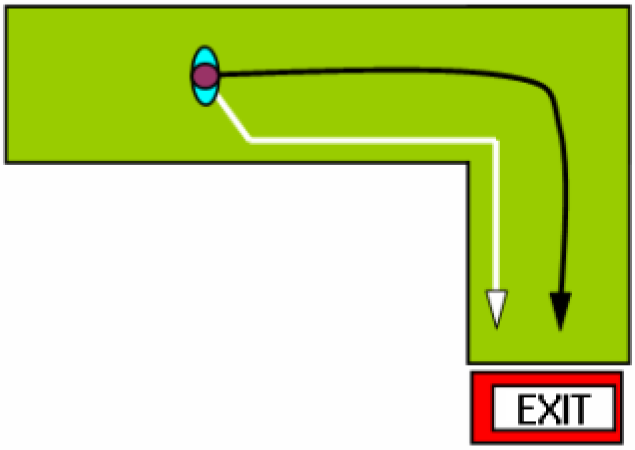

Figure 7-7: Paths that are adopted by an individual given the standard EXODUS behaviour (white line) and the enhanced wall proximity method (black line).

NOTE:
This behaviour is enabled by flagging the "Wall Proximity" option in the Behaviour tab of the BEHAVIOUR OPTIONS dialogue box in SIMULATION mode.

7.2.1.7 (vii) Adaptive Exit Selection In Relation to Congestion.

This algorithm determines whether an individual will redirect away from a particular *External Exit* due to the level of congestion around that exit (i.e. the exit towards which they are heading). The use of this function is dependent upon the user having generated visibility zones for the exits establishing the information to which the evacuees have access. This will require the association of a sign to the exit in question in order for the visibility catchment area (i.e. the area from which the exit/sign can be seen) to be accurately established. The method by which this can be achieved is outlined in CHAPTER 8: and Chapters 3 and 5 in the User Guide.

On running the simulation, the agents will not then simply move towards their current exit irrespective of the level of crowding around it, but will instead take into consideration the other options available to them according to the information at their disposal.

During the simulation, an individual will head towards their exit until they fall into the visibility catchment area calculated for the sign located near to (and is associated with) the exit indicating the existence of the exit (i.e. they are able to 'view' the conditions of the exit). At this point the algorithm will be accessed. However, if the agent has previously redirected (or have temporarily been prevented from redirecting due to a decision within the algorithm) or if they have been attributed with a target exit which has to be maintained (i.e. the *Maintain Target Exit* switch has been enabled), then the algorithm will progress no further. If the exit is being used and there is relatively minor congestion around it (i.e. if the exit is deemed to be acceptable), then no alternatives need be considered. The algorithm also ends if the individual is within 2m of their current exit, as the individual is then assumed to be close enough to their exit for them not to wish to redirect.

NOTE:
It is very important that the user correctly locates and associates signs in order to accurately produce the visibility catchment areas that drive the redirective behaviour. This is described in more detail in CHAPTER 8: *and Chapters 3 and 5 in the User Guide.*

Otherwise, once the agents are in sight of the associated sign indicating the existence of the exit to which they are heading (or their nearest exit if they are simply following the potential map) then they may make comparisons between the exits that are available to them. At this stage, if the agent is not entirely surrounded by others and they are sufficiently driven and have a low patience attribute then they will consider redirection. If the agent is surrounded by others then they will temporarily be prevented from redirection, although this condition may change (see Figure 7-8). This is intended to simulate both the difficulty that these people would have in seeing the conditions at the other exits as well as the difficulty that they would have in redirecting away from their current exit.

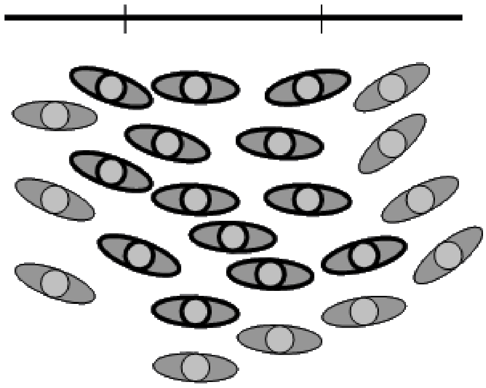

Figure 7-8: The occupants with bold edges are not able to redirect as they are completely surrounded.

The method of navigation used will have an impact upon the redirective behaviour, largely because of the assumptions on which these techniques are based. If the *Local Familiarity* method or the *Target Exit* method is used then the exit towards which the individual is heading (i.e. the exit specified as their target) is examined. If the *Potential Map* system is used then the individual has no target specified but instead is entirely dependent upon the potential map. This map can be biased according to the seed potentials associated with the exits. In this case the individuals will not have a specified exit associated with them and there is no reliable means by which to determine the exit towards which they are heading. The agents in these cases will make the calculation according to the exit nearest to them (still being dependent upon its visibility). The exits known by the individual are then cycled through. This knowledge may be specified (i.e. the individual will be using the *Local Familiarity* system and will have a list of exits associated with them or unspecified (i.e. they will be using the *Target Exit* or the *Potential Map* system, which will assume a complete knowledge of the exits available). In addition, the *Local Familiarity* system can be influenced by the usage of the exits. As with the general adoption of exits under this system, exits deemed to be used in *Emergency* conditions, will only be considered when the *Extreme* behavioural regime is in effect.

The exits are then separated according to whether they are visible or not visible (i.e. whether a sign is located at the exit and the individual falls inside the visibility catchment area of this sign). The visible exits are examined according to the time taken to arrive at them and the time that the crowd around it takes to evaporate (Exits 1 and 2 in Figure 7-9). This will be based on the size of the exit and the assumed flow rate at that exit, which in this instance is assumed to be 1.33 occ/m/s (the HMSO flow rate). The evaporation of this congestion will then be calculated according to:

$$t_{congestion} = N/(f * w) \qquad \textbf{(23)}$$

where N is the number of individuals in the surrounding area of the exit, f is the assumed unit flow rate and w is the width of the exit.

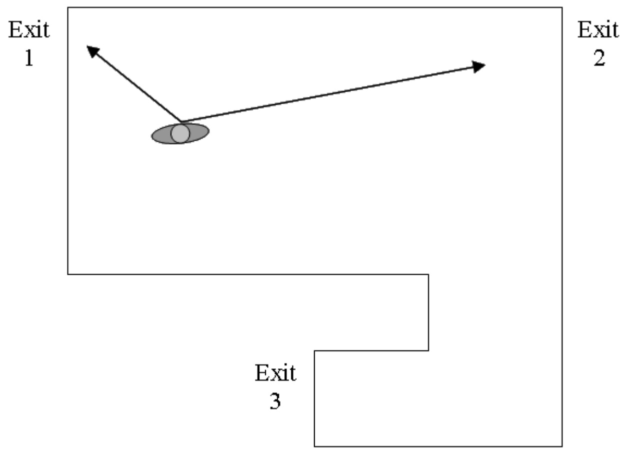

Figure 7-9: Only Exits 1 and 2 are visible while Exit 3 is deemed not to be visible.

The non-visible exits (those that either have not been associated with a sign or where the individual falls outside of the visibility catchment area produced by the associated sign) are only examined for the time to cover the distance to them as no other information is available to the agent (Exit 3 in Figure 7-9). Therefore the estimated time of arrival at a visible exit will be:

$$\max(t_{arrival,}\, t_{congestion}) \qquad \textbf{(24)}$$

whereas the estimated arrival at a non-visible exit is simply:

$$t_{arrival} \qquad \textbf{(25)}$$

Where:

$t_{arrival}$ is the time to cover the distance between their current location and the exit, and
$t_{congestion}$ is the time for the congestion around an exit to evaporate.

The exits are then compared according to the predicted time of arrival. Only those that will be arrived at more quickly than their current exit are examined. Preference is given to the visible exits (i.e. these are examined first and separately from other exits). If there are no visible exits that can be arrived at more quickly than the current exit, then the non-visible exits are interrogated. The adoption of non-visible exits will be less likely to be adopted in *Extreme* mode, based on the assumption that individuals are less likely to wander back into the structure when they can see an available and active exit, even if it is congested. In both cases (*Normal* or *Extreme*), a comparison will be made against the *Drive* of the individual.

If the agent is surrounded by others or they fail one of the attribute comparisons then they will be attributed with a temporary attribute preventing them from redirecting. This is examined during each time frame (through comparison between their *Drive* and a probability factor).

This algorithm therefore allows for the sub-optimal adoption of routes. If, for instance, an agent is in a position where they are able to redirect and are aware of two alternative exits, one of which is visible and one of which is not, if the time calculated to arrive at the visible exit is greater than the arrival time at the current exit, then it is disregarded (see Figure 7-10). The non-visible exit is then considered (Exit 3). In this case, only the time to arrive at the exit is considered as no other information is available. If this time is less than the arrival time at the current exit, then there is a possibility that this exit may be adopted. If it is adopted then the agent has redirected without considering the congestion that might be evident at their new non-visible exit, which may be far in excess of their current exit. In this case, Exit 3 has significant congestion at it and would not provide a reduction in the overall evacuation time.

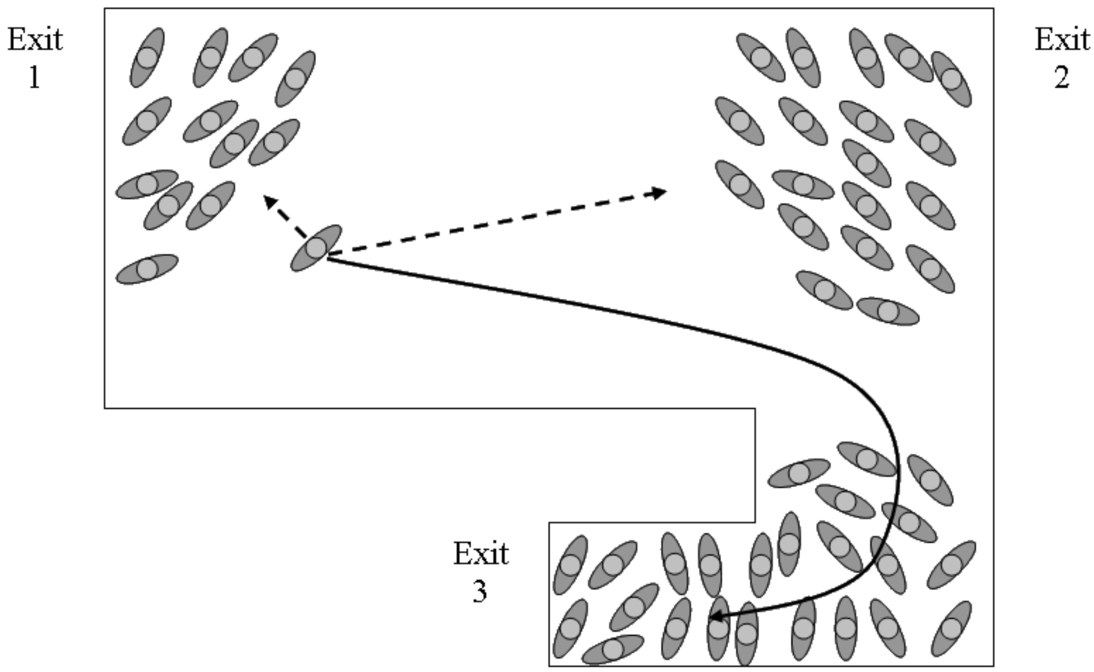

Figure 7-10: Occupant redirects the non-visible exit that has an elevated level of congestion.

As mentioned previously, the method of navigation used will have an impact upon the redirective behaviour. This will now be briefly discussed.

7.2.1.7.1 (a) Local Occupant Familiarity

If the *Local Familiarity* method or the *Target Exit* method is used then the exit towards which the individual is heading (i.e. the exit specified as their target) is examined. If *Local Occupant Familiarity* is utilised, then the occupant's *OEK* is interrogated to determine whether alternative egress routes exist. If this is the case then the occupant will head towards the nearest exit with which he/she is familiar. In addition, the *Local Familiarity* system will be influenced by the usage of the exits. As with the general adoption of exits under this system, exits deemed to be used in Emergency conditions, will only be considered when the *Extreme* behavioural regime is in effect.

7.2.1.7.2 (b)Potential Map Option Selected

If the *Potential Map* system is used, the occupant will be assumed to be familiar with all of the exits within the structure and will examine the exits according to the *Potential* of the exits. This behaviour is consistent with the assumptions on which the potential map is based, such as the occupant's global familiarity with the exits of the enclosure. If the potential map system is used then the individual has no target specified but instead is entirely dependent upon the potential map. This map can be biased according to the seed potentials associated to the exits. In this case the individuals will not have a specified exit associated with them and there is no reliable means by which to determine the exit towards which they are heading, given the nature of the potential map. The individuals in these cases will make the calculation according to the exit nearest to them (still being dependent upon its visibility).

7.2.1.7.3 (c) Target Door Option Selected

If the *Target Door* system is used, then the occupant *will not* redirect if they are forced to maintain this exit, as it assumed that the occupant has either a limited understanding of the structure (limited to a single exit), or is following an instruction. It is only if the target is lost, due to the occupant becoming impatient, that the occupant will revert to the potential map system and may then redirect according to the conditions described above.

7.2.1.7.4 (d) Occupant Itinerary List (OIL)

If the occupant is performing a task from an itinerary list, the potential redirective behaviour will supersede the task behaviour. Therefore, if the occupant decides to redirect, then their list of tasks will be lost and, according to the familiarity system in use, the occupant will redirect towards their next choice of exit. However, the user has the ability to force the occupant to follow their *Itinerary Lists* by enabling the *Maintain Itinerary* option in the *BEHAVIOUR CONTROL* dialogue box.

7.2.2 Environmental interaction:

Here we consider the impact of the atmospheric conditions resulting from fire.

NOTE:

The behaviour interaction associated with conditions of smoke, heat and toxic gas is only available with level C of buildingEXODUS.

7.2.2.1 (i) Smoke and Temperature trigger.

Under non-fire conditions occupants will not begin to actively take part in the evacuation until their *Response Time* has elapsed (see Section 2.3.7.2 part 3). However, if the person is made aware of the danger of the fire through his/her perception of the local temperature or smoke concentration, he/she will begin to evacuate before his/her *Response Time* has elapsed. This behaviour can either be controlled through the setting of the *Triggering Temperature* and *Triggering Smoke Concentration* attributes in the Toxicity Sub-Model in the SCENARIO mode (see CHAPTER 6:), or by the setting of the *Response Time* overrides of a zone in the HAZARD mode (see the User Guide, Chapter 5).

7.2.2.2 (ii) Smoke interaction.

It is known that a person's *Walk Rate* decreases with increasing smoke concentration [17,18]. EXODUS links the smoke concentration (i.e. both irritant and non irritant) with the *Mobility* attribute (see Section 3.1, part 2). As the smoke concentration increases the *Mobility Degradation Factor* decreases (see Table 3-2) and this in turn decreases the persons *Mobility* and hence travel speed (see Section 3.1, part 3). This is also dependent upon the nature of the irritant products represented within the environment.

In a similar manner, smoke concentration can also affect an agent's *Agility* attribute. As the smoke concentration increases the *Agility Degradation Factor* also decreases and this in turn decreases the person's *Agility* (see Section 3.1, part 4). However, this effect only occurs with irritant smoke. Although non irritant smoke affects an agent's *Mobility* it has no corresponding affect on their *Agility*. This is because non irritant smoke only causes visual obscuration. While this obscuration may affect an agents ability to move within a structure (i.e. their *Mobility* and *Travel Speed*) it is felt unlikely to affect their ability to negotiate obstacles (i.e. their *Agility*).

NOTE:
In versions of buildingEXODUS prior to v6.0 an agents exposure to non-irritant smoke would also have reduced their corresponding Agility value, meaning that their ability to negotiate obstacles would previously have been reduced.

7.2.2.3 (iii) Inefficient movement within smoke

Under experimental conditions Jin found that when encountering smoke, in addition to a reduction in travel speed (see CHAPTER 7:) the evacuee movement became increasingly inefficient, with evacuee's 'staggering' along a smoke-filled corridor [17,18]. This was due to the visual obscuration caused by the dense, irritant environmental conditions. An additional behaviour noticed by Jin was that in smoke conditions, occupants tended to use the walls to assist them in navigation [17,18]. Both these behaviours have been included in EXODUS [55].

The staggering behaviour is controlled through the setting of the *Smoke Stagger* option on the *BEHAVIOUR OPTIONS* dialogue box. This function operates independently of the physical impediment provided by smoke that is described in CHAPTER 3:. It should be possible for the user to enable these functions simultaneously or independently according to the user's needs.

This function only affects occupant behaviour if the occupant is not situated on stairs, amongst seats or is not adjacent to an exit (i.e. the occupant is not located on a node that is connected to an exit). Under these conditions it is assumed that the occupant's ability to navigate is not further impaired. As described later, the algorithm also has no impact upon the occupant movement if the conditions have forced the occupant to crawl, which supersedes any other form of behaviour.

Effectively, as the extinction coefficient of the smoke increases, so the likelihood of inefficient movement increases; the possibility of the occupant straying from his optimal path is directly linked to the environmental conditions. Instead of the nodal attractiveness being entirely dependent upon its distance from a particular destination or the nodal potential differences, the extinction coefficient of the smoke also has an impact upon the attractiveness of the node, allowing non-optimal decisions to be made.

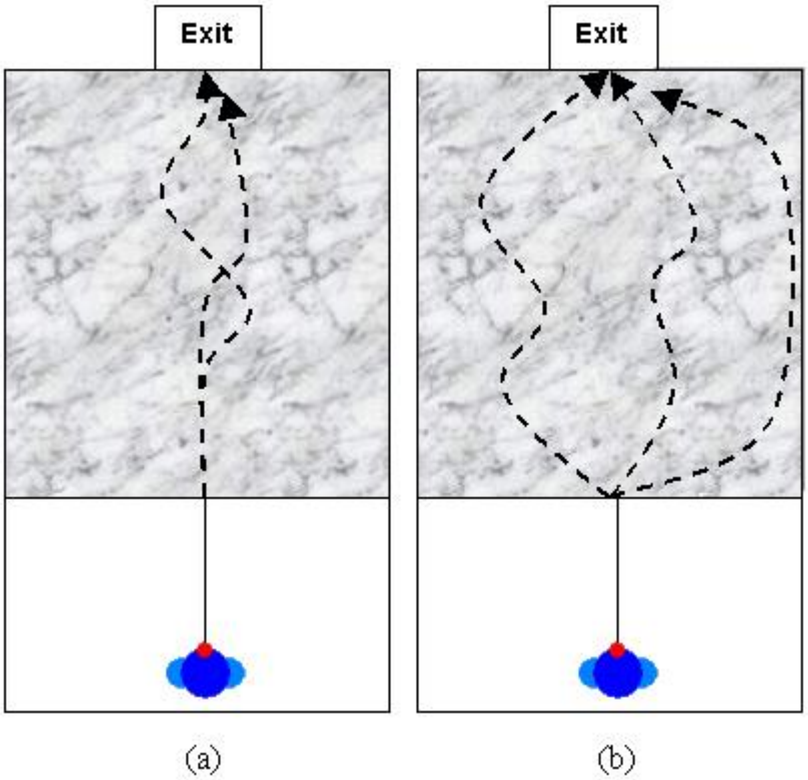

Figure 7-11: In (a) the environment conditions are severe enough to cause only minor inefficient movement. In (b) the conditions worsen, causing more extravagant staggering. During this process the occupant comes into contact with the edge of the boundary, along which he continues until it no longer provides an advantage. In both cases several possible example paths are provided. It should also be noted that prior to entering the deteriorating environment, the occupant adopts a normal path.

The occupant initially prioritises the nodal options according to their attractiveness in relation to their target. Once completed the attractiveness of the node is biased in relation to the environmental conditions. This biasing is weighted according to the severity of the conditions and is then multiplied by a random value (between 0 and 1.0), to incorporate stochastic processes. Therefore although the list of nodes may contain the same nodal entries representing the nodal options available to the occupant, they may appear in a different order. The random element allows the path adopted by the occupant to fluctuate between simulation runs. Therefore, as the occupant passes through a smoke-filled environment his identification of nodes as viable options becomes inaccurate, forcing him to adopt node off his optimal line of egress, and causing a stagger.

The probability of this movement occurring is bounded, with smoke levels influencing occupant behaviour between extinction coefficients of 0.1/m and 0.5/m. The lower limit is derived from the original Jin experiment, while the upper limit represents the stage at which buildingEXODUS assumes that the occupant is forced to crawl if using simplified model (i.e. irritant model not enabled) (see CHAPTER 4: and CHAPTER 5:).

The *Stagger Behaviour* may also cause the occupants to temporarily avoid nodes that have a relatively high smoke concentration. If an alternative node exists, which satisfies the rules of movement according to the behavioural regime and that has a lower smoke concentration, it will appear more attractive than an otherwise equivalent node of higher smoke concentration. This behaviour cannot be guaranteed, as the implementation of this behaviour is dependent upon the potential map or the distance map generated within the structure. Therefore, the occupant's avoidance of smoke-filled environments, where alternatives exist, is strongly dependent upon the location of the smoke barrier, the location of the occupant and the nature of the surrounding environment.

The second behaviour implemented concerns the use of walls as navigation aids. Within EXODUS, occupants will prefer nodes adjacent to walls if available. Only those nodes that have a reduced level of connectivity (implicitly representing the walls of the structure) and are closer to the occupants target than their present location are considered in this behaviour. Once engulfed in a smoke-filled environment, occupants will prefer nodes that have reduced connectivity (representing edge or wall nodes) rather than centrally positioned fully connected nodes. This behaviour is dependent upon the geometry to some extent, as the occupant will only prefer 'wall' nodes if they provide an advantage to the occupant and if they are connected to the occupant's present location.

NOTE:
The toxicity or irritancy of the smoke is not considered in the occupants stagger response, only the concentration of the smoke.

7.2.2.4 (iv) Occupant redirection in response to smoke

The interaction of the occupant with the environment is not simply limited to physiological considerations. Of equal importance is the effect that the environment can have upon the occupant's decision-making, especially on the egress route that might be adopted. A great deal of investigative work has been conducted concerning the occupant decision-making process in relation to smoke [63-70]. Where practical, this information and data has been used in the development of this behavioural feature.

The behaviour implemented in EXODUS and described here represents the occupant response when confronted with a perceived barrier of smoke.

NOTE:
The smoke re-direction behaviour represents a considerable simplification of an extremely complex and incompletely understood behaviour. As a result a number of influential factors have been omitted.

This behaviour is enabled through flagging the *Redirection* option in the *BEHAVIOUR OPTIONS* dialogue box in *SIMULATION* mode.

NOTE:
It is recommended that smoke-redirection be used in conjunction with Extreme behaviour.

Initially, for the behaviour to be activated, the occupant has to be located in a relatively clear environment. This is determined as being when the occupant is located on a node with an extinction coefficient less than 0.1/m [17,18]. If this is not the case then the occupant is assumed to be 'engulfed' in smoke, which may trigger other environmental behaviours (see Section 7.2(2)).

If the occupant satisfies this condition, it is determined whether the occupant is confronted with a smoke barrier. For any analysis of this type to occur the occupant must be *adjacent* to the smoke hazard (i.e. the occupant only analyses those nodes connected to the node presently occupied). The smoke barrier is only influential if it lies closer to the occupant's target than the occupant's present position. Therefore if all of the nodes that are situated sufficiently closer to the occupant's target have an extinction coefficient greater than 0.1/m, then it is assumed that a barrier has been formed and the behaviour continues (see Figure 7-12).

NOTE:
The occupant does not have a long-range vision capability and so will not react to a distant, yet potentially visible smoke barrier.

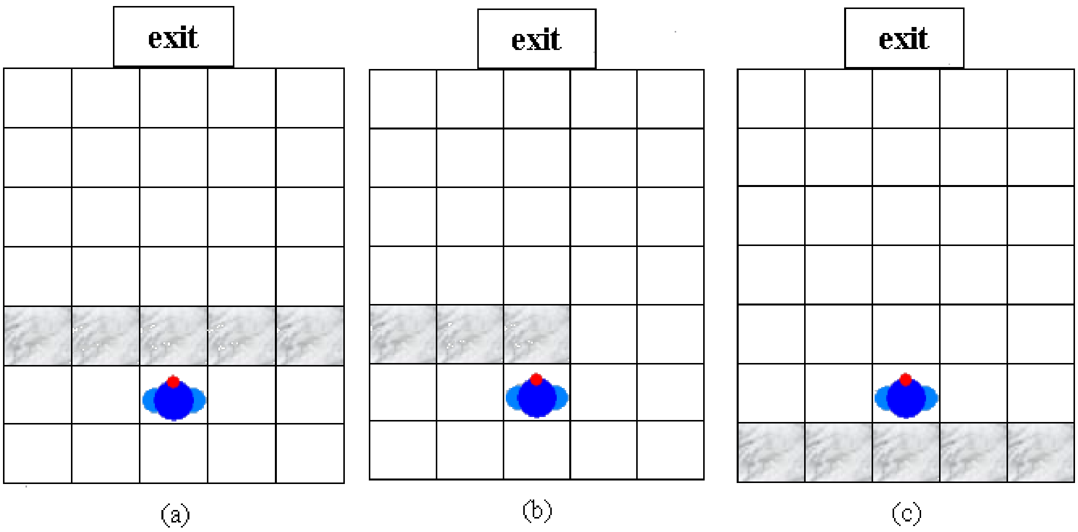

Figure 7-12: (a) The occupant is faced with a barrier of smoke and may redirect. (b) A smaller barrier exists, allowing the occupant passage. (c) Although a barrier exists is does not affect the occupant's path to his exit and is therefore ignored.

If this is not the case, then the occupant will not perceive the environment as presenting a barrier to his egress route and will continue onwards irrespective of the conditions. Given that the barrier is perceived, the visibility afforded by the environment is calculated according to the Jin calculations [17,18]. The default position adopted by the model is to use the following Jin formulation

$$v = \frac{\lambda}{K} \quad \textbf{(26)}$$

Where:

v is the visibility (m) afforded by the environment,
K is the extinction coefficient (m^{-1}), and
λ represents a dimensionless constant.

The default value of this constant is 2.0. This is the most conservative estimate of the occupant's visual ability, given the environmental conditions. Several other formulations are also provided by Jin [17,18], (with λ=3.0, 4.5 and 8.0 [17,18]), as well as formulations by other researchers [17,18]. These formulations tend to refer to the impact of different forms of signage or to the tenability limits provided by the environment. Given these variations, the user is able to adjust the λ value, by altering the *Visibility Coefficient* in the BEHAVIOUR OPTIONS dialogue box.

The extinction coefficient of the smoke analysed is calculated as being the average derived from the nodes that form the barrier immediately ahead of the occupant.

If the visibility afforded by the smoke barrier is greater than his expected travel distance to his target exit, then the occupant *automatically* maintains the current egress route. This prevents occupants redirecting once they are close to their exit or when they are 'in sight' of their exit. This calculation utilises the visibility formulation defined above. If the occupant is not within the calculated visible range of his exit, the occupant must decide his next course of action. Only two options are available to the occupant, either *continuing through the smoke barrier* or *re-directing away from it*.

Prior to any examination of the occupant's decision, alternative routes need to exist, in order for redirection to be contemplated. The method used to represent exit familiarity will have a significant effect upon the nature of occupant redirection.

NOTE:
In reality, occupants may not only redirect towards other exits, but may also move towards areas of refuge. This behaviour is not considered in the current implementation.

NOTE:
It is recommended that the Smoke Redirect option be used in conjunction with the Smoke Stagger behaviour. This is to guarantee that occupants once confronted by complex environmental zones have exhausted their potential clear egress routes by attempting to navigate around the hazard prior to the redirection process.

7.2.2.4.1 (a) Local Occupant Familiarity

If *Local Occupant Familiarity* is utilised (see Section 7.1(c)), then the occupant's *OEK* is interrogated to determine whether alternative egress routes exist. If this is the case then the occupant will head towards the nearest exit with which he/she is familiar, (irrespective of the exit type). An exit will not be adopted if it entails the occupant passing through the smoke barrier previously identified as preventing safe passage. Therefore alternative exits, which entail moving through the adjacent smoke barrier, are discounted.

NOTE:
The use of the Local Familiarity system can have an important impact upon the occupant's ability to redirect, as they do so, on the basis of their knowledge of the structure. In the current software release, occupants have no ability to search for previously unknown exits.

TIP:
It is recommended that the Extreme behavioural regime should be enabled when using Smoke Redirection, as it should be assumed that occupants encountering significant levels of smoke would be aware of the seriousness of the incident.

7.2.2.4.2 (b)Potential Map Option Selected

If the potential map system is used, then once the smoke barrier is identified, the occupant will move towards the *next most attractive* exit that does not entail movement through the same barrier of smoke. This behaviour is consistent with the assumptions on which the potential map is based, such as the occupant's global familiarity with the exits of the enclosure.

7.2.2.4.3 (c) Target Door Option Selected

If the *Target Door* system is used, then the occupant *will not* redirect, as it assumed that the occupant has either a limited understanding of the structure (limited to a single exit), or is following an instruction. It is only if the target is lost, due to the occupant becoming impatient, that the occupant will revert to the potential map system and may then redirect according to the conditions described above.

NOTE:
If the Maintain Target Exit option is enabled the occupant is prevented from redirecting due to smoke conditions.

7.2.2.4.4 (d) Occupant Itinerary List (OIL)

If the occupant is performing a task from an itinerary list when encountering a barrier of smoke, the potential redirective behaviour will supersede the task behaviour. Therefore, if the occupant decides to redirect away from the barrier of smoke, then their list of tasks will be lost and, according to the familiarity system in use, the occupant will redirect towards their next choice of exit. However, the user has the ability to force the occupant to follow their *Itinerary Lists* by enabling the *Maintain Itinerary* option in the *BEHAVIOUR CONTROL* dialogue box.

NOTE:
If the Maintain Itinerary switch, available in the Route tab of the BEHAVIOUR CONTROL dialogue box, is enabled, then the occupant will complete their tasks, irrespective of the environmental conditions.

Finally, given that a barrier is perceived and alternative routes exist, the likelihood of the occupant redirecting is determined. This is dependent upon the data-set selected by the user, from which the likelihood of redirection is taken. The user has the option to select between the data-sets of Bryan [64], Wood [63], an average of these two data-sets, or a user defined data-set (see Figure 7-13).

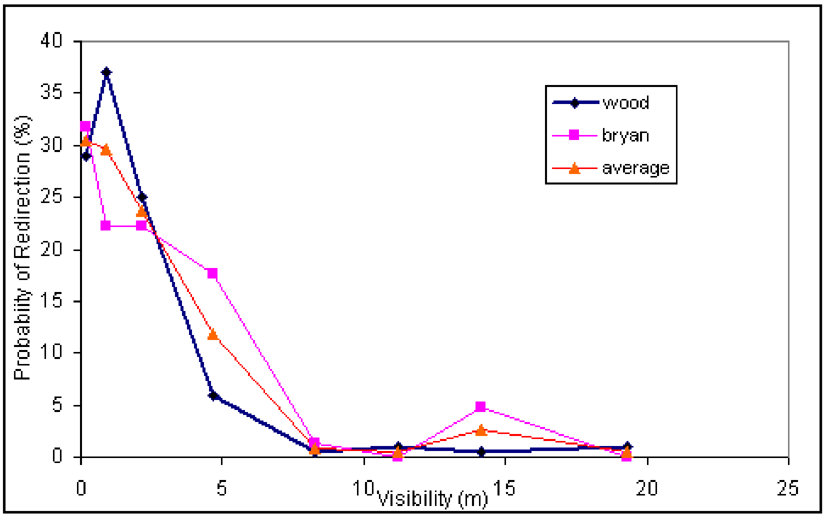

Figure 7-13: Representation of the smoke-redirection data-sets provided within the buildingEXODUS model based on the data of Bryan [64] and Wood [63].

The likelihood of the occupant moving through smoke is extracted from the work of Bryan [64] and Wood [63]. Bryan analysed data from 335 incidents, involving 584 occupant interviews originally conducted by fire personnel. The incidents all occurred in residential enclosures in the US [35,56]. Wood investigated 952 fires in the United Kingdom, questioning 2,193 occupants who were involved in residential and industrial incidents [35,56]. Although limited, these data-sets provide some of the best indications of occupant behaviour when confronted with smoke-filled environments.

Furthermore, Bryan and Wood attempted to quantify their findings and provide information concerning the proportion of people electing to redirect when faced with a given smoke concentration. From Wood's data, 88% of the sample claimed to have been involved in situations involving smoke. Of these, 40% did not necessarily have to move through the smoke. The remaining 60% encountered a smoke barrier that needed to be negotiated to reach their primary object. Of these people a proportion turned back, according to the density of the smoke (see Figure 7-13). Bryan found that 62.7% moved through smoke and 18% redirected, again according to the density of the smoke [27,34,36,64]. The data provided in Figure 7-13 relates to the conditions faced by those occupants that *decided to redirect* away from smoke.

NOTE:
The Wood and Bryan smoke redirection data should be considered at best only indicative of an occupant's response to smoke. The data suffers a number of defects that limits its applicability. For instance, smoke concentrations are crudely estimated, potentially important factors such as age, gender and occupant experiences are ignored, and the occupant response to the nature of the structure is not identified. This is discussed in more detail in the User Guide, Chapter 6 and the Application Manual, Chapter 5.

The Bryan and Wood data-sets have been distilled into composite-linear functions. This method is used to limit the number of mathematical assumptions required in the formulation of the data used. Additionally, data was converted into metric measurements and an upper limit was provided for the data-set to allow a consistency in the function generation. It should also be recognised that these functions are only one aspect of the assessment made by the agent when redirecting away from a smoke barrier; for instance, the agent's exit familiarity, the agent's prior actions, the presence of other agents, etc., all influence the outcome of the function.

Therefore, the impact of the function will not be identical across scenarios, or between multiple runs of the same scenario.

An average of the Wood and Bryan figures has been calculated to provide a general indication of occupant behaviour when faced with smoke. The user is also able to provide data concerning the likelihood of occupant redirection. The user is able to provide eight data-points representing the probability of redirection, given the visibility afforded to the occupant by the environment. Once the data-points have been provided, the model conducts a linear interpolation between the data-points to allow intermediate figures to be evaluated (see the User Guide, Chapter 6 and the Application Manual, Chapter 5). It should be borne in mind that this probability represents one aspect of the redirection process and that numerous other factors (e.g. familiarity, response of those deemed to be socially significant, etc.) will all have an influence. Given these facts, the numbers redirecting will be largely dependent upon the scenario.

These figures provide a likelihood of redirection given the visibility afforded to the occupant by the environment; the decision to redirect is stochastically based. If the simulation is repeated, potentially different results may be generated. If the occupant decides to redirect, the alternative exit identified previously is adopted as a new target. This action is then stored by the occupant as an act of redirection in relation to the smoke.

Once the decision to redirect away from smoke has been made the occupant is prevented from doing so again. An occupant who has previously redirected will automatically maintain their egress route and will not redirect once confronted by smoke. This rule is enforced to prevent continual redirection. However, the occupant may redirect in relation to other circumstances such as an exit becoming unavailable, etc.

Once a new target has been adopted, the occupant's movement will be subjected to the local and global rules that apply to general occupant movement.

7.2.2.5 (v) Toxicity interaction.

EXODUS links the *FIN* and *FIC* attributes to the *Mobility* and *Agility* attributes (see Section 3.1 part 2 and Section 3.1 part 4 respectively). As the *FIN* and the *FIC* increases, the *Mobility Degradation Factor* decreases in turn decreasing the person's *Mobility* and also *Travel Speed* (see Table 3-1). Similarly, As the *FIN* and the *FIC* increases, the *Agility Degradation Factor* also decreases in turn decreasing the person's *Agility*.

7.2.3 Terrain interaction:

7.2.3.1 (i) Free-space.
While travelling through a region defined as *Free-Space*, a person is allowed to travel at his/her maximum *Fast Walk Rate* (see CHAPTER 3:, Table 3-3 and Table 3-6) if unhindered.

7.2.3.2 (ii) Seats.
While travelling within a seat row (seat to seat), an occupant is allowed to travel at his/her maximum *Walk Rate* (see CHAPTER 3: and Table 3-3 and Table 3-6) if unhindered. Under *Normal* behaviour, occupants travelling within seat rows will choose to travel to the end of the seat row and wait (if necessary) to join the general flow.

If travelling between seat rows (over seats), an occupant is allowed to travel at his/her maximum *Leap Rate* (see Section 3.1 part 3 and Table 3-3 and Table 3-6). An occupant is given the option of leaping over a seat only if his/her *Wait Counter* attribute (see CHAPTER 3:) exceeds his/her *Patience* attribute (see CHAPTER 3:) and his/her *Agility* attribute (see CHAPTER 3:) exceeds the arc *Obstacle* attribute (see Section 2.4). *Seat* jumping may be completely forbidden by the user, by de-selecting the *Seat Jumping* option in the *Behaviour* tab of the *BEHAVIOUR OPTION* dialogue box in *SIMULATION* mode.

Occupants will not usually enter seat-rows to bypass crowds unless exhibiting *Extreme* behaviour.

7.2.3.3 (iii) Obstacles.
Obstacles can be represented within EXODUS through *Boundary Nodes* and through the obstacle values placed on arcs connected the nodes. Each approach will be discussed in turn.

7.2.3.3.1 (a) Nodes.
Obstacles modelled using *Boundary Nodes* should not be confused with the *Obstacle* attribute in arcs. They are intended to model obstacles that can be negotiated by all occupants irrespective of their *Agility* attribute. Occupants will attempt to avoid *Boundary* nodes in their path by veering around them. Any occupants that are forced to travel over *Boundary* nodes are slowed down to their *Walk-Speed* (see CHAPTER 3:). *Boundary* nodes are usually placed around immovable objects such as tables in a restaurant or along walls. Figure 7-6 displays an example of this type of behaviour. In this figure, occupant 2 starts in the restaurant area and takes a winding path to the exit passing the tables on the way. In this geometry, the tables are surrounded by a halo of boundary nodes thus forcing the occupants to slow down while moving in this area. A further discussion on the use of *Boundary* nodes can be found in the User Guide, Chapter 3.

Boundary nodes can be generated automatically through flagging the *Auto. Boundary Nodes/ Generate* option on the *Tools/Generate* menu in *Geometry* mode. This is done on the basis of nodal connectivity. The software searches for all *Free-space* nodes that have a missing connection to a node vertically or horizontally, and when located, the node is converted to a *Boundary* node. Once in place, the *Boundary* nodes may be manually converted to any other form of node if required. This feature is considered an aid to constructing complex geometries.

It is recommended that *Boundary* nodes be imposed around the boundary of an enclosure and around internal obstacles. This improves the representation of the occupant's behaviour by simulating the propensity of occupants to maintain a distance between themselves and the structure and the potential reduction in travel speed that may occur when in close proximity to immobile structures.

NOTE:
It is recommended that the automatic generation of Boundary nodes feature only be used in geometries that have NOT had significant levels of arc manipulation.

NOTE:
If the user is dissatisfied with the conversion process during the automatic generation of Boundary nodes, then the Undo function can be used to remove the action.

7.2.3.3.2 (b)Arcs.

Another method by which an obstacle may be represented concerns the *Obstacle* value of the *Arc*. By default all arcs have an *Obstacle* value of 0 indicating that no obstacle is present (i.e. no obstacle exists between the two nodes connected by the corresponding arc). Any arc having a *Obstacle* value >0 is assumed to represent an obstacle being present. The larger the *Obstacle* value the larger the obstacle posed to occupants. The ability of occupants to traverse arcs that have an associated *Obstacle* value (i.e. >0) will be dependent upon the node types of the nodes connected by the arc and also the *Agility* and *Mobility* of the occupant. If the occupants *Agility* value is greater than or equal to the arc *Obstacle* value then the occupant is deemed capable of traversing the obstacle (i.e. arc). If the occupants *Agility* value is less than the arc *Obstacle* value the occupant may still be able to traverse the arc if both the types of the nodes connected by the arc are *Free-space* or *Stair,* and the *Mobility* value of the occupant is ≥ 1.

Hence, in summary an occupant can traverse an arc with an assigned *Obstacle* value if:

1) Their *Agility* value ≥ the arc *Obstacle* value, **OR**
2) The types of the nodes connected by the arc are *Free-space* or *Stair* **AND** the *Mobility* value of the occupant is ≥ 1.

If occupants are capable of traversing the arc containing the obstacle then the arc *Obstacle* value will affect their travel speed. The manner in which the arc *Obstacle* value affects the travel speed of the occupant is dependent upon the types of node connected by the arc (i.e. the types of node between which the occupant is moving, see Table 7-3). It is important to note that the manner in which an arc *Obstacle* value reduces occupant travel speed is arbitrary and is merely intended to represent an increased level of difficulty in traversing the arc.

NOTE:
When an occupant is incapacitated through heat or toxic gases, the obstacle values associated with the arcs connected to the node of incapacitation is increased by 1.

Table 7-3: Impact of arc obstacle value on the occupant travel speed

Arc Obstacle Value	Resulting Occupant Speed on Free-Space	Resulting Occupant Speed on Stair
0.0	*Fast Walk*	*Stair Speed*
1.0	*Walk*	*Stair Speed/2.0*
2.0	*Walk/2.0*	*Stair Speed/3.0*
x (where x>2.0)	*Walk/x*	*Stair Speed/(x+1)*

If the nodes involved are of a different type (e.g. *Seat* nodes) then the *Obstacle* value may prevent progress. Under these circumstances, an occupant will only be able to travel across an arc if his/her *Agility* attribute is greater than the arc *Obstacle* value.

7.2.3.4 (iv) External Exits.

External exits connect the building interior with the exterior and are modelled using Exit Nodes. They have a profound impact on occupant behaviour as they are used to control the overall potential distribution, hence influencing global behaviour, and can be used to regulate the local behaviour of people passing over them. For this reason, each of the exit node attributes affecting flow performance are discussed in turn.

NOTE:
An External Exit will always be assumed to flow out of the structure. These occupants attracted to the exits are then assumed to be leaving the building.

7.2.3.4.1 (a) External Exit Potential.

The nodal potential distribution is *grown* from each external exit and increases with each step from the seed exit (see Section 2.3). The exit potential is used as the starting value for the map in the vicinity of the seeded exit. If each exit exerts an equal influence on occupants they should each carry an identical exit potential. The exit potential is given a default value of 100. Exits with differing attraction qualities can be represented by setting different exit potentials. The smaller the exit potential the more attractive it becomes (see CHAPTER 7:). Exit potentials can be set in *SCENARIO MODE* (see the User Guide, Chapter 5). A similar capability is available for internal exits (see the User Guide, Sections 3.2.1.7 and 5.3.3).

Attribute	: *Exit Potential.*
Range	: 0 - indefinite.
Default	: 100
Influenced by	: None.
Influences	: Potential map, global behaviour.
Used in level	: A, B and C
Note	: The Exit Potential attribute is a measure of the degree of familiarity occupants may have with an exit.

7.2.3.4.2 (b)External Exit Attractiveness

During the use of Localised Occupant Familiarity system, there are two means by which the occupant can be attributed with an exit list. Firstly, the exit list can be manually defined by the user (see the User Guide, Section 4.5). Secondly, the occupant's exit list can be automatically generated. The likelihood of adoption of an exit is then dependent upon a user-specified probability known as the *Exit Attractiveness*. The external *Exit Attractiveness* is a global parameter i.e. the value set for a particular exit is the same for each occupant. External exits are assigned to the *OEK* of each occupant's list through a stochastic process. For each occupant, each exit is examined in turn, through the generation of a random number. If the generated random number is less than or equal to the *Exit Attractiveness*, the occupant adds the exit to his *OEK* list. Therefore the higher the *Exit Attractiveness* the more likely that the exit will be familiar to an occupant. A *Exit Attractiveness* figure of 100 will guarantee that the occupant population will be familiar with the exit.

Attribute : *Exit Attractiveness.*
Range : 0% - 100%
Default : 100%.
Influenced by : None.
Influences : *Local Occupant Familiarity* List.
Used in level : A, B and C
Note : The Exit Familiarity attribute reflects the probability of an occupant being familiar with an exit.

7.2.3.4.3 (c) External Exit Usage

When using the *Localised Occupant Familiarity* system, a means is implemented to determine the occupant's initial target. The usage of the external exits of which the occupant is aware, will affect the adoption of an exit as a target. This reflects the general propensity of occupants not to blindly use exits according to their proximity alone, but also being influenced by the day-to-day usage of the exit and the circumstances in which they are placed. Under the *Normal* behavioural regime exits in *General Use* will be given priority over *Emergency* exits. Under the *Extreme* behavioural regime, no prioritisation occurs.

Attribute : *Exit Usage.*
Range : *General Use* and *Emergency*.
Default : *General Use*.
Influenced by : None.
Influences : *Local Occupant Familiarity*.
Used in level : A, B and C
Note : The *Exit Usage* attribute reflects the normal use of a particular exit.

7.2.3.4.4 (d) External Exit Availability.

During a simulation a particular external exit may become open or closed, once or several times. This may have an impact on the potential map depending on whether the exit is active or inactive (see Section e). The opening/closing times can be set by the user as a list of *event times* that alternately open and close the exit (see the User Guide, Section 5.3.2). Multiple event times should be avoided where possible, as each occurrence of an exit opening/closing requires a recalculation of the potential map that may slow the simulation significantly. Exit availability is set in *SCENARIO MODE* (see the User Guide, Chapter 5).

Attribute	: *Exit Event Times.*
Range	: 0 - indefinite (seconds).
Default	: 0 seconds.
Influenced by	: None.
Influences	: Exit availability, potential map, global behaviour.
Used in level	: A, B and C
Note	: Indicates when the exit availability changes, i.e. opens or closes.

7.2.3.4.5 (e) External Exit Activity.

External exits may be *active* or *inactive*. This will have an influence on the behaviour of the occupants, which is dependent on the nature of the global behaviour being implemented i.e. *Potential Map, Occupant Exit Knowledge* or *Target Exits*. Each case is examined in turn below.

7.2.3.4.5.1 (i) Potential Map

An external exit may be set to be *active* or *inactive*. An *active* exit will influence the potential map - occupants will be attracted to it - whereas an *inactive* exit has no effect on the potential map - occupants will ignore the exit. An exit may be any combination of open/closed and active/inactive, although some basic rules apply:

- Exits initially closed that open, become *active*.
- Exits initially open that close, become *inactive*.
- Occupants will be attracted to *active* exits, regardless if the exit is open or closed.
- Occupants will ignore *inactive* exits, regardless if the exit is open or closed.

7.2.3.4.5.2 (ii) Occupant Exit Knowledge

If the occupant is attributed with an exit list, the impact of the availability and activity of exits will largely reflect that described in the previous section. The occupant's initial targets are calculated normally as described in CHAPTER 7:. An occupant will then move towards the nearest *active* exit that is familiar to them, according to the priorities afforded by the behavioural regime.

If an exit becomes active during an evacuation and becomes the occupant's nearest viable exit, it is selected according to the preferences outlined in CHAPTER 7: (i.e. if the *Normal* regime is applied, exits in constant use are preferred to *Emergency* exits, while under the *Extreme* regime, no preference is applied). If an exit becomes unavailable during a simulation, the occupant is forced to examine his/her exit list to determine the next nearest viable exit. Again this is conducted using the preferences outlined in CHAPTER 7:. If no other exits exist, then the occupant may become trapped within the geometry.

NOTE:
The user should take care when designing the evacuation scenario not to generate unrealistic situations, i.e. leaving the occupant without any viable exit route. The OEK should represent the entire exit awareness of the occupant. Unlike the application of Target Exits, OEK will not default to the underlying Potential Map system if the occupant is left without an exit, as this would present a significant compromise to the underlying assumptions of the system.

7.2.3.4.5.3 (iii) Target exits

Occupants with a specified *Target Exit* have the ability to open the Target Exit if it is found closed, according to the attributes of the *Target Exit*. Table 7-4 outlines the relationship between door attributes and the action performed by such an occupant. The occupant will open the door upon arrival if the door is closed and no opening time is scheduled for the door. Note that there is no delay time associated with the opening process. If an opening time *is* specified for a target exit, occupants targeted to it will wait at the door until it becomes open according to the opening time. This is a particularly useful feature when using the software to reconstruct known events.

If at any point, a *Target Exit* becomes unavailable, i.e. deleted, disconnected or Off, then any occupants targeted to the door will follow the potential map, i.e. the *Target Exit* is reset.

Table 7-4: Door attribute and occupant action relationship for occupants with a target door specified.

Target door Availability	Target door Opening Times	Occupant Action at target door
Open	Specified	Exits
	Unspecified	Exits
Closed	Specified	Queues
	Unspecified	Opens Door
Off	-	Follow potential map

NOTE:
The activity of an exit has no impact upon its attractiveness to an occupant, if the occupant has had the exit specified as a target exit.

Note that the active attribute can only be set initially in *SCENARIO MODE* (see the User Guide, Chapter 5).

Attribute : *Activity.*
Range : Active and Inactive.
Default : Active.
Influenced by : None.
Influences : Exit availability, potential map, global behaviour.
Used in level : A, B and C
Note : The Exit Activity determines whether or not occupants are attracted to the exit.

7.2.3.4.6 (f) External Exit Width.

The width of an external exit is defined by the *number* of arcs connected to it. Generally, two arcs should be used per metre width of exit. Fractional widths should generally be rounded down in order to produce conservative results. The user sets the exit width in *GEOMETRY MODE* (see the User Guide, Section 3.5). Exit widths are important as they specify the maximum number of people that can occupy an exit opening at any one time.

7.2.3.4.7 (g) Unit flow rate (UFR) Applied to External exits.

The UFR attribute for an external exit is a measure of the number of occupants per metre of exit which can pass through the exit per second (occupants/metre/second). The flow rate through the door (occupants/second) is calculated by multiplying the exit width (see Section f) by the UFR assigned to the exit. This is used to specify a maximum allowable flow rate for the external exits. If a UFR is specified for an exit, the actual flow rate determined by buildingEXODUS during the course of a simulation cannot exceed this maximum value. Thus, the buildingEXODUS predicted flow rate achieved by the exit is effectively CAPPED by specifying a unit flow rate. The UFR of an external exit is defined in *SCENARIO MODE* (see the User Guide, Section 5.3.1).

The UFR assumes that a linear relationship exists between exit width and flow rate. According to the literature [19,28-30] this assumption is valid for exits with widths in excess of 1 metre. However, different authors have published different gradients [19,28-30] for the relationship. buildingEXODUS offers a choice of these values from the literature, or the possibility of defining one explicitly. Rather than define a single value for the UFR at each exit, buildingEXODUS uses a range of values. At each exit an upper and lower limit is specified and occupants passing through the exit are randomly assigned a maximum UFR between the specified limits. In the User Guide, Section 5.3.1 explains this process in more detail. It should be noted that this process defines the maximum flow rate that an external exit can achieve at the instant an occupant passes through the exit. It does not mean that the exit will actually achieve this flow rate. In practice, as the number of people using an exit increases to the point that the exit is working at its full capacity (i.e. the exit has achieved its nominal maximum flow rate), further increases in occupant density around the exit will tend to decrease the flow rate as the number of conflicts predicted by buildingEXODUS increases.

Table 7-5 lists the available UFR for standard exit types.

Table 7-5: Exit Unit Flow Rate options in buildingEXODUS

	Unit Flow Rate (occ/m/s)	
Source	**Min**	**Max**
User-Defined	999	999
HMSO [28]	1.33	1.33
HANKIN [30]	1.46	1.46
POLUS [29]	1.25	1.58
FRUIN [19]	1.33	2.0

NOTE:

By specifying a very large UFR for an exit through the User-Defined option (e.g. 999 occ/m/s), the CAP on the buildingEXODUS predicted flow rate is effectively removed. Under these conditions the exit is described as being in a "free-flow" condition. In free-flow situations, buildingEXODUS can produce UFRs of around 2 occ/m/s, which appear large when compared with the HMSO approved rate, but are in line with the upper end of the Fruin measured flow rates [19] and other experimental observations [37,44,45,50].

NOTE:
The flow capability of an exit is dependent on the exit type. For example free-swinging exits have different capabilities to turnstile exits and hence will have a different range of exit unit flow rates. Unfortunately, the current literature does not suggest appropriate gradients for these exit types. However, if data were available it is possible to specify the appropriate unit flow rates using the user-defined option.

Attribute : *Exit Unit Flow Rate.*
Range : 0 - indefinite (occupants/m/sec).
Default : 1.33 occupants/m/sec.
Influenced by : Exit type.
Influences : Exit flow rate.
Used in level : A, B and C
Note : The UFR determines the maximum flow rate of people through an external exit per metre of width.

7.2.3.4.8 (h) Extreme Behaviour at External Exits.

Extreme behaviour at external exits will have an influence on the behaviour of the occupant, however this is dependent on the type of global behaviour employed, i.e. *Potential Map*, *Occupant Exit Knowledge* or *Target Exit*. Each case is described below in turn.

7.2.3.4.8.1 (i) Potential Map

Occupants are prevented from exhibiting *Extreme* behaviour (see CHAPTER 7:) if they are within 2.5 metres of a functioning External exit. Within this distance, occupants will crowd around the exit and wait their turn to exit. If however the occupant is further than 2.5m from their initial exit, they will be able to become *Extreme*. The occupant is then able to move away from their current exit and may fall into the potential well of a nearby exit. *Extreme* behaviour effectively prevents occupants from queuing and may allow them to use other exits situated nearby (see Table 7-6) if the potential map dictates that these other exits are sufficiently attractive.

Table 7-6: Extreme behaviour at exits and resulting possible redirection, when the potential map system is used.

Potential Map	***Normal* Behaviour**	***Extreme* Behaviour**
Occupant within 2.5m of an exit	No redirection possible, however, occupant may fall under the potential well influence of a nearby exit and thus be attracted to it.*	Occupant unable to become *Extreme*. No redirection possible, however, occupant may fall under the potential well influence of a nearby exit and thus be attracted to it.*
Occupant outside 2.5m of exit	No redirection possible, however, occupant may fall under the potential well influence of a nearby exit and thus be attracted to it.*	Occupant is able to become *Extreme*. No redirection possible, however as occupant can go extreme, it is more likely that he may fall under the potential well influence of a nearby exit and thus be attracted to it.*

**it should be noted that as the exits become further apart, so the occupant is progressively less likely to fall into an alternative potential well.*

7.2.3.4.8.2 (ii) Occupant Exit Knowledge

As in the previous section, the exit location and the distance of the occupant from his target exit is influential, as is the behavioural regime applied (see Table 7-7).

In all cases the occupant must become impatient prior to considering redirection. If the *Normal* behavioural regime is applied, then redirection is only available to non-emergency exits once the occupant is impatient. When the *Extreme* behavioural regime is applied, the occupant is able to redirect if the occupant is between 2.5-5m of their initial exit and have become impatient. Their ability to do so is dependent upon the existence of an exit within 5m of their current exit. If this exit appears in the occupant's exit list, then they are able to redirect towards it. If it is not initially in the occupant's exit list, then the occupant is assumed to *become* aware of the exit, add the exit to their exit list and then redirect towards it (see Figure 7-14). The use of the exit does not affect the likelihood of exit adoption. This is intended to reflect the greater desire of the occupants under the extreme behavioural regime to move to safety.

Table 7-7: Extreme behaviour at exits and resulting possible redirection, when the *OEK* system is used.

OEK		
	***Normal* Behaviour**	***Extreme* Behaviour**
Exits separated by greater than 5m and occupant within 2.5m of an exit	No redirection	Occupant unable to become *Extreme*. No redirection.
Exits separated by greater than 5m and occupant further than 2.5m of an exit	No redirection	Occupant is able to become *Extreme*, but redirection will not occur.
Exits separated by less than 5m and occupant located between 2.5 - 5m of an exit	Redirection to other exit if it is a non-emergency exit and if occupant is closer to other exit (as determined by distance map). If this exit is not known to the occupant, the occupant will adopt this exit in his OEK.	Occupant is able to become *Extreme*. Redirection will occur to exit within 5m of initial target exit if occupant becomes *Extreme*. If this exit is not known to the occupant, the occupant will adopt this exit in his OEK.
Exits separated by less than 5m and occupant is within 2.5m of exit	No redirection.	Occupant is unable to become *Extreme*, redirection will not occur.
Exits separated by less than 5m and occupant further than 5m of an exit	No redirection.	Occupant is able to become *Extreme*. However, redirection will not occur.

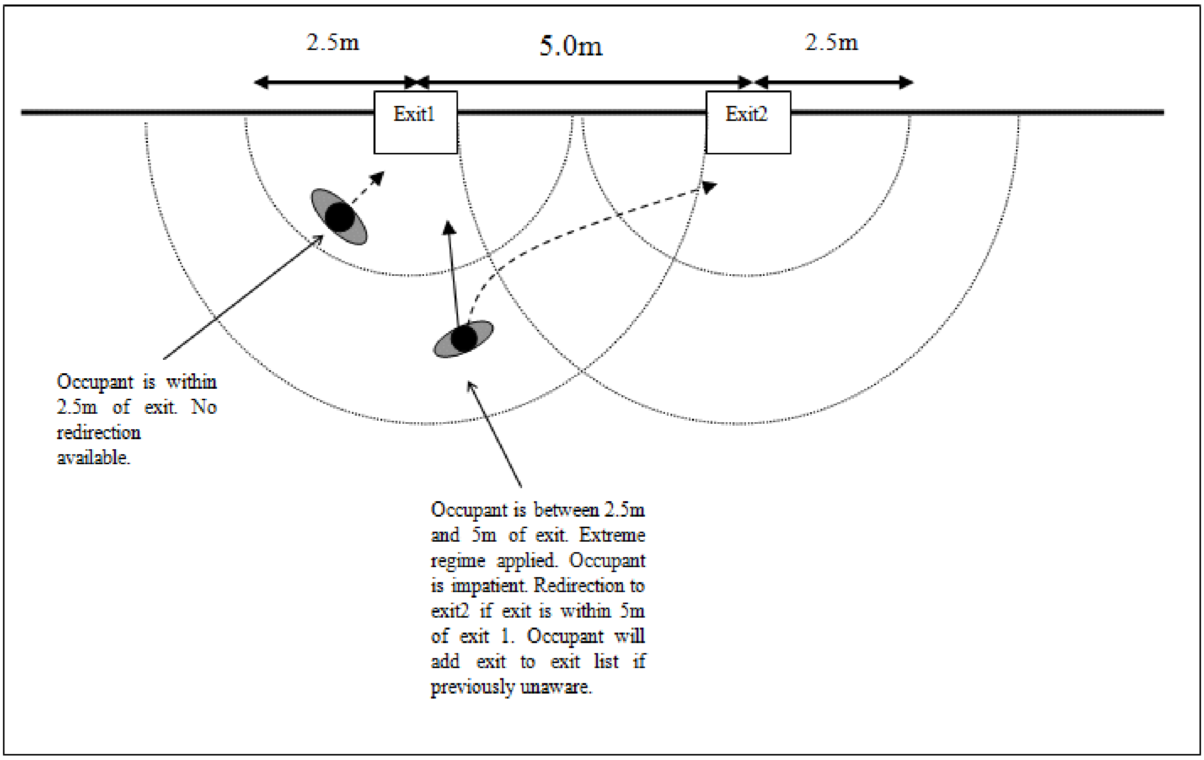

Figure 7-14: Potential redirection from target exit when *Extreme* behaviour and *OEK* options are in operation.

NOTE:
It should be noted that the redirection referred to in relation to the OEK system involves the occupants adopting a new target, rather than relying upon falling into the catchment area of a new exit, as is the case with the use of the Potential Map system.

NOTE:
Occupant redirection when using the Potential Map or OEK options are not directly comparable, as they reflect the underlying differences in the mechanisms used. The redirection evident when using the OEK system is due to the occupant adopting the alternative exit, whereas when the Potential Map system is used, the redirection is dependent on the nature of the potential map.

7.2.3.4.8.3 (iii) Target Exit

If the occupant has a Target Exit, then they will continue moving towards their *Target Exit* until they become impatient due to congestion. At this stage two options are available. If the *Maintain Target Exit* option has been flagged in the *BEHAVIOUR OPTIONS* dialogue box (which it is by default), then the occupant will continue on towards their current exit irrespective of the congestion around it. If this option is not flagged then the occupant will revert to the underlying navigation being used at the time. This will either be the potential map system (described in (i)) or the *OEK* system (described in (ii)).

NOTE:
It should be noted that the Target Exit system can be used simultaneously with the Potential Map system and the Local Familiarity system.

7.2.3.4.9 (i) Conflict Resolution at External Exits.

When an occupant occupies an external exit node, no further conflicts can occur and the occupant is considered "safe". However, the conflict resolution process applies around external exits so that occupant-occupant interactions influence the exiting procedure. As a result, arch-like structures can be seen to form and collapse at crowded exits and exit apertures may not be fully occupied, even under crowded conditions.

7.2.3.4.10 (j) The Exiting Procedure at External Exits.

In buildingEXODUS the flow capability of an external exit is a function of the exit width (see Section f) and the unit flow rate (see Section g). The maximum flow rate capability through the exit (occupants/second) can be calculated by multiplying the exit width (metres) by the unit flow rate (see Section g) assigned to the exit (occupants/metre/second).

As an occupant reaches an exit node, he/she is delayed at the exit according to the specified unit flow rate. As the unit flow rate decreases at an exit, the delay for each occupant increases. Under free-flow conditions (see Section g), the delay at the exit is infinitesimal, i.e. the occupants are not delayed. Under such conditions, the flow rate of the exit is simply dependent upon the actual number of occupants passing through the exit.

7.2.3.5 (v) Internal Exits.

Within buildingEXODUS internal doorways can be represented using three different means (a) simply as restrictions to the flow caused by a reduced number of nodal connections, (b) using the *Internal Exit* option or (c) using *Attractor/Discharge* nodes.

Using the first method, internal exits are simply considered to be constrictions to occupant flow and are essentially modelled as free space. This is the simplest (and the default) method and the majority of internal exits are usually represented in this manner. Using this approach, no manipulation of internal exits is required and internal exits behave as if under free-flow conditions.

However, if it is considered necessary to regulate the flow through certain internal exits additional manipulation is required and either the *Internal exit* or *Attractor/Discharge* nodes are required. It is however recommended that the *Internal exit* method is always used. The *Internal Exit* option is preferred as it is the more versatile of the two techniques, allowing internal doorways to have all of the functionality of *external exits*.

The use of *Attractor/Discharge* nodes in controlling the flow through internal exits is also discussed in Section (3)(vi). Although a secondary means by which to model internal exits, *Attractor/Discharge* nodes are a useful tool in manipulating occupant movement within a geometry, as they allow the flow rate to be capped and also allow the local biasing of the potential map. A general discussion on the detailed use of A*ttractor/Discharge* nodes is presented in the User Guide, Chapter 5 and the Application Manual, Chapter 2. Some additional discussion on the effects of limiting the flow on internal exits can be found in the User Guide, Chapter 3.

Internal Exits are created by selecting the group of nodes that are to form the internal exit and selecting the *Internal Exit* option from the *Generate* sub-menu of the *Tools* menu. This provides the selected nodes with the collection of properties necessary to define the exit, namely those required to describe an external exit (Exit width, Flow rate, etc.). This functionality includes the ability of the user to control exit flow rates, the exit opening times, the exit's potential and the width of the exit. Once defined, the use of the exit is recorded in the relevant simulation file, in an identical manner to external exits.

NOTE:
It is recommended that all internal exits that are required to be controlled in some manner are represented using the Internal Exit node as it allows a consistent representation of exits throughout the geometry and provides an increased level of information to the user.

NOTE:
Users are advised not to attempt to substitute free-space nodes for Internal Exit nodes. If the user wishes to remove an Internal Exit from a geometry the exit should be selected and then deleted using the DELETE button on the dialogue box provided.

NOTE:
It is recommended that the user places Internal Exit nodes on either side of the internal opening. The occupants will incur a delay on moving from a free-space node to an Internal Exit node, irrespective of the direction of movement.

NOTE;
It should be remembered that whilst using the Local Familiarity system, the delay incurred while passing through the Internal Exit is recorded and will affect the occupant's movement. However, the alteration of the potential map made by the user will have no effect on the direction of the occupant's movement as the Local Familiarity system is dependent upon the distance map and not the potential map system. Alternative means are available to manipulate occupant movement within the structure.

Each of the key *Internal Exit* node attributes affecting flow performance are discussed in turn below.

7.2.3.5.1 (a) Internal Exit Width.

To remain consistent with the calculation of the width of *External Exits*, the calculation of *Internal Exit* Width is *lane-based*. This approach provides a consistent basis for behaviour in and around an internal opening. Each lane is assumed to carry a single occupant at any one time and is assumed to be 0.5m in width. The user must take this into consideration when calculating the impact of an imposed flow rate cap applied to the internal opening. This is covered in great detail in the User Guide, Chapter 5.

Several nodes may be required to model the opening of an internal exit. Where buildingEXODUS has automatically positioned *Free-Space* nodes in the exit (through the use of the *Node Flood* tool), the user may not be satisfied with the results and may wish to manually alter their position (see the User Guide, Section 3.3.3), to more accurately represent the actual internal exit's construction. This process occurs in *GEOMETRY MODE* (see the User Guide, Chapter 3).

NOTE:
It is recommended that users check that the appropriate number of nodes has been placed in the internal opening.

7.2.3.5.2 (b) Internal Exit UFR.

The UFR attribute for an *Internal Exit* functions in the same manner as the UFR for an *External Exit*, i.e. it limits the maximum flow of occupants through the exit. The maximum flow capability (measured in occupants per second) can be calculated by multiplying the UFR (occupants/metre/second) by the width (m). It should be noted that the UFR for the entire *Internal Exit* is defined by setting the UFR on any one of its constituent nodes. This is based on the assumption that *Internal Exits* are a single entity. Again it should be remembered that the movement of occupants through an *Internal Exit* is assumed to be lane-base, i.e. that the occupant is assumed to occupy 0.5m of the exit. This assumption should be taken into consideration when calculating an appropriate UFR setting.

Attribute	: *Unit Flow Rate (UFR).*
Range	: 0 - indefinite (occupants/m/sec).
Default	: 999.
Influenced by	: Exit type.
Influences	: Exit flow rate.
Used in level	: A, B and C
Note	:The UFR determines the maximum flow rate of people through the constriction, per metre of width.

It should be noted that the following attributes associated with *Internal Exits* (described in (c), (e) and (f)) are designed to be used with the *Potential Map* guidance system. The results that might be produced when using these attributes, which affect the generation of the *Potential Map* may not be predictable when used with other systems (i.e. the *Local Familiarity* or the *Target Exit* systems). It is reasonable for the user to utilise the Internal Exits to delay the movement of the Evacuees through the application of a flow rate, without applying the changes to the generation of the potential map. There is also an associated computational expense with the recalculation of the potential map due to the changing availability of internal exits.

7.2.3.5.3 (c) Internal Exit Potential

It is possible to modify the local *Potential Map* within a room through the use of an Internal Exit node (see the User Guide, Chapter 5). This allows the *Potential Map* within a room to be set to achieve the desired local behaviour, for example making one internal exit more attractive than another. The *Potential Map* is developed in a similar manner to that for exits (see Section (3)(iv) part (a)). A seed potential is placed on the *Internal Exit* node and the potentials are "grown" from this point, thereby modifying the *Potential Map*. The *Potential Map* is first determined by calculating the global potentials as specified by the seed potentials on the exit nodes. The *Potential Map* is then modified by the seed potentials specified on the *Internal Exit* nodes. However, these potentials are only considered if the *Local Potentials* option is set in the *BEHAVIOUR OPTIONS* dialogue box, see the User Guide, Section 6.4.3. Smaller Internal Exit potentials exert a greater influence on the potential map. The seed potential for Internal Exits are specified in *SCENARIO* mode (see the User Guide, Chapter 5).

Once created, the *Internal Exit* is a single entity. Therefore once the potential of a constituent node within an *Internal Exit* has been altered this will affect the whole exit (i.e. all of the other constituent nodes). This therefore allows comparison *between* Internal Exits rather than *between constituent nodes* of an Internal exits. This is to allow each of the nodes within the exit to appear equally desirable to the occupants.

Attribute : *Potential.*
Range : 0 - indefinite.
Default : 0
Influenced by : None.
Influences : Potential map, global behaviour.
Used in level : A, B and C
Note : *Internal Exit* potentials modify the local potential map and only take effect if the *Local Potentials* option is set in the *Rule Base* menu (see the User Guide, Section 6.4.3).

7.2.3.5.4 (d) Internal Exit Availability.

During a simulation a particular *Internal Exit* may become open or closed, once or several times. This may have an impact on the local potential map depending on whether the exit is active or inactive (see Section e). The opening/closing times can be set by the user as a list of *Event Times* that alternately open and close the exit. Multiple event times should be avoided where possible, as each occurrence of an exit opening/closing requires a recalculation of the potential map that may slow the simulation significantly. Exit availability is set in *SCENARIO MODE* (see the User Guide, Chapter 5).

Attribute : *Exit Event Times.*
Range : 0 - indefinite (seconds).
Default : 0 seconds.
Influenced by : None.
Influences : Exit availability, potential map, global behaviour.
Used in level : A, B and C
Note : Indicates when the exit availability changes, i.e. opens or closes.

7.2.3.5.5 (e) Internal Exit Activity.

Internal Exits may be *active* or *inactive*. This will have an influence on the behaviour of the occupants which is dependent on the nature of the global behaviour being implemented i.e. *Potential Map, Occupant Exit Knowledge* or *Target Exits*. Each case is examined in turn below.

An *Internal Exit* may be set to be *active* or *inactive*. An *active* exit will influence the potential map - occupants will be attracted to it - whereas an *inactive* exit has no effect on the potential map - occupants will ignore the exit. An exit may be any combination of open/closed and active/inactive, although some basic rules apply:

1) Internal Exits initially closed that open, become *active*.
2) Internal Exits initially open that close, become *inactive*.
3) Occupants will be attracted to *active Internal Exits*, regardless if the Internal exit is open or closed.
4) Occupants will ignore *inactive Internal Exits*, regardless if the *Internal Exit* is open or closed.

NOTE:
It is recommended that the user limits the number of changes in the availability of the exit as each time it changes the Potential Map is recalculated, slowing down the simulation.

7.2.3.5.6 (f)Direction of Use

Internal Exits allow occupants to move across them in both directions. This is particularly useful if the user is interested in examining numerous circulation patterns within a structure. However, this does not represent *simultaneous* bi-directional flow. The user may control the direction in which occupants are attracted towards *Internal Exits,* through the use of the *Direction* attribute.

NOTE:
The Direction attribute for Internal Exits is only accessible in SCENARIO mode.

The *Direction* attribute determines the manner in which the *Internal Exit* affects the spreading of the *Potential Map*. By altering the direction of occupant attraction from *Forward* to *Backward* (or visa versa) the direction in which the occupants use an *Internal Exit* can be manipulated.

NOTE:
Simultaneous bi-directional flow through an Internal Exit can occur through the use of a combination of navigational systems. However it cannot occur through the influence of the potential map alone.

Occupant movement is calculated according to the occupant's distance from *Internal Exits* and from *External Exits*. This calculation takes into consideration the direction of the *Internal Exit* and whether the *Local Potential* option is enabled in *SIMULATION* mode. If the *Local Potential* option is enabled occupants are attracted to their nearest exit point that flows away from them, irrespective of whether it is an *Internal* or an *External Exit*. Therefore if the occupant is nearest to an *Internal* exit that is flowing away from them, they will pass through this exit prior to any other *External Exit*s. If the nearest exit point is an *External exit*, the occupant will ignore any other *Internal* exits and will move directly to it. If the *Local Potential* option is disabled then the *Internal Exits* will have no impact upon the *Potential Map*.

NOTE:
All External Exits are assumed to flow away from the occupants out of the structure. Thus it is not necessary to specify a Direction attribute at these exit points.

By changing the *Direction* of flow of an *Internal Exit*, its catchment area and therefore the occupants that will use the exit can be influenced (see Figure 7-15).

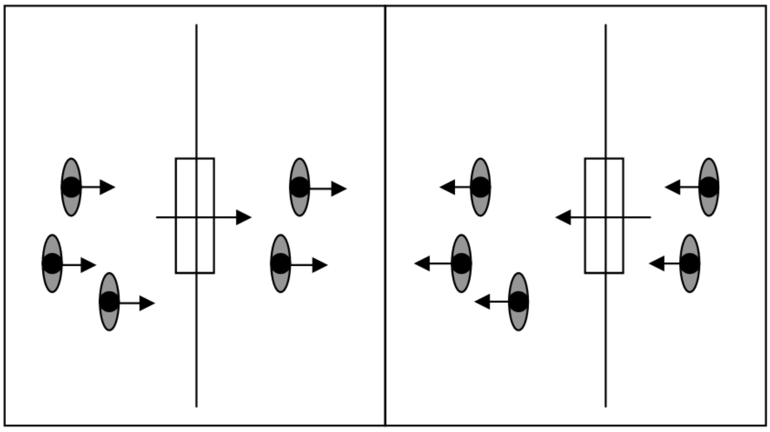

Figure 7-15: The alteration of the Direction of the Internal Exit affects the movement of the surrounding population.

This reversal of the flow direction may change which occupants eventually use the exit, as shown in Figure 7-15.

Once the geometry has been created and the simulation started the model makes a pass at estimating the direction of flow for the *Internal Exits*. The initial direction of flow through the *Internal Exit* will be towards the nearest active *External Exit*. In certain geometries this may not accurately reflect the movement required by the user. For example in Figure 7-16 the initial direction of the flow for the two *Internal Exits* is demonstrated. However, this may not satisfactorily reflect how the user wishes the resident population to behave. To correct this, the flow direction of the northern *Internal Exit* is reversed (see Figure 7-16). Once this has been performed the resident occupants will then flow 'out' of the compartment and will then be sensitive to the biasing applied to the *Internal Exits* (i.e. the actual use of the *Internal Exits* by the resident occupants can be regulated by altering the relative biasing of the *Internal Exits*).

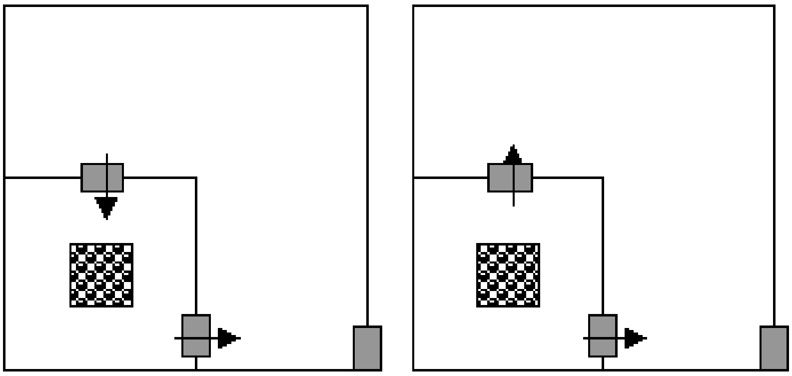

Figure 7-16: (left) Initial pass of the model. (right) User has altered flow direction of northern internal exit.

Therefore, when the usage of *Internal Exits* does not satisfactorily reflect the user's expectations the user should manipulate the *Direction* of flow for the *Internal Exits*. It should also be remembered that the *Direction* of flow affects occupant movement in conjunction with the *Potential* biasing attached to the exits.

7.2.3.5.7 (g) Extreme behaviour at Internal Exits.

Occupants are able to go *Extreme* at and around *Internal Exit* nodes. *Extreme* behaviour at *Internal Exits* will only influence the occupant's ability to occupy nodes of a higher potential/greater distance (see CHAPTER 7:).

7.2.3.5.8 (h) Conflicts at Internal Exits.

As the occupant is expected to continue evacuating once through the *Internal Exit*, rather than to have completed the evacuation, the occupant will experience conflicts before, during and after occupying an *Internal Exit* node.

7.2.3.6 (vi) Attractor/Discharge Node Representation of Internal Openings

Attractor/Discharge nodes are generally used to control the local *Potential Map* at a desired location in order to better control occupant movement in difficult areas of the geometry or areas where the *Potential Map* has not satisfactorily reproduced occupant movement. A general discussion on the detailed use of these nodes is presented in the User Guide, Chapter 5. However, in addition some limited functionality of *Internal Exits* can be represented in the software through a combination of *Attractor/Discharge* Nodes.

Due to the generality of their application, there are a number of differences in the uses and generation of *Attractor/Discharge* nodes in comparison with *Internal Exit* nodes. Firstly, a group of *Attractor/Discharge* nodes are not considered as a single entity, allowing both UFR and *Potential* levels to alter across them. This is particularly important when manipulating occupant movement along corridors or around corners, allowing the user greater control over the performance of specific nodal locations. Secondly, diagonal arcs must not be used to connect *Attractor/Discharge* nodes. This is due to the structures used in their construction and also the manner in which the nodal width and the UFR interact. This is further explained in the User Guide, Chapter 5.

In this section the use of *Attractor/Discharge* nodes in controlling the flow through an opening is discussed. Some additional discussion on the effects of limiting the flow on *Internal Exits* and in other areas can be found in the User Guide, Section 3.2.1.11.

7.2.3.6.1 (a) The Width of the Internal Opening using Attractor / Discharge nodes.

Several *Attractor/Discharge* nodes may be required in order to model an opening inside the geometry. The actual width of each node is defined by the arcs emanating from it (see Section 2.5). In order to correctly control the flow rate through an internal opening, buildingEXODUS requires an *accurate* knowledge of the width of each *Attractor* node used to model the exit. Where buildingEXODUS has automatically positioned free-space nodes in the exit (through the use of the *Node Flood* tool), these must be manually changed to *Attractor* nodes (see the User Guide, Chapter 3) after which buildingEXODUS will automatically estimate the internal exit width. This process occurs in *GEOMETRY MODE* (see the User Guide, Section 3.3.3).

NOTE:
When using Attractor/Discharge nodes to specify an internal exit, it is recommended that the user check that the appropriate number of nodes have been placed in the Internal Exit. It is also important to check that the arc lengths are correct as the width of an exit is an important factor in determining the overall flow capability of an exit.

NOTE:
The user has the ability to replace the width of an Attractor node using the Substitute function, overriding any other automatic calculation concerning the width of the nodes involved. This capability is available in Scenario mode and is described in the User Guide, Chapter 5.

7.2.3.6.2 (b) The Application of UFR to Attractor nodes.

Unlike the specification of *External / Internal Exits*, a menu selection of *standard* UFRs is not provided for *Attractor/Discharge* nodes, only a *User Defined* option is available. This is because the manner in which the user may apply *Attractor/Discharge* nodes is not restricted and therefore the standard UFRs may not be appropriate. However, the capability of having a maximum and a minimum UFR is maintained, allowing the user to provide their own flow rates. In all other respects, this feature is identical to that found on the *Internal Exit* nodes, see (3iv) part (g) above for further discussion. It should be noted that when using *Attractor/Discharge* nodes to represent internal openings, the UFR for the internal exit is set on the *Attractor* nodes and is required to be set on each of the *Attractor* nodes representing the opening. As with *Internal Exits* the delays incurred by an evacuee when crossing an Attractor node will be the same irrespective of the use of the *Potential Map* or the *Distance Map*.

Attribute	: *Unit Flow Rate (UFR).*
Range	: 0 - indefinite (occupants/m/sec).
Default	: 999.
Influenced by	: Exit type.
Influences	: Exit flow rate.
Used in level	: A, B and C
Note	: The UFR determines the maximum flow rate of people through the constriction, per metre of width.

7.2.3.6.3 (c) Attractor Node Potential

It is possible to modify the local *Potential Map (*within a room, compartment or corridor) through the use of *Attractor* nodes. However, the node *Potential* must be set for *each* of the *Attractor* nodes. This allows the potential map within a room to be set to achieve the desired local behaviour. The *Attractor* node potentials will only be effective if the *Local Potential* option is set in the *BEHAVIOUR OPTIONS* dialogue box, see the User Guide, Chapter 6 and if the *Potential Map* is used during the simulation.

Attribute	: *Potential.*
Range	: 0 - indefinite.
Default	: 0
Influenced by	: None.
Influences	: Potential map, global behaviour.
Used in level	: A, B and C
Note	: The potential attributed to the Attractor node modifies the local potential map and only take effect if the *Local Potentials* option is set on the *BEHAVIOUR OPTIONS* dialogue box from the *Rule Base* menu (see the User Guide, Chapter 6).

7.2.3.6.4 (d) Extreme behaviour at Attractor/Discharge nodes.

Occupants are able to become *Extreme* at and around *Attractor/Discharge* nodes. *Extreme* behaviour functions normally around *Attractor/Discharge* nodes (as described in CHAPTER 7:). The conditions surrounding the internal opening may generate congestion, causing an increase in the extent of *Extreme* behaviour.

7.2.3.6.5 (e) Conflicts.

The conflict resolution process applies in and around *Attractor/Discharge* nodes. This means that occupant-occupant interactions have a strong influence on the flow rates achieved within them.

7.2.3.7 (vii) Occupant selection of Transit Nodes

The *Transit Node* functionality allows different components to be represented within the model, including travelators, corridors, metered gates, stairways, escalators and lifts/elevators. Each *Transit Node* is a single entity, but represents components that are typically made up of many nodes. These components are typically considered as single entities in practice as they influence occupant performance in a specific manner. For instance, a staircase is not considered to be a collection of individual steps and lanes, but instead merely a single entity.

NOTE:
While a Transit Node is a single entity, within buildingEXODUS the space represented by the Transit Node is modelled as if it was meshed with fine nodes. Thus agent-agent interaction is simulated within the Transit Node. For example, if sufficient space was available on a flight of stairs, a faster moving agent could overtake a slower moving agent. Furthermore, if an agent is incapacitated at some point on a flight of stairs represented by a Transit Node, part of the stair occupied by the incapacitated agent can be blocked, effectively impeding the motion of other agents.

Transit nodes enable the user to manage the local performance of occupants on each component, and also (if required) cap the achievable performance. They also allow information to be collected on a component-basis; i.e. flow through a stair, number of people who used the stair, etc. Previously, the user would have had to insert *Census* components or collect data manually.

When faced with alternatives routes (i.e. *Transit Nodes* in parallel leading from one level to another), occupants will need to choose between components; e.g. select between stairs and escalators when moving between levels. Within the model, it is assumed that occupants will make this decision based on:

(1) the conditions evident as they approach the components, or
(2) the requirements of the user.

This functionality therefore allows approaching occupants to select between components that have been identified as part of a *Transit Node* group. This explicitly identifies alternatives to move between the regions/levels within the geometry.

NOTE:
It is important that when producing Transit Node groups, that the user ensures that they are genuine alternatives for the same route; i.e. that they start and end on equivalent levels. If this is not the case then the model may produce unexpected results.

The user is able to define the area associated with each *Transit Node* from which the conditions can be assessed (e.g. congestion levels, clearance times, etc.). This is defined by modifying the *Catchment* attribute associated with each *Transit Node*. Agents have access to the conditions at the *Transit Node* when located within the *Catchment* area of the *Transit Node*. When a *Transit Node Group* is defined, the areas associated with each of the *Transit Nodes* are grouped, forming a superset of all of the areas (see Figure 7-17); i.e. information on any of the grouped *Transit Nodes* can be gathered when entering this area. This is a necessary simplification to allow the alternative routes to be compared given that an agent has entered this (superset) area and wishes to use one of the *Transit Nodes.*

Components can be broadly categorised into those that provide *Continuous Service* and those that provide *Intermittent Service*. Continuous service components include travelators, corridors, ticket gates, stairways, and escalators. These devices, providing there is no crowding, do not require agents to wait to use them: they are continually available. However, intermittent service components, including lifts/elevators, can require agents to wait for the device to service a given area, i.e. they are intermittently available. Occupant selection of which transit node to use depends on whether the *Transit Nodes* provide *Continuous* or *Intermittent Service*. The following two sections describe the selection method for the *Continuous* and *Intermittent* service *Transit Nodes*. It is important to note that transit nodes groups that include both *Continuous* and *Intermittent* service components require the user to manually assign agents to use a given device. However, it is generally advised that users should not mix the two types of *Transit Nodes* within groups as they use different agent selection systems (these are described in the following sections).

Continuous Service

For groups of *Transit Nodes* that provide continuous service, there are three methods available to choose between the *Transit Nodes* available [124,157].

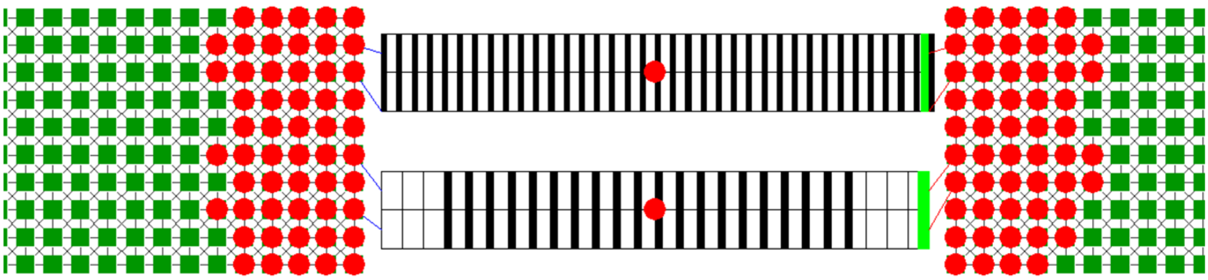

Figure 7-17: Escalator and stairway transit node group with the catchment area displayed (red circles).

In the **first approach** the user merely states the proportion of agents that will use each component. This approach is most appropriate for scenarios where lower levels of crowding occur at the entrance to the components, or where the user has specific information on the expected use of the components. This approach relies on the user providing data that is representative of the scenario being considered. Using this approach, when an agent enters the catchment area they are assigned a component based on the assigned probabilities and will use this component regardless of the conditions faced.

The **second approach** relies on each agent making a decision as to which device to use based on their perception of local conditions. This decision is based on the assumption that the main motivation of the agent is to reduce their overall travel time through the selection of a particular component. This approach is most appropriate for scenarios involving motivated agents, where high levels of congestion occur at the entrance to the component (where there is the potential and desire for agents to optimise the travel times by moving between components), and/or when the user is not aware of the expected use of the components. This is determined by the agent initially estimating the time it would take to travel to the component (i.e. distance / projected travel speed), the estimated congestion that will be encountered before reaching the component (current wait time of people directly between the agent and component) and the projected time to traverse the component (based on the distance to be travelled and the agent's component travel speed). The agent then selects the component within the *Transit Node Group* that provides the shortest expected travel time and moves towards the selected component. This assessment is made using the following function:

$$C_j = \sum_{i=1, i \neq j}^{n} w_i + \left\{ \frac{d_j + p_C}{v_j} \right\} + t_C \tag{27}$$

Where:
C_j is the expected time using transit node C by occupant j,
p_C is the user-defined penalty associated with component (see the User Guide, Chapter 5),
n is the number of agents located between agent j and component C,
i is the agent index associated with n,
w_i is the time that occupant i has spent waiting in their current position,
d_j is the distance from agent j to the component, and
v_j is the fast walking speed of agent j.
t_C is the expected time to traverse component C.

The user-defined penalty value (p_C) is associated with each transit node in a transit node group. It allows a user to make a given transit node more or less attractive to agents (using the second approach) relative to other transit nodes in the group. Increasing the penalty value decreases the attractiveness of the associated transit node by increasing the agents perceived travel distance to the transit node. For example, if a user defined a penalty of 5, this would add an extra 5 metres onto the calculated travel distance to the transit node (i.e. making it seem further away). The default value of this penalty is 0 for all transit nodes (i.e. all transit nodes irrespective of type have no additional bias for agent usage). A more detailed description of this behaviour is provided in the User Guide, Chapter 5.

The equation calculating the expected travel time to use a transit node (C_j) is dependent on the nature of the component given such that t_C is set to the following:

$$t_{corridor} = \frac{l_c}{v_j} \tag{28}$$

$$t_{staircase} = \frac{l'_C}{v_j} \tag{29}$$

$$t_{escalator} = \frac{l'_C}{(v_j + v_C)} \tag{30}$$

$$t_{travelator} = \frac{l_c}{(v_j + v_C)} \tag{31}$$

Where:

l_c is the length of the component (including any landings that may be present),
v_j is the speed of agent j in the desired direction,
l'_c is the diagonal length of component C, and
v_c is the speed of the component.

While in the catchment area and prior to boarding the desired transit node, the agent will reassess their choice when the elapsed time is 10% greater than expected; i.e. where the agent has spent more time reaching the transit node than predicted. Therefore

$$Experienced\ Wait\ Time\ >\ 1.1\ * w_i\ +\ \left\{\frac{d_j + p_c}{v_j}\right\} \quad \textbf{(32)}$$

At this point, the agent will again perform the calculations described above to reassess the conditions, and select a route accordingly.

The **third approach** incorporates both the fixed proportion system and the shortest time system; the user provides a threshold that determines which of the two selection methods are employed according to the levels of congestion in the catchment area. Human factors data analysis studies [124, 157] suggest that during low levels of crowding at the base of an escalator/stair, individual personal factors are of primary influence for device selection. However, during high levels of crowding the primary influence for device selection are the local crowd congestion levels. The selection system requires the user to specify a congestion *Threshold* value for the transit node catchment area. This congestion threshold value is set to between 0-100 and represents the proportion of nodes in the transit node catchment area occupied by agents who are targeting the transit node group. While the occupancy levels in the catchment area are below the specified threshold, agents entering the catchment area select the component based on the fixed proportions defined by the user (Approach 1). When the levels of congestion reach or exceed the congestion threshold, agents instead attempt to optimise their travel time (i.e. employ Approach 2). This hybrid approach allows the component selection to be sensitive to the scenario: suboptimal/probability-based component selection assumed at low levels of congestion and adaptive component selection at higher levels of congestion. The default congestion *Threshold* value is 0 for a transit node group so the second approach is used by default.

In some situations, the configuration of *Transit Nodes* connecting the same two levels may be complex. For instance, a single escalator may connect two levels, while alternatively three flights of stairs are required to move between the same two levels (see Figure 7-18).

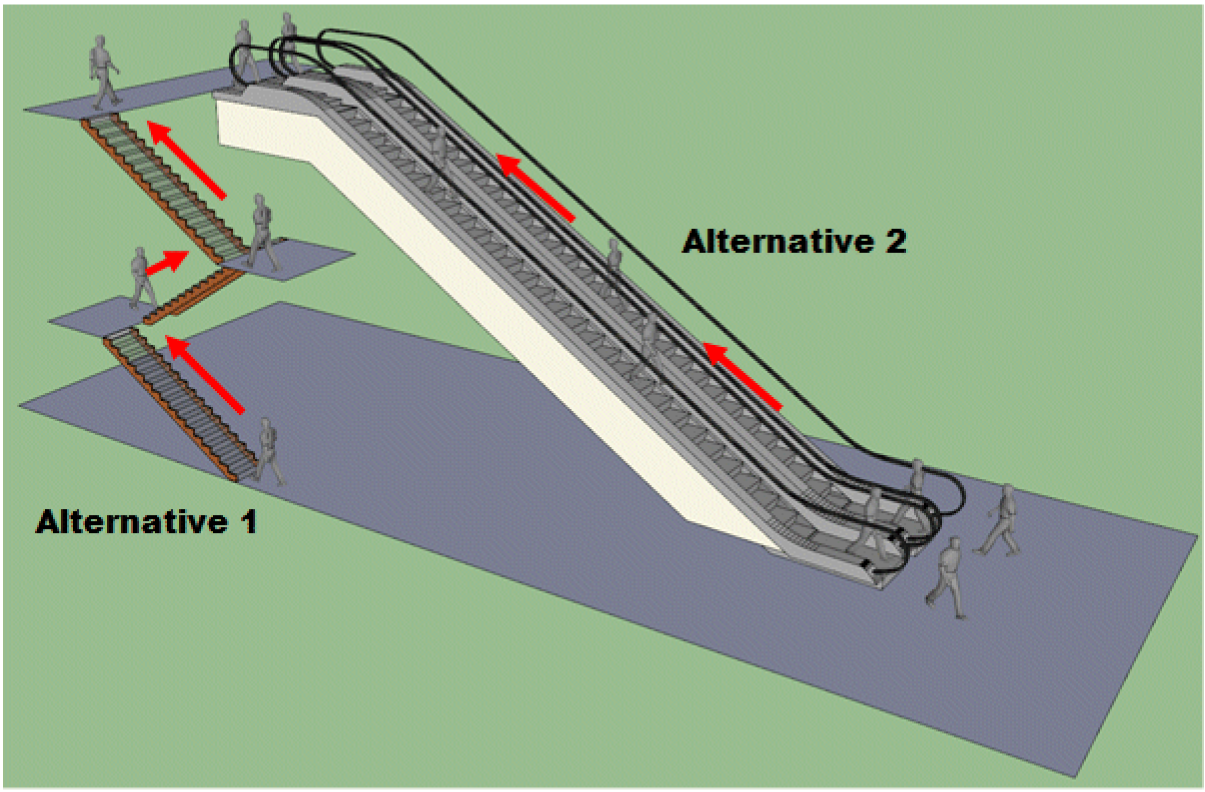

Figure 7-18: Three flights of stairs and an escalator.

The *Transit Node Group* functionality allows this situation to be addressed. The user is able to group together *Transit Nodes* that appear in series as part of the same route, such that they then form a single alternative, rather than distinct options. These will then be a sub-group, compared against other alternatives. In our example, the escalator would be one alternative (Alternative 2 in Figure 7-18), while the group of the stair flights would be another (Alternative 1 in Figure 7-18). One calculation would then be made when assessing Alternative 1 that incorporated the journey along the group of flights to reach the next level.

$$C_j = \sum_{i=1, i\neq j}^{n} w_i + \left\{\frac{d_j + p_C}{v_j}\right\} + \left\{\frac{l'_{C1}}{v_j} + \frac{l'_{C2}}{v_j} + \frac{l'_{C3}}{v_j} + \frac{l'_{C4}}{v_j}\right\} \quad \textbf{(33)}$$

Intermittent Service

At present only *Lift Shaft Opening* transit nodes (used to define a lift/elevator system) provide an intermittent service within buildingEXODUS.

There are four key decision points within the buildingEXODUS agent lift model which determine whether an agent will use a lift/elevator (see Figure 7-19):

(1) Consideration and decision to use a lift,
(2) Initial lift area assessment,
(3) Lift wait behaviour, and
(4) Lift redirection.

Each of these is explained in the following sections. The model has been used to explore the impact of lift strategies and human factors upon an evacuation in past studies [124, 168, 169].

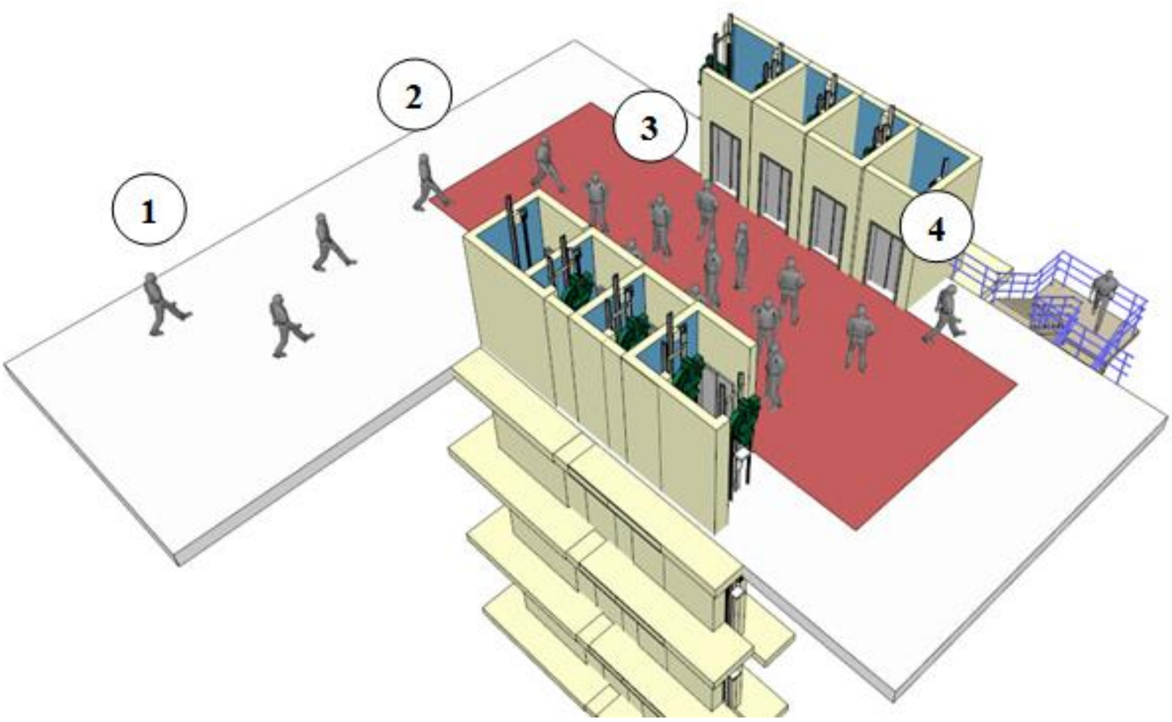

Figure 7-19: Agent Lift Model - Key Decision Points

Lift Shaft Opening transit nodes on the same floor/level can be grouped to form lift banks that have a combined catchment areas (see Figure 7-20). These define the lift waiting area where agents, upon entering, assess the local conditions and decide whether or not to wait for a lift (see decision points 2-4 mentioned above).

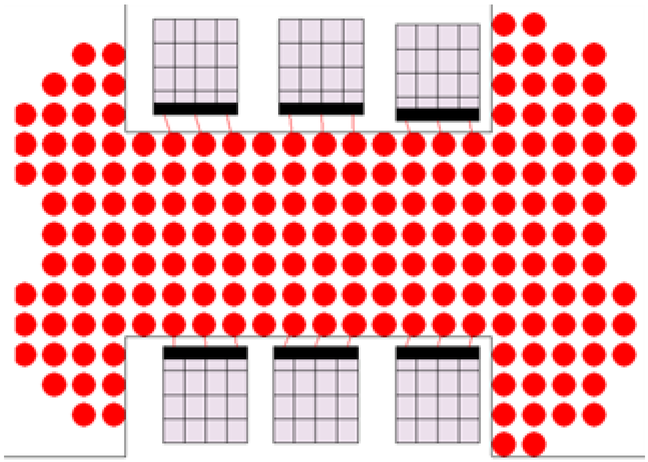

Figure 7-20: Lift waiting area defined by a transit node group catchment area (displayed in red circles)

The default values for each of the key decision points included in the model are based on data collected via an international lift/stair usage survey conducted by FSEG [124, 166, 167]. Based on the survey data, agents on progressively higher floors are:

(a) more likely to use a lift,
(b) likely to be more tolerant to higher levels of congestion, and
(c) are likely to be prepared to wait longer for a lift than those agents on lower floors.

An overview of the data and how it is employed within the model is presented in the following sections.

Consideration to use a lift

Data collected in the online survey [124, 166, 167] suggests that levels of familiarity with a building (e.g. layout, routes, exits etc.) influence peoples' consideration to use a lift during an evacuation. The results suggest that the more familiar someone is with a building the less likely they would be in using a lift during an evacuation. More specifically, if people were familiar with a building then only approximately 1 in 3 would consider using a lift during an evacuation. If people were unfamiliar with a building then approximately 1 in 2 people would consider using a lift during an evacuation. These results also highlight peoples' reservations about using a lift during an evacuation, with at least half of all people, irrespective of levels of familiarity, not prepared to use a lift during such situations. It is suggested this is due to most buildings not allowing lifts to be used during an evacuation.

In consideration that a building would employ suitable training which would sufficiently assure people that lifts were safe to use during an evacuation, the default parameters in the lift model assume that all agents would at least consider using a lift. The percentage of agents that would consider using a lift on each floor can be configured to the values obtained in the survey (see the User Guide, Section 5.17.4) or to user specified values. It should be noted that this relates to whether agents will **consider** using a lift and not whether they will **actually** use a lift; this is described in the following sections.

Decision to use a lift

Agents that would consider using a lift during a simulation, will then decide whether to move towards a lift or stair based on a proportional system. This represents a subset of the percentage of agents that would consider using a lift (those agents that would not consider using a lift will only choose to use stairs). By default the percentage of agents that choose to use a lift is defined according to a regression formula based on the floor the agents are initially located (see Figure 7-21, [124, 166, 167]). It is possible to alter the formula or manually prescribe a different percentage of agents that will use a lift on each floor during a simulation.

Agents that choose to use a lift will move towards the lift bank containing the chosen *Lift Shaft Opening* transit node (agents select *Lift Shaft Opening* transit nodes to move towards as these represent openings where the lift can be accessed). By default agents who choose to use a lift will randomly choose to move towards a lift bank (containing the *Lift Shaft Opening* transit node). This means approximately an even number of agents from a given floor should elect to use each of the available lift banks on the service floor (i.e. no single lift bank should be disproportionally used by agents).

Agents will only move towards lifts assigned a *Shuttle Floor* procedure if the agent's initial floor is part of the lift's sequence. Similarly, by default agents will only move towards lifts assigned a *Sky Lobby* procedure if the agent's initial floor is part of the lift's sequence, or the floor in which the sky lobby is located falls between the agents initial floor and their exit (i.e. if an agents exit is located below them they will only consider moving down to a sky lobby, not up). Agents can move up to lifts assigned *Sky Lobby* procedures, but only if the number of floors the agent needs to travel up in order to reach the sky lobby floor is:

(a) less than or equal to a user defined number of floors (as defined by the *Sky Lobby Up Floors* variable, see Chapter 6 of the User Guide), and

(b) less than the number of floors that they would traverse in travelling down to either another sky lobby lift or an exit (i.e. agents will not travel up to a sky lobby lift if there is an alternative down option that is closer (or as close) in terms of the number of floors they are required to traverse).

By default the number of floors agents are prepared to travel up in order reach the lift (i.e. the *Sky Lobby Up Floors* variable, see Chapter 6 of the User Guide) is set to zero, thereby implying that agents are not prepared to move up in order to reach a sky lobby lift.

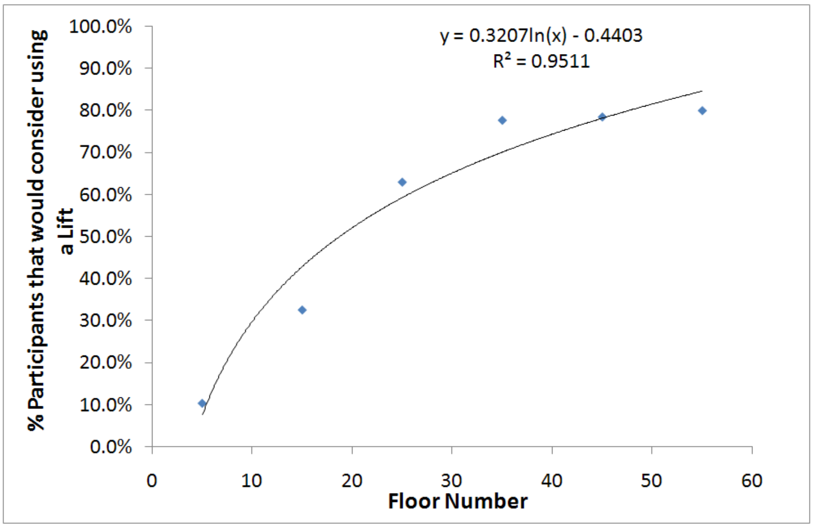

Figure 7-21: Percentage of Participants that would consider using a lift for each floor if they were familiar with a building

Once a simulation has started agents do not change their device selection or redirect to another lift/stair after the initial choice is made. Once the agent's response time has expired they move towards the chosen *Lift Shaft Opening* transit node. Agents that do not elect to use a lift, either evenly distribute between the available stairs or use the stairs which forms part of the route to their closest exit (the selection of which method is dependent on the defined settings within the *Behavioural Control* dialogue box, see the User Guide, Section 6.6.4.5).

NOTE:
In addition to the automated approach, users can explicitly specify which lift bank an agent will use via assigning them a Lift Bank task (see the User Guide, Chapter 4). When assigning Lift Bank tasks to agents users are required to manually define both the Congestion Threshold (occ/m^2) and Wait Time (s). As a result, these values are not automatically assigned in the same manner that Lift Bank tasks assigned via the automated approach are (i.e. according to a probability distribution based on the floor the lift they are targeting is located).

Initial lift area assessment
When an agent enters into a lift waiting area, they will decide if the levels of congestion are too high in the area. This is achieved by agents being assigned a *Congestion Threshold* value (measured in occ/m^2) based on a probability distribution according to the floor on which the lift they are targeting is located [124, 166].

Figure 7-22: Online lift/stair survey data - Percentage of participants on different floors that would redirect from a lift waiting area due to a given level of congestion

Floor Range Location	Percentage of participants that would redirect from a lift waiting area due to a given level of congestion %							
	Congestion doesn't matter	0.13 occ/m^2	0.5 occ/m^2	1.0 occ/m^2	1.5 occ/m^2	2.0 occ/m^2	2.5 occ/m^2	2.5 occ/m^2+
2-10	14.3	21.4	42.9	78.6	85.7	85.7	85.7	85.7
11-20	11.4	15.9	31.8	59.1	86.4	86.4	88.6	88.6
21-30	7.1	4.8	25.0	63.1	83.3	90.5	92.9	92.9
31-40	10.1	3.0	19.2	47.5	77.8	86.9	89.9	89.9
41-50	9.1	3.0	14.1	37.4	65.7	82.8	90.9	90.9
51-60	11.0	3.0	12.0	32.0	57.0	72.0	84.0	89.0

If the agent's assigned *Congestion Threshold* is exceeded then the agent will redirect to the stairs. If the *Congestion Threshold* is not reached or exceeded then they will choose a location to wait in the lift waiting area near one of the lifts (i.e. within the transit node group catchment area).

If a lift is open in a lift waiting area, an agent who has chosen to redirect to the stairs due to the congestion levels can still elect to use a lift whilst in the lift waiting area (defined as the transit node group catchment area, see Figure 7-20). This allows agents who are required to travel through a lift waiting area whilst redirecting to the stairs to still be able to use a lift if they have the opportunity. This is based on the assumption that such agents will not walk past a lift that they have the option to board.

Agents who initially choose to use a lift but arrive in the lift waiting area after all lifts in the bank have serviced the agent's floor will redirect to use the stairs. This represents the influence of either dynamic signage or a communication system informing the agents that the lift has already serviced their floor and that they should use the stairs instead. This system prevents agents from waiting indefinitely for a lift that will not arrive.

Lift wait behaviour
Upon entering the lift waiting area, providing the agent does not redirect due to congestion, the agent is assigned a *Lift Wait Time*. This defines the maximum time the agent is prepared to wait to use a lift. By default, the *Lift Wait Times* are assigned based on a probability distribution according to the floor on which the lift they are targeting is located. These probability distributions are based on the international lift/stair usage survey [124, 166] and can be seen in see Figure 7-23. These distributions range from 0-60 minutes. A given percentage of agents will also be assigned to wait for as long as is necessary to use a lift (i.e. as long as a lift in the bank is servicing the agent's floor they will wait).

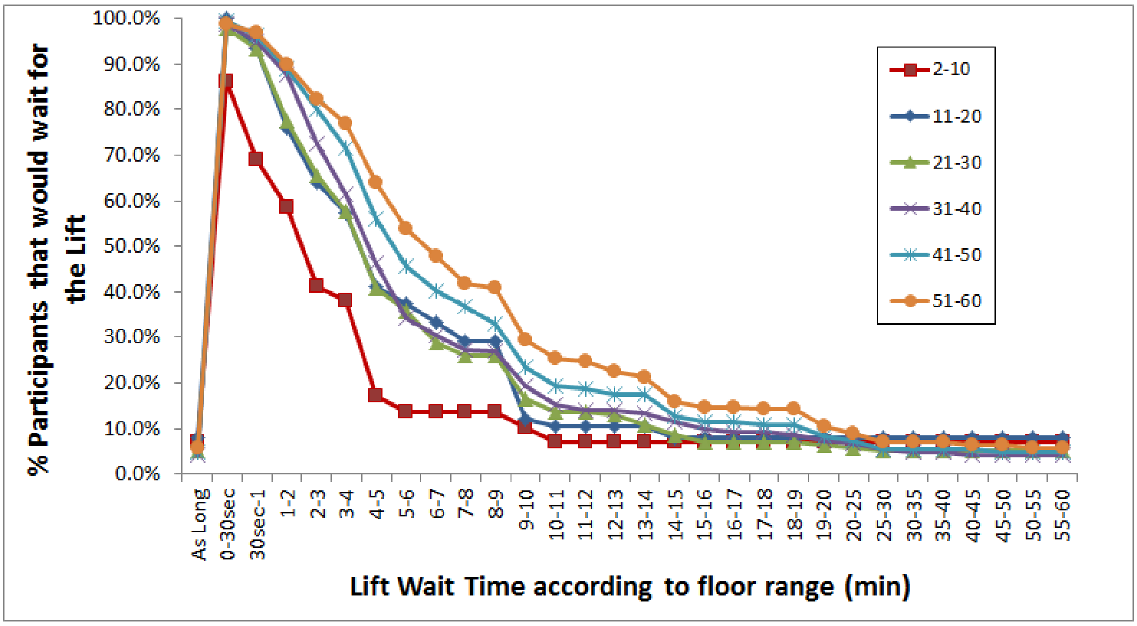

Figure 7-23: Cumulative frequency distribution of maximum lift wait times for each floor range

The *Lift Wait Time* starts when the agent enters into the lift wait area. If the *Lift Wait Time* expires before a lift in the bank has serviced the floor (i.e. the agent has waited that length of time for a lift) the agent will redirect to the stairs. If their *Lift Wait Time* has not expired and a lift in the bank starts servicing the floor then their *Lift Wait Time* will be increased by half of their original wait time. This is intended to represent an assumption that agents who see a lift servicing their floor are prepared to wait longer in anticipation that they will be able to board a lift soon. The *Lift Wait Time* is only increased once and does not affect agents who would 'wait for as long as it takes'.

Agents that decide to wait for a lift are randomly assigned to wait outside one of the *Lift Shaft Opening* transit nodes within the lift bank targeted by the agent. Randomly assigning agents in this manner ensures that waiting agents are evenly spread throughout the lift waiting area, and hence that a disproportionate number of agents will not wait outside any one lift. Once a *Lift Shaft Opening* transit node has been randomly assigned to an agent they will then move towards it. Once an agent is both within 3m of their target *Lift Shaft Opening* transit node and within its corresponding catchment area (as defined by its *Catchment* attribute) the agent will be deemed to have reached their assigned wait location and will commence milling. The milling process involves the agent occasionally randomly moving while waiting for a lift. Agents waiting for lifts are free to ‘mill’ anywhere within their lift banks corresponding waiting area (i.e. the area representing the combined catchment areas of each *Lift Shaft Opening* transit node within the lift bank). Hence while milling, agents will not move outside of the lift waiting area.

When a lift door opens in a lift bank, the nearest agents who are waiting in the lift waiting area move to use the lift. Only the number of agents that can fit inside the lift (derived from the lift *Capacity*) attempt to board the lift. As a result, there is no competition for lift boarding and the boarding process is orderly. If multiple lifts open their doors simultaneously at a given floor, the agents select the nearest open lift that is not oversubscribed. Agents who could not board their nearest lift will assess if they can board their second nearest lift using the same process. This process is repeated until all spaces in all open lifts have been allocated to waiting agents. The agent’s *Lift Wait Time* ends when they decide to either board an open lift or redirect to the stairs: not when they physically leave the lift waiting area (i.e. board the lift).

Having boarded a lift agents will then travel in it to its designated drop off floor, where upon they will alight. Upon exiting the lift agents will then search for another available lift to use as a means of egress. It is important to note that having exited a lift agents will not make a decision on whether or not to use another lift in the same manner as their initial decision to use a lift (see the *Decision to Use a Lift* section). Hence, the decision on whether to use another lift is (unlike the agent's decision to initially use a lift) **not** determined according to the percentage of agents choosing to use a lift on the floor on which the agent is currently located. Therefore, if an agent has initially decided to use a lift it is assumed that they will continue to seek to use lifts throughout their egress. Agents who identify an available lift will then once again move towards the designated lift/lift bank. As a result, during egress agents can search for and target lifts several times in succession and therefore evacuate from structures using multiple lifts.

Figure 7-24 shows the route adopted by an agent evacuating from a structure via three lifts (i.e. *Lift_A*, *Lift_B* and *Lift_C*). Here the points at which the agent actively sort available lifts and subsequently targeted them are shown via the blue dots. Hence the agent initially chooses to use a lift at the start of the simulation based upon their initial floor and the percentage of agents choosing to use a lift on that floor (see the *Decision to Use a Lift* section). Having decided to use a lift the agent then identifies *Lift_A* as an available lift and then sets about using it. Having alighted from *Lift_A* the agent then again sets about identifying an available lift. Once *Lift_B* has been identified as an available lift the agent then in turn sets about using it. Alighting from *Lift_B* the agent then identifies *Lift_C* as an available lift, before in turn using it to egress. Having alighted from *Lift_C* the agent is then unable to identify any available lifts, and hence continues their evacuation by walking to the nearest available exit point. It is important to note that within each lift waiting area the agent could have chosen to redirect from the lift (either as a result of excessive congestion or wait times or as a result of the lift no longer servicing the floor).

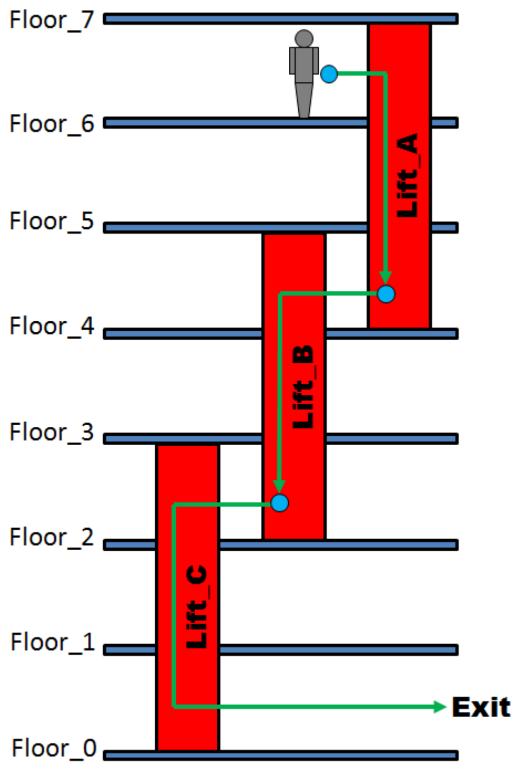

Figure 7-24: An agent evacuating from structures using multiple lifts (points at which the agent sort out and targeted lifts shown as blue dots)

Lift redirection
Agents who redirect to the stairs are either randomly assigned to use given stairs or follow the potential map to their nearest exit. The selection system used is determined according to the behavioural options specified in *Simulation Mode* (see the User Guide, Section 6.6.4.5). The random assignment redirection system reduces the likelihood of any stair being oversubscribed and causing a bottleneck during a scenario. If this is not used then it may result in a disproportionate number of agents electing to use a given stair, as agents will typically all choose to use the stair which leads them to the nearest exit.

If an agent redirects to the stairs from a lift bank but is still inside the transit node group *Catchment* area, they are still able to board a lift should one become available (i.e. an agent will not walk past an open lift if they are able to board it).

It is important to note that once an agent has redirected from a lift (either as a result of excessive congestion or wait times or as a result of the lift no longer servicing the floor) they will not consider using a lift again during their egress. Consequently the agent will be required to evacuate via the stairs. Figure 7-25 shows the routes available to an agent initially located on *Floor_2*. Here the valid routes available to the agent are shown in green. The agent can initially either catch *Lift_A* down to *Floor_1* (and then catch *Lift B* down to the ground floor) or redirect away from the lift. If the agent redirects away from *Lift_A* they will be required to evacuate the structure solely via the stairs. In this instance this will mean the agent descending the stairs first to *Floor_1* and then descending the stairs further to *Floor_0*, before then moving to the exit. It is important to note that having redirected away from *Lift_A* the agent will NOT consider using *Lift_B* (i.e. the red route, see Figure 7-25), even though it may provide a faster means of egress.

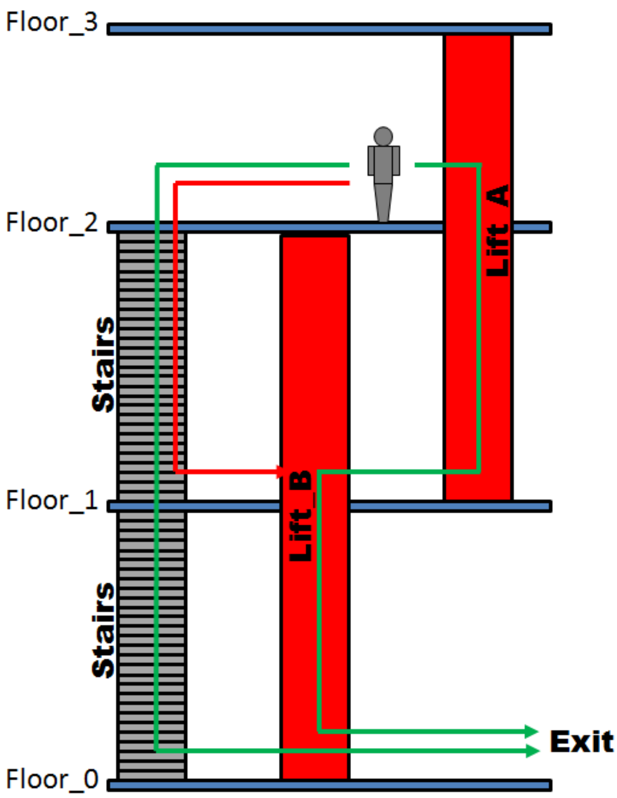

Figure 7-25: Evacuation routes from a structure via lifts or as a result of redirection from them (valid routes shown in green, invalid routes shown in red)

7.2.3.8 (viii) Staircases

Staircases in buildingEXODUS are used to connect floors. The generation of staircases is a semi-automatic process, the user must define where a staircase is required and provide some basic geometry information. Staircases are created using *Transit Nodes*.

7.2.3.8.1 (a) Staircase geometry

Two approaches are available to define the staircase geometry, namely *Standard* and *Effective Width* methods. The user is free to choose the preferred approach. The first simply utilises the maximum amount of space (width) available on the staircase, regardless of the presence of handrails or other behavioural considerations. The second utilises the effective width model [19,31] to take account of handrails and the fact that people try to maintain a small distance from the edge of the staircase when moving on a staircase. Figure 7-26 summarises the two approaches. The staircase geometry is specified in *GEOMETRY MODE* (see the User Guide, Section 3.7).

In either case, the default space across a step occupied by each occupant is 0.76m. Thus the number of lanes across the width of the staircase is the width available divided by 0.76m. In the case of the *Standard* approach the useable width of the staircase is the entire available width of the staircase. In the *Effective Width* approach the useable width of the staircase takes into account handrail spacing and boundary layer width (also known as effective width [19, 31]). If no handrails are present, only the boundary layer width is still taken into account.

NOTE:
The default lane width on staircases is 0.76m. This increase in node width is intended to represent body sway on the stair.

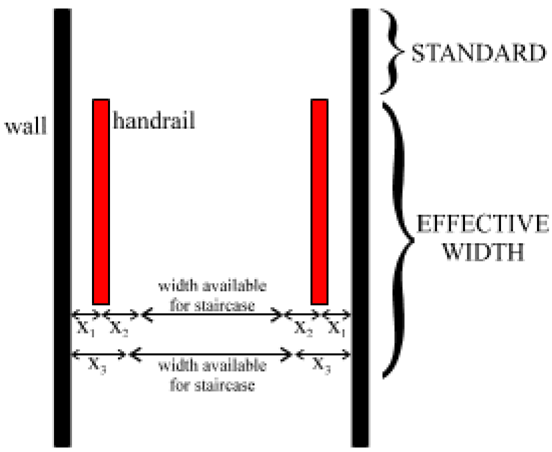

Figure 7-26: Stair case models available in buildingEXODUS

When defining a staircase, the user must specify which approach is required. If the effective width model is chosen then the presence of handrails must be registered (see the User Guide, Section 3.7.4). In total the user has three options:

a) Enable the presence of handrail (X_1 on Figure 7-26)
b) Enable the presence of handrail and the effective width (X_1+X_2 on Figure 7-26)
c) Enable only the effective width with no handrail (X_3 on Figure 7-26)

In each case, the values X_1, X_1+X_2 or X_3 are subtracted from the overall staircase width to produce the width available. When no handrails are present, only X_3 is subtracted. Table 7-8 provides default values for each of these parameters.

Table 7-8: Parameter values for Figure 7-26

Parameter	Value (m)	Explanation
X_1	0.1 (each side)	distance from wall to centre of handrail
X_2	0.0875 (each side)	effective width from centre of handrail
X_3	0.15 (each side)	effective width with no handrail

7.2.3.8.2 (b)Stair Potentials / Distances

Transit Node staircases have an internal potential/distance map which is independent from the global map used by other buildingEXODUS nodes. Each riser, within a *Transit Node*, is treated by the occupants as being *equipotential* (or *equidistant* from their current exit point). In other words, while on a riser, the occupant expresses no desire to occupy a particular position simply on the basis of that location's potential or distance map (see Figure 7-27). The localised map is then created according to the location of the exits in relation to the staircase. Therefore moving towards an exit - either up or down the stair - is considered desirable.

This localised map is affected by all of the other behavioural rules designed to apply on staircases.

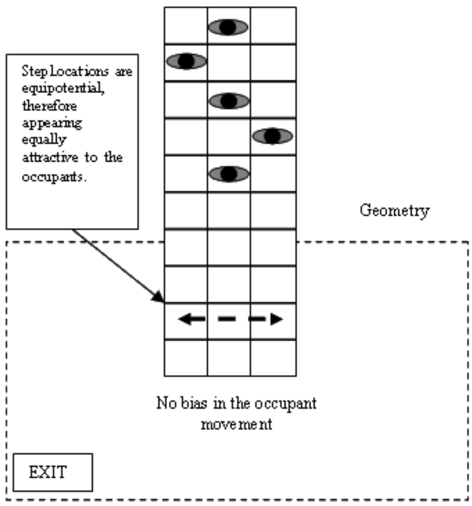

Figure 7-27: Constituent nodes within a step appear to be equipotential, generating no nodal bias on the basis of potential.

7.2.3.8.3 (c) Staircase behaviour

The behaviour of occupants on staircases is different to that of their behaviour in open space. This section describes the behaviour rules that apply to staircases.

(i) General Behaviour

A stair in buildingEXODUS represents a single flight of stairs. The maximum number of occupants on any one step is dependent upon the number of unit lanes modelled per step. The maximum number of occupants on a flight of stairs is equal to the number of lanes times the number of steps. However, the achieved occupancy depends on the behavioural regime selected during the simulations. There are two behavioural regimes available to describe staircase behaviour which will affect the usage level of a flight of stairs: *Packed* and *Staggered* (default) behaviour, see Section c part ii for further information. Note that several general rules apply to both regimes:

No sideways movement is allowed on a step.
Occupants will attempt to use handrails if they are present and thus occupy the space next to the handrails if possible (see (iv)).
Occupants attempt to maintain one step spacing between their current position and the person in front.

NOTE:
If the staircase is in excess of 5m in width, then the occupants will not attempt to move towards the handrails.

These are the fundamental behavioural rules that are applied to occupant movement on staircases irrespective of the navigational system used (i.e. *Potential map* or *Distance map* systems) or the behavioural regime applied to stair travel. *Transit Nodes* allow performance to be established on an individual level, or through capping the attainable flow levels. The behaviours described below relate to the application of the *Individual Model*; i.e. when performance is assessed at the individual level (see the User Guide, Chapter 5).

(ii) Staggered / Packed Behaviour
There are two basic behavioural regimes that may be applied to staircases: the *Packed* and *Staggered* (default) behavioural regimes. These control the level of crowding that an occupant will tolerate whilst moving along a staircase through controlling the space each occupant maintains around themselves whilst moving. Figure 7-28 indicates the main difference between the *Packed* and *Staggered* staircase behaviour regimes. It should be noted that staircases with centre handrails are essentially treated as separate staircases in this respect.

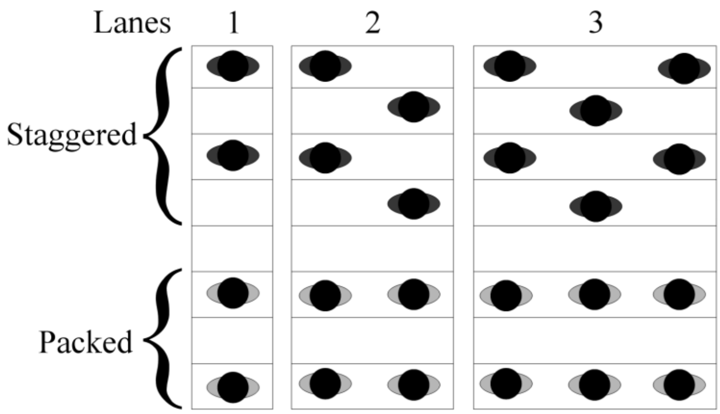

Figure 7-28: Packed and Staggered occupant behaviour on staircases

In the *Packed* regime, occupants are allowed to use any location available to them across a riser, but attempt to keep one riser spacing from the occupant in front. This could allow staircases to become fully loaded with no free space available. In the *Staggered* regime, occupants attempt to maintain one riser spacing and one occupant per riser for narrow staircases (one or two lanes wide) or keep one lane spacing on risers in wide staircases (two lanes or more). From Figure 7-28 it is apparent that the level of congestion that is maintained during the *Packed* regime can be expected to be greater than during the *Staggered* regime. These regimes are then based around several simple rules that when combined can generate significantly different behaviour (see Table 7-9).

Table 7-9: Packed and Staggered behavioural rules used on staircases

Riser Spacing Requirements	PACKED	STAGGERED
Riser Spacing in the direction of movement	A location will only be adopted if there is an empty space ahead of it (i.e. a riser ahead of the occupant is still maintained once the desired location is occupied).	A location will only be adopted if there is an empty space ahead of it (i.e. a riser space ahead of the occupant is still maintained once the desired space is occupied).
Riser Spacing in the direction perpendicular to direction of movement	No additional rules are in place.	A location will only be adopted if there is an empty lane space either side of it (i.e. a lane space is maintained on either side of the selected location, across the riser).

The *Staggered / Packed* regimes are determined in *Simulation* mode. The default setting is *Staggered*, which was originally derived from the work of Fruin and Pauls [19,34]. However, the user has the ability to set the spacing required for the *Staggered / Packed* regimes on each *Transit Node*. This allows them to represent staircases that are known to be used in a different manner from the rest in the geometry. The user can therefore provide local values (via the *Stair* dialogue box in *Scenario Mode*) that influence how that particular staircase is used. This is consistent with the approach adopted in other *Transit Nodes*, such as *Escalators*.

(iii) Overtaking Behaviour

In the following discussion, the terms *overtaking* and *passing* are used. It is important that the differences between these two behaviours is understood. An agent is said to *pass* another agent if the agent passes a slower moving agent located in *another lane*. The *passing* manoeuvre does not involve the agent making a lane change. *Overtaking* differs from *passing* in that the *overtaking* manoeuvre requires the faster moving agent to make a lane change in order to pass the slower moving agent (i.e. the slower moving agent is initially in the same lane as the overtaking agent).

An agent may *pass* a slower moving agent under either behavioural regime so long as they are not in adjacent lanes. If agents are in adjacent lanes and the *Staggered* behavioural regime is in force, a faster agent can pass another only if the passing agent becomes impatient or/and the difference in their travel speeds is greater than 50% of the slower moving agent.

Agents will consider an *overtaking* manoeuvre either if the faster moving agent becomes impatient or if the difference in travel speeds is greater than 50% of the slower moving agent. The decision to make a lane change occurs when there is a one-riser gap between the two agents. A lane change can only occur if the way ahead in the adjacent lane is clear. If this is not the case, the agent will not attempt to *overtake*. In making the lane change it is possible in some circumstances that the faster moving agent will “cut up” a slower moving agent who is travelling behind him in the adjacent lane. This may result in the single riser spacing rule being violated. Once in the new lane, the *overtaking* agent follows the *passing* rules and all other appropriate stair rules.

(iv) Stair Edge Preference.

Individuals will attempt to occupy space next to handrails. This can be very important when trying to reproduce results that are in accordance with various prescribed rates (see the Application Manual, Chapter 2).

NOTE:
This will only occur if the individual is within 2.5m of the handrail. This is to prevent individuals travelling large diagonal distances to reach the handrails.

7.2.3.8.4 (d) Speed on staircases

Agent travel speeds while on stairs are based on data from Fruin [19]. The travel speeds on staircases vary according to the *Age* and *Gender* of the occupant. A full list of stair travel speeds can be found in Table 3-4, CHAPTER 3:.

NOTE:
The Fruin specified stair travel speed [19] is an indication of the MAXIMUM likely travel speeds. It should be remembered that the travel speed of an occupant across staircases can also be affected by the congestion. Therefore delays may be incurred due to congestion that may force the occupants to be static for small periods of time.

NOTE:
The travel speeds attained on staircases may be affected by the orientation of the structure.

Transit Nodes also allow performance to be established through capping the attainable flow levels. The SFPE formulation is used to estimate the achievable flow rates that can be attained on stairs [151]. This formulation can be seen in the equation below.

$$F_s = (1 - 0.266 * D) * kD \quad \textbf{(34)}$$

Where:

Fs = Specific flow (occ/m/sec)
D = Population density (occ/m^2)
k = Constants for evacuation speed (m/s)

The achievable flow rate calculated using the SFPE method is sensitive to the density levels assumed. The achievable flow rate is constrained to prevent unrealistic levels being generated [151]. The value of the k constant is varied to reflect the riser and tread size being represented. This produces a series of parabolas, each reflecting the changing flow rates achievable for specific step designs. These parabolas have maximum values that range between 0.94-1.16 occ/m/sec [151]. It is these maximum values that are used within buildingEXODUS to cap the flow rates that can be attained. These rates are applied across the entire width of the component, irrespective of the number of lanes. This is in line with the assumption that the flow through a stair is proportional to the width of stair, rather than purely on a lane-based calculation.

7.2.3.9 (ix) Escalators

The data underlying the occupant use of escalators was derived from two sources: an EU-sponsored project called AVATARS [124,144,158,160,161,162,164], which involved collecting escalator usage data in the Barcelona metro (Spain); and observations made in the London (UK) Underground [124]. Both data-sets were produced from observations made in rail/underground/metro terminals, given the frequent use of escalators. Data was collected using video cameras, and the conditions/dimensions/configuration of the escalators and associated staircases were recorded to ensure that they were representative of such configurations elsewhere. The data collected informed four elements of escalator performance: flow rates, travel speeds when walking, likelihood of walking, and standing position when not walking. These were derived from the results produced in one or both of the source data-sets.

The default escalator flow rates, presented in Table 7-10 were derived from the UK data-set and reflect an upper bound (i.e. collected during peak conditions) of a 1.0 m wide escalator having a vertical drop of 3.65m and horizontal speed of 0.5 m/s. These are employed when the *Flow Model* is enabled (see the User Guide, Chapter 5).

Table 7-10: Default Escalator flow data. Data collected by FSEG staff at a London Underground station. [124]

Direction	Max occ/m/s [occ/m/min]
Down	0.783 [47.0]
Up	1.25 [75.0]

Alternatively, the *Individual Model* can be enabled, allowing performance to be established as individuals arrive and traverse the component. The default walker (horizontal) travel speeds were derived from the UK data-set (collected by FSEG staff at a London Underground station) [124] and are assigned to the occupants on an individual basis and then employed when they reach the component. These are categorised according to direction of travel (i.e. either up or down) and gender. The number of data points is also provided (see Table 7-11).

Table 7-11: Default Escalator horizontal travel speeds - Individual Model.

Direction	Gender	Speed (m/s)	Frequency
Down	**Female**	0.75 [0.40-1.44]	112
	Male	0.84 [0.38-1.67]	247
	Average/Total	0.82 [0.38-1.67]	359
Up	**Female**	0.67 [0.41-1.02]	126
	Male	0.72 [0.32-1.51]	325
	Average/Total	0.71 [0.32-1.51]	451

The UK data-set was based on occupants traversing a moving escalator. However, no data was collected concerning travel speeds on static escalators. As such, if an escalator's speed is set to 0 m/s (i.e. the escalator is static), then agents will use their stair walker speed (derived from the Fruin data) to traverse the escalator [124]. The user should refer back to Section 3.1 for more information on occupant travel speeds.

If the *Flow Model* is enabled then the default range of travel speeds over the escalator in each direction are assumed to be the same as the corresponding average individual values shown in Table 7-11. Consequently, the default range in the *Down* direction is assumed to be 0.38-1.67m/s, while the default range in the *Up* direction is assumed to be 0.32-1.51m/s.

The percentage for occupants opting to walk while on the escalator was derived from the Spanish data-set [124,144]. This percentage is categorised merely according to the direction of movement (see Table 7-12).

Table 7-12: Default Escalator probability of walking.

Direction	% of Occupants	Frequency
Up	27.0	159
Down	21.0	77

Given that an occupant has decided not to walk, the standing location was derived from the Spanish data-set [124,144]. This data is categorised according to the side of the escalator on which the occupants stand and the direction of movement (see Table 7-13).

Table 7-13: Default Escalator probability of standing.

Direction	% of Occupants (Left)	Frequency	% of Occupants (Right)	Frequency
Up	10.9	47	89.1	384
Down	26.9	72	73.1	196

Finally, buildingEXODUS also represents the delay incurred as an individual steps onto the escalator. By default, the data shown in Table 7-14 is used. This was derived from the work of Fruin [19]. These values are insensitive to direction, but sensitive to gender, in accordance with the original data.

Table 7-14: Default Escalator Entry Delay.

Gender	Min. Delay (sec)	Max. Delay (sec)
Male	0.95	1.16
Female	1.06	1.18

7.2.3.10 (x) Travelators

The data underlying the occupant use of travelators was derived from the CIBSE guide [125] and from the work of Fruin [19].

The default travelator flow data was derived from the CIBSE guide [125]. This data is employed when the *Flow Model* is enabled. The original data was generated with a population density of 2.0 occ/m^2 and a travelator speed of 0.5m/s and equates to 1.0 occ/m/s (60 occ/m/min). Little relevant data is available describing travel speeds on travelators. Therefore, the buildingEXODUS default *Fast Walk* speeds (1.2-1.5m/s) are used by default.

Alternatively, the *Individual Model* can be employed. Little data is available to support the likelihood of occupants standing on a particular side of the escalator or choosing to walk or stand. Therefore, no bias is assumed in selected a side of the elevator to use, and an arbitrary value of 20% is assumed to represented the percentage of walkers present. The same entry delays are assumed as for the use of escalators (see Table 7-14).

7.2.3.11 (xi) Metered Gates

Transit nodes can also be used to represent pedestrian metered gates that allow access from one part of a structure to another. Such gates are found in a wide variety of settings including mass transit stations, office buildings, libraries, and generally, in any structure where limited or restrictive access is required between separate areas. During the normal operation of the structure metered gates are usually operated using either a ticket or a card and are set to a *Closed* state. However, during an emergency, these can be set to an *Open* or free flow state to ease the flow and facilitate the evacuation of the structure.

Within EXODUS two main types of metered gate can be represented, namely turnstiles and ticket gates (see Figure 7-29). In addition, two movement models are available that govern the behaviour of the agents when using each type of metered gates (i.e. ticket gates or turnstiles). The first is a *Flow* based model and the second is an *Individual* based model. Each of these types of model will now be discussed:

Turnstile

Ticket Gate

Figure 7-29: Common types of Metered Gate

Within the *Flow* based model the flow through the device is defined (per person) according to a random uniform value between a range of defined minimum and maximum flows (occupants/sec). When using the *Flow* based model the flow achieved through the device is assumed to be uniform and is imposed on the entire length of the device. A typical metered gate configuration is shown in Figure 7-30. In heavy use conditions, namely when there is a steady and continuous flow of individuals through the gate, the flow imposed through the device is a capping flow (i.e. the achieved flow can be lower than the maximum imposed flow but it can never exceed it). Hence when specifying minimum and maximum flow values, the actual achieved flow under heavy constant usage will typically be a value midway between the minimum and maximum imposed values.

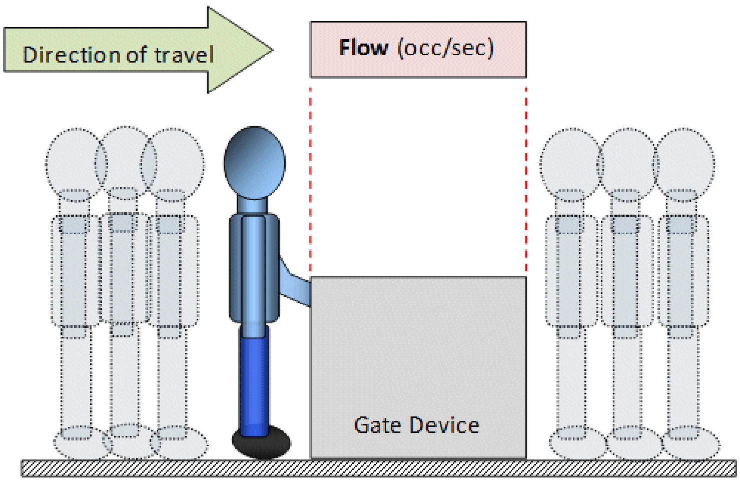

Figure 7-30 Representing a metered gate via the Flow based model

Within the *Individual* based model the time each agent takes to traverse a metered gate is comprised of both their corresponding *Entry Delay* and a *Travel Time* (i.e. T_1 and T_2, see Figure 7-31(a)). Using this approach the *Entry Delay* is assumed to be the time (in seconds) between the pedestrian inserting and retrieving their ticket from the machine (i.e. T2 – T1 in Figure 7-31(b)). Similarly, the *Travel Time* is assumed to be the time (in seconds) between the pedestrian collecting their ticket and exiting the device by placing their first foot past the metered gate (i.e. T3 – T2 in Figure 7-31(b)). Within EXODUS, both the *Entry Delay* and *Travel Time* for a given metered gate are defined by either uniform or user defined probability distributions. Each agent traversing the metered gate is then randomly assigned both an *Entry Delay* and a *Travel Time* according to their corresponding distributions. In this manner the total time taken for the agent to traverse the metered gate can be defined. If no *Entry Delay* data is available for a given metered gate then the time taken to traverse the entire metered gate can be defined purely by a *Travel Time* distribution (i.e. the *Entry Delay* and *Travel Times* are effectively *Combined*).

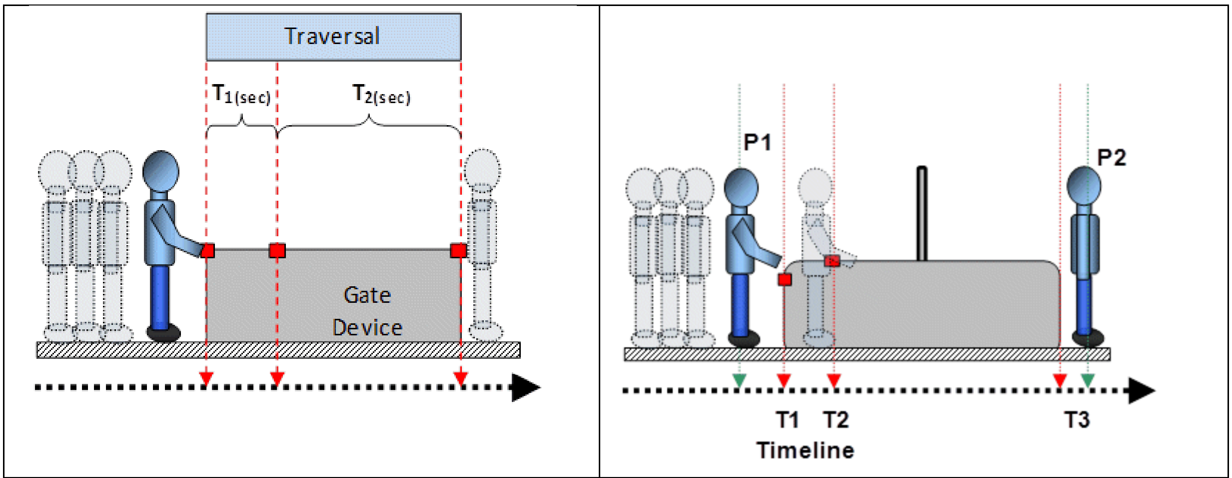

(a) (b)

Figure 7-31 Representing a metered gate via the Individual based model

Within EXODUS each type of metered gate (irrespective of whether it is defined via the *Flow* or *Individual* based model) is assumed by default to have a capacity of one agent (i.e. only one agent can operate the gate at any one point in time and the gate can only be operated again once the agent currently using the gate has passed through the gate).

Figure 7-32 below shows default data sets available for each type of metered gate available for both the *Flow* and *Individual* based models. For each metered gate and movement model combination the name of the data set is stated, indicating whether it was derived from an *Open* or *Closed* metered gate, as well as the location of the device (i.e. station or library). In addition the source of the data is also provided.

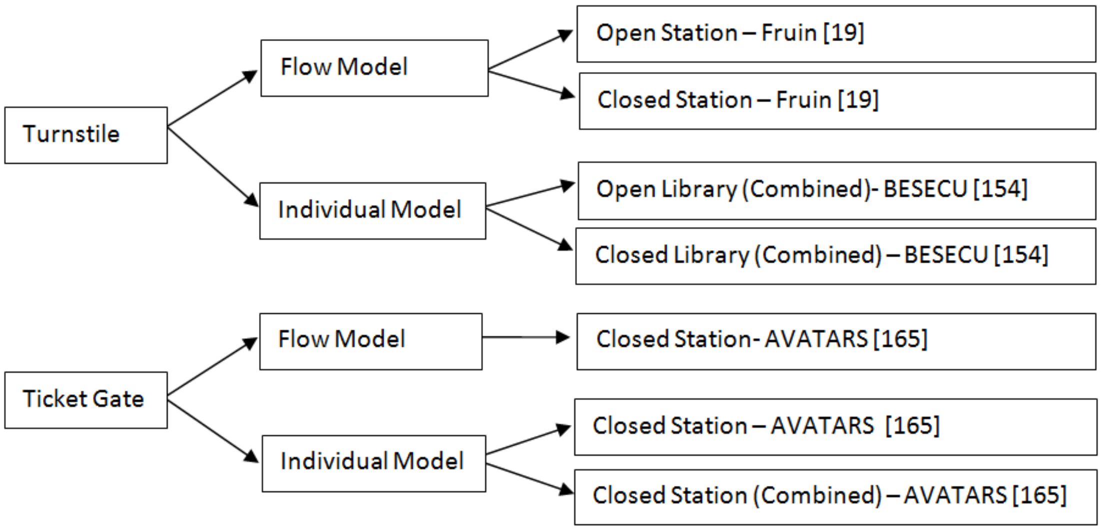

Figure 7-32 The metered gate types and the default data available for each movement model type (i.e. Flow and Individual).

In total, the user can select from four different default data sets defining the performance of turnstiles. Two data sets (i.e. *Open Station* and *Closed Station*) correspond to the performance of turnstiles via the *Flow* model as derived by Fruin [19]. A further two data sets (i.e. *Open Library* and *Closed Library*) correspond to the performance of turnstiles by the *Individual* model via data derived from a turnstile located in a Turkish Library collected as part of the BESECU project [154].

Similarly, the user can also select from a total of three different default data sets defining the performance of ticket gates. One data set defines the performance of ticket gates via the *Flow* model (i.e. *Closed Station*), while two data sets define the performance of ticket gates via the *Individual* model (i.e. *Closed Station* and *Closed Station (Combined)*). All data relating to ticket gates (i.e. for both the *Flow* and *Individual* model) were derived from metro ticket gates within a Spanish underground station collected as part of the AVATARS project [124,144,158,160-165].

Within EXODUS selecting any of the seven available data sets will in turn automatically set the metered gate's operating characteristics to match the selected device. Hence in the case of data sets relating to the definition of metered gates via the *Flow* based model, the length of the metered gate and its corresponding minimum and maximum flow rates will be automatically set. Conversely, for data sets relating to the definition of metered gates via the *Individual* based model, the length of the metered gate and its corresponding *Entry Delay* and *Travel Time* distributions will be automatically set. In addition to the seven available default data sets representing experimental data collected from equipment located and installed in existing structures the user also has the option to incorporate *User Defined* data to represent other types of metered gates. Hence the user is able to define/import both *Flow* and *Individual* based data. By default all *Metered Gate* transit nodes are assumed to be *Closed* turnstiles defined by the *Flow* based model (i.e. the *Closed Station* Fruin data set [19], see Figure 7-32).

Each of the default types of metered gate (i.e. turnstile/ ticket gate) which can currently be represented within EXODUS will now be discussed:

Turnstile Data
(a) Open/Closed Station Turnstile Gate
The *Flow* data for the *Open/ Closed Station Turnstile* gate is derived from data collected by Fruin [19] and is presented in Table 7-15. This consists of the minimum and maximum flow data imposed on the model.

Table 7-15: Flows achieved in different operational modes of a registering turnstile [19]. The Open and Closed operated options represent default settings in EXODUS.

Device Type: Registering Turnstile	Flow rate (occ/min)	Flow rate (occ/s)
Open Station Turnstile (Free Admission)	**40 - 60**	**0.67 - 1.00**
Closed Station Turnstile (Card/coin operated turnstile)	25 - 50	0.42 - 0.83
Ticket Collector	25 - 35	0.42 - 0.58

The *Open Station Turnstile* represents a free admission turnstile or *Open* gate, while the *Closed Station Turnstile* represents a card/ticket/coin operated gate or *Closed* gate. There are at least two types of open gate operation, one in which the gate barrier is permanently open and another in which the gate barrier can be operated without the use of a ticket or card. The former is typically found in mass transit stations and allows free but restricted movement, while the later is typically found in office buildings and libraries and offers free and somewhat more restricted movement. Within EXODUS the operating characteristics for both *Open* and *Closed Station Turnstiles* are set using default data (see Table 7-15). In both cases, once selected the length of the metered gate (i.e. turnstile) is by default automatically set to 1m. Thus on passing through an *Open/Closed Station Turnstile*, the agent will have travelled 1m.

Figure 7-33: Turnstile examined in the work by Fruin [19]

The flows presented in Table 7-15 were derived from the turnstiles shown on Figure 7-33. The default data used in EXODUS for the *Closed Station Turnstile* corresponds to the Fruin data for coin operated gates [19]. It is noted that these particular gates are not in common use in European mass transit stations and the use of coin operated ticket gates are now rare all over the world. Modern gates instead often require pedestrians to insert and collect a ticket or simply to swipe a card on a reader.

While the flows through this type of gate will not be identical to that for a card/ticket operated gate, both data sets will include a degree of delay as the pedestrian locates the coin/ticket/card and also a delay time while the pedestrian inserts/swipes the coin/ticket/card. For the ticket/card operated gate, there may be an additional delay as the pedestrian retrieves the ticket/card which is not present in the coin operated gate. Thus the flow rates specified for the *Closed Station Turnstile* gate should be considered indicative (at best) rather than definitive of modern card operated gates. However, the range of flow rates used for the *Open Station Turnstile* gate is considered appropriate for this type of gate.

The *Closed Station Turnstile* option would normally be used for circulation applications while the *Open Station Turnstile* option would normally be used in an emergency evacuation application.

To apply this data within EXODUS the *Turnstile* gate type and *Flow* movement model should be selected on the transit node. The user can then select either the *Open Station* or *Closed Station* options.

(b) Open/Closed Library Turnstile Gate
The *Individual* data for the *Open/ Closed Library Turnstile* gate relates to a turnstile installed in the Izmir Yuksek Teknologi Entitusu library in Izmir, Turkey which was collected as part of the BESECU project [154]. This turnstile operates in two modes. In the normal use of the building the users are required to swipe their ID card over the card reader to unlock the turnstile (i.e. *Closed*). When a card is correctly swiped the user hears an audio signal indicating that the card was accepted. The turnstile then unlocks allowing the user to traverse it. In contrast during emergencies the turnstile is set to free flow allowing free traversal without the use of an ID card (i.e. *Open*). The physical dimensions of the turnstile are shown in Figure 7-34. It has a length of 1.06m and a minimum width bounded by a pillar at 0.59m.

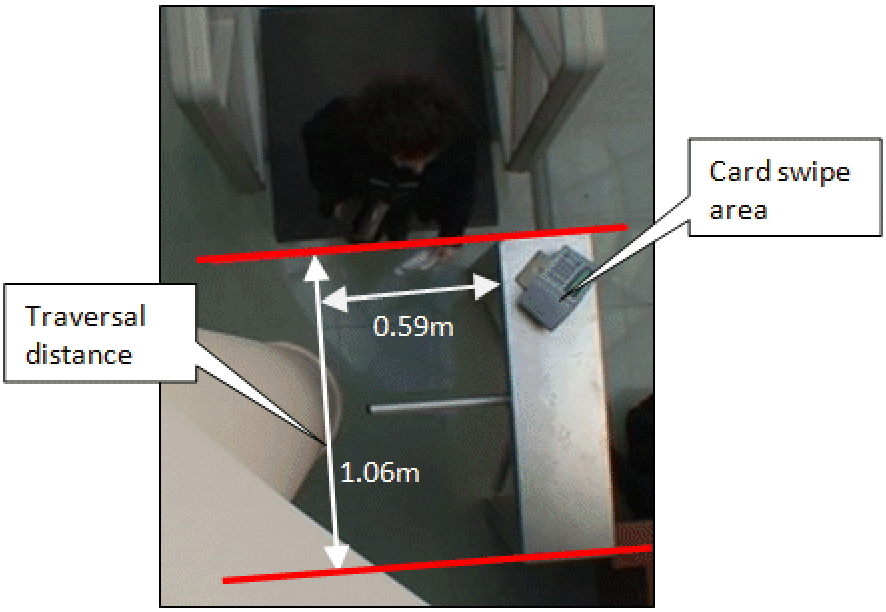

Figure 7-34: Turnstile in the Izmir Yuksek Teknologi Entitusu library, Izmir

Video footage of the turnstile being used during both circulation (i.e. *Closed*) and evacuation (i.e. *Open*) from the structure was collected. The video footage that was collected during the normal operation (i.e. circulation) of the structure amounted to approximately 46 minutes. During this time 60 people used the turnstiles. The observed *Travel Time* distribution limits were between 2 sec and 8.5 sec peaking between 3.5sec and 4sec. The video footage that was collected during the evacuation of the library amounted to just over 3.5 minutes during which 55 people used the turnstile. The observed *Travel Time* limits were between 1 sec and 8 sec with the most frequent traversal time between 2 and 2.5 sec. The data collected for the turnstile is shown in Figure 7-35 and Table 7-16.

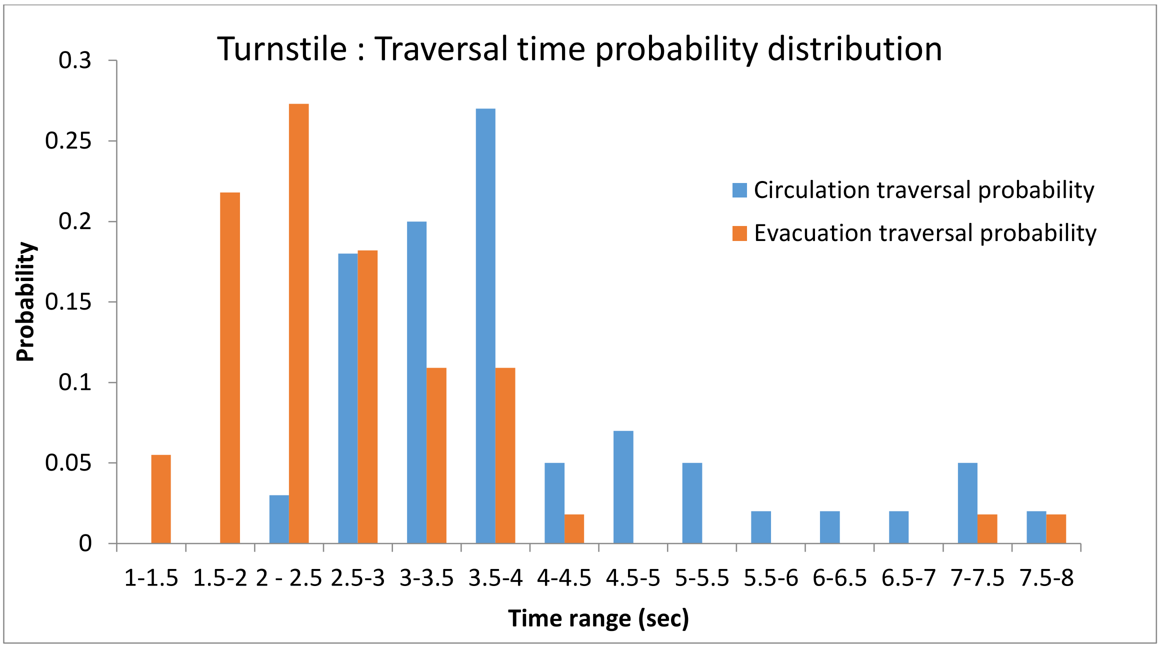

Figure 7-35: Traversal time distribution for turnstile use during circulation (i.e. Closed) and evacuation (i.e. Open)

Table 7-16: Turnstile traversal time frequency

	Circulation (i.e. Closed)		Evacuation (i.e. Open)	
Time range (sec)	**Freq**	**Probability**	**Freq**	**Probability**
1 - 1.5	0	0.00	3	0.055
1.5 - 2	0	0.00	12	0.218
2 - 2.5	2	0.03	15	0.273
2.5 -3	11	0.18	10	0.182
3 - 3.5	12	0.20	6	0.109
3.5 - 4	16	0.27	6	0.109
4 - 4.5	3	0.05	1	0.018
4.5 - 5	4	0.07	0	0.000
5 -5.5	3	0.05	0	0.000
5.5 - 6	1	0.02	0	0.000
6 - 6.5	1	0.02	0	0.000
6.5 - 7	1	0.02	0	0.000
7 - 7.5	3	0.05	1	0.018
7.5 - 8	1	0.02	1	0.018
8.0 - 8.5	2	0.03		

To apply this data within EXODUS the *Turnstile* gate type and *Individual* movement model should be selected on the transit node. The user can then select either the *Open Library* or *Closed Library* options.

Ticket Gate Data

Within EXODUS the default data for *Ticket Gates* corresponds to data collected from a group of ticket gates present in several metro stations in Barcelona as part of the AVATARS project [124,144,158,160-165]. The passengers operate these ticket gates in two phases. In the first phase the passenger inserts a ticket through a slot located at the front of the gate system (i.e. T1, see Figure 7-31(b)). If the ticket is not valid it is returned to the passenger. If however, the ticket is valid it reappears through a different slot located nearer the centre of the gate structure. The second phase starts from the moment the gate user retrieves the ticket which opens the gate and traverses the gate (i.e. T2, see Figure 7-31(b)). This phase ends when the passenger traverses the length of the gate and completely clears the device (i.e. T3, see Figure 7-31(b)). The physical dimensions of the ticket gates analysed are shown in Figure 7-36.

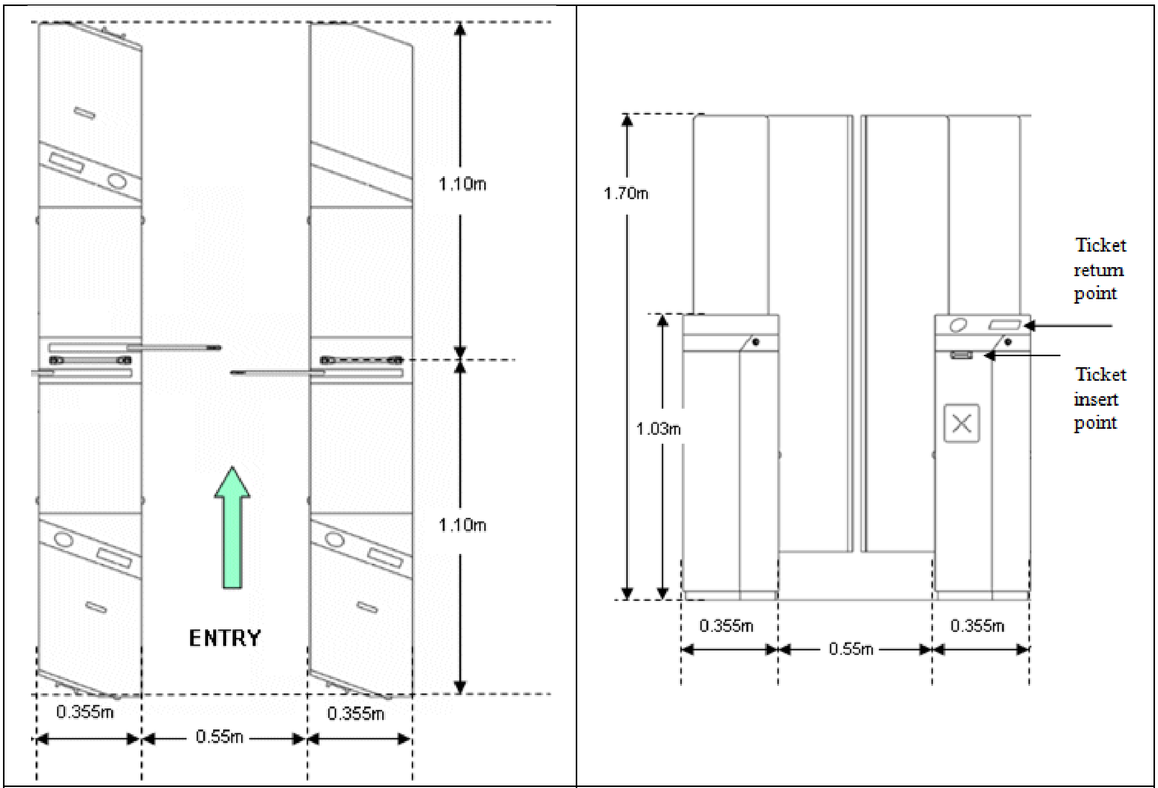

Top view	Front view

Figure 7-36: The ticket gate dimensions present within the Barcelona metro station.

The data was collected during the morning and evening rush hour on three separate days over a total duration of 28 minutes. In total 125 people were recorded using the gates over a series of 9 batches of pedestrians arriving from trains (see Figure 7-37).

Figure 7-37: A row of ticket gates within a Barcelona metro station.

Since the process of using and traversing the ticket gates consists of two phases two sets of data were produced. The first set corresponds to the *Entry Delay* time representing the time between the individual inserting and retrieving their ticket. This can be seen as T_1 within Figure 7-31(a) or alternatively the difference between T1 and T2 within Figure 7-31(b) (i.e. T2 – T1). The second set corresponds to the *Travel Time* representing the time between the individual collecting their ticket and them placing their first foot past the ticket gate (i.e. exiting the device). This can be seen as T_2 within Figure 7-31(a) or alternatively the difference between T2 and T3 within Figure 7-31(b) (i.e. T3 – T2).

The data collected from the ticket gates is shown in Figure 7-38 and Table 7-17 representing the probability distributions for the *Entry Delay* and *Travel Time*. The separation of the two data sets (*Entry Delays* and *Travel Times*) suggests that the *Entry Delay* of a person is independent of that person's traversal time. Since this is not known the model also allows the user to select a combined data set that represents the time from the moment the subject inserted the ticket into the slot and validated the ticket up to traversing and clearing the length of the gate.

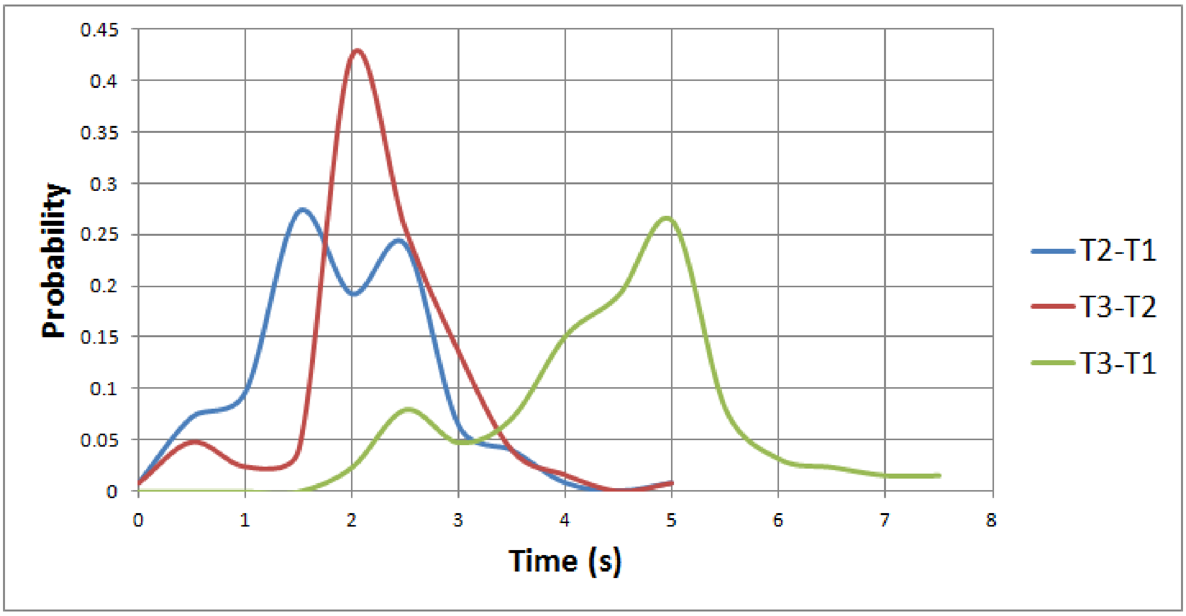

Figure 7-38: Ticket barrier usage time probability distribution, showing the corresponding Entry Delay (T2-T1, blue), Travel Time (T3-T2, red) and combined time (T3-T1, green).

Table 7-17: Ticket gate entry delay, travel time and combined probability distributions

Time range (sec)	Entry delay Probability	Travel time Probability	Combined time Probability
0.0-0.5	0.008	0.008	0.000
0.5-1.0	0.072	0.048	0.000
1.0-1.5	0.096	0.024	0.000
1.5-2.0	0.272	0.040	0.000
2.0-2.5	0.192	0.424	0.024
2.5-3.0	0.240	0.256	0.080
3.0-3.5	0.064	0.136	0.048
3.5-4.0	0.040	0.040	0.072
4.0-4.5	0.008	0.016	0.152
4.5-5.0	0.000	0.000	0.192
5.0-5.5	0.008	0.008	0.264
5.5-6.0			0.080
6.0-6.5			0.032
6.5-7.0			0.024
7.0-7.5			0.016
7.5-8.0			0.016

To apply this data within EXODUS the *Ticket Gate* gate type and *Individual* movement model should be selected on the transit node. The user can then either select the *Closed Station* option in order to model the *Entry Delay* and *Travel Time* via separate distributions, or alternatively select the *Closed Station (Combined)* option in order to combine the two distributions into a single distribution which is then used to model the agents total traversal time of the entire ticket gate.

Analysis of the Barcelona metro station ticket gate data also enabled flow rates through the devices to be calculated. The flow rate for each batch of 9 pedestrians was calculated individually which in turn enabled estimates of the minimum and maximum flow rates through the devices to be derived (see Table 7-18).

Table 7-18: The minimum and maximum flow rates observed through the Barcelona metro station ticket gates.

Device Type:	Flow rate (occ/min)	Flow rate (occ/s)
Closed Station Ticket Gate	12.7 – 16.6	0.211 – 0.277

To apply this data within EXODUS the *Ticket Gate* gate type and *Flow* movement model should be selected on the transit node. The user can then select the *Closed Station* option.

User Defined

In addition to the default *Turnstile* and *Ticket Gate* data the user also has the ability to define additional metered gates by manually defining their operating characteristics. The user can define the performance of a metered gate either via the *Flow* or *Individual* based model. If the metered gate is defined via the *Flow* based model then the user is required to define both the metered gate's length (m) and the flow rate obtainable through the device (i.e. by defining the minimum and maximum flow rates (occ/s)). Conversely, if the metered gate is defined via the *Individual* based model then the user is required to define either both the *Entry Delay* and *Travel Time* distributions, or alternatively merely a *Travel Time* distribution defining the total time for the agent to traverse the entire device. In addition the user should also define the length of the metered gate. It is important to note that when the metered gate is defined via the *Individual* based model its length is not used in the calculation of the time taken for the agent to traverse the gate. Despite this it is important to define the length of the metered gate in order to both accurately represent the distance traversed by the agent and hence their corresponding speed through the device.

CHAPTER 8: SIGNS AND THE VISIBILITY CATCHMENT AREA CONCEPT

buildingEXODUS allows the user to introduce signs into a structure. The user has the ability not only to examine the areas from where each sign, or combination of signs, is visible but also to assess the impact that signage has on the behaviour of the agents during circulation or evacuation scenarios. The visibility of a sign is determined by the *Visibility Catchment Area* (VCA) [103,104]. The VCA of a sign represents the physical extent to which a sign is visible within a structure: the area from which it can be seen. By placing a series of signs within a structure a complete signage system can be represented, along with their respective catchment areas. The implemented signage system can influence the behaviours and decisions taken by agents that detect the signs.

Placing *Sign* objects (see the User Guide, Section 3.3.1.19) within the geometry and then calculating the VCA for each sign enables the user to examine the areas from which specific signs are visible, as well as determining the combined catchment area of the signage system as a whole. This combined catchment area represents a footprint within which the agent population can receive information from one or more signs. Of more significance, agents located outside the VCA are not able to receive information from the signs present in the structure.

There are two possible strategies available to determine the VCA of a sign. The most obvious approach is to focus on the calculation of the VCA from the occupant viewpoint. Using this approach the VCA is determined according to whether an individual is able to ‘see’ the sign. The alternative approach is based around the locations from where each sign can be ‘seen’ by an agent. Using this strategy the VCA is determined according to the location of the sign in context with the surrounding configuration and this is interpreted by the population during the simulation. While both methods can produce identical results relating to the level of visibility and the VCA produced, the latter method is not only more computationally efficient (as there are generally fewer signs than people), it can also be used to determine a range of additional information and statistics relating to the positioning and effectiveness of a signage component. This method has therefore been adopted within buildingEXODUS.

The decision as to where to place a sign can be made before the interior of the structure is complete; i.e. before furnishings and fittings are introduced. Thus, while the signage system could be designed to be compliant to building regulations, when the structure has its full complement of furnishings and fittings, the signage system may become less effective due to these additions as they can obscure signs reducing their visibility. This may also occur when a different internal layout is implemented after a refurbishment. Examples of where this frequently occurs include supermarkets, furniture warehouses, department stores, libraries. Furthermore, occupants might lose sight of the exits and signs due to the internal complexity and configuration of the enclosure. For example, airport terminals are usually populated with a large amount of signage for normal circulation and retail concession advertising in addition to those required for evacuation.

In addition to these physical issues there may be a psychological component affecting the recognition of signs. The number of visible signs can create an information overload that may inhibit an individual’s ability to observe a particular sign. In effect, the clutter of signs creates visual “noise” through which it is difficult to discern the required information. So while the sign can be visible, the signal from the sign is overlooked. Other psychological aspects may come into play such as the attentiveness and the state of mind of the individual. There may also be physiological components to the visibility of signs, for example the nature of the eyesight of the observer. Furthermore, environmental factors such as the presence of fire products can have a significant impact on the visibility of signs.

The VCA of a sign attempts to address the *physical* visibility of the sign in relation to the surrounding structure that defines the region over which the sign is visible - the psychological and physiological aspects are dealt with by the behavioural model; e.g. the occupant's attentiveness (i.e. whether they notice and absorb the information on the sign) and the probability of an individual acting upon seeing a sign (i.e. whether they act in accordance with the information provided). *In essence, the VCA of a sign defines the region of space over which it is possible to visually receive information from the sign without guaranteeing that an individual who is located within the VCA will actually see the sign.*

The VCA of a sign is determined by taking into account the obstruction caused by other objects or structural features. This accounts for the location and height of the sign and obstructions and assumes an average observer height. The actual size of the sign - and by implication the lettering on the sign - is indirectly also considered (see below). By default the height of the observer is set to 1.75m (this can be altered by the user) and the height where the sign is assumed to be located is set by default at 2.2m (this can also be altered by the user).

The VCA algorithm used within buildingEXODUS uses a line of sight search methodology to determine the free space that has visibility access to the sign in question. The space that has visibility access to the sign is marked and then grouped in a zone representing the VCA that is associated to the sign. In Figure 8-1 the highlighted area corresponds to the VCA of sign 'S' located above an exit.

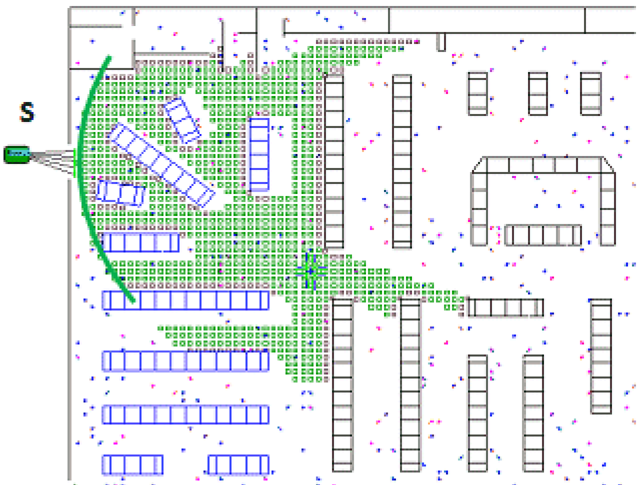

Figure 8-1: VCA of sign S, the highlighted area indicates the space that has visibility access to the sign S assuming an average observer height of 1.75m and a cut off distance of 30m. The curved line highlights the circular nature of the VCA.

To determine the level of visibility the algorithm uses the central point of the lower edge of the sign and a point in space at a height equal to the assumed average occupant height (see Figure 8-2).

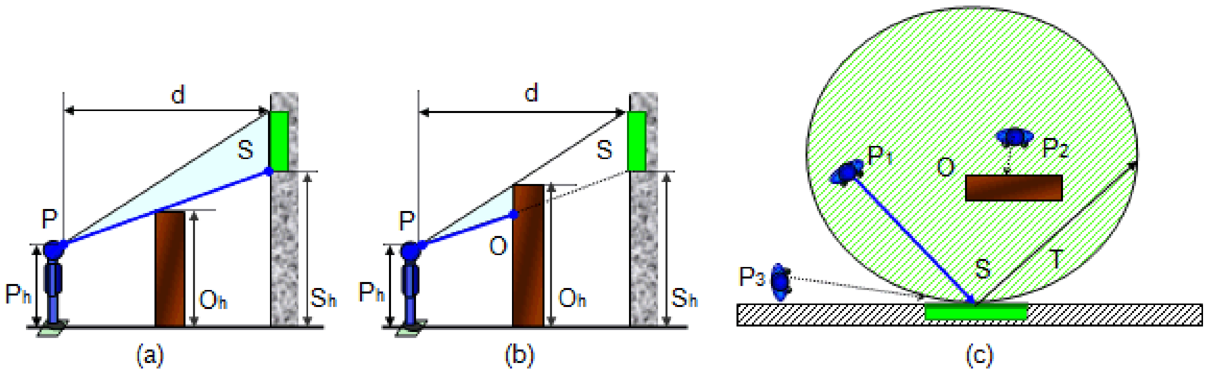

Figure 8-2: VCA is dependent on relative location of occupant, obstructions and sign. (a) The occupant is able to see the lower edge of sign S, while in (b) the occupant cannot see the base of the sign. In (c) only occupant P1 is able to see the lower edge of sign S.

The relative heights between observer, sign and obstacle can, and in many cases will, create a non-visibility "shadow" region behind some obstacles. Figure 8-3 shows the relation between observer, obstacle and exit sign in determining the visibility of a particular exit sign. In this example, person P2 is located too close to the obstruction O in order to see the sign S located above the exit. However, person P1 is able to see the sign as this person is situated far enough from the obstruction to get a clear view of the sign.

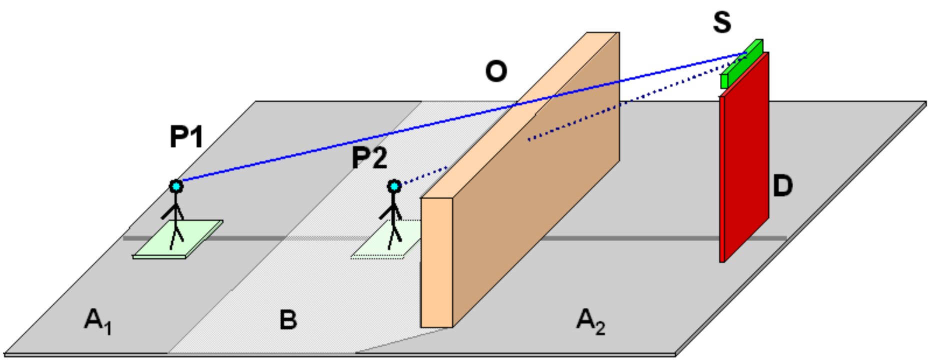

Figure 8-3: Exit sign visibility behind an obstacle

By default every boundary line in buildingEXODUS is assumed to represent a wall that does not allow the observer to see over or through it. This is because buildingEXODUS does not import 3D information from third party geometry files (i.e. DXF, IFC, FDS or SMF), other than the existence of the line itself (e.g. the opacity/height of the line) and thus has to assume that each imported line represents a solid wall that expands from the floor up to the ceiling. Therefore, the user has to manually assign the heights of the obstacles or walls in buildingEXODUS if these do not reach up to the ceiling. This is done in *Geometry Mode*. If obstacle or wall heights have been specified, they will be taken into account when determining the VCA of a sign. During this process only regions that have direct visibility access to a sign are marked as being part of the sign's VCA.

Additionally, when the VCA is calculated the observation angle is taken into account to determine the shape and size of the VCA. It has been shown theoretically and experimentaly [126] that there is a strong dependency on observation angle and signage VCA resulting in the VCA being circular in shape and tangent to the central point of the sign (see Figure 8-2(c)). The observation angle is defined as *the angle subtended by the observer's line of sight to a normal line bisecting the surface of the sign*. An observation angle of 0° means that the observer is viewing the sign straight on (i.e. the line of sight from the observer to the centre of the sign is perpendicular to sign surface). At this angle the observation distance will be the maximum distance as defined by the regulation or the user defined value. However, as the observation angle increases the observation distance will decrease. This relationship between observation distance and observation angle defines the circular nature of the VCA of signs.

In the current implementation, five factors are considered when determining the physical extent of the VCA of a sign:

- The height of the elements on the sign which can be the letters, pictographs, symbols or the height of the sign itself (these influence the maximum termination distance)
- The height at which the sign is placed (measured from the floor to the central point of the lower edge of the sign)
- The assumed height of the average observer (by default 1.75m)
- The height and locations of obstructions
- The agent's observation angle

However, in reality, the actual physical extent of the VCA depends on other factors such as the nature of the sign (reflective, self illuminating), the environmental conditions that both the observer and sign are exposed to and the physiological characteristics of the observer. The sign's physical properties that influence the level of visibility include: lettering colour and nature of the sign's illumination. Environmental conditions include the level of lighting within the compartment, whether or not smoke is present and whether crowding affects the visibility of the sign. In the present implementation, many of these contributory factors are not included for the sake of simplicity and the VCA is simply terminated according to the five factors described in the bulleted list above.

The maximum extent of a VCA can be modified by the user to represent local regulations or the nature of the signage system being simulated, or set by the user to any arbitrary value. For example, the NFPA Life Safety Code Handbook suggests that when the letters on an exit sign have a height of 15.2 cm it is legible for a distance of up to 30m when viewed in good visibility conditions [105]. To extend the visibility of a sign the letter height can be increased. Similarly, British Standard BS 5499 [127-129] defines the methodology for determining the maximum legibility distance of a sign. According to BS 5499 the maximum visibility distance is determined by the following formula:

$$D = Z * h \qquad \textbf{(35)}$$

Where:

D is the maximum viewing distance,
Z is a distance factor, and
h is the height of the element on the sign (letter, pictograph, sign surface).

Table 8-1 lists the suggested distance factors Z for various types of signs and luminance levels as defined by various revisions of BS5499 and BS5266. Similarly ISO 3864-1:2002 suggests two values for Z for reflective and self-illuminated signs.

In both NFPA Life Safety Code and BS 5499 it is suggested that the relationship between the element height and the visibility distance is linear. Therefore, as the element on the sign increases in size so does the visibility distance increase.

Table 8-1: Evolution of BS5499, BS5266 and ISO standards that define various termination distances based on the type of sign (reflective, self illuminated) and various lighting conditions

Standard	Definition of h	Distance factor Z, applicable conditions and notes
Legacy code Fire Precautions Act and BS 5499 Used for escape route signs	Height of uppercase letters on sign	Z=250, a safety factor of 2 is included
BS 5499-1:2002 Used for safety signs other than escape route signs	Height of the sign excluding white border	Z=120 for externally illuminated signs with a minimum vertical illuminance of 50 lux at the sign Z=225 for text only signs
BS 5499-4:2000 Used for escape route signs	Height of the graphical symbol	For externally illuminated signs: Z=95, Vertical illuminance at sign 5 lux Z=170, Vertical illuminance at sign 100 lux Z=185, Vertical illuminance at sign 200 lux Z=200, Vertical illuminance at sign 400 lux For internally illuminated signs: Z=150, Mean luminance of white contrast colour 10.0 cd/m^2 Z=175, Mean luminance of white contrast colour 30.0 cd/m^2 Z=200, Mean luminance of white contrast colour 100.0 cd/m^2 Z=215, Mean luminance of white contrast colour 200.0 cd/m^2 Z=230, Mean luminance of white contrast colour 500.0 cd/m^2 Note: The options for full level of lighting conditions for both externally (Z=200) and internally illuminated (Z=230) signs have been incorporated into buildingEXODUS model
BS 5266-7:1999 Used for escape route signs	Height of the sign excluding white border	Mains failure conditions Z=100 for externally illuminated signs Z=200 for internally illuminated signs Note: The options for mains failure conditions for both externally (Z=100) and internally (Z=200) illuminated signs have been incorporated into the buildingEXODUS model
ISO 3864-1:2002 Used for all safety signs [130]	Height of the sign excluding white border	Z=100 for externally illuminated signs, a minimum incident illuminance of 50 lux on the sign surface is required Z=200 for transilluminated signs, an average luminance of the contrast colour greater than 500 cd/m^2 is required

Within buildingEXODUS four visibility distance options are available to the user (see Table 8-2). These options include: the NFPA standard for full level of lighting; the BS5499-4:2000 standard for full level of lighting for reflective and self illuminated signs (Z=200 and Z=230 respectively); the BS5266-7:1999 standard for mains failure for reflective and self illuminated signs (Z=100 and Z=200 respectively) and a user defined model where the user can specify an arbitrary maximum visibility distance for a particular sign.

Table 8-2: Four options to determine sign maximum visibility distance available within buildingEXODUS

Model	Use	Sign Type	Maximum visibility distance
NFPA	Full level of lighting	Reflective	30.0m per 15.2cm of letter height
BS 5499-4:2000	Full level of lighting	Reflective	13.2m per 6.6cm of graphical symbol height
		Self illuminating	15.2m per 6.6cm of graphical symbol height
BS 5266-7:1999	Mains failure, emergency lighting available	Reflective	6.6m per 6.6cm of sign height excluding white border
		Self illuminating	13.2m per 6.6cm of sign height excluding white border
User defined	User specific	N/A	User specifies maximum visibility distance in metres

Once the VCAs are calculated the software can display the VCA of a particular selected sign on the floor plan of the structure as a shaded region. The model also returns statistics such as the total floor area covered by the signage system and the floor area covered by individual signs.

NOTE:
Within the current implementation of buildingEXODUS the physical extent of a sign's VCA is not directly affected by the level of lighting or the presence of smoke within the geometry. It is up to the user to specify an appropriate termination distance.

As mentioned earlier, the observation angle has an impact in defining the extent of the VCA: as the observation angle increases the viewer has to be closer to the sign in order to resolve the information relayed by the sign. This is due to the fact that viewer must be able to maintain an *angular separation* (i.e. angle φ in Figure 8-4) that will keep the sign legible [126]. The minimum angular separation (φ_{min}) which can be resolved (i.e. make out the individual elements on the sign) by the human eye is is assumed within the model to be constant. This constant is calculated using the maximum viewing distance for viewing signs with an observation angle of 0^o (i.e. straight on) as specified in the NFPA Life Safety Code Handbook [105]. For signs with lettering of 15.2 cm height, the maximum viewing distance is 30m. This produces a φ_{min} of 0.29^o. The angular separation of a sign is dependent on the size of the sign (or more correctly the size of the letters on the sign), and the distance of the observer from the centre of the sign and the observation angle (i.e. θ in Figure 8-4). An observation angle of 0^o (i.e. viewing the sign straight on) provides the maximum angular separation as well as the maximum ovservation distance wheras an observation angle of 90^o (i.e. viewing the sign side on) results in an angular separation (φ) of 0^o, effectively making the sign invisible to the observer. Therfore, to maintian the minimum angular separation ($\varphi_{min} = 0.29^o$) when viewing a sign at greater than 0^o angles ($\theta > 0$) the observer will have to be closer to the sign. This relationship defines the circular shape of the VCAs. This theoretical model has been confirmed through experimental analysis using test subjects [126].

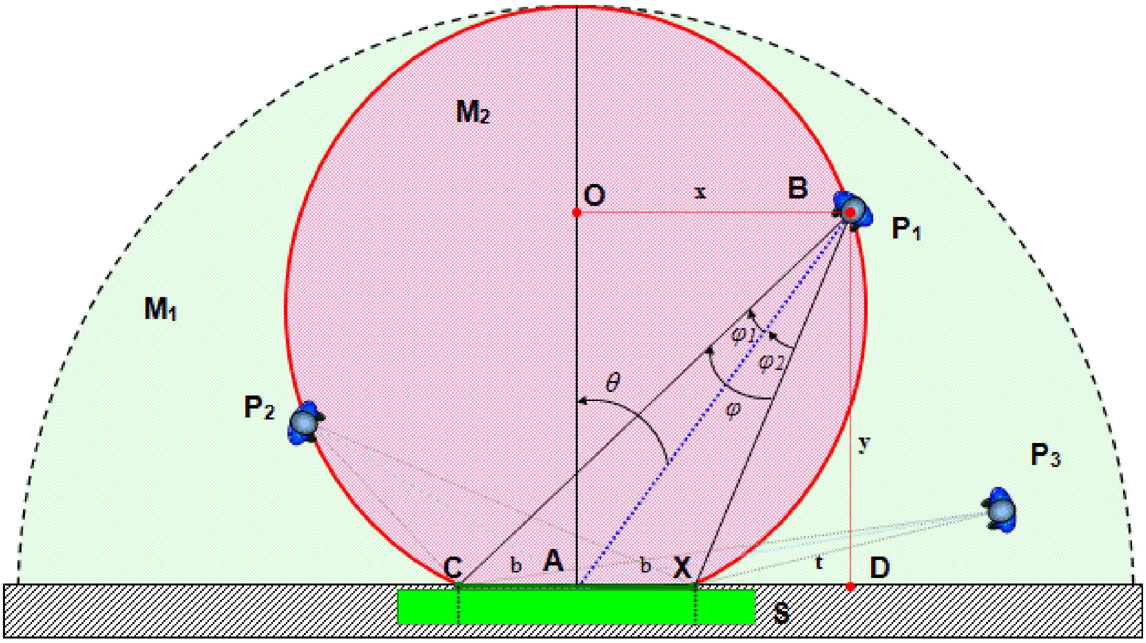

Figure 8-4: The composition of a sign object. Area M_1 defines the VCA as implemented by the first VCA implementation in EXODUS which was based on the assumption that the observation distance is independent of observation angle. Area M_2 defines the improved circular VCA implementation within EXODUS.

Thus, for an observer to be able to resolve a sign at the maximum observation distance, the observation angle should be such that the angular separation of the individual elements making up the sign are greater than or equal to φ_{min} and this can be achieved only if the observer moves closer to the sign than the maximum observation distance at 0^o degrees.

NOTE:
Within the current implementation of buildingEXODUS the physical size and shape of the VCA of a sign depends on its height from the floor, the height of the elements on the sign, the height and location of obstructions, the assumed height of the average observer and the observation angle

Once the VCA for a sign (or systems of signs) has been determined, buildingEXODUS is able to display the results on a floor plan of the structure (see Figure 8-5). buildingEXODUS also produces statistics such as the total floor area covered (i.e. the overall *Visibility Catchment Area*) by the signage system and the floor area covered by individual signs (see the User Guide, Chapter 5).

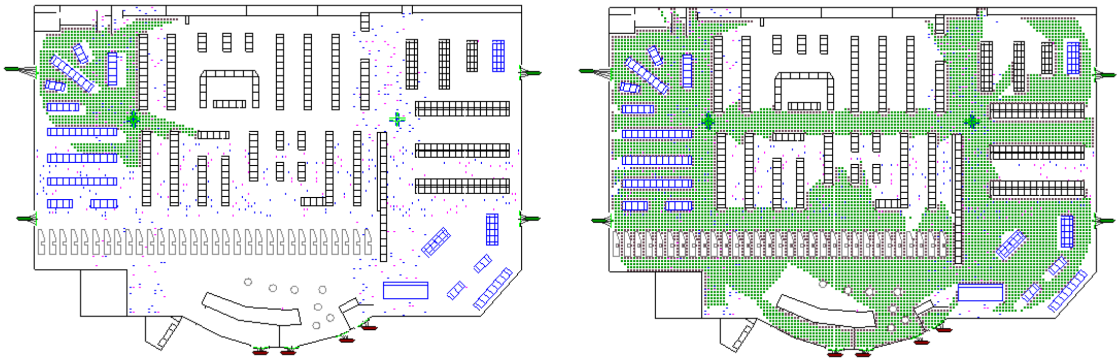

(a) **(b)**

Figure 8-5 (a) visibility catchment area of a single sign (b) visibility catchment area of all present signs combined

Two behavioural models within buildingEXODUS utilise signage during the evacuation simulations:

- *Adaptive Exit Selection* behaviour
- *Signage Model*

Adaptive Exit Selection behaviour enables agents to select which exits to use based on the congestion observed in the vicinity of the exits. The *Signage Model* allows for more advanced behaviours to be modelled. More specifically it allows the agents to interact with the signage system to gain directional information regarding points of interest within the structure or receive instruction from purposely placed signs. The *Signage Model* thus incorporates two main variations: *Signage Driven Behaviour* and *Agent Driven Behaviour* that determine the manner in which agents interact with the signs, use the signs, the information provided and the method used to link the signs together with their targets. Generally *Signage Driven Behaviour* allows agents to receive instructions or prompts from signs and *Agent Driven Behaviour* allows agents to seek locations within the structure that they wish to reach.

8.1 Impact of VCA on Adaptive Exit Selection

The VCA functionality is utilised by the *Adaptive Exit Selection* process in relation to observed congestion around an exit. By placing a sign above or immediately next to an exit or in the immediate vicinity (i.e. within a couple of metres) of the exit and then associating this sign and hence VCA with the exit, the agents are able to observe the congestion around the exit i.e. within the VCA associated with the exit. It is important to note that agents can only observe the congestion around an exit (i.e. within its VCA) when they themselves are also within the VCA. If the levels of congestion are sufficiently high at their current exit the agent may choose to use an alternative known exit (see Section 7.2) given the assessed conditions. If the agents fall outside of the VCA, then a similar judgement might be made, without however taking into account the congestion around the exit in question. The algorithm represents the impact that information can have upon the adaptive process.

NOTE:
Adaptive Exit Selection works only with signs that are placed above or immediately next to an exit (known as zero order signs - see Section 8.2) or are in the immediate vicinity of an exit (known as first order signs - see Section 8.2) and are directly associated with the exit. These signs are also known as zero or first level signs respectively.

8.2 Signage Model

The *Signage* sub-model allows agents to interact with signs, navigate within a structure to previously unknown locations and to gain knowledge of previously unknown exits, all this while they are circulating within or evacuating from a structure. The model thus allows an agent to interact with an isolated sign or a chain of signs, and also can represent the numerous factors that influence this interaction.

EXODUS classifies signs (referred to as the order) within the wayfinding system according to their relationship with the sign target. This classification includes *zero* order, *first* order and *high* order signs [104]. A *zero* order sign is located immediately above or next to the object it is intended to identify; e.g. an exit sign located above a door. A *first* order sign refers to a sign that directly points to the target object, but is not located adjacent to that object; e.g. a sign in a corridor that points directly to an exit, but is not adjacent to the exit. To be effective (both in reality and in the model), it is preferable that *first* order signs lie within the VCA of the object to which they are pointing. Finally, a *higher* order sign refers to a sign that does not directly point to the target object, but leads the observer to an area where another sign in the signage system exists that is closer to the target object. A *higher* order sign points to another sign of lower order (i.e. a *fifth* order sign will point to a *fourth* order sign, etc). Figure 8-6 shows the relationship between signs of various order level. It should be noted that the schematic in Figure 8-6 may include numerous turns and obstacles which are not shown for clarity.

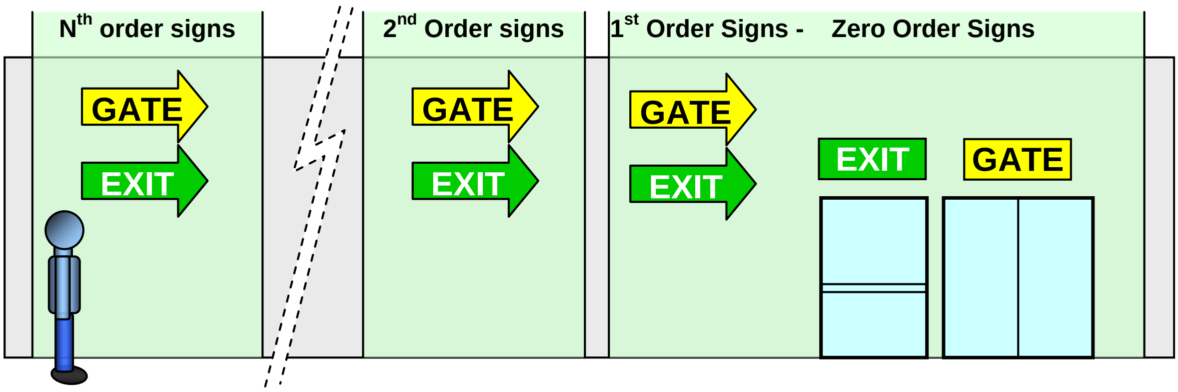

Figure 8-6 Signage system involving zero, 1st, 2nd and higher order signs

In reality, a well designed signage network would consist of a series of signs that are:

- in appropriate locations
- in close proximity to each other or in view of each other
- direct the occupant towards their target intuitively
- reinforce the observer’s confidence in the signage system

Such a system would allow the occupant to follow the instructions relayed by the signs, and thus move from sign to sign until the target object is reached.

Two variations of the signage interaction model are available. The first is known as the *Signage Driven Behaviour* that instructs agents to follow the information provided who had no previous intent to use the signage system. In this system the agents are not looking for a specific target location, but once they (coincidentally) observe a sign it provides information that the agents then use (i.e. follow) instructing them to follow a particular path. The second variation is known as the *Agent Driven behaviour*. It refers to a system where agents deliberately seek out signage in order to guide them to specific locations they wish to visit, or exits they wish to reach. In this system agents are assumed to have a predetermined desire to reach a target, but do not necessarily know how to do so.

In both variations the (chain of) signs can potentially influence the evacuation or circulation behaviour of the agents by allowing them to adopt new paths and also to utilise previously unknown exits. Both systems can operate simultaneously within a simulation, assuming that the signs are appropriately configured, and are then used according to the agent's attributes; i.e. whether their use of signage is coincidental (i.e. *Signage Driven* behaviour) or pre-determined (i.e. *Agent Driven behaviour*).

For *Signage Driven* behaviour the signs need to point directly to either another sign or to an exit. Agents that have not been assigned a target exit or itineraries can then utilise the information provided by these signs. This requirement assumes (i.e. represents) that the agent has *some* understanding of the routes suggested - they have some prior understanding of the routes available, with the signs reinforcing their knowledge of the existence of objects/exits within the structure.

For *Agent Driven* behaviour the signs point to an eventual target location or object. This information is coupled with an associated direction that leads the agents towards this target (or towards another sign in a chain of signs that continues pointing towards their eventual target). The associated direction prompts the agent to move in the general direction indicated by the sign (in reality) without guaranteeing their arrival at the intended target. This is to represent the fact that the agent is not necessarily aware of the location of the target and is therefore moving in the general direction of the intended target (e.g. falling within the VCA of the intended sign). In order that agents can utilise this system they must be given itineraries that indicate that they have to find their desired target using signage.

In the next sections, the key components of the *Signage Model* are described. Sections 8.2.1 and 8.2.2 address how an occupant may initially encounter a sign (whether they use *Signage* or *Agent Driven behaviour*) and under which circumstances each system can affect an agent's behaviour. Section 8.3 addresses how an agent interacts with the sign. Section 8.3.6 addresses movement between the signs.

8.2.1 Signage Driven Behaviour

The *Signage Driven Behaviour* system represents the unplanned use of signage by an occupant - where the occupant coincidentally encounters the information provided and then makes use of this information. It assumes that the occupant has a basic familiarity with the structure in question. The purpose of this system is two-fold:

a) To allow occupants to learn and use previously unknown, but closer exits.
b) To allow occupants to follow arbitrary paths defined by signs. The use of non optimal paths is possible.

This system provides instructions to agents that *happen* to fall in the VCA of a sign and it can direct occupants towards another sign or to an exit. In some circumstances this system can guide agents to previously unknown and closer exits. *Agents that navigate purposefully in the structure by either having a target exit or following an itinerary are not influenced by this system and ignore the signs completely.*

The signage system incorporates a series of signs linked together in a simple list fashion where one sign points *directly* to another sign, until a final sign which can lead agents to an arbitrary location within the structure or point directly to an external exit. Figure 8-7 shows the relationship between a sign and its target information which can be either another sign or an exit in which case by definition it would be the last sign in the chain of signs.

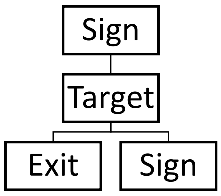

Figure 8-7 A sign points to a target that can be either another sign or an exit

Figure 8-8 depicts the results of the interaction between agents and the signage system depending on whether the agents have a full knowledge of the available exits or not. In both cases if the agents see a *second* or *higher* order sign they will follow its instructions by following the indicated path. However, in the case when the agent has limited knowledge of the available exits, 0^{th} or 1^{st} order signs can influence agent exit behaviour. Agents with reduced awareness of the available exits can learn new exits and can chose to use them provided they are located closer than the exit they were originally heading for.

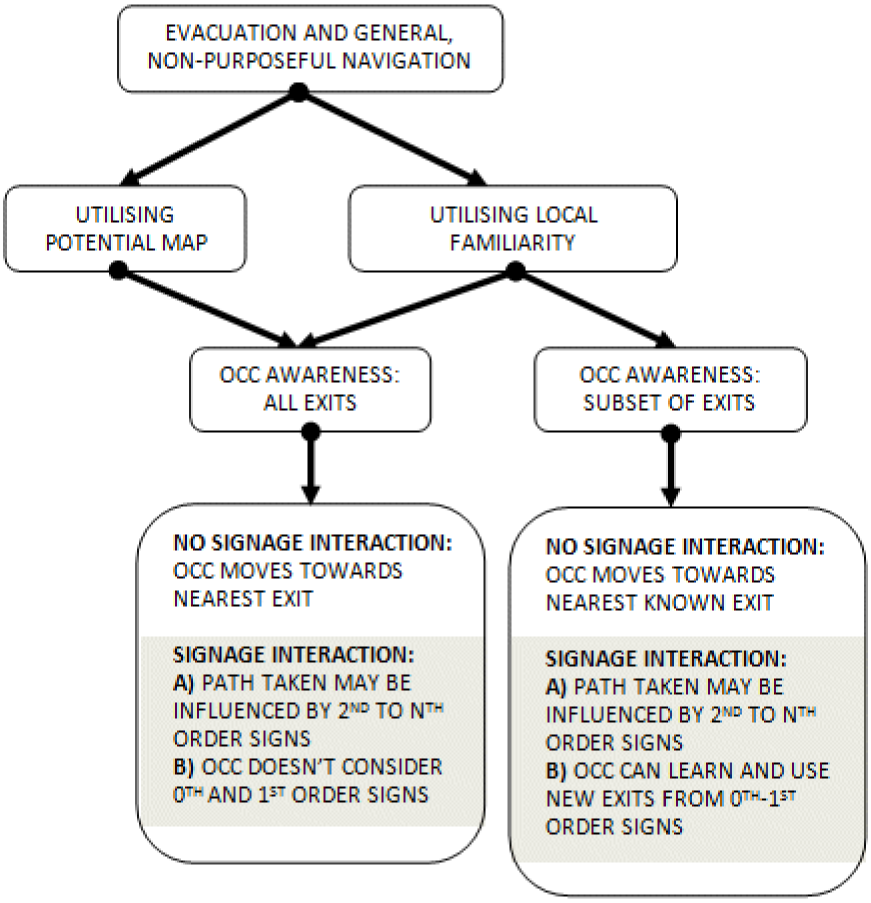

Figure 8-8: Signage Driven Behaviour: the highlighted sections indicate the outcome of occupant interaction with signage

The two outcomes of *Signage Driven Behaviour* are now described. These two outcomes do not require specific user actions for them to function - they are a consequence of the behaviours included and described in Figure 8-8.

Allowing agents to learn and use previously unknown but closer exits

When using the *Potential Map* system it is assumed that all the agents have complete knowledge of all the available exits and so using this approach it is not possible to differentiate between exit knowledge levels within the population. However, when using the *Local Familiarity* system it is possible to represent differing levels of exit knowledge within the population, i.e. some agents can be aware of fewer exits than are available. This exit knowledge influences the impact of *Signage Driven Behaviour*. When an agent is aware of all of the available exits (i.e. either the *Potential Map* system is used or the *Local Familiarity* is used with the agent being aware of all of the exits), the agent can follow instructions only from second or higher order signs influencing the adopted route. Once they reach the end of the chain of signs (or a *second* order sign), the agent will then move towards their nearest exit.

When an agent has limited awareness of the available exits (i.e. the *Local Familiarity* is used with the agent being aware of a sub-set of the exits), the agent can follow the instructions from *second* or *higher* order signs influencing the adopted route. The presence of *zero* or *first* order signs can also influence the exit usage by allowing the agent to use a previously unknown, but closer exit; i.e. in effect learning new exits that are of use to them.

Once the agent reaches the end of the chain of signs, the agent will then move towards the exit indicated by the sign they have followed provided it is the closest exit. If the signage chain ends at an arbitrary location within the geometry and not on an exit, then upon reaching the last sign in the chain the occupant will then progress towards their nearest known exit.

This allows agents that have a reduced familiarity with the structure's exits to learn new exits. If this takes place the new exits are added to the agent's *Occupant Exit Knowledge* (OEK) list. If the agent is evacuating and not following an itinerary or moving towards a target exit then when a new exit is learned the agent will use it provided it is closer than the exit they were originally heading for. Only 0^{th} and 1^{st} order signs can influence the agents in the described manner. *Second* or *higher* order signs cannot directly affect exit usage however they can alter exit usage indirectly by guiding agents to a location where an alternative exit exists.

Figure 8-9 provides an example of this behaviour. Here it is assumed that agent P is aware of only Exit 1 and is unaware of the existence of Exit 2. As agent P moves towards Exit 1 and passes through the VCA of sign S there is a chance of seeing sign S and therefore becoming aware of the existence of Exit 2. If this takes place then agent P may choose to redirect and evacuate via Exit 2 which is much closer than Exit 1.

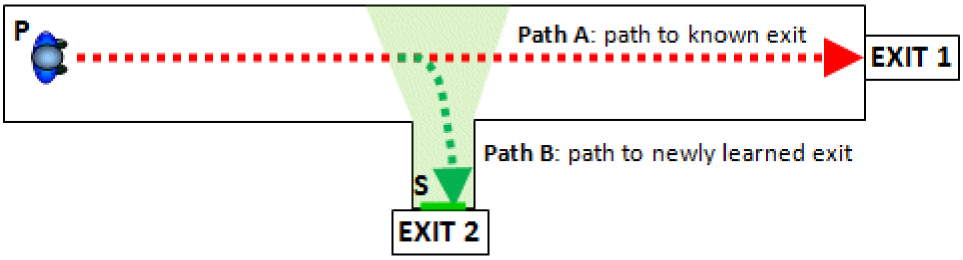

Figure 8-9: Agent P moves towards known Exit 1 however, as the agent sees sign S chooses to redirect and evacuates via Exit 2

Allowing agents to follow arbitrary paths defined by signs with the possibility of following sub-optimal paths

This allows the agents to adopt arbitrary paths within the structure by following the instructions of the signage system. Only *second* or *higher* order signs can influence the agents in the described manner. The navigational system used (either *Potential Map* or *Local Familiarity*) has no impact in this behaviour. If the agents see the sign they will follow its instruction by adopting the path indicated by the signage chain. When the agents are evacuating or are engaged in non-purposeful navigation (i.e. the user has not assigned them with a target exit or itineraries), then this system can instruct the agents to alter their path and use alternative, unplanned paths, possibly non-optimal paths. An example of where this system can be used is presented in Figure 8-10. Here the normal route of agent P in the structure shown would be to move straight ahead (Path A). However, an obstruction blocks his path. As a temporary measure the operators of this structure decided to redirect the occupants to use an alternative route sign posted by signs S1, S2 and S3. By seeing these signs agent P can follow the alternative path (Path B) and hence avoid the obstruction.

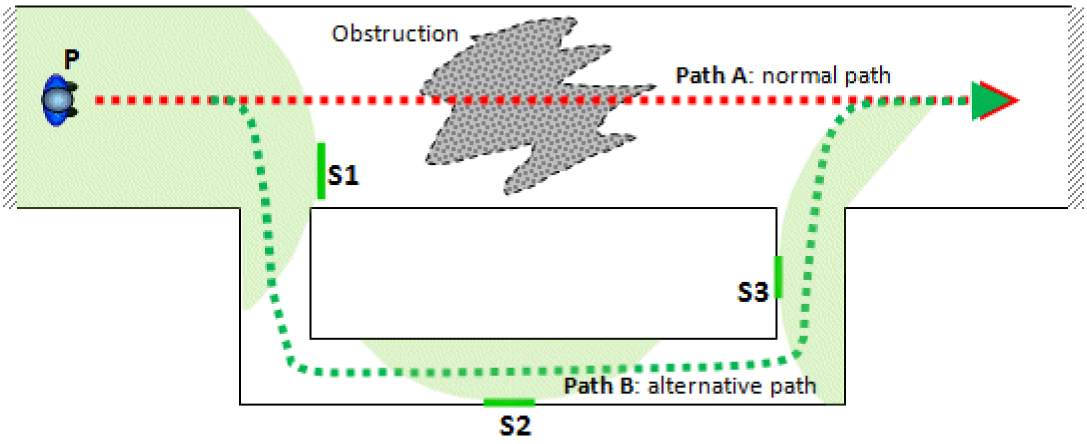

Figure 8-10: Agent avoids obstruction by following Path B that is sign posted by signs S1, S2 and S3

8.2.2 Agent Driven Behaviour

In *Agent Driven Behaviour*, the agents intend to visit or use specific locations or exits, but are completely unfamiliar with the routes to reach them; i.e. they know where they want to end up but do not know how to accomplish this. For instance, the agent may have been instructed as part of a procedure to reach a refuge location. This requires the use of the *Find via Signage* itinerary (see the User Guide, Chapter 4). The agent has no specific knowledge of the route to the refuge; however, they have been told to follow the available signs to reach the refuge in question. The agents purposely start the process of moving towards their goal by following the signage system.

The *Agent Driven Behaviour* is applicable when the agent has a predetermined goal; i.e. when the agents have an itinerary of targets to visit without knowing their exact location. Figure 8-11 depicts the likely outcome of the interaction between agents and the signage system when the *Agent Driven Behaviour* is utilised. It should be noted that whether the *Local Familiarity* or *Potential Map* system is employed is not relevant as the agents are assigned specific itineraries for the functionality to be enabled.

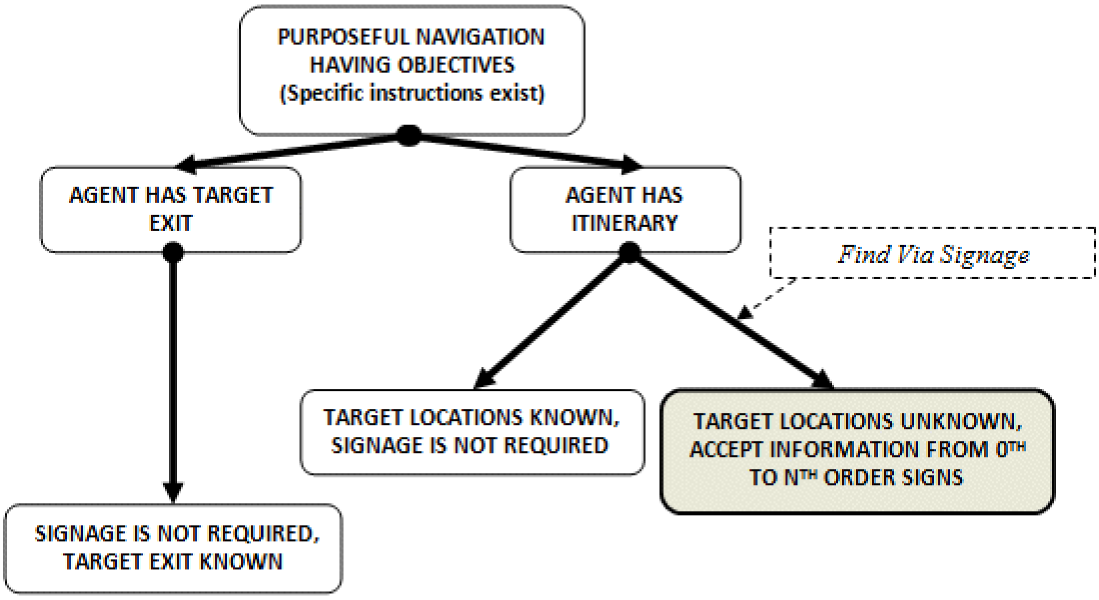

Figure 8-11: Greyed box indicates when the *Agent Driven Behaviour is active*. Note that the *Find via Signage* itinerary should be assigned.

In the model, the *Agent Driven Behaviour* system utilises a signage system that incorporates a series of signs each of which provides information for one or more targets along with a general direction towards these targets. For each target a sign should provide two pieces of information: a target specified as a nodal location within the structure or an exit - *this is the location where the agent wants to end up*; and a nodal location that determines the agent's direction of movement that they need to follow to reach their target. This represents:

(1) where the sign is actually intended to lead the agent,
(2) the fact that the guidance provided is general and is reliant on the agent moving in the general direction indicated - in effect, the process is imperfect.

These information pairs (i.e. eventual targets and directions) are assigned by the user for each sign by selecting at least one target object (e.g. a nodal location or an exit) and a corresponding direction towards that target object (reflecting the direction indicated by the sign). It should be noted that to ensure that an agent reaches the next intended link in the 'information chain' (e.g. sign), the direction node should preferably lie beyond the VCA of the next sign in the signage system. This allows a route to the VCA of the sign to be located (see Figure 8-12).

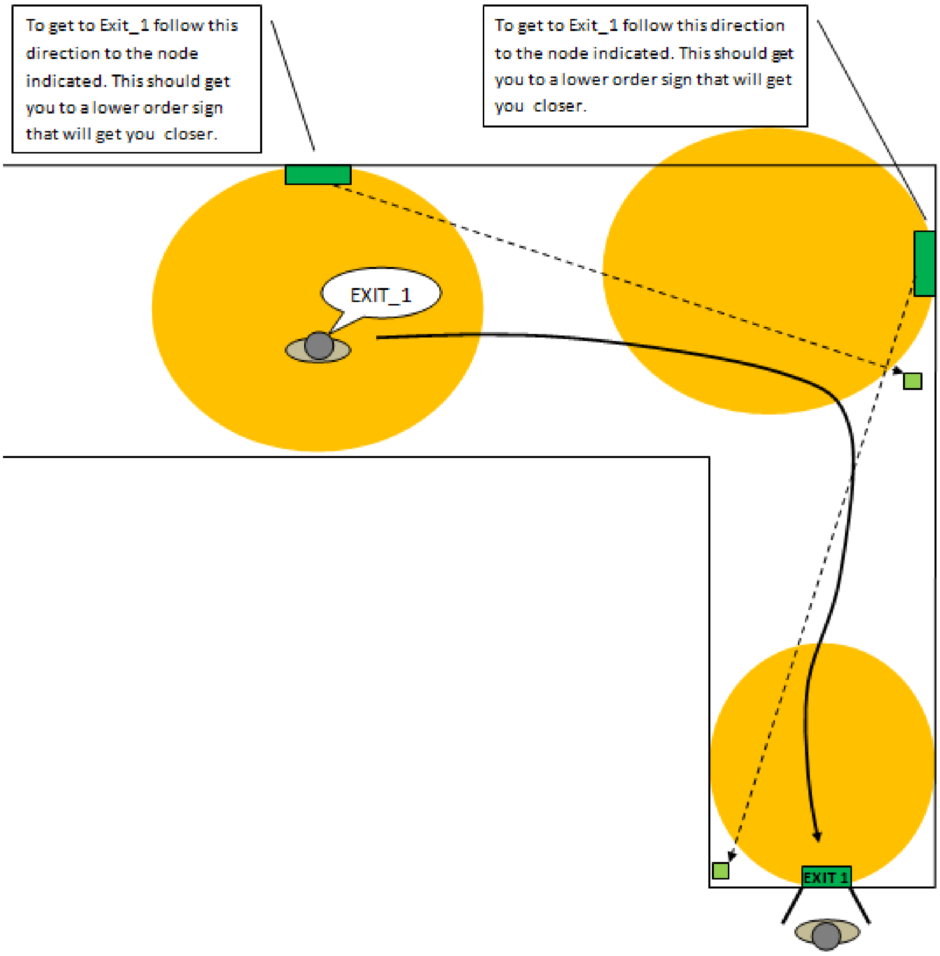

Figure 8-12: Positioning of route nodes.

A more complex example is presented in Figure 8-13, where the interaction between the agent, the real world information provided by the signs, and the buildingEXODUS representation is shown. In this example the signage system consisting of signs S0 to S3 all pointing to Exit 1 (i.e. Exit 1 is the target of all signs). Each sign is indirectly linked with its next lower order sign by indicating a general location to reach the sign; i.e. a node is indicated by the user describing the direction of the movement indicated on the sign.

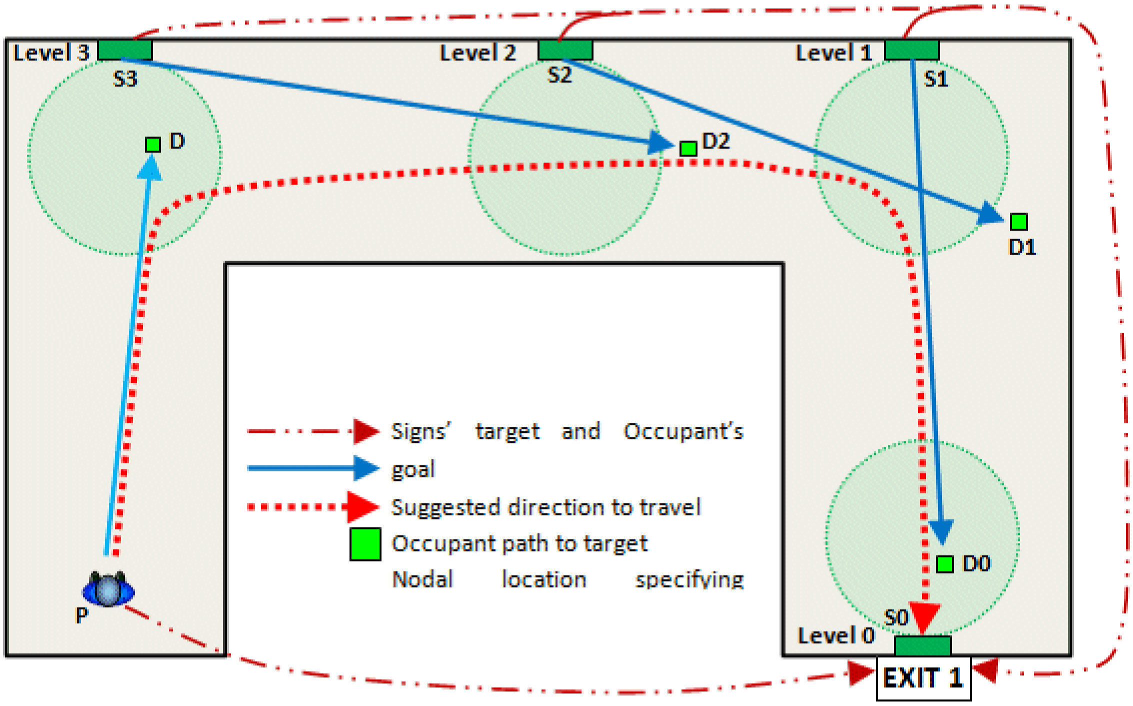

Figure 8-13: Movement produced by the Occupant Driven System.

A sign can provide information leading to several targets; i.e. numerous target exits/locations each of which requires a nodal location indicating the short-term direction that moves the agent on a path that will eventually get them to their objective (see Figure 8-14). The value and use of this information depends on the itineraries of the occupants in the VCA of the sign; i.e. their desired target point.

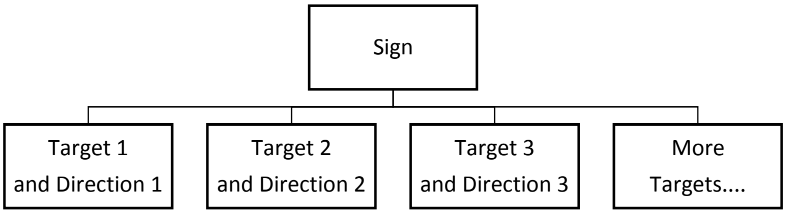

Figure 8-14: A sign may be associated with multiple targets, each of which should have an associated initial direction of travel

An agent may have several tasks to complete as part of their itinerary - several of which may involve them moving to a sign to receive new information and be led towards target locations (which may also incur an associated delay as part of the task completion). The information pairings on the sign used by occupants are in the same format - a long-term target location and a short-term general direction (see Figure 8-15).

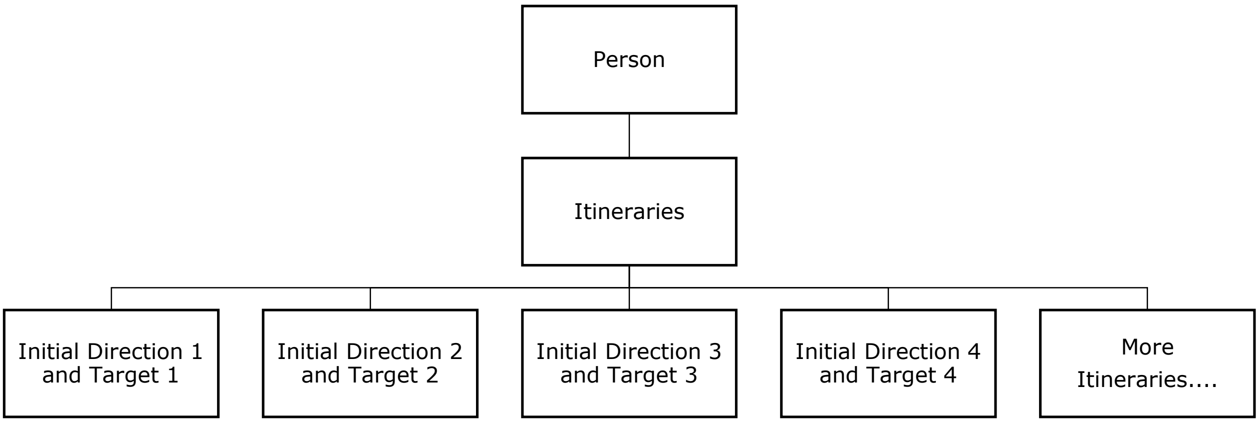

Figure 8-15 A person can be assigned multiple targets.

An agent will have an initial instruction to move towards a sign in order to start receiving information from the signage chain. Once they are in the VCA of the sign, they will consider the information available and attempt to obtain additional directional information to guide them towards their target location. Depending on the nature of the sign, this may involve them going through several pieces of information provided by the sign (e.g. a multidirectional sign) to find something relevant to their current target. Once this direction is found, they will head off in the assigned direction. **Other information provided by the sign is not considered - only information that addresses the agent's final target is considered.**

An agent may receive information from the same sign several times during a simulation provided it relays information on multiple targets that the agent wishes to visit.

The use of signage allows relatively complex emergency or non-emergency procedures to be represented. However, for an agent to be able to utilise this system there must exist at least one sign that provides information related to the agent's desired target or targets. **Furthermore, this transfer of information takes place only for one objective (i.e. target) at a time. The agent can only seek and register information from a sign relevant to their current task (i.e. target).** When an agent wants to visit multiple targets the agent addresses them one at a time.

Figure 8-16 is an example showing the relationship between an agent and a sign that provides information that the agent wishes to use. In this example, agent P1 wants to visit a coffee shop and then visit the rest room (in this order). In this case, agent P1 is considered to have two tasks, task 1 to visit the coffee shop and task 2 to visit the rest room. It is assumed that this agent does not initially know the locations of either of these targets. However, this information is provided by sign S1. If agent P1 is observant and sees sign S1 then this agent will acquire information regarding their location. However, since the transfer of information takes place for one target at a time agent P1 will first try and acquire information for task 1 (coffee shop) and once this task is completed will try and acquire information for task 2 (rest room). In addition, sign S1 also provides information for the duty free shop however, agent P1 is not interested in this piece of information and thus does not register it. The arrows in Figure 8-16 indicate the information required and received by agent P1 from the sign S1. Note that agent P1 first receives information related to task 1 and then once completed for task 2 (this assumes that the agent will be physically capable of observing sign S1 once task 1 is completed).

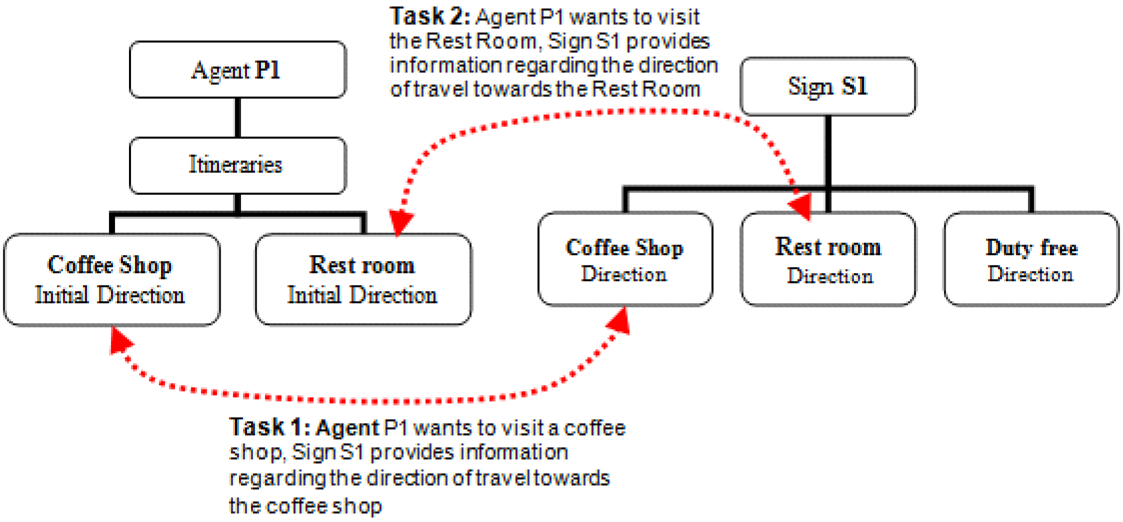

Figure 8-16: Agent P1 wants to visit a Coffee Shop and the Rest Room. This information is provided by Sign S1.

Note that due to the way the navigational model functions the agents must be given an initial direction of travel by the user (i.e. direction towards a sign), along with their long-term target. An agent's initial direction is similar in concept to the direction of travel incorporated within the signs (e.g. indicated by an arrow on the sign). It requires the generation of a *Find Via Signage* itinerary where a corresponding node represents the agent's initial direction of travel. A user specified direction of travel is only necessary for the initial movement towards the first observable sign. The directional information on the sign will thereafter provide the direction of travel to observant agents. If such an initial direction is not given then the agents would not know in which directions they should start moving in order to find their long-term target. Once an agent observes a sign the signage chain will provide further information regarding the path to the agent's desired target.

Figure 8-17 depicts another example where this system has been utilised. Agent P1 wishes to visit room R1 and then exit the structure using Door 1. The signage system (comprising signs S1 to S5) provides the directional information to this agent - effectively pushing them towards their goal in stages: signs S1and S2 guiding the agent to R1, and then S3, S4 and S5 guiding them to Door1. The other agents within this structure wish to evacuate the structure using Door 1. The signage system provides information on the location of Door 1 through the path defined by signs S1, S3, S4 and S5 to the exit Door 1 (S2 provides information about room R1 only).

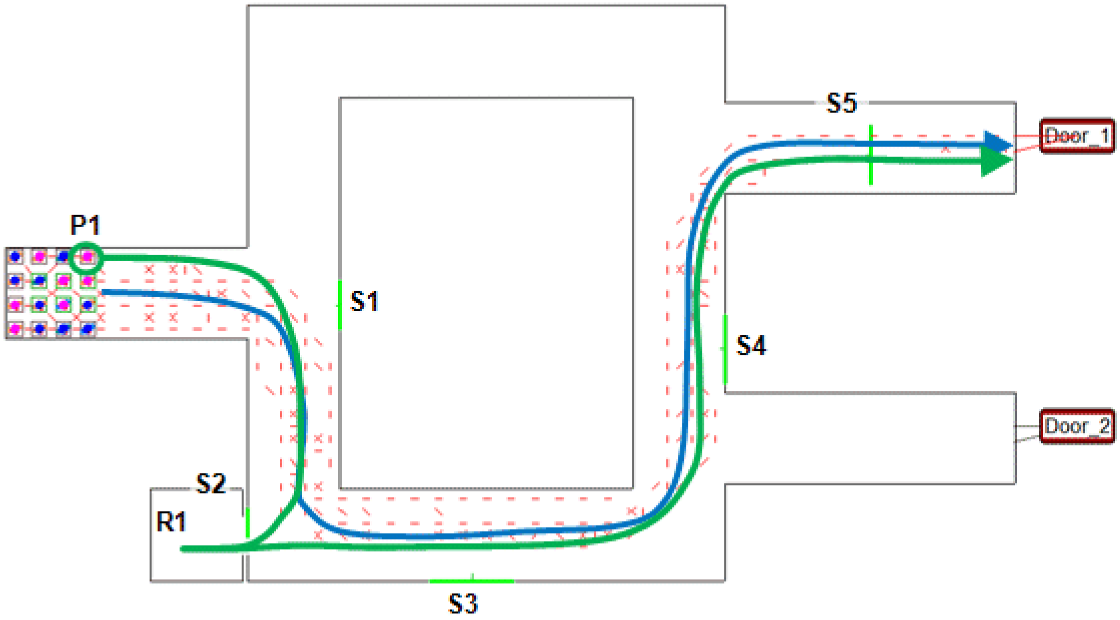

Figure 8-17: Complex routes produced by use of Find via Signage itinerary.

NOTE:
The Agent Driven behaviour system assumes that the agents know where they want to go but do not know how to get there. Agents that observe a sign can receive information regarding the direction they should travel in order to find their desired target.

8.3 Sign Interaction

Irrespective of which of the two signage interaction systems are used within the simulation, the agents will need to interact with individual signs and the model will need to determine the outcome of this interaction. Ultimately this interaction will determine whether the agent observes and uses the information provided by the sign. Three methods are provided to determine whether or not the agent observes and uses the signage information: *Predictive*, *Empirical* and *User Defined*.

Currently little data is available on how individuals actually interact with signage systems. The *Predictive Approach* represents an initial understanding of how people interact with a signage system. In order to better understand the interaction between individuals and signage FSEG conducted a study involving a series of experimental trials [126]. This study resulted in a revision of the *Predictive Approach* that led to the *Empirical Approach*. The *User Defined* approach is based on the *Empirical Approach* however it allows the user to specify user defined values for detection and compliance probabilities associated with this approach. The three available methods are described in the following sections.

8.3.1 Predictive Model

The *Predictive Model* is a hypothetical model which determines whether or not the agent detects a sign and then follows the sign's directional information. The model represents three aspects of the agent interaction with a sign:

- Whether or not it is possible for an agent to physically see the sign
- Given that it is possible to see the sign, does the agent actually detect the sign
- Given that the agent has detected the sign does the agent then follow the sign's guidance instructions

In the first step, the agents must be located within the VCA to be physically able to see the sign. However, this alone does not necessarily determine whether or not the agents can physically see the sign. For instance, if an agent has their back to the sign, they cannot physically see the sign even though they are in the sign's VCA. To determine whether or not it is possible for an agent to see a sign it is thus necessary to take into consideration the agent's angle of approach relative to the sign. This is known as the *visibility probability*. This is determined by an arbitrary sigmoid function (see Section 8.3.4) that determines the *visibility probability* as a function of angle of approach towards the sign. This function is based on the assumption that if the agent is moving in a direction that is coincident to the direct line of sight of the sign, the greater will be the probability that the agent will be able to physically see the sign (see Section 8.3.4). The user has the option of using the sigmoid function to determine the visibility probability or specifying a single user defined arbitrary visibility probability value (0% to 100%).

Once it is determined that the agent can physically see the sign it is necessary to determine whether or not the agent detects the sign. Within the model it is arbitrarily assumed that there is a 50% chance per step that the agent is attentive and so is able to detect the sign while within the VCA of the sign. As this determination is made each step the agent completes while within the VCA of the sign, the longer the agent is moving within the VCA the higher the chances are that they will be attentive and able to see the sign. The default *detection probability* of 50% means that the agent has a high degree of attentiveness and so has a high probability that he will detect the sign. The user has the option of using the default detection probability (50%) or specifying an arbitrary visibility probability (0% to 100%).

If the agent detects the sign, it is necessary to determine whether or not the agent will follow the directional information provided by the sign i.e. given the sign was detected, does the agent use the information provided. In the *Predictive Model* this is by default set to 100%, although it can be modified by the user (0% to 100%).

The following pseudo-code corresponds to the *Predictive Model* and determines whether a person will detect and eventually follow the instructions provided by the sign or not.

Table 8-3: The pseudo-code for the *Predictive Approach*

```
a)  IF Person is in a VCA THEN
b)  IF Person will see the sign (visibility probability as a function of angle of approach) THEN
c)  IF Person is attentive (50% detection probability each step in the VCA for this to be true) THEN
d)  IF Person will act on information received from sign (100% compliance probability to be true) THEN
e)  Person learns new exit and uses it, if appropriate, OR Person follows sign's instructions
```

The probabilities that are used in statements (b), (c) and (d) have been determined arbitrarily and indicate the default values used by the model, and can be modified by the user.

If the criteria listed in Table 8-3 are met, then the agent is deemed to have observed the sign and has chosen to act on the information received from the sign. Figure 8-19(a) depicts a representation of this algorithm in a flow chart form.

8.3.2 Empirical Model

The *Predictive* approach represents a hypothetical model of how people may interact with signage systems. In order to better understand the interaction between individuals and signage systems FSEG conducted a series of experiments involving individuals who had no knowledge of the egress routes from a particular building [126]. The results from this experiment provided data on how people perceived different levels of emergency signs during an evacuation situation. The data generated from these experiments was used to develop an alternative model to the *Predictive Model* called the *Empirical Model*.

The experimental trials were conducted in the area shown in Figure 8-18. It consisted of several large, relatively wide corridors with access to offices or windows on either side. Given the size of the corridors some parts are used for general circulation. These corridors end in four exits, two that lead to other corridors and parts of the structure (Exits D1, D5) and two that lead to staircases to the outside of the structure (Exits D2, D4). The test subjects were asked to evacuate as quickly as possible by any viable means. The test subjects evacuated individually and were instructed that there was an emergency in the building however, they were not instructed to follow the emergency exit signs. The starting location of the subjects is shown on the diagram as point P. Once a subject reached an exit (D1, D2, D4 or D5) the experiment would end and, the test subject having passed three emergency exit signs on their way to an exit. Based on this experimental work the probability of detecting a sign and the probability of accepting the directional information provided by the sign was determined.

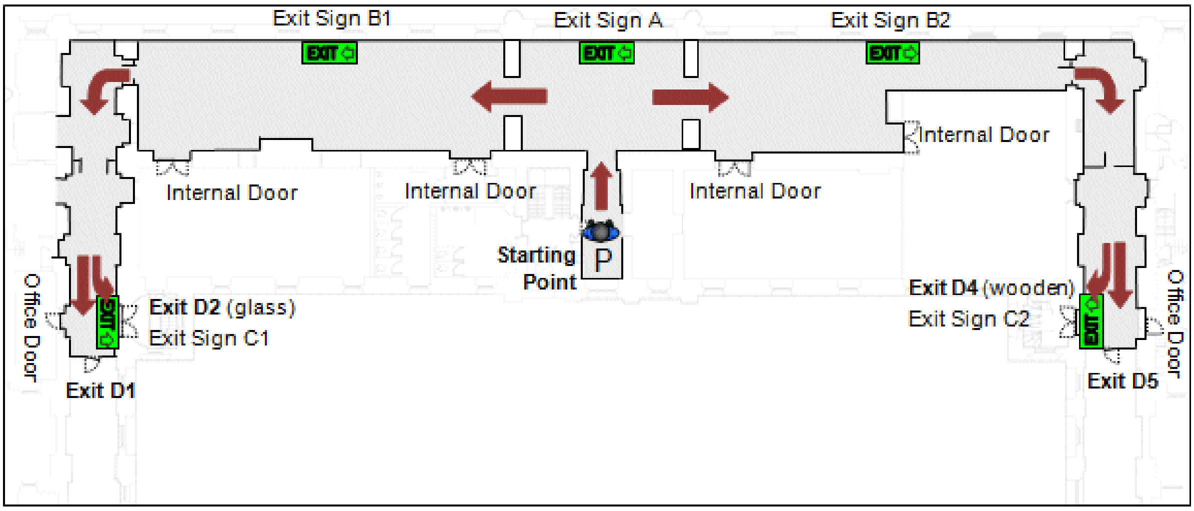

Figure 8-18: The experimental area that was used to determine the effectiveness of signage

The following pseudo-code represents the *Empirical Model* and determines whether an agent will observe and eventually follow the instructions provided by the sign or not. The probabilities that are used in statement (b) have been determined arbitrary and are the same as that used in the *Predictive Model* statement (b). However, the probabilities that are used in statements (c) and (d) have been derived from the experimental work [126]. This approach is recommended, given that it is based on experimental data.

Table 8-4: The pseudo-code for the *Empirical Approach*

a) IF Person is in a VCA THEN
b) IF Person will see the sign (visibility probability as a function of angle of approach) THEN
c) IF Person is attentive (30% to 38% cumulative detection probability for this to be true) THEN
d) IF Person will act on information received from sign (94% to 97% compliance probability to be true) THEN
e) Person learns new exit and uses it, if appropriate, OR Person follows sign's instructions

As with the *Predictive Model*, the algorithm is activated when an agent is within the VCA of a sign. The algorithm establishes whether the agent is physically able to see the sign based on the angle of approach towards the sign (*Visibility probability*). This is determined by either the sigmoid function (as in the *Predictive Approach*, see Section 8.3.4) or by a user defined arbitrary value between 0% and 100%. The algorithm then checks whether the person detects the sign. In contrast with the *Predictive Model* this is determined cumulatively across the projected path that the agent will make, ignoring the influence of the sign, within the VCA. This means that the *detection probability* is capped at the value assigned by the model, which by default is between 30% and 38% for each agent. Finally, the algorithm determines whether the agent acts on the information, given that they have received it and is determined by a *compliance probability*. The experimental study indicated a variance between the *compliance probability* for the first sign that an individual encounters and the *compliance probability* for all subsequently encountered signs. The *compliance probability* range for the first sign was calculated from 94% to 97% whereas for all subsequent signs it was calculated at 100%. In the *Empirical Model* therefore, the *compliance probability* is set by default between 94% and 97% for the first sign an agent encounters and at 100% for all other encountered signs.

If all of the criteria listed in Table 8-4 are met, then the agent is deemed to have observed the sign and chosen to act on the information received from the sign. Figure 8-19 (b) depicts a representation of this algorithm in a flow chart form.

8.3.3 User Defined Model

Within the model the user is also able to select the *User Defined* option. In terms of the algorithm used, it is based on the *Empirical Model* however, it allows the user to access and specify user defined values for probabilities used in this approach, i.e. the *detection* and *compliance probability*. The user can therefore specify the *detection* range by specifying different minimum and maximum probabilities. The user can also specify the *compliance probability* upon receiving information from the signs by specifying different minimum and maximum probabilities for the first observed sign and a single probability for all subsequently observed signs. Using this approach the user is therefore able to customise the method used, especially where the probabilities used are not supported by empirical data.

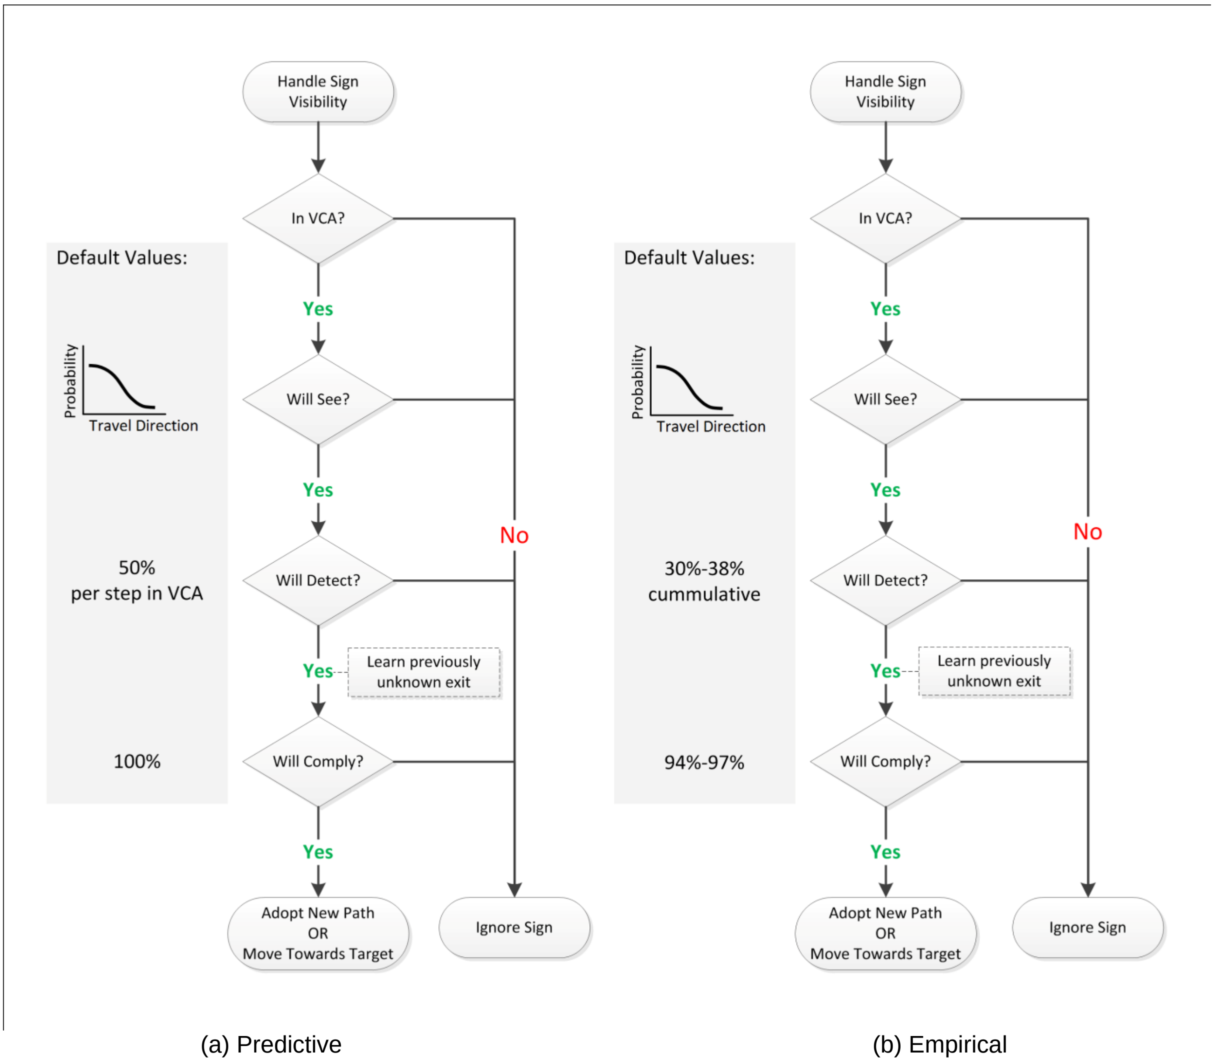

Figure 8-19: Representation of the algorithms that determine whether an agent will act on information received from a sign or not.

The various behaviours above are now described.

8.3.4 Detecting the Sign – Visibility Probability as a function of angle of approach

While a VCA defines the region from where it is physically possible to see a sign, it does not take into consideration the physical orientation of the person relative to the sign. Thus even if a person is within the VCA of a sign it does not mean that they will be physically able to see the sign, for example this would be the case if their back was turned towards the sign. Thus while occupants located outside the VCA will definitely not be able to physically see the sign, occupants located within the VCA will have a probability that they will be able to physically see the sign based on the direction they are facing. This is represented within the model by resolving the relative orientation of both the sign and the agent.

The relative orientation between the agent and the sign is examined only when the agent is within the VCA of the sign. Once this is the case the angle between the agent's travel vector and the location of the sign is determined. A visibility probability is then introduced that is dependent on the relative angle between the location of the sign and the agent's travel vector. A hypothetical relationship linking the visibility probability and the orientation angle is proposed (see Figure 8-20). It is assumed that this relationship is non-linear and that the smaller the angle between the agent's travel vector and the sign's location the higher the probability that the occupant will be able to see the sign. Furthermore, when the agent is directly facing the sign they will certainly be physically able to see the sign and when they have their backs to the sign they will certainly not be physically able to see the sign without turning around. In addition, due to peripheral vision and the likelihood that the agents will move their heads or eyes slightly as they walk, the probability of detection is small but non zero at an angle of 90^{o} and diminishes to zero at 180^{o}. This hypothetical probability distribution is displayed in Figure 8-20 (a) where the orientation angle varies from 0^{o} for a head on approach to 180^{o} for when the agent has their back to the sign and is walking away. The hypothetical probability distribution displayed in Figure 8-20(a) represents a step-wise approximation of the sigmoid function mentioned in the previous section and shown on the flow charts of Figure 8-19 (a) and (b).

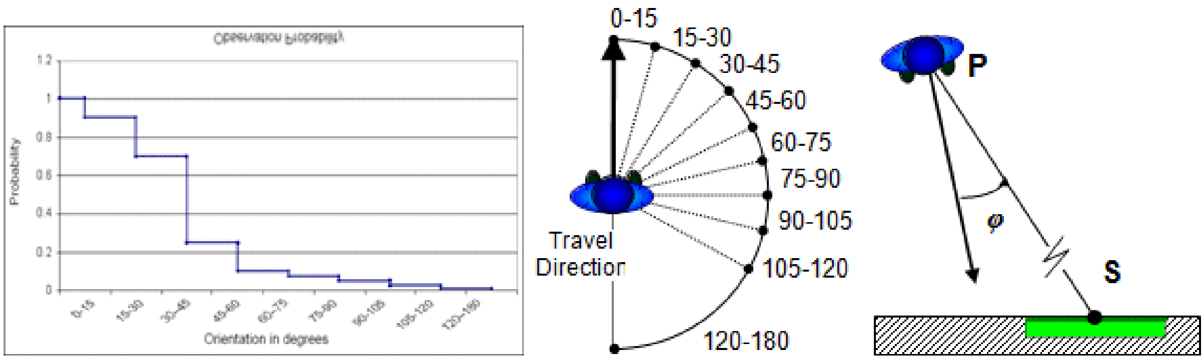

(a) (b) (c)

Figure 8-20 a) Function describing the probability of seeing the sign, b) Space around an agent is split into angular zones, c) Angle φ between direction of travel and sign object

Both systems that handle visibility within the model described in Sections 8.2.1.1 and 8.2.1.2 use the methodology described here to determine the probability of an occupant being physically able to see a sign while within the VCA. The user can override this by specifying a single probability of seeing the sign within the range of 0% to 100%.

8.3.5 Agent behaviour on detecting a sign – Detection and Compliance Probability

The agent must detect (i.e. recognise and correctly interpret the information provided by the sign) and then determine whether they act upon the information provided [103,126]. In reality, this is an extremely complicated process sensitive to physical, psychological and social factors. In the model, detection and interpretation of the sign by the agent is determined by the *detection probability*. In the *Predictive Approach* whether or not an individual recognises and interprets the sign is expressed as a probability which is tested every time an agent moves to a new node within the VCA of the sign (i.e. *detection* is predicted). In the *Empirical Approach*, irrespective of the size of the sign's VCA, the likelihood that the agent will detect the sign is set to a value within a range determined by the original experimental results. In the *User Defined* Approach the likelihood of the agent detecting the sign is set to value within a user defined range and is applied irrespective of the size of the sign's VCA.

If the agent has recognised and correctly interpreted the sign's message (i.e. they are deemed to have detected the sign) it is necessary for them to decide whether to act upon the information received. In the model it is assumed that on correctly interpreting the sign and on deciding to follow the direction of the sign the agent will move towards the target to which it points. Depending on the system used (*Signage* or *Agent Driven behaviour*) this target can be a sign, another exit or an arbitrary location within the structure. In *Signage Driven Behaviour* the agents will follow the signs even though they may guide them to use non optimal paths. If an agent is located in an overlap region of two or more VCAs then that agent will follow the instructions of the first observed and registered sign as determined by the *Predictive* or *Empirical* model. In *Agent Driven Behaviour* the agents will follow the signage provided that it leads them towards the location they want to visit. Agents with pre-assigned tasks (not related to signage) or target exits will ignore all signage information. If an agent is located in an overlap region of two or more VCAs corresponding to exit signs or other signs located within the geometry they will follow only those instructions that lead them to a location they want to visit.

8.3.6 Occupant Behaviour while moving between signs

If an agent using the signage system observes and follows the instructions of an N-order sign they will progress by moving towards the (N-1)-order sign (see Figure 8-21), and so on until the agent reaches the zero-order sign and, therefore, their destination.

Figure 8-21 Signage chain leading to the target location or exit indicated by the 0 order sign.

In a successful signage system the agent will be directed from sign to sign until the target location is attained. However, not all signage systems are well-designed and agents relying on a signage system may lose track of the next sign; i.e. not fall into the associated VCA of the next sign in the system. If the agent fails to detect the next sign some or all of the following non optimal behaviours may be observed; *Searching Behaviour*, *Backtracking Behaviour*, and *Fail Safe Behaviour* [104, 153]. These behaviours are briefly explained in turn. These behaviours can be exhibited when both the *Signage Driven Behaviour* and *Agent Driven Behaviour* is used.

While no precise data is available on actual occupant behaviour during signage assisted navigation, it is hypothesised that these behaviours can be exhibited at various stages during navigation towards their target locations (see Figure 8-22) [153]. Within the model these behaviours are exhibited as a function of the agents' expected travel distance to observe the *next* lower order sign. It is assumed that the agents expect the signs to be placed at regular intervals within the structure. They then make an assessment of their progress in relation to these distances; i.e. their expectations.

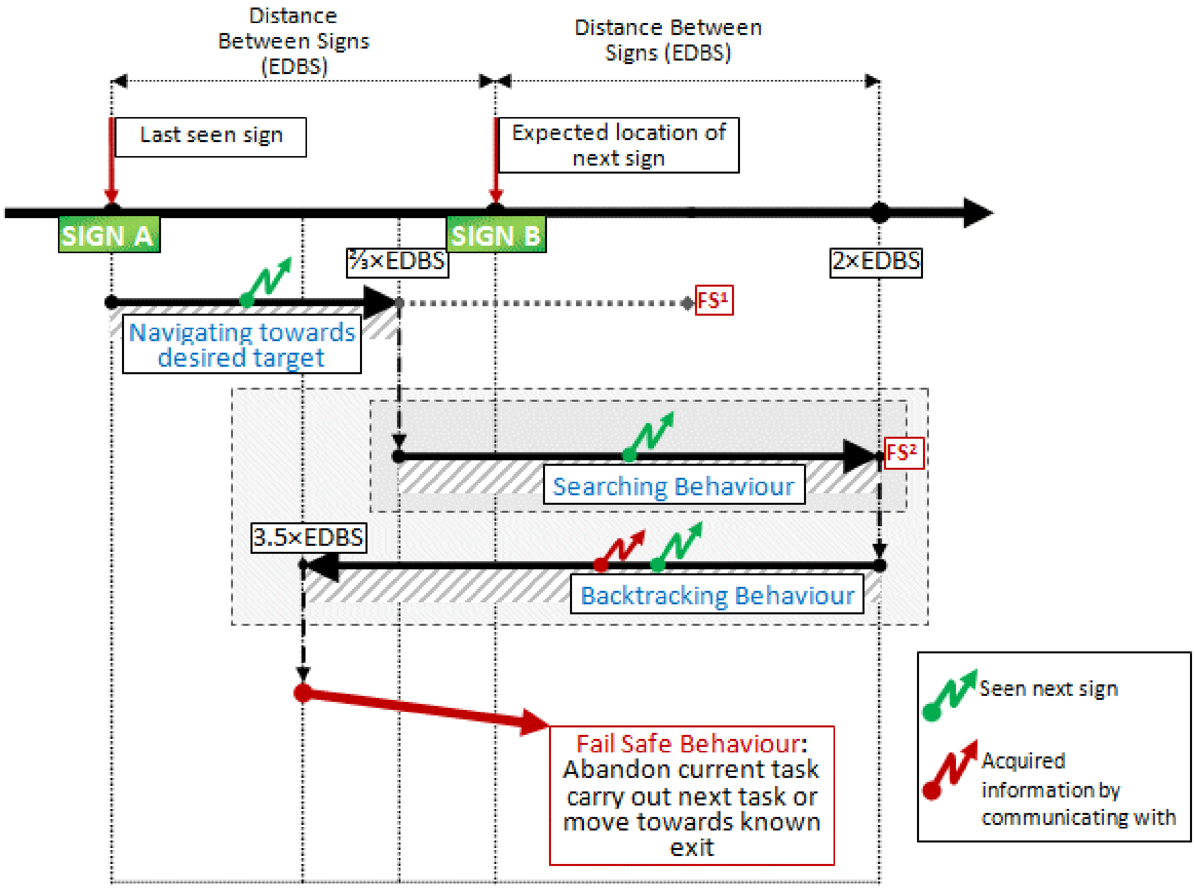

Figure 8-22: The timeline describing recovery behaviours

Within the model the *Expected Distance Between Signs (EDBS)* represents the distance within which an agent is expected to see the next sign. By default this value is initially set to 30m. The initial value of the *EDBS* can be modified by the user. As the simulation progresses and the agents observe new signs the value of the *EDBS* adapts to the agent's experiences. Each time an agent observes a sign that is part of a chain of signs they re-evaluate their expected distance between signs. The following expression shows how the *EDBS* value is re-evaluated for each agent that has encountered a new sign:

$$EDBS = 0.9 * CurrentEDBS + 0.1 * DLE \qquad \textbf{(36)}$$

Where *CurrentEDBS* is a function of the travelled distance between the agent's initial location and first observed sign (distance d1 in Figure 8-23) plus the travelled distances between all subsequently observed signs (distances d2 to d5). This excludes the travelled distance from the last two seen signs, represented by d6 in Figure 8-23. The *DLE* (*Distance Last Experienced*) value represents the distance an agent travelled between the last two seen signs from which the agent gained information.

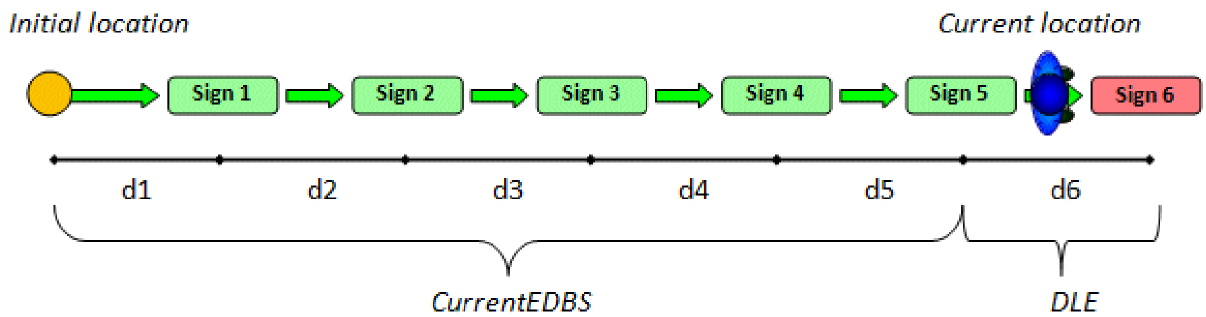

Figure 8-23: The expected distance between signs is calculated as a weighted average of CurrentEDBS and DLE

The EDBS value is renewed each time an agent observes a new sign. As a consequence an agent's expectation can change as that agent observes new signs. The shorter the distances at which an agent observes the signs the higher the expectation will be that the next sign will be observed at a relatively short distance. This formula produces a weighted average of the current expectation and experience. This method allows the *Searching*, *Backtracking* and *Fail-safe* behaviours to take place at distances representative of the sparseness of the experienced signs in the signage system.

The *DLE* value does not represent only the distance between the last two seen signs but it can also include the distance travelled while backtracking. The *DLE* value is influenced by the source of information used:

1. If the agent receives information from the targeted sign or from another agent then *DLE* excludes any backtracking travel distance and represents the distance between the last two seen signs
2. If the agent receives information from a sign other than the targeted sign then *DLE* may include backtracking travel distance if applicable. Therefore, any additional distance incurred by missing a sign is included in the *DLE* distance.

Searching Behaviour

An agent is assumed to commence *Searching* once they have travelled a distance that should have brought them within the VCA of the next sign. By default this value is set to 2/3xEDBS. During this behaviour the agent continues moving in the general direction of travel in accordance with the information received from the previously observed sign up to a distance of 2xEDBS. In certain circumstances an agent can engage in *Searching* behaviour prior to travelling the required distance of 2/3xEDBS. This can take place when the agent that has a *Find via Signage* task (i.e. follows the signage system using the *Agent Driven Behaviour* model) reaches the direction node before walking a distance of 2/3xEDBS. In this case the agent will engage in *Searching* behaviour early, and will move in the general direction of travel as it was prior to reaching the direction node.

Backtracking Behaviour

When an agent has commenced *Searching* and has travelled a distance of 2xEDBS from the last seen sign and failed to detect the next sign or target exit then *Backtracking* is enabled. An agent *Backtracking* will change travel direction and head back towards the location of the last known sign for a travel distance of 1.5xEDBS or to their start location. This behaviour represents a verification stage in the agent's behaviour. The agent goes back to the last location where they received information from a sign and tries again to navigate towards the target. If the user has enabled communication between agents then the *Backtracking* agent can receive information from other more knowledgeable agents within the structure regarding their target location; once new knowledge is acquired the agent will start moving in accordance with the new information.

Fail-Safe Behaviour

If an agent has reached the end of backtracking behaviour and still has found neither their target nor another relevant sign or acquired information from the surrounding population they will commence *Fail Safe* behaviour. Under these circumstances, the agent will give up the search for their target and will carry out the next task in their itinerary list or will exit the structure via the nearest available known exit if they have reached the end of their itinerary list. *Fail-safe* behaviour is exhibited at 3.5xEDBS.

The user is able to disable *Searching* or *Backtracking* behaviour (from the *Behavioural Options* dialogue box in *Simulation Mode*). If *Searching* is disabled then by implication *Backtracking* is disabled too; however, *Backtracking* can be disabled independently from *Searching*. If *Searching* is disabled and the agents fail to observe the next sign then *Fail-safe* behaviour is initiated at 1.5xEDBS (point FS^1 in Figure 8-22). If *Backtracking* is disabled and the agents failed to observe the next sign during searching behaviour then *Fail-safe* behaviour is initiated once *Searching* ends at 2xEDBS (point FS^2 in Figure 8-22).

CHAPTER 9: LIFTS/ELEVATORS

Lifts/elevators in buildingEXODUS are created using *Transit Nodes* in a two stage process. Initially a lift shaft is created using a series of identical *Lift Shaft Opening* transit nodes spanning multiple floors. Once the lift shaft is defined via a set of *Lift Shaft Opening* transit nodes, a lift and its associated properties can then be associated with the shaft.

(a) Lift Shaft

Lift Shaft Opening transit nodes span each of the floors within a geometry which define the vertical path and distance a lift moves between each floor (as defined by the floor height, see Figure 9-1). The dimensions of the shaft (i.e. length and width) on each respective floor can be defined within the model by altering the attributes of each *Lift Shaft Opening* transit node. Once a *Lift Shaft Opening* transit node has been associated with a lift, if its dimensions are changed then the dimensions of **all** the other *Lift Shaft Opening* transit nodes associated with the same lift are also changed accordingly. It is important to note that the size of the door for a given lift shaft is assumed to be equal to the width of its associated *Lift Shaft Opening* transit nodes.

It is important to note that the transit nodes do not represent the lift, but define the lift shaft that the lift moves within, hence being called *Lift Shaft Opening* transit nodes. An example of a 3 floor geometry with a single lift shaft represented using 3 *Lift Shaft Opening* transit nodes, where the lift is currently on *Floor 0* can be seen in Figure 9-1.

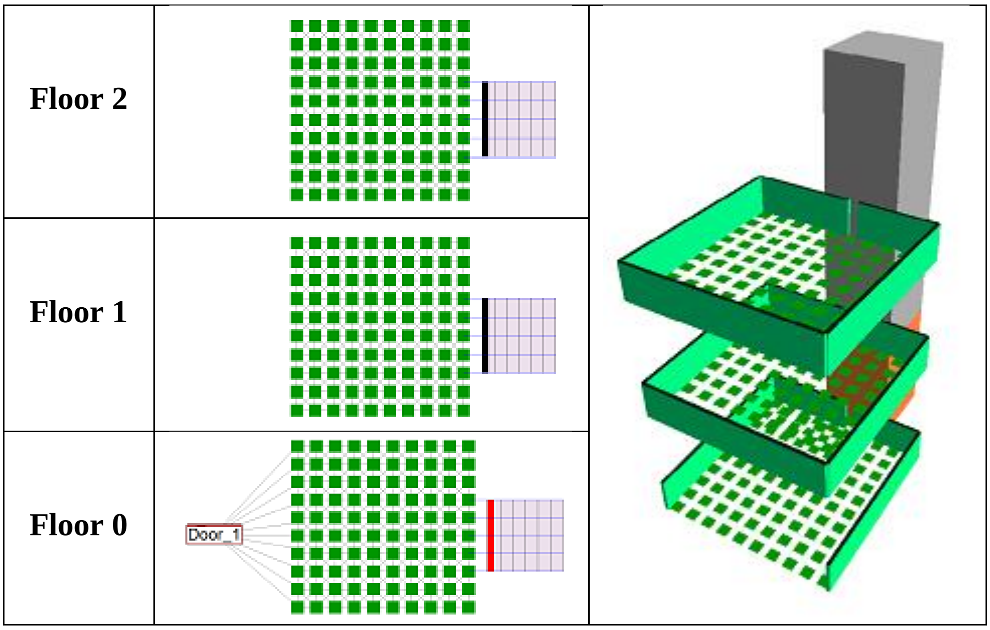

Figure 9-1: Transit node lift shaft example

(b) Lift Car

Only a single lift can be associated with each lift shaft. There are a number of attributes that define the movement, delay times and capacity of a lift within the model, these are listed in Table 9-1. The default kinematic and delay time values for a lift are set according to those recommended by CIBSE Guide D [125]. Within buildingEXODUS the default kinematics for a lift (i.e. maximum speed and acceleration) are derived from CIBSE Guide D and correspond to a lift with a maximum travel distance of 50m. Since lift deceleration rates are not provided in CIBSE Guide D, the default deceleration rate of a lift within buildingEXODUS is therefore assumed to be the same as the rate of acceleration.

Table 9-1: buildingEXODUS lift attributes and default values

Attribute	Default	Description
Title	Lift#	The name of the lift.
Current Floor	-	The *Lift Shaft Opening* transit node and *Floor* that the lift is currently at.
Start Floor	-	The *Lift Shaft Opening* transit node and *Floor* that the lift starts at the beginning of the simulation.
Status	Closed	The status of the door (i.e. *Open* or *Closed*).
Direction	-	The direction the lift is travelling in (i.e. *Up*, *Down* or *Waiting*)
Capacity	13	The maximum number of agents that can simultaneously occupy the lift.
Occupancy	-	The number of agents currently inside the lift.
Start Delay (s)	0.0	The delay time before a lift begins servicing its floor sequence at the beginning of a simulation.
Is in service	Yes	Defines if a lift is in service or not.
Max Speed (m/s)	2.5	The maximum speed the lift can travel at.
Acceleration (m/s^2)	0.8	The rate at which the lift accelerates towards the maximum speed.
Deceleration (m/s^2)	0.8	The rate at which the lift decelerates from the maximum speed.
Opening Time (s)	3.0	The time the lift door takes to open (agents cannot board a lift until the door is fully opened).
Closing Time (s)	4.0	The time the lift door takes to close (agents cannot board a lift once the doors have begun to close).
Dwell Time (s)	3.0	The time the lift doors remain open after fully opening once no agents in the transit node catchment area are targeting the lift bank.
Sensor adjusted Dwell Time (s)	0	The time the lift doors remain open after an agent boards the lift and no agents in the transit node catchment area are targeting the lift bank.
Motor Delay (s)	0.5	The time it take a lift motor to start moving the lift after the doors have fully closed.

CIBSE Guide D [125] recommends a variety of lift kinematic values according to the maximum travel distances a lift can make (see Table 9-2).

Table 9-2: CBISE Guide D [125] typical lift kinematic values

Lift travel (m)	**Maximum speed (m/s)**	**Acceleration/Deceleration (m/s^2)**
<20	<1.0	0.4
20	1.0	0.4-0.7
32	1.6	0.7-0.8
50	2.5	0.8-0.9
63	3.0	1.0
100	5.0	1.2
120	6.0	1.2
>120	>6.0	1.2

CIBSE Guide D [125] also recommends a variety of lift door opening/closing times depending on the width and type of lift door used (see Table 9-3).

Table 9-3: CBISE Guide D [125] typical lift door closing and opening times

Door type	**Closing and opening times (s) for stated door width (m)**					
	Closing		**Opening (normal)**		**Opening (advanced)**	
	0.8m	1.1m	0.8m	1.1m	0.8m	1.1m
Side	3.0	4.0	2.5	3.0	1.0	1.5
Centre	2.0	3.0	2.0	2.5	0.5	0.8

The default *Opening* and *Closing* times for lifts within buildingEXODUS therefore correspond to a conventional side opening door 1.1m in width. These recommended values maybe more appropriate to use than the default values for different height buildings with lift systems of different characteristics.

Previous prototype versions of the lift model have included the representation of Jerk (i.e. the change in acceleration/deceleration as the lift moves towards constant acceleration/deceleration) [124, 168, 169]. After careful review this was considered unnecessary to represent due to the increase in computational overhead, increased algorithmic complexity, and marginal benefit provided with regards to accuracy. Regarding the kinematic attributes, if a user is unaware of the rates of acceleration/deceleration then the model can function by only using the maximum speed (i.e. by setting the acceleration/deceleration to 0 m/s^2).

When a lift changes its state or moves, the colour of the entrance to the associated *Lift Shaft Opening* transit node changes to indicate whether the lift door is open, closed or the lift is not at the floor in question (see Figure 9-2).

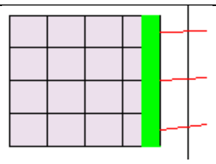	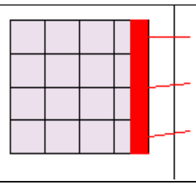	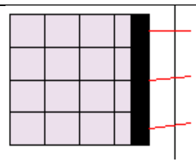
Open (Green)	Closed (Red)	Lift is not at the floor (Black)

Figure 9-2: Lift Shaft Opening transit node door status colour coding

(c) Lift Motion Controls

Lifts can be controlled by a user specifying a sequence of floors a lift will service prior to a simulation being run. This method allows the user to easily specify and alter the lift evacuation strategy by giving priority to agents on certain floors during the evacuation. The system assumes that the lifts are either autonomously controlled by a computer system or manually controlled by an operator i.e. the system is able to detect the number of people in the lift and lift waiting areas on the specified floors either by sensors or visual inspection.

There are three types of sequence that can be defined:

- *Floor*-sequence
- *Shuttle-floor* sequence
- *Sky-lobby* sequence.

Floor Sequence

Using a *Floor*-sequence, a user explicitly specifies the sequence of floors a lift will service during the simulation. The assigned lift will serve each of the floors specified in the sequence irrespective of whether agents are waiting for the lift on each of the floors.

Shuttle Floor Sequence

Using a *Shuttle-floor* sequence, a user specifies a paired-sequence of pick-up/drop-off floors where the lift could pick agents up from and shuttle them to. This shuttle process repeats until there are no more agents in the catchment area of the *Lift Shaft Opening* transit node on the pick-up floor. The process is then repeated for the next pair of pick-up/drop-off floors in the sequence (i.e. once the lifts have dropped off agents and there are no more agents waiting, the lifts will move to the next pair of pick-up/drop-off floors). If the lift is on a pick-up floor and is ready to leave (i.e. no more agents are waiting to board the lift and it is not full to capacity), it will directly move to the next pick-up floor to collect more agents (i.e. it will not drop-off agents first with a partially full lift).

The ability of lifts assigned *Shuttle-floor* sequences to return to pick-up floors after they have been cleared (i.e. once all the agents waiting in the lift catchment area on the pick-up floor have been collected) is controlled via the *Shuttle Floor Return* option (see Chapter 6 of the User Guide).

If the *Shuttle Floor Return* option is disabled then once a given pick-up floor has been cleared the lift will then move onto the next pick-up/drop off pair within the sequence. Consequently, once a given pick-up floor has been cleared the lift will therefore not return to it, even if agents subsequently turn up on that floor hoping to catch the lift. In this case, the agents arriving within the catchment area of a lift on a floor already cleared will be deemed to have missed the lift and will therefore redirect to use the stairs. This represents the influence of either dynamic signage or a communication system informing the agents that the lift has already serviced their floor and that they should use the stairs instead. This system prevents agents from waiting indefinitely for a lift that will not arrive.

However, if the *Shuttle Floor Return* option is enabled then lifts assigned *Shuttle-floor* sequences can return to previously cleared floors. As a result, agents arriving within the catchment area of a lift on a previously cleared pick-up floor will not redirect to the stairs, but will instead commence waiting for the lift in the conventional manner. The pair of pick-up/drop floors corresponding to the previously cleared floor on which an agent has arrived will then be added back into the lifts *Floor Requests* list (see Chapter 5 of the User Guide). The lift will then return to the previously cleared floor in order to collect the agents waiting for the lift. Once all the agents on that floor have then been collected the lift will then move onto the next pair of pick-up/drop off floors etc. Hence, when the *Shuttle Floor Return* option is enabled pairs of pick up/drop off floors can be both dynamically added (i.e. as agents arrive within the lift catchment areas of previously cleared floors) and removed (i.e. as all the agents waiting for the lift on a given pick-up floor are collected) from lifts corresponding *Floor Request* lists throughout the simulation. By default the *Shuttle Floor Return* option is disabled (i.e. lifts assigned *Shuttle-floor* sequences will not return to previously cleared floors).

Sky Lobby Sequence
Using a *Sky-lobby* sequence, a user only specifies a pair of floors defining a single pick-up (sky-lobby) and drop-off floor where lifts pick agents up from and shuttle them to. Unlike the *Shuttle floor*-sequence system outlined previously the lift will always return to the pick-up floor to collect more agents irrespective of whether any are waiting.

Whether agents seek to use a lift assigned a *Sky-lobby* sequence is dependent upon the agent's initial floor, the pick-up floor of the lift and the number of floors the agent is prepared to travel up in order reach the lift. The number of floors an agent is prepared to travel up in order to reach a sky lobby lift is controlled via the *Sky Lobby Up Floors* variable (see Chapter 6 of the User Guide). By default the *Sky Lobby Up Floors* variable is set to zero, thereby implying that agents are not prepared to move up in order to reach a sky lobby lift. As a result, an agent will only move towards a lift assigned a *Sky-lobby* sequence if the agent's initial floor is the same as the lift's pick-up floor, or the floor on which the sky lobby is located (i.e. its pick-up floor) falls between the agent's initial floor and their exit (i.e. if an agent's exit is located below them they will only consider moving down to a sky lobby, not up).

If the value of the *Sky Lobby Up Floors* variable is greater than zero then agents can additionally travel up to a sky lobby lift. However agents can only do this if the number of floors they are required to traverse in or to reach the sky lobby lift is less than or equal to the value of the *Sky Lobby Up Floors* variable. In addition, agents will still only travel up to the sky lobby lift if the number of floors that they are required to traverse is less than the number of floors they would traverse in travelling down to either another sky lobby lift or an exit (i.e. agents will not travel up if there is an alternative down option that is closer (or as close) in terms of the number of floors they are required to traverse).

It is important to note that if no floor sequence is defined for a given lift then it will be assumed not to represent a viable egress route during a lift evacuation. Consequently agents will not be automatically assigned itinerary tasks to use the lift as part of a lift evacuation. However, agents who have been manually assigned to the lift (i.e. via the user manually assigning them *Lift Bank* tasks within their itinerary) can still use the lift if the *Lift Bank* task is followed by subsequent tasks within the agent's itinerary (i.e. if the *Lift Bank* task is not the agent's last task). In these instances, when the agent enters the waiting area (i.e. catchment area) they will consider the floors serviced by the lift. If the target of their next task (i.e. the task immediately following the *Lift Bank* task) is on a floor directly serviced by the lift (i.e. if a *Lift Shaft Opening* transit node on the target floor is defined within the lifts *Nodes* tab, see the User Guide, Chapter 5) then the agent will deem that the lift can take them to where they need to go. In this case, the agent will then be assumed to call the lift. The lift dispatch algorithm will then register the request and set about sending a lift from the designated lift bank to the appropriate floor. In this manner, not defining a floor sequence for a given lift means that the lift will not be used as part of a lift evacuation, but that it can instead be used to dynamically model the circulation of agents within the structure.

(d) Lift Door Controls

Once a lift has opened its doors on a floor, the doors will remain open for the specified *Dwell Time* before closing. However, whilst an agent is within a lift waiting area and is targeting one of the lifts in the lift bank, all open lifts that are not full to capacity will keep their doors open and not begin their *Dwell Time*. This is in anticipation that another agent might board the lift. This assumes that the lift car doors are kept open either by people already inside the lift or staff controlling the lift. Once there are no more agents left targeting a lift in the bank or the lift is full to capacity the *Dwell Time* begins. This means that lifts have an increased chance of being filled to their maximum capacity when collecting agents.

(e) Lift Kinematics

The kinematics of a lift is defined by its acceleration, deceleration and maximum velocity. There also are a number of other factors that influence lift's actual journey such as friction, 'wear and tear' of machine parts, etc. However, these are not considered within the current lift model. As such, the kinematics of the lift model should be considered "ideal". Specifying the cars acceleration, deceleration and maximum velocity is sufficient to determine the location of a lift at any point during its journey. The time at which the lift passes each respective floor between its original location and its destination is determined using a series of formulae.

Considering a lift's acceleration, deceleration and maximum velocity, there are three types of journey a lift can make:

A) a lift reaches maximum velocity,
B) a lift just reaches maximum velocity but immediately decelerates, or
C) a lift does not reach maximum velocity.

Figure 9-3 below shows an example of each of the three different types of journey that can be made by a lift in terms of velocity and time.

Condition A: Lift reaches maximum velocity	**Condition B: Lift just reaches maximum velocity**	**Condition C: Lift does not reach maximum velocity**
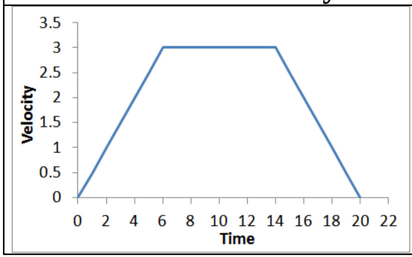	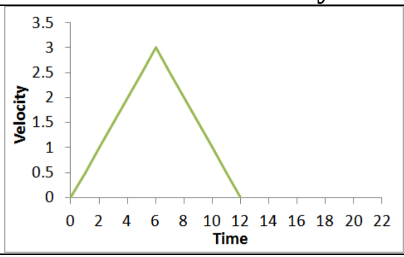	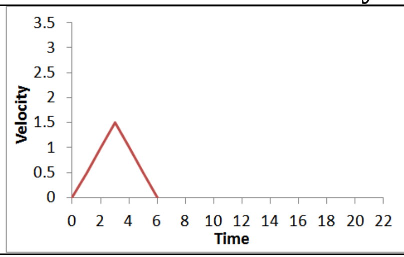

Figure 9-3: Graph showing the velocity against time for the three types of journey a lift can make (Acceleration = 0.5m/s^{2,} Deceleration=0.5m/s^2, Max velocity = 3.0m/s)

For each journey type there are a series of key time points when a kinematic variable changes according to the required travel distance. These are the time a lift reaches its maximum velocity (if it does reach its maximum velocity) and the time a lift needs to decelerate before stopping. These time points along with associated travel distances are used to determine which formula can be used to calculate the time a lift arrives/passes a given floor. These are also used to display the lift on each floor within a simulation and so represent the animation of a lift moving. This is achieved by changing the state of the *Lift Shaft Opening* transit node as shown in Figure 9-2. The following sections describe the criteria for deciding which journey type a lift will make in addition to the associated formula for calculating the time points (i.e. the time a lift passes each floor on its journey (t_i)).

Specified kinematic parameters

V = maximum velocity (m/s)
a = rate of acceleration (m/s^2)
d = rate of deceleration (m/s^2)

Defined distance parameter

D = total travel distance (m)
D_i = total travel distance to floor ‘i’ (m)

Calculated distance parameters

$D_a = \frac{V^2}{2a}$ = total distance travelled whilst accelerating to maximum velocity (m)
$D_d = \frac{V^2}{2d}$ = total distance travelled whilst decelerating from maximum velocity (m)
$D_v = D - (D_a + D_d)$ = total distance travelled whilst travelling at maximum velocity (m)

Calculated time parameters

$t_a = \sqrt{\frac{(2*D_a}{a}}$ = total time to accelerate to maximum velocity (s)
$t_d = \sqrt{\frac{(2*D_d}{a}}$ = total time to decelerate from maximum velocity (s)
t_i = total time to travel to floor ‘i’ (s)

Travel time formula according to velocity phase for each journey type

Condition A: Lift reaches maximum velocity
For a lift to travel at its maximum velocity then the total travel distance (D) must be greater than the combined travel distance of the acceleration phase (D_a) and deceleration phase (D_d).

Condition statement: $D > D_a + D_d$

Travel distance condition	Velocity Phase	Travel Time Formula
$D_i \le D_a$	Acceleration	$t_i = \sqrt{\frac{(2*D_i}{a}}$
$D_i > D_a$ and $D_i \le (D - D_d)$	Max velocity	$t_i = \sqrt{\frac{(2*D_a}{a}} + \frac{D_i - D_a}{V}$
$D_i > D - D_d$	Deceleration	$t_i = \sqrt{\frac{(2*D_i}{a}} + \frac{D-(D_a+D_d)}{V} + \frac{V - \sqrt{(V^2 - (2d*(D_i - D_a - D_v))}}{d}$

Condition B: Lift only just reaches maximum velocity but immediately decelerates
For a lift to just reach (but not maintain) its maximum velocity then the total travel distance (D) must be equal to the combined travel distance of the acceleration phase (D_a) and deceleration phase (D_d).

Condition statement: $D = D_a + D_d$

Travel distance condition	Velocity Phase	Travel Time Formula
$D_i \le D_a$	Acceleration	$t_i = \sqrt{\frac{(2*D_i}{a}}$
$D_i > D_a$	Deceleration	$t_i = \sqrt{\frac{(2*D_a}{a}} + \frac{V - \sqrt{(V^2 - (2d*(D_i - D_a))}}{d}$

Condition C: Lift does not reach maximum velocity
For a lift to not reach its maximum velocity then the total travel distance (D) must be less than the combined travel distance of the acceleration phase (D_a) and deceleration phase (D_d).

Condition statement: $D < D_a + D_d$

$V = \sqrt{\frac{D}{\frac{1}{2a}+\frac{1}{2d}}}$ = maximum velocity reached (m/s)

$D_a = \frac{1}{2a} * \frac{D}{\frac{1}{2a}+\frac{1}{2d}} = \frac{D}{1+\frac{2a}{2d}}$ = total distance travelled whilst accelerating to maximum reached velocity (m)

Travel distance condition	Velocity Phase	Travel Time Formula
$D_i \leq D_a$	Acceleration	$t_i = \sqrt{\frac{(2*D_i}{a}}$
$D_i > D_a$	Deceleration	$t_i = \frac{1}{a} * \sqrt{\frac{D}{\frac{1}{2a}+\frac{1}{2d}}} + \frac{V - \sqrt{(V^2 - (2d*(D_i - D_a))}}{d}$

(f) Lift Dispatch Algorithm

In the case when a lift has no defined floor sequence it is assumed the lift can only be used for circulation (i.e. not used for evacuation). As a result, agents will not be automatically assigned to use the lift when a lift evacuation is being performed and therefore the lift with no pre-defined floor sequence will be ignored. However, agents who have user defined itineraries can use such lifts if the task within their itinerary instructing them to use the lift is followed by other additional tasks (i.e. if the task instructing the agent to use the lift is not the last task within their itinerary, see User Guide, Chapter 4). In this instance, upon entering the waiting area (i.e. catchment area) of a given lift bank the agent will consider three things, namely:

1) Are lifts within the current lift bank answering service calls?
2) Can a lift within the current lift bank take the agent to their desired destination?
3) Is the congestion within the lift bank waiting area less than the agent's defined *Congestion Threshold*?

When considering whether lifts within the current lift bank can take the agent to where they want to go the agent considers the location that the next task within their itinerary instructs them to move to (i.e. the task immediately following the task telling them to use the lift). If the floor on which the target of their next task is located is serviced by lifts within the current lift bank then the agent will deem that a lift will take them to where they wish to go. If any of the three above criteria are not met (i.e. if lifts are not in service, do not take the agent to where they want to go or the congestion within the lift bank waiting area exceeds their defined congestion threshold) then the agent will abandon their use of the lift and redirect to the stairs as an alternative means of achieving their next task.

If all of the three above criteria are met then the agent will next consider whether they need to call a lift. If a lift within the lift bank is currently on the agent's floor, its doors are open and it is travelling in the required direction (i.e. up or down) then the agent will simply board the lift (i.e. no lift request is required). If however no lifts within the lift bank meet these criteria then the agent will instead call a lift. It is important to note that if an agent calls a lift they are assumed to automatically do this upon entering the lift banks corresponding catchment area. Hence the movement of the agent to the lift control panel and the subsequent pressing of the request button is not explicitly modelled.

Upon requesting a lift both the floor on which the request was made and the required direction of travel (i.e. up or down) are sent to the lift dispatch algorithm. It is important to note that the agent's target floor is not sent to lift dispatch algorithm as part of the initial lift request. Once the request has been made the lift dispatch algorithm then sets about determining which lift within the current lift bank is best placed to service the call. To achieve this, the boarding floor and direction of travel are inserted into the current floor request sequence of each lift. An

Estimated Time to Destination (ETD) algorithm is then used to determine the time for each lift to service the boarding request.

The estimated time taken for each lift to service the boarding request (i.e. ETD_i) is based upon the time taken for the lift to travel from its current location to each of the floors within its floor request sequence up to and including the boarding floor. In each case the time taken for the lift to move between the floors in its floor request sequence is based upon the corresponding distance between the floors and the kinematics of the lift (i.e. acceleration, deceleration, max speed etc.). Similarly, the time spent at each interim floor is also taken into consideration. Hence at each floor the time for the lift to open its doors, dwell, close its doors and the time taken for the motor to start (i.e. motor delay) are all taken into consideration. The estimated time for a lift to service a specific floor is calculated based on the following equation:

$$ETD_i = T_i(C_i, R_{i,1}) + L_{i,o} + \sum_{j=1}^{k-1} (T_i(R_{i,j}, R_{i,j+1}) + L_{i,o} + L_{i,d} + L_{i,c} + L_{i,md})$$

Where:

ETD_i = the estimated time for lift i to service the agent's boarding request.
R_i = the set of *n* floor requests for lift i
k = the position in the set R_i of the agent's boarding floor
$L_{i,o}$ = the opening time of lift i
$L_{i,c}$ = the closing time of lift i
$L_{i,d}$ = the dwell time of lift i
$L_{i,md}$ = the motor delay of lift i
$T_i(x,y)$ = the time for lift i to travel between floors x and y
C_i = the current floor of lift i

The lift which services the boarding request in the shortest possible time is then assigned the boarding floor request. In this manner, a single boarding request results in only a single lift being dispatched, even in lift banks comprising two or more lifts.

It is important to note that the occupancy level of the lift (i.e. the number of agents in the lift) is not taken into consideration when determining which lift to dispatch to a given boarding request. This is because it is assumed that the lift itself (and hence the lift dispatch algorithm) is not aware of either the current occupancy of the lift or the exact number of people who intend to board or alight on any given floor. As a result, having requested a lift, agents waiting for it may find that it turns up on their floor fully loaded, and hence they are unable to board it. In these cases, having been unable to board the lift the agents will issue another boarding request, thereby resulting in the lift dispatch algorithm allocating another lift to service the call.

Having issued a boarding request the agent will then set about waiting for a lift within the lift bank waiting area. If the amount of time spent waiting for the lift by the agent exceeds their designated *Wait Time* then the agent will abandon their use of the lift and redirect to the stairs as an alternative means of achieving their next task. It is important to note that if an agent abandons their use of a lift as a result of excessive waiting then their boarding request still remains. Hence the lift allocated to service the agent's boarding request will still stop at the floor (i.e. since the lift is unaware that the agent no longer wishes to board).

If a lift stops on the agent's floor travelling in the required direction (i.e. up or down) then the agent will attempt to board it. It is important to note that agents will not board lifts travelling in the opposite direction to that in which they want to travel. Assuming the agent can board the lift (i.e. assuming that the lift is not full to capacity) then upon entering the lift the agent is assumed to issue an alighting floor request (i.e. the agent tells the lift which floor they want to get off at). It is important to note that alighting requests are assumed to occur automatically upon the agent entering the lift. Hence the movement of the agent to the control panel within the lift and the subsequent pressing of the alighting floor button is not explicitly modelled. Furthermore, the behaviour of the agent's that are within a lift is not explicitly modelled. As a result no movement calculations or other behavioural aspects of an agent's existence within the lift is considered. Once the alighting request has been made by the agent their corresponding alighting floor is automatically inserted into the lift's floor request sequence. Agents will then remain within the lift until it reaches their designated alighting floor, where upon they will alight and then set about achieving their next task within their itinerary.

Using this approach the commonly accepted rules for lift behaviour defined by Closs [170] are enforced:

1) A lift car may not stop at a floor where no passenger enters or exits (i.e. a lift car may not stop at a floor unless it has been specifically requested to do so by a passenger),
2) A lift car may not pass a floor at which a passenger wishes to exit,
3) A passenger may not enter a lift car carrying passengers and travelling in the reverse direction to his required direction of travel,
4) A lift car may not reverse its direction of travel while carrying passengers.

CHAPTER 10: VALIDATION AND VERIFICATION

Validation is an essential step in the continual development and acceptance of evacuation modelling. While no degree of successful validation will prove an evacuation model correct, confidence in the technique is established the more frequently it is shown to be successful in as wide a range of applications as possible.

While the term "validation" is often used its meaning is often misinterpreted. Here we take validation to mean the systematic comparison of model predictions with ***reliable*** information. The information used for validation purposes may comprise experimental data, numerical data, or experiential insight or a combination of these sources. Depending on the nature of the data, the validation may comprise [36]:

- **(i) Component testing:** routine checking of major software sub-components
- **(ii) Functional validation:** checking model capabilities and inherent assumptions are compatible with intended use
- **(iii) Qualitative verification:** compare predicted human behaviour with informed expectations
- **(iv) Quantitative verification:** detailed comparison of model predictions with reliable experimental data

As the number and variety of evacuation models increase it becomes essential to provide a discriminating basis of comparison. Success at a wide range of standard 'validation' exercises provides one means to this end. However, to date, little effort has been invested in the systematic comparison of various evacuation models with common experimental data.

The lack of a convincing quantitative validation history is due for the most part to the scarcity of suitable experimental benchmark evacuation data. The majority of evacuation trials are not conducted for model validation purposes but to demonstrate the suitability of a building/ design/ staff procedures and training or to gauge compliance to a regulation or standard. In most of these cases insufficient data is recorded to allow a detailed "validation" exercise.

The "ideal data set" would require staging an evacuation exercise using a representative target population within the structure subjected to as realistic a scenario as possible. Both the structure and the population must be characterised in full. This means providing a full description of the structure including all relevant dimensions, presence and nature of internal obstructions such as seating, nature of the alarm system and a description of the evacuation procedures employed. A profile of the participating population – including a description of age, gender, level of physical ability, familiarity with the structure and procedures, etc must also be provided. The starting locations of each of the individuals must be known as well as their paths to exit. The development of pinch points and crowding should be recorded, in particular, when and where this occurred and how long it lasted. It should be possible to derive the time history for each individual including their response time and their time to exit. In addition to the number of people to use an exit and the total evacuation time, the flow time and time dependent flow rate for each exit should be recorded. The participants should also be followed up with an appropriate questionnaire designed to reveal their experiences through the evacuation. Several evacuation trials should also be performed, ideally using different individuals. Information missing from the data set will result in the necessity of the modeller to make assumptions and estimations that can adversely influence the fidelity of the model predictions.

NOTE:
Ideal validation data as described above has been generated by FSEG as part of the EU FP7 projects BeSeCu [154] for building evacuations and also SAFEGUARD [155] for passenger ship evacuation. These data sets were being used to test both buildingEXODUS and maritimeEXODUS. The results show very good agreement between the validation data and the EXODUS predictions.

The EXODUS suite of software has undergone - to varying degrees - and continues to undergo each of the four forms of validation specified above. This has involved direct comparison of model predictions with historic experimental data, comparisons of "blind" model predictions with experimental data and comparing the nature of predicted human behaviour with expectations. In the following section, the attempts of the EXODUS developers to follow the validation protocol outlined above are briefly described. Interested readers are referred to documents [5,36–42,48-50,55,71,80] for further details.

10.1 Component Testing of EXODUS

The EXODUS development group employ standard software engineering principles with regard to the development and testing of software. Furthermore, once a new version of the software is produced it is alpha tested internally by the EXODUS applications group and beta tested by the EXODUS distributor network and advanced users prior to general release. The alpha testing involves checking that the user interface, movement, behaviour, hazard and toxicity sub-models within EXODUS are functioning as expected (see for example reference [15]). Furthermore, cases run using previous versions are repeated with the latest version to ensure that similar results are generated [71]. The beta testing involves a similar exercise. In addition, beta testers make use of the software on "live" applications, providing field testing for the software.

10.2 Functional Validation of EXODUS

Functional validation of EXODUS has essentially been performed via two routes. Firstly, numerous peer-reviewed publications have been produced describing the details and operation of the EXODUS software (for example [1-3,5-12,38,39,50,80]). In addition, the EXODUS software has been presented and discussed at many academic conferences concerning fire safety engineering. Both of these approaches have opened the principles of the EXODUS software for comment by the professional and academic fire engineering and fire science communities.

The second route concerns the publication of the EXODUS User Guide and Technical Manual. A version of the manual exists for the buildingEXODUS and maritimeEXODUS software products. The technical guide is a full account of the capabilities and limitations of the software. It provides a comprehensive account of the assumptions employed with references to 47 related publications. In addition a detailed testing document has been produced that details the testing process and test cases that EXODUS undergoes through its development [71].

10.3 Qualitative Verification of EXODUS

The third form of model validation concerns comparing the nature of predicted human behaviour with informed expectations or trends in behaviour noted from actual experimental results. A number of such cases have been produced using both the airEXODUS (for example, see references [10,15,39,42]) and buildingEXODUS (for example, see references [1,3,37,38,41,48,50]) software.

A form of independent qualitative validation of the airEXODUS evacuation model has been undertaken by NIST [42]. This work involved the study of a range of sensitivity studies based on airEXODUS predictions of evacuation from an aircraft geometry.

Examples of qualitative validation of the buildingEXODUS evacuation model are provided in the simulation studies of the Milburn House evacuation [37] and the Tsukuba pavilion evacuation [38]. These simulations demonstrate that plausible model predictions are produced in these applications that are in line with the trends observed in the associated experimental evacuations.

Another example is provided by the application of buildingEXODUS V2.0 to the simulation of evacuation from a hypothetical supermarket/restaurant complex (see the Application Manual, Chapter 8). In particular the behaviour of the population during the evacuation is noted. The ability of the buildingEXODUS behaviour sub-model to display the type of behaviours described in these examples demonstrates the qualitative capability of the model.

Two more recent examples include the simulation of the Rhode Island Disco fire [156] and the 2001 World Trade Centre evacuation [134] using buildingEXODUS V4.07. In both these examples buildingEXODUS was used to simulate real world incidents. The first case included coupling to the SMARTFIRE CFD fire simulation of the Rhode Island fire [156]. In this case buildingEXODUS was shown to be able to produce a realistic reconstruction of the actual incident, including predicting the approximate number of fatalities. In the second case, buildingEXODUS was used to reconstruct the evacuation of the North Tower [134]. The buildingEXODUS simulation produced a realistic prediction of the evacuation time for the North Tower and was also used to examine the impact of the fire fighters ascending the stairs of the North Tower on the evolving evacuation and to predict the outcome of a fully occupied North Tower.

10.4 Quantitative Verification of EXODUS

Wherever possible, EXODUS predictions have been quantitatively compared with data generated from *REPEATED* experimental trials. This has been done in an attempt to account for the variability in human behaviour. With a few exceptions, this data has been generated from research in the aviation industry. Comparisons of this type can be viewed as providing some justification for the buildingEXODUS methodology.

As a first attempt at this type of validation, a pre-release version of the airEXODUS model was compared with a selection of the Cranfield Trident Three experiments [21]. In this series of experiments, competitive evacuations were performed from a Trident Three aircraft cabin section, comprising 12 rows of seats organised six abreast and parted by a single aisle. The airEXODUS model was found to correctly predict the trends found in evacuation times [10,15].

More recently, predictions produced by airEXODUS were compared with a series of evacuation trials involving a B-737 mock-up [43]. The results shown in Figure 10-1 represent a subset of these trials. The straight lines mark out the outline of the experimental envelope. The envelope represents the outer bounds of the data generated by four repeated experimental evacuation trials. The stepped lines represent four airEXODUS predictions for this configuration. Clearly the airEXODUS simulations fall within the variation observed in the experiment [39].

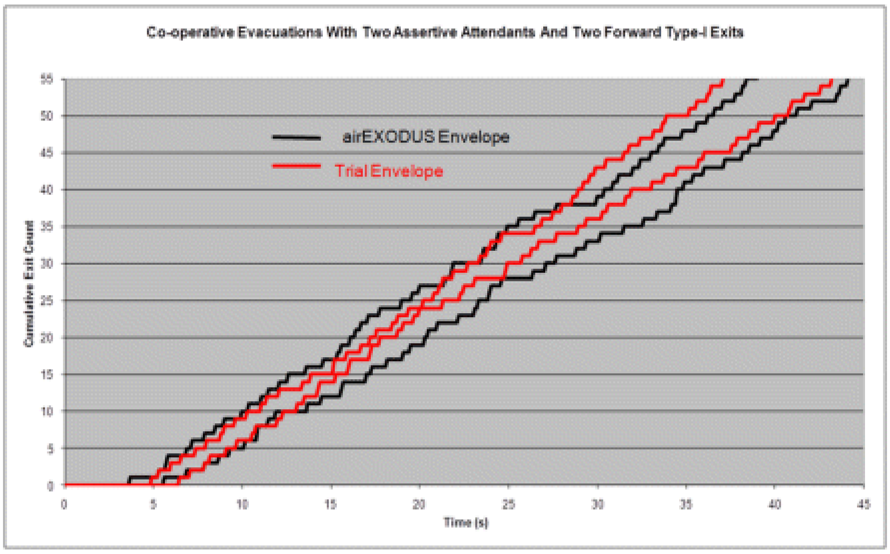

Figure 10-1: Comparison of the results obtained from the B737 Simulations and the airEXODUS predictions

The above cases represent comparisons of the airEXODUS model with historic data. In the second series of validation exercises, predictions from the airEXODUS model were compared with data generated from the certification trial for the B767-304ER aircraft (seating 351 passengers). However, the model predictions were performed and reports submitted to the UK CAA and US FAA several weeks prior to the actual certification exercise taking place [39,49].

Unlike in the experimental evacuation trials, only a single certification trial is performed. As this does not allow for the variability in human performance a number of model predictions were performed to cater for a range of possible outcomes. Two types of scenario were investigated, each specific case within each scenario being repeated a number of times. The first scenario involved passengers heading towards the exit that is deemed optimal. An optimal selection of exits may necessitate some passengers using an exit that is not necessarily their closest exit. These cases give an indication of the best times that can be achieved by crew and aircraft during the trial assuming all goes well. A number of sub-optimal cases were also run. These cases give an indication of times that may be achieved if problems are encountered during the trial. Scenarios investigated included late opening of exits and inefficient crew performance resulting in poor passenger distribution between the available exits. In total, 321 evacuation simulations for the B767-304ER were produced using airEXODUS.

In order to specify various levels of sub-optimal performance a parameter was defined which enables the level of optimal performance in the simulations to be compared with the optimal performance achieved during the trial. Using this parameter it was determined that performance achieved by the passengers/crew/aircraft on the day was near optimal. The optimal airEXODUS predictions for the average evacuation time were within approximately 2% of the measured time [39,49]. Furthermore, general trends in passenger flow behaviour predicted by airEXODUS appear to have been corroborated by actual events, for instance, the passenger split within the cabin predicted by airEXODUS was achieved in the actual trial.

Direct quantitative validation of the buildingEXODUS model has been limited due to the lack of suitable experimental data, however several studies have been undertaken using available data [5,37,38,40,41,48,50,80].

For the purposes of this document some of the results from the comparisons with the Stapelfeldt [44] data will be described. Interested readers are referred to references [37,50] for a more complete discussion. This validation is based upon the experiments conducted by Stapelfeldt [44] and reported by Paulsen et al [45]. This involved placing 100 police cadets in a compartment containing a single exit. Four compartment exit widths were considered namely, 0.75m, 0.80m, 1.50m and 1.60m. The population density around the exit was measured at 4 persons/m^2.

The data reported in Paulsen et al [45] concerns the evacuation of the room under 'normal' conditions, without simulated 'panic'/increased motivation to evacuate, or the effects of hazards.

Within buildingEXODUS a 3.0m x 8.5m enclosure and a population of 100 males was specified. This achieved a population density of 4 persons/m^2 around the exit. Here the results for the 1.5m exit using a population with a 50% range in the 'DRIVE' attribute and allowing the software to predict the flow rate through the exit will be briefly described.

In the Paulsen report [45], the experimental evacuation time for the 1.5m exit is quoted as 30 seconds. This is apparently for a single trial. Five repeat runs of buildingEXODUS produced evacuation times of 28.8, 29.1, 29.6, 31.1, 31.4 seconds with a mean time of 30.0 seconds, resulting in a difference of 0%.

buildingEXODUS was also used to simulate the evacuation of one of the pavilions at the Tsukuba world exposition [38]. As in the previous example, it should be noted that this data-set suffers from a range of defects and is thus not ideal for quantitative validation purposes [38]. As a result the data set was primarily used for qualitative validation. One of the difficulties with the data was that on the day of the evacuation trials it rained. This provided the evacuees with a disincentive to vacate the pavilion. As a result, considerable crowding occurred at the exit, slowing the actual evacuation. For validation purposes it is essential to be fully aware of this situation in order that measures can be taken to account for this behaviour. Another difficulty was that the response time for the audience was not explicitly measured.

However, accepting the limitations of the available data and including a mechanism that creates the observed crowding by the exit, it was found that buildingEXODUS was able to produce reasonable agreement with the data [38] *under certain conditions*. As part of the study, a sensitivity analysis of the response time distribution was undertaken. Several ranges of response time distributions were used that broadly corresponded to the described situation. For example, assuming an occupant response time distribution of 0 to 90 seconds and by extending the solution domain to include the outside region by the exit (in order to include the exterior crowding), buildingEXODUS produced a mean total evacuation time of 160.5 seconds. The predicted evacuation times varied from 156.8 to 164 seconds. This compared with a measured mean evacuation time of 158.6 seconds, with a range of measured evacuation times varying from 152 to 166 seconds (results from four repeat drills). Interested readers are referred to [38] for a full account of the simulations.

Recently, the manner in which EXODUS is able to integrate response time data and subsequently simulate an evacuation was examined [106]. This was based on the evacuation of a university structure, where data describing the overall evacuation performance and the response time distribution of the evacuees was collected. This was then introduced into the model which was then used to simulate the trial evacuation. The results produced were impressive, to the extent that enough confidence was established, enabling the model to be used to examine procedural changes in the structure. In a similar manner, the buildingEXODUS model was used to examine the Gothenburg tragedy, where a fire consumed a disco, causing 63 fatalities. In conjunction with the Swedish National Testing and Research Institute, who provided experimental data describing the environmental decline, buildingEXODUS simulated the outcome of the evacuation. There was a favourable comparison between the simulated and actual results produced, especially given the limitations of the experimental results [107].

Independent quantitative validation of buildingEXODUS and three other evacuation models, has also been provided by the VTT laboratory in Finland [41,48]. This involved comparing model predictions with observed evacuation times for a theatre containing some 600 people. The models were found to agree quite well with each other and with the observed evacuation times. In this example the model developers were not involved in the validation modelling process, nor were they even aware of the nature of the example being simulated. The model developers were however available to provide answers to technical questions raised by the investigators and to provide advice on model usage when asked. In addition, the maritimeEXODUS model has been subjected to scrutiny by Monash University in relation to the application of the model to the maritime IMO regulations [109]

buildingEXODUS has also been applied to simulate hospital evacuation trials [5,40,80]. Model predictions were compared with data generated from trial evacuations of a hospital ward [47]. Unfortunately, (as with most experimental trials) some important data required for model validation purposes was not collected during making detailed comparison impossible. However, bearing these limitations in mind, buildingEXODUS was able to produce good agreement with the collected experimental data.

NOTE:
Quantitative validation involving evacuation from a two floor library has recently (August 2011) been undertaken based on data generated by FSEG as part of the EU FP7 project BeSeCu [154]. The simulations show that buildingEXODUS V4.07 is capable of reproducing the experimental results for the total evacuation time and for the evacuation history, as measured for each of the two external exits. In addition, another set of evacuation data generated from the same project (Polish library experiment) shows that buildingEXODUS V4.07 is capable of reproducing the evolving crowd density in the vicinity of a main exit, the exit history and total evacuation time. Both these data sets will be published in the near future. For the latest information concerning these data sets, please consult the FSEG web site publication pages (http://fseg.gre.ac.uk/fire/pub.asp).

There have also been a number of validation projects relating to the other variants in the EXODUS software suite of models [109-118]. Although the buildingEXODUS model was not directly involved in these validation tests, a number of the components tested are shared between the models in the suite of software. Therefore, indirectly, the buildingEXODUS model benefited from the validation cases of the maritime and aviation versions of the model. In addition to this a number of other publications are available in relation to the EXODUS suite of models (see Chapter 12).

NOTE:
Quantitative validation involving large passenger vessels has recently (August 2011) been undertaken based on data generated by FSEG as part of the EU FP7 project SAFEGUARD [155]. The simulations show that maritimeEXODUS is capable of reproducing the

experimental results for the total assembly time and for the assembly time for each assembly station. This data set will be published in the near future. For the latest information concerning these data sets, please consult the FSEG web site publication pages (http://fseg.gre.ac.uk/fire/pub.asp).

The validation of any model is an on-going, continual activity. Hence, the validation of the buildingEXODUS model is by no means complete, further quantitative validation will take place when suitable data is available. The other forms of validation described in this document continue.

CHAPTER 11: REFERENCES

1. M Owen, E R Galea and P J Lawrence, The EXODUS evacuation model applied to building evacuation scenarios, Journal of Fire Protection Engineering, Vol.8(2), pp.65-86, 1996.

2. E R Galea, M Owen and P J Lawrence, Computer Modelling of Human behaviour in aircraft fire accidents, Toxicology, vol 155, 1-3, pp 63-78, 1996.

3. M Owen, E R Galea and P J Lawrence, The EXODUS evacuation model applied to building evacuation scenarios, Fire Engineers Journal, pp 26-30, vol 56 No. 183, July 1996.

4. S Prew, Computerised evacuation modelling, Cabin Safety Update, pp 1-4 Vol.2 No. 3, 1996.

5. E Galea, P Lawrence, L Filippidis, Extending the Capabilities of the buildingEXODUS Evacuation Model to Cater for Hospital Evacuations, Proceedings of the Third International Conference on Fire Research and Engineering, 4-8 October 1999, Chicago, USA, pp 39-50, 1999.

6. E R Galea, The Role of Evacuation and Fire Modelling in the Development of Safer Aircraft, Proceedings International Conference on Cabin Safety Research, pp 145-160, Atlantic City, Nov. 1995, Report: DOT/FAA/AR-95/120 pub: Office of Aviation Research, Washington, 1995.

7. E R Galea, M Owen, and P J Lawrence, Emergency Egress from Large Buildings under Fire Conditions Simulated Using the EXODUS Evacuation Model, Proceedings of the 7th International Fire Science and Engineering Conference: Interflam'96, pp 711-720, St John's College, Cambridge, England, March 1996, compiled by C Franks and S Grayson, Published by Interscience communications Ltd, ISBN 0 9516320 9 4, London, UK, 1996.

8. M Owen, E R Galea and P Lawrence, The EXODUS evacuation model applied to building evacuation scenarios, Proc International conference, Fire Safety By Design, ISBN 0 9527398 3 6, Vol.3, pp.81-90, Compiled by G W Butler, 1995.

9. E R Galea and M.Owen, Predicting the Evacuation Performance of Mass Transport Vehicles, 'Fire Safety on Ships', Developments into the 21st Century. ISBN 0-907206-57-3, IMarE Conference, Vol 106,2, 239 – 247, Lond 1994.

10. E R Galea and J M P Galparsoro, A Computer Based Simulation Model for the Prediction of Evacuation from Mass Transport Vehicles, Fire Safety Journal Vol 22 pp 341 - 366, 1994.

11. E R Galea, A Computer Based Simulation of the Evacuation of Passengers under Hazardous and Non-Hazardous Conditions from Wide and Narrow Body Aircraft, Invited Paper, Proc. 11th Annual Int. Aircraft Cabin Safety Symp., Southern California Safety Institute, Calif. USA, pp 259 – 280, 1994.

12. E R Galea, J M P Galparsoro and J Pearce, EXODUS: A simulation Model for the Evacuation of Large Populations from Mass Transport Vehicles and Buildings under Hazardous Conditions. CIB W14, Int. Symp. and Workshop Engineering Fire Safety in the Process of Design. Univ. of Ulster, Part 3 pp 11-25, 1993.

13. E R Galea, How University Professor's Prototype Egress Model Helps Evacuation Design, Fire, p.23, July 1993.

14. E R Galea, J M Perez Galparsoro and J Pearce, A Brief Description of the EXODUS Evacuation Model, Proceed, 18th Int Conf on Fire Safety. San Francisco, USA, Vol.18, pp.149-162, 1993.

15. E R Galea and J M Perez Galparsoro, EXODUS: An Evacuation Model for Mass Transport Vehicles. UK CAA Paper 93 006 ISBN 0 86039 543X, 1993.

16. E.R, Galea, G Sharp, P Lawrence, Investigating the Representation of Merging Behavior at the Floor Stair Interface in Computer Simulations of Multi-Floor Building Evacuations, Journal of Fire Protection Engineering, Volume 18, pp. 291-316, 2008. Publisher SAGE. DOI:10.1177/1042391508095092.

17. Jin, T. Visibility through fire smoke, J of Fire and Flammability, 9, 135-155, 1978.

18. Jin, T. and Yamada, T Experimental Study of Human behaviour in smoke filled corridors. In: T. Wakamatsu, Y. Hasemi, A. Sekizawa, P. Seeger, P. Pagni and C. Grant (Eds.) Fire Safety Science Proceedings of the Second International Symposium, Hemisphere Publishing Company, New York, 561-570, 1989.

19. J.Fruin. Pedestrian planning and design, Metropolitan Association of Urban Designers and Environmental Planners, New York, 1971

20. D.A.Purser, Modelling time to incapacitation and death from toxic and physical hazards in aircraft fires. AGARD No 467, pp 41-1 - 41-12, 1989.

21. H.Muir, C.Marrison and A.Evans, Aircraft evacuations: the effect of passenger motivation and cabin configuration adjacent to the exit. CAA Paper 89019, ISBN 0 86039 406 9, 1989.

22. Galea E R, On the field modelling approach to the simulation of enclosure fires, Journal of Fire Protection Engineering, vol 1 (1), pp 11-22, 1989.

23. Purser, D.A., Toxicity Assessment of combustion products. In: C.L.Beyler (Ed) SFPE Handbook of Fire Protection Engineering, National Fire Protection Association, Quincy M.A., 1-200 - 1-245, 1988.

24. Purser, D.A., Interactions between behaviour patterns and physiological impairment in escape from fire. In: C.A.Franks (Ed.) Interflam'93, Interscience communications Ltd, London, 579 – 593, 1993.

25. Purser, D.A., Toxicity Assessment of combustion products. In: J.L.Linville (Ed) SFPE Handbook of Fire Protection Engineering 2nd Edition, National Fire Protection Association, Quincy M.A., 2-85 - 2-146, 1995.

26. Speitel, L.C., Analytical Method for water vapour collection and analysis in aircraft cabin fires. DOT/FAA/CT - TN93/33, US Department of Transportation, 1993.

27. Speitel, L.C., Toxicity Assessment of Combustion Gases and Development of a Survival Model, DOT/FAA/AR-95/5, US Department of Transportation, 1995.

28. HMSO, The Building Regulations 1991 - Approved Document B, section B1, HMSO Publications, London

29. Polus A, Schofer J L, Ushpiz A, Pedestrian flow and level of service, Journal of Transportation Engineering, PROC ASCE, 109, pp46-47, 1983

30. Hankin B D, Wright R A, Passenger flows in subways, Operational Research Quarterly, Vol 9, pp81-88, 1958.

31. Pauls J L, Building Evacuation: Research Findings and Recommendations, chapter 14 Fires and Human Behaviour, D.Canter (ed.), pp251-275, John Wiley and Sons Ltd, London

32. Shields T J, Dunlop K E, and Silcock G W H, Escape of Disabled People From Fire. A Measurement and Classification of Capability For Assessing Escape Risk, BRE Report 301, ISBN 1 86081 0675, 1996.

33. Ando K, Ota H, and Oki T, Forecasting The Flow Of People, Railway Research Review, (45), pp 8-14, 1988.

34. SFPE Handbook of Fire Protection Engineering 2nd Edition, Editor: J.L.Linville, National Fire Protection Association, Quincy M.A., 2-85 - 2-146, 1995.

35. Gwynne, S. , Galea, E.R., Owen, M. and Lawrence, P.J., Escape As A Social Response, Published by the Society of Fire Protection Engineers, 1999.

36. Galea,E.R, A General Approach to Validating Evacuation Models with an application to EXODUS, Journal of Fire Sciences, Vol 16, Technomic Publishing Co., November / December, 1998.

37. Gwynne, S., Galea, E. R., Owen, M., Lawrence, P. J., and Filippidis, L., Validation of the buildingEXODUS Evacuation Model. CMS Press Paper 98/IM/29, ISBN 1899991298, 1998.

38. Gwynne, S., Galea, E. R., Owen, M., Lawrence, P. J., and Filippidis, L., A Systematic Comparison of Model Predictions Produced by the buildingEXODUS Evacuation

Model and the Tsukuba Dataset. Journal of Applied Fire Science, Volume 7 Number 3, pp 235 – 266, 1997/1998.

39. Owen, M., Galea E. R., Lawrence, P., and Filippidis, L., The Numerical Simulation of Aircraft evacuation and its application to aircraft design and certification, Aeronautical Journal, June/July pp 301- 312, 1998.

40. Galea, E.R., Owen, M., Lawrence, P.J., Confidential Report prepared for the NHS, Mar 1997.

41. Weckman, H., Calculation and simulation of evacuation from buildings, VTT report 1890, ISBN 951-38-5195-8, 1998.

42. Bukowscki, R., Analysis of the airEXODUS aircraft evacuation model, NIST report in preparation. 1998.

43. Muir, H.C., Cobbett, A., Influence of Cabin Crew Behaviour in Emergency Evacuations At Floor Level Exits, Report Prepared for CAA/FAA, CAA Paper 95006: Part A, (FAA No.DOT/FAA/AR-95/52), CAA, London, ISBN 0 86039 6495, 1996.

44. Stapelfeldt, J-P Angst- und Paknikstande aus der Sicht des Brandschutzes (Conditions of anxiety and panic from the viewpoint of fire protection), vfdb-Zeitschrift 2,41,1986

45. Paulsen, T., Soma, H., Schneider, V., Wiklund J., Lovas, G., Evaluation of simulation models of evacuation from complex spaces, SINTEF REPORT, STF75 A95020, ISBN 82-595-8583-9, June 1995.

46. Butler, G.W., The Factors Involved in Evacuation and the Extent to Which Efficient Management Can Influence These, The Proceedings Interflam '93, Interscience Comms Ltd, London, 1993.

47. Appleton, I., Quiggin, P., Hackney Hospital Evacuation Exercises Fire Research Station Contribution An Evacuation Model. Note N 22/78, March 1978.

48. Weckman, H., Lehtimaki, S., and Mannikko, S. Evacuation of a theatre: Exercise Vs Calculations., Human Behaviour in Fire – Proceedings of the First Int Symp, Ed: J.Shields, ISBN 1 85923 103 9, pp 479 - 488, 1998.

49. Galea E.R., Gwynne S., and Lawrence, P., Using Computer Simulation to Predict the Evacuation Performance of Passenger Ships, Paper 00/IM/61, CMS Press, ISBN 189999162X, 2000 .

50. Gwynne,S., Galea, E.R., Owen, M., Lawrence, P.J. and Filippidis, L., A Comparison of Predictions from the buildingEXODUS Evacuation Model with Experimental Data, pp711-721, Human Behaviour in Fire: Proceedings of the 1st International Symposium, Ed: Shields,J., University of Ulster, ISBN 1859231039, TextFlow Ltd, 1998.

51. Predtechenskii, V.M. and Milinskii, A. I., Planning For Foot Traffic Flow In Buildings, Published For The National Bureau Of Standards, Amerind Publishing Co., 1978, Translated From The Russian Publication Which Appeared In 1969, STROIZDAT Publishers, Moscow, 1969.

52. Hartzell, G. E., “Engineering analysis of hazards to life safety in fires: The fire effluent toxicity component.” Human Behaviour in Fire – Proceedings of the First Int Symp, Ed: J.Shields, ISBN 1 85923 103 9, pp 621 - 629, 1998.

53. Fire Safety Engineering in buildings, Part 1 Guide to the application of fire safety engineering principles, DD240: Part 1: 1997, Publisher: BSI., ISBN 0 580 27952 9, 1997.

54. Gwynne,S. and Galea, E.R.,A Review of the Methodologies Used in the Computer Simulation of Evacuation from the Built Environment, Published by the Society of Fire Protection Engineers, 1998.

55. Gwynne,S., Galea, E.R., Owen, M., Lawrence, P.J. and Filippidis, L., Modelling Occupant Interaction with Fire Conditions using the buildingEXODUS Evacuation Model, CMS Press Paper 00/IM/54, ISBN 1899991557, 2000.

56. Gwynne, S., Galea, E., R., Lawrence, P.J. and Owen, M., An Investigation Of The Aspects Of Occupant Behaviour Required For Evacuation Modelling, Journal of Applied Fire Science, Vol 8(1), pp 19-59, 1998/1999.

57. Gwynne,S. Galea, E., R., Lawrence, P.J.,Owen, M. and Filippidis, L., Adaptive Decision-Making in buildingEXODUS, Journal of Applied Fire Science, 8 (4), pp265-289, 1999.

58. Gwynne,S., Galea, E.R., Owen, M., Lawrence, P.J. and Filippidis, L.,A Review of the Methodologies used in Evacuation Modelling, FIRE AND MATERIALS, 23,pp383-388, Wiley & Sons, 1999.

59. Munro, J., and Edwards, C., Clinical Examination, 8th edition, Churchill Livingstone, Edinburgh, ISBN 0443040796, 1990.

60. A Technical Reference For CFAST: An Engineering Tool For Estimating Fire And Smoke Transport, NIST Technical Note 1431, Building And Fire Research Laboratory, January 2000.

61. NIST Website: http://cfast.nist.gov

62. Purser, D.,A., Inaugural Lecture Series, University of Greenwich, ISBN 1861661177, 1999.

63. Wood, P.G, Survey Of Behaviour In Fires, Edited D. Canter pp 83 - 95). Fires And Human Behaviour (2nd Edition), Ed. D. Canter, Fulton, pp205-234,1990.

64. Bryan, J., Behavioural Response To Fire And Smoke, The SFPE, Handbook Of Fire Protection Engineering(2nd Edition), Ed:Dilenno, P.J., Beyer, C.L.,Custer, R.L.P.,Walton, W.D.,Watts, J.M.W.,Drysdale, D., Hall, J.R., National Fire Protection Association,Quincy,Ma, pp (1-241)-(1-262), 1996.

65. Jones, B.K. And Hewitt, J.,A., Leadership And Group Formation In High-Rise Building Evacuations, Fire Safety Science-The Proceedings For The First International Symposium, Hemisphere Publishing Corp., Washingtion, pp513-522, 1986.

66. Brennan, P., Smoke Gets In Your Eyes:The Effect Of Cue Perception On Behaviour In Smoke, Proceeding For The International Conference On Fire Science And Engineering, Asiaflam'95,Eds: Chow,W.K.Et Al, pp187-197, ISBN 0951632078, 1995.

67. Brennan, P., Impact Of Social Interaction On Time To Begin Evacuation In Office Building Fires:Implications For Modelling Behaviour, Proceedings Of The Seventh International Fire Science And Engineering Conference, Interflam'96, Ed:Franks,C. and Grayson, S., pp701-710, ISBN 0951632094, 1996.

68. Proulx, G. And Fahy,R., A Study Of The New York World Trade Centre Evacuation, Proceedings Of The Third International Symposium On Fire Safety Science, Eds: Cox, G. And Langford, G., pp199 -209, Elsevier, London, 1991

69. Johnson, N.,R., Feinberg, W.,E. And Johnston, D.,M., Microstructure And Panic:The Impact Of Social Bonds On Individual Action In Collective Flight From The Beverly Hills Supper Club Fire, From Disaster, Collective Behaviour And Social Organisation, (Ed:Dynes, R., And Tierney, K.)

70. Best, R.,L., Reconstruction Of A Tragedy: The Beverly Hills Supper Club Fire, Southgate, Kentucky, May 28, NFPA, 1977.

71. Gwynne, S., buildingEXODUS Functionality Testing Document (Draft), University of Greenwich, 2000.

72. Proulx, G., Lessons Learned On Occupants' Movement Times And Behaviour During Evacuation Drills, Interflam 96, pp1007-1011, 1996.

73. Tong, D. And Canter, D., The Decision To Evacuate: A Study Of The Motivations Which Contribute To Evacuations In The Event Of A Fire, Fire Safety Journal , 9, pp257-265, 1985.

74. Tong, D. And Canter, D., Informative Warnings: In Situ Evaluations Of Fire Alarms, Fire Safety Journal , 9, pp267-279, 1985.

75. Geyer, , Bellamy, L., Max-Lino, R.,Harrison, Bahrami, Z. And Modha, B., An Evaluation Of The Effectiveness Of The Components Of Information Fire Warning Systems, In Sime J.,(Ed) Safety In The Built Environment, E And Fn Spon, pp36-47, 1988.

76. Proulx, G.,Sime, J., To Prevent Panic In An Underground Emergency: Why Not Tell People The Truth?, Fire Safety Science- 3rd Symposium, Elsevier, Appl. Sci., NY, pp843-853, 1991.

77. Gardner, J., Building Design And Human Behaviour In Fires, Fire Engineer, pp18-20, July, 1996.

78. Fire Research Station, Experimental Programme To Investigate Informative Fire Warning Characteristics For Motivating Fast Evacuation, Br172, Building Research Establishment, UK, 1990.

79. Proulx, G., Evacuation Time and Movement in Apartment Buildings, Fire Safety Journal, 24, pp229-246, Elsevier, 1995.

80. Galea, E.R., Lawrence, P., Owen, M., and Filippidis, L., Adapting the buildingEXODUS Evacuation Model for Hospital Specific Evacuation Scenarios, Proceedings of the 8th International Fire Science and Engineering Conference: Interflam'99, Vol. 2, pp 1247-1252, Edinburgh, Scotland, June 29 - July 1st, published by Interscience Communications Ltd, London, UK, 1999. ISBN 0 9532312 3 2, 1999.

81. Hymes, I., Boydell, W., and Prescott, B., Thermal radiation: Physiological and pathological effects, IChemE Major Hazards Monograph, 1996.

82. Hymes, I., Brearley, S., Prescott, B.L. and Zahid, M., The prognosis of burn injury victims. SRD/HSE/600, AEA Technology, 1993.

83. Purser, D.A. "Modelling Time To Incapacitation And Death From Toxic And Physical Hazards In Aircraft Fires", Agard No 467, pp41-1 - 41-12, 1989.

84. Purser, D., "People And Fire", Inaugural Lecture Series, 17th February, University Of Greenwich, ISBN 1-86166-117-7, 1999.

85. Purser, D.A., "Toxicity Assessment Of Combustion Products", The SFPE, Handbook Of Fire Protection Engineering (1st Edition), Ed: Beyer, C.L. National Fire Protection Association, Quincy, Ma, pp(1-200) - (1-245), 1988.

86. Purser, D.A., "Toxicity Assessment Of Combustion Products", The SFPE, Handbook Of Fire Protection Engineering (2nd Edition), Ed:Dilenno, P.J., Beyer, C.L.,Custer, R.L.P.,Walton, W.D.,Watts, J.M.W.,Drysdale, D., Hall, J.R., National Fire Protection Association,Quincy,Ma, pp (2-85)-(2-146),1996.

87. Purser, D., “Human Tenability”, Technical Basis for Performance-Based Fire Regulations, Engineering Foundation Conference, 7-11 January, San Diego, United States, 2001

88. Taylor S., Petridis M., Knight B., Ewer J., Galea E.R. and Patel M. K. “SMARTFIRE: An Integrated Computational Fluid Dynamics code and Expert System for Fire Field Modelling”, Fire Safety Science – Proceedings of the Fifth International Symposium, pp 1285-1296, 1997.

89. Ewer J., Galea E.R., Patel M.K., Taylor S., Knight B. and Petridis M., “SMARTFIRE: An Intelligent CFD Based Fire Model, Fire protection Engineering”, vol 10, no 1, pp 13-27, 1999.

90. Taylor S, Galea E, Patel M K, Petridis M, Knight B and Ewer J, “SMARTFIRE: An Intelligent Fire Field Model”, Proceedings Interflam 96, Cambridge, UK, pp 671-680, 1996.

91. Wang, Z., Jia, F., Galea, E.R., Patel, M.K. and Ewer, J., Simulating one of the CIB W14 round robin test cases using the SMARTFIRE fire field models, Fire Safety J., vol. 36, 2001, pp. 661-677.

92. Ewer, J., Jia, F., Grandison, A., Galea, E.R., and Patel, M.K.,. SMARTFIRE V3.0 User Guide and Technical Manual, Fire Safety Engineering Group, University of Greenwich, UK, 2002."

93. “Smartfire Verification and Validation Report”, Software Version 2.01, Report Version 1.01, Fire Safety Engineering Group, University of Greenwich, Revision Date 25/05/99.

94. Galea E.R., “On the field modelling approach to the simulation of enclosure fires”, Journal of Fire Protection Engineering, vol 1 (1), 1989, pp 11-22.

95. “Combustion Fundamentals of Fire”, Editor: Cox G., Academic Press, 1995.

96. Kumar S., Gupta A.K. and Cox G., “Effects of Thermal Radiation on the Fluid Dynamics of Compartment Fires”, Fire Safety Science - Proc. of the Third Intl. Symp., pp 345-354, 1991.

97. Raithby G.D., Chui E.H., “A finite volume method for predicting a radiant heat transfer in enclosures with participating media”, Journal of Heat Transfer, vol 112, pp 415-423, May 1990.

98. Lewis M.J., Moss M.B. and Rubini P.A., “CFD Modelling of Combustion and Heat Transfer in Compartment Fires”, Fire Safety Science, Proc. of the 5th Int. Symp., pp 463-474, 1997.

99. Pantakar S.V., “Numerical Heat Transfer and Fluid Flow”, Intertext Books, McGraw Hill, New York, 1980.

100. Grandison A.,Galea E.R., and Patel, M.K., “Fire Modelling Standards/Benchmark Report on Phase 1 Simulations”, CMS Press London UK, Report 01/IM/72, ISBN 1899991735, June 2001.

101. Grandison A.,Galea E.R., and Patel, M.K., “Fire Modelling Standards/Benchmark Report on Phase 2 Simulations”, CMS Press London UK, Report 01/IM/82, ISBN 1899991824, Oct 2001.

102. NIST Website: http://fire.nist.gov/fire/fires/fires.html

103. Filippidis L, Galea E R, Lawrence P and Gwynne S, Visibility Catchment Area of Exits and Signs, INTERFLAM 2001, Proc 9th Int Conf, Vol. 2, p1529-1534.

104. Filippidis L, Galea E R, Gwynne S and Lawrence P, Simulating The Interaction of Pedestrians with Wayfinding Systems, Pedestrian and Evacuation Dynamics 2003, Proc 2nd Int Conf, p 39-50.

105. National Fire Protection Association, Life Safety Code Handbook, National Fire Protection Association, Quincy, MA, 1997.

106. Park J., Gwynne S., Galea E.R., and Lawrence P., “Validating the buildingEXODUS Evacaution model using Data from an Unannounced Trial Evacuation.”, Proceedings of the 2nd International Pedestrian and Evacuation Dynamics Conference, Ed: E.R.Galea, CMS Press, Greenwich, UK, ISBN 1904521088, pp295-306, 2003.

107. Jiang, H, Gwynne S., Galea E.R., Lawrence P., Jia F and Ingason H., “The use of Evacuation simulation, Fire Simulation and Experimental Data in Forensic Fire Analysis.”, Proceedings of the 2nd International Pedestrian and Evacuation Dynamics Conference, Ed: E.R.Galea, CMS Press, Greenwich, UK, ISBN 1904521088, pp341-354, 2003.

108. Galea E.R., Filippidis L., Gwynne S., Lawrence P.J. and Blackshields D,"Predicting the Performance of Passenger Ships using Computer Simulation", Proc Fire Safety and Survivability, Applied Vehicle Technology (AVT) Panel Symp, NATO Research and Technology Organisation, Aalborg Denmark 23-26 Sept 2002, RTO-MP-092, (AVT-088), CD, paper 33, 14 pages, published November 2003.

109. Galea E.R., Blake S.J., Lawrence P.J., and Gwynne S., "The airEXODUS Evacuation Model and its Application to Aircraft Safety", Proc Fire Safety and Survivability, Applied Vehicle Technology (AVT) Panel Symp, NATO Research and Technology Organisation, Aalborg Denmark 23-26 Sept 2002, RTO-MP-092, (AVT-088), CD, paper 31, 15 pages, published November 2003.

110. Galea E.R., Blake S., Gwynne S. and Lawrence P.,“The use of evacuation modelling techniques in the design of very large transport aircraft and blended wing body aircraft”. The Aeronautical Journal, Vol 107, Number 1070, pp 207-218, 2003.

111. E.R.Galea, S. Gwynne, D. Blackshields, P.Lawrence and L.Filippidis, "Predicting the evacuation performance of passenger ships using computer simulation", Proceedings of the 9th International Fire Science and Engineering Conference: Interflam '01, Vol. 2, pp 853-864, Edinburgh, Scotland, Sept 17-19 2001, published by Interscience Communications Ltd, London, UK, 2001. ISBN 0 95323129 1 (vol2).

112. Galea E.R., Lawrence P, Gwynne S., Filippidis L., Blackshields D., Sharp G., Hurst N., Wang Z., and Ewer J ., Proc 2nd Int Pedestrian and Evacuation Dynamics Conference, Ed: E.R.Galea, CMS Press, Greenwich, UK, ISBN 1904521088, pp159-172, 2003.

113. Galea, E R, Gwynne, S., Lawrence P., Blackshields, D., Ewer, J., Wang, N., Hurst, N. and Mawhinney,N., "The Application of Fire and Evacuation Simulation in Ship Design". COMPIT 2003, 2nd International EuroConference on Computer and IT Applications in the Maritime Industries, Hamburg, pp55-69, 2003.

114. Galea, E R, Gwynne, S., Lawrence P., Blackshields, D., Ewer, J., Wang, N., Hurst, N. and Mawhinney,N., "Fire Safety Engineering and Ship Design Using Advanced Fire and Evacuation Simulation", 1st International Conference – Fire on Ships: Prevention, Detection and Control on Ships and Offshore Structures, 11-12/3/03.

115. Galea E.R., Filippidis L, Gwynne S, Lawrence P, Sharp G, Blackshields D., "The Development of an Advanced Ship Evacuation Simulation Software Product and Associated Large Scale Testing Facility for the Collection of Human Shipboard Behaviour Data", Proceedings Int Conf on Human Factors in Ship Design and Operation, The Royal Institution of Naval Architects, London, pp37-50, 2002.

116. Gwynne S, Galea E.R., Parke J, Hickson J., "The Collection and Analysis of Pre-Evacuation Times from Evacuation Trials and their Application to Evacuation Modelling". Fire Technology, Kluwer Associates, US, pp173-195, vol 39, number 2, 2003.

117. Gwynne S, Galea E.R., Lyster C, Glen I., "Analysing the Evacuation Procedures Employed on a Thames Passenger Boat using the maritimeEXODUS Evacuation Model". Fire Technology, 2003. Kluwer Associates, US, pp225-246, vol 39, number 3, 2003.

118. Blake, S. J., Galea, E. R., Gwynne, S., Lawrence, P. J., and Filippidis, L., "Examining The Effect Of Exit Separation On Aircraft Evacuation Performance During 90-Second Certification Trials Using Evacuation Modelling Techniques". The Aeronautical Journal of the Royal Aeronautical Society, pp1-16, January 2002.

119. Carroll, L.J., Koss, L.L., and Brumley, A.T., “Evacuship Analysis Based on IMO Test Cases and Other Benchmark Comparisons”, Proceedings of the 2nd International Pedestrian and Evacuation Dynamics Conference, Ed: E.R.Galea, CMS Press, Greenwich, UK, ISBN 1904521088, pp185-196, 2003.

120. So, A., Lai, T., and Yu,.J. “Lift Logic”, Fire Prevention Fire Engineers Journal, pp26-28, August, 2003.

121. “Floating ASSETS”, Fire Prevention Fire Engineers Journal, pp43-45, August, 2003.

122. http://jfe.sagepub.com/content/11/2/88.full.pdf+html

123. Engineering Guide for Predicting 1st and 2nd Degree Skin Burns from Thermal Radiation, Society of Fire Protection Engineers, Bethesda, MD, 2000.

124. Kinsey, M.J, Vertical Transport Evacuation Modelling, PhD Thesis, University of Greenwich, London, UK 2011.

125. Transportation systems in buildings, CIBSE Guide D:2005, Chartered Institution for Building Service Engineers (CIBSE), London, UK, 2005

126. Hui X, Filippidis L, Gwynne S, Galea E.R., Blackshields, D., and Lawrence P., Signage Legibility Distances as a Function of Observation Angle, Journal of Fire Protection Engineering, Vol 17, No1, pages 41-64, 2007. DOI: 10.1177/1042391507064025.

127. BS 5499-1:1990, Fire safety signs, notices and graphic symbols. Specification for fire safety signs, ISBN 0 580 18830 2

128. BS 5499-1:2002, Graphical symbols and signs — Safety signs, including fire safety signs — Part 1: Specification for geometric shapes, colours and layout, ISBN 0 580 38258 3

129. BS 5499-4:2000, Safety signs, including fire safety signs — Part 4: Code of practice for escape route signing, ISBN 0 580 33205 5.

130. ISO 3864-4:2011, Graphical symbols -- Safety colours and safety signs, 2011.

131. Hui Xie, L. Filippidis, E. R. Galea, D. Blackshields, P. J. Lawrence, "Experimental Study of the Effectiveness of Emergency Signage". Proceedings of the 4th International Symposium on Human Behaviour in Fire, Robinson College, Cambridge, UK, 13-15 July 2009, pp. 289-300, ISBN 978-0-9556548-3-1.

132. Al-Sharif, L., “Escalator Handling Capacity: Standards Versus Practice”, Elevator World, 1996.

133. E.R.Galea and A.J.P.Dixon, Collection and Analysis of Emergency Services Data Relating to the Evacuation of the World Trade Centre Towers of 11 September 2001 (7/2005), Office of the Deputy Prime Minister, 2005.

134. Galea,E.R, Sharp, G., Lawrence, P.J., Holden, R., Approximating the Evacuation of the World Trade Center North Tower using Computer Simulation, Journal of Fire Protection Engineering, Vol 18 (2), 85-115, 2008. DOI:http://dx.doi.org/10.1177/1042391507079343.

135. Gwynne, S.M.V., Optimizing Fire Alarm Notification for High Risk Groups: Summary Report, Report for The Fire Protection Research Foundation, MA, USA, June 2007.

136. Purser, D., (2005), Comparisons Of Evacuation Efficency And Pre-Movement Times In Response To A Sounder And Two Different Voice Alarm Messages, Hartford Environmental Research.

137. G.Proulx, (1998), The Impact Of Voice Communication Messages During A Residential High-Rise Fire, Human Behaviour In Fire:Proceedings Of The First International Symposium, 31 August-2 September 1998, ISBN 1859231039, pp265-271.

138. G.Proulx, (1999), Occupant Response During a Residential High-Rise Fire, Fire and Materials, v. 23, no. 6, pp. 317-323.

139. Proulx, G.,(1996), Lessons Learned On Occupants' Movement Times And Behaviour During Evacuation Drills, Interflam 96, pp1007-1011.

140. Proulx, G., (1995), Evacuation Times And Movement Times In Apartment Buildings, Fire Safety Journal, 24, pp229-246, Elsevier,0379-7112.

141. Fahy,R. and Proulx, G., (2001), Towards Creating a Database on Delay Times to Start Evacuation and Walking Speeds for Use in Evacuation Modelling, Human Behaviour in Fire, Second International Symposium on Human Behavior in Fire, Boston, MA, March 2001, pp. 175-184, Interscience Communications, London.

142. Lord, J., Moore, A., Meacham, B., Fahy, R., Proulx, G., Uncertainty in egress models and data: investigation of dominant parameters and extent of their impact on predicted outcomes – initial findings, NRCC-47351, NRC-CNRC, 2004.

143. T.J. Shields, *Fire and Disabled People in Buildings*, Report BR 231, Building Research Establishment, Garston, 1993.

144. EU FP6 project AVATARS N° TST4-CT-2005-012462 (http://cordis.europa.eu/fetch?CALLER=FP6_PROJ&ACTION=D&DOC=7&CAT=PROJ&QUERY=0130bce1d89c:8f2b:5839a1d6&RCN=74788)

145. G. Proulx, J. Latour and J. MacLaurin, Housing Evacuation of Mixed Abilities Occupants, Internal Report No. 661, National Research Council of Canada, Ottawa ON, July 1994.

146. G. Proulx, D. Tiller, B. Kyle and J. Creek, Assessment of Photoluminescent Material During Office Occupant Evacuation, Internal Report 774, National Research Council of Canada, Ottawa ON, April, 1999.

147. G. Proulx, J.C. Latour, J.W. McLaurin, J. Pineau, L.E. Hoffman, and C. Laroche, Housing Evacuation of Mixed Abilities Occupants in Highrise Buildings, Internal

148. G. Proulx, A. Kaufman and J. Pineau, Evacuation Time and Movement in Office Buildings, Internal Report No. 711, National Research Council of Canada, Ottawa ON,March 1996.

149. Gwynne, Steven, Kuligowski, Erica, and Nilsson, Daniel, Representing Egress Behaviour in Engineering Terms, SFPE Conference on Performance-Based Design, Lund University, 2010.

150. International Maritime Organization, Interim Guidelines for Evacuation Analyses for New and Existing Passenger Ships, MSC/Circ.1033, June 2002.

151. Gwynne, S.M.V. and Rosenbaum, E.R., “Employing the Hydraulic Model in Assessing Emergency Movement,” SFPE Handbook of Fire Protection Engineering, 4th Edition, DiNenno et al. (eds.), National Fire Protection Association, Quincy, MA, 2009, pp3-373 – 3-396.

152. Gwynne, S.M.V., Conventions in the Collection and Use of Human Performance Data, NIST GCR 10-928, National Institute of Standards and Technology, February 2010.

153. Filippidis, L., Lawrence, P., Galea E.R., Blackshields, D. “Simulating the Interaction of Occupants with Signage systems”, Proceedings of the 9th IAFSS Symposium Karlsruhe, Germany, 2008, ISNN 1817-4299, pp 389-400. DOI:10.3801/IAFSS.FSS.9-389.

154. Galea, E.R., Sauter, M., Deere, S.J., Filippidis, L., “Investigating the Impact of Culture on Evacuation Behavior – A Turkish Data-Set”, Proceedings of the 10th IAFSS Symposium, University of Maryland, USA, 20-24 June 2011.

155. Galea, E.R., Brown, R.C., Filippidis. L., and Deere. S., “Collection of Evacuation Data for Large Passenger Vessels at Sea”, Pedestrian and Evacuation Dynamics 2010. 5th International Conference. Proceedings. March 8-10, 2010, Springer, New York, NY, Peacock, R.D., Kuligowski, E.D., and Averill, J.D., Editor(s), pp 163-172, 2011.

156. Galea E.R., Wang, Z., Veeraswamy, A., Jia, F., Lawrence, P., and Ewer, J. “Coupled Fire/Evacuation Analysis of the Station Nightclub Fire”. Proceedings of

9th IAFSS Symposium Karlsruhe, Germany, 2008, ISNN 1817-4299, pp 465-476. DOI:10.3801/IAFSS.FSS.9-465.

157. Kinsey, M.J, Galea, E.R, Lawrence, P.J, (2009), ‘Extended Model of Pedestrian Escalator Behaviour Based on Data Collected within a Chinese Underground Station’, Proceedings of the Human Behaviour in Fire Conference, pp173-182.

158. Kinsey, M.J, Galea, E.R, Lawrence, P.J, (2008), ‘Modelling Pedestrian Escalator Behaviour’, Proceedings of the Pedestrian and Evacuation Dynamics Conference, pp689-696.

159. Nelson, H.E, MacLennan, H.A., (1996) “Emergency Movement” in The SFPE Handbook of Fire Protection Engineering, 2nd ed (P.J DelNenno et al), National Fire Protection Association, Quincy, MA, pp 3-286-3-295.

160. A. Tortajada, J. Quintana, D. Martinez, FGC, "D2.1 Case Studies Data Collection", Report for EU project AVATARS, FP6 N°TST4-CT-2005-012462, March 2006.

161. "D6.3 Final Evaluation Report", Restricted report by BH, BMT, UOG, USAL, ATM and FGC, for EU project AVATARS, FP6 N°TST4-CT-2005-012462, August 2008.

162. "D4.2 Test Results and Stage Two Recommendations inc Appendix 1", Report by FGC, BMT and UOG, for EU project AVATARS, FP6 N°TST4-CT-2005-012462, April 2007.

163. "Video Analysis Instructions for Ticket Barrier and Ticket Machine usage data", Internal report by L. Filippidis, University of Greenwich, for the EU project AVATARS FP6 N°TST4-CT-2005-012462, May 2006.

164. Private Communication from Diego Martinez (FGC) to Lazaros Filippidis (University of Greenwich) regarding Spanish turnstile data analysis, email including Excel datasheets representing the use of turnstile gates from four underground stations in Barcelona, for the EU project AVATARS FP6 N°TST4-CT-2005-012462, 06 November 2006, 10:56am.

165. Kinsey, M.J., Galea, E. R., Analysis and Simulation of Human Factors Associated with Using Ticket Gates in a Spanish Underground Station. CMS Press Paper 13/IM/162, ISBN 978-1-904521-80-8, 2013.

166. Kinsey, M.J, Galea, E.R, Lawrence, P.J, "Human Factors Associated with the Selection of Lifts/Elevators or Stairs in Emergency and Normal Usage Conditions", Fire Technology, Vol 8, Issue 1, pp3-26, DOI: 10.1007/s10694-010-0176-7, 2010.

167. Kinsey, M.J, Galea, E.R, Lawrence, P.J, "Stairs or Lifts? – A study of Human Factors associated with Lift/Elevator usage during Evacuations using an online Survey", Proceedings of the Pedestrians and Evacuation Dynamics Conference, National Institute of Standards & Technology (NIST), US, pp627-636 , 2010.

168. Kinsey, M.J, Galea, E.R, Lawrence, P.J, "Investigating the Use of Elevators for High rise Building Evacuation through Computer Simulation", Proceedings of the Human Behaviour in Fire Conference, Robinsons College, University of Cambridge, UK, pp85-96, 2009.

169. Kinsey, M.J, Galea, E.R, Lawrence, P.J, "Modelling Human Factors And Evacuation Lift Dispatch Strategies", Proceedings of the 5th International Symposium on Human Behaviour in Fire 2012, UK, ISBN 978-0-9556548-8-6, pp386-397, 2012.

170. Closs, G.D., "The computer control of passenger traffic in lift systems", PhD Thesis, The Victoria University of Manchester, 1970.

171. SFPE Handbook of Fire Protection Engineering 5th Edition, Editor: M.J. Hurley, Springer-Verlag New York, ISBN 978-1-4939-2564-3, DOI: 10.1007/978-1-4939-2565-0, 2015.

172. ISO 13571: 2012(E), Life-threatening components of fire – Guidelines for the estimation of time to compromised tenability in fires, International Standard, 2012.

173. David A. Purser and Jamie L. McAllister, Assessment of Hazards to Occupants from Smoke, Toxic Gases, and Heat, Chp 63, The SFPE Handbook of Fire Protection Engineering, 5th Edition, M.J. Hurley (ed.), pp 2308 – 2428, 2016, DOI 10.1007/978-1-4939-2565-0_63.

174. Rosepark Nursing Home. *Sheriff's inquiry report and appendices*. http://www.scotcourts.gov.uk/opinions/2011FAI18.pdf

CHAPTER 12: PUBLICATIONS

A complete and up to date list of all EXODUS publications can be obtained from the website http://fseg.gre.ac.uk/fire/pub.asp

A complete and up to date list of all EXODUS publications available from CMS PRESS can be obtained from the website http://fseg.gre.ac.uk/fire/press.asp