An Object-Oriented Simulation (Crisp II) for Fire Risk Assessment

J. N. FRASER-MITCHELL
Fire Risk Unit, Fire Research Station, Borehamwood, Herts, UK

Abstract

Risk assessment methods are briefly reviewed and the advantages of simulation are discussed. FRS is developing a zone model (CRISP II) of the complete fire system for this purpose. CRISP II is written in object-oriented fashion for maximum flexibility. The types of objects currently modelled are items of furniture, hot and cold gas layers, vents, walls, rooms, alarms, fire brigade and occupants of the building. The occupants exhibit the most complex object behaviour. Their actions in a domestic dwelling are governed by a table of rules. Each action requires the person to move to a specified room, followed by a time delay until the action is complete. The action may be abandoned before completion if room conditions become untenable, or the person is overcome by toxic products. In the final section of the paper it is shown how the model may be used for Monte-Carlo studies of risk in domestic houses.

Keywords: Risk Assessment, Simulation, Human Behaviour, Monte-Carlo Studies.

List of symbols

A_i	area of burning item (m^2)
C_d	discharge coefficient of vent
dt	length of simulation timestep (s)
FED_X	fractional effective dose of species X
h_F	height of the flame (m)
h_i	height of smoke:air interface in room i (m)
L_{vap}	latent heat of vapourisation/pyrolysis ($kJ\ kg^{-1}$)
$\dot{m}''_p$	rate of pyrolysis per unit area ($kg\ m^{-2}\ s^{-1}$)
N	number of samples used to estimate the mean
p_F	probability that a fire is fatal
P(h)	pressure at height h ($N\ m^{-2}$)
$\dot{Q}$	fire's total heat output rate (kW)
$\dot{q}''_{ext}$	heat flux from sources other than flames ($kW\ m^{-2}$)
R	radius of the fire (m)
R_{max}	maximum fire radius (m)
t	simulation time (s)
U(0,1)	random number with uniform distribution between 0 and 1
V_{CO2}	a multiplication factor in FED calculation
w	vent width (m)
Y_X	yield of species X (kg of X per kg of fuel)
Z	height above virtual origin of plume (m)

z_0	height of virtual origin of plume above burning surface (m)
ΔH_{comb}	heat of combustion per kg of O_2 consumed (kJ kg^{-1})
$\Delta P(h)$	difference between room pressure on either side of vent, at height h (N m^{-2})
Φ	configuration factor
λ_r	fraction of fire's heat output as radiation
μ	mean value of a population
ρ_{hot}	density of hot layer (kg m^{-3})
σ_μ	standard deviation of estimates of the mean
σ_{pop}	standard deviation of the population
τ_{grow}	pyrolysis growth time (s)

RISK ASSESSMENT METHODS.

Risk assessment is growing in importance as many countries seek to replace their existing prescriptive building codes with performance–based ones. In order to demonstrate that a building satisfies the performance requirements, some probabilistic estimate of the fire risks is required. Risk assessment in industry is well–established, and uses a number of techniques to predict the consequences of a component's failure, and hence calculate the risk (the product of cost and probability of failure). Attempts have been made to apply these techniques to fires [1,2].

Points Schemes [3,4] would ideally be a simple attempt to correlate fire statistics with parameters such as size of building, fire load, presence of sprinklers, alarms, etc. It is difficult to assign numerical values to some of the parameters, so the correlations will probably be rather weak, leading to large uncertainties in the final measure of risk. The method is also not applicable to novel buildings/techniques, since the statistical data required is not available. In practice, the points awarded for various parameters may be subjective, based on expert opinion.

A more sophisticated approach is based on the use of **State Transition Models** (eg.event trees), with probabilities assigned to each event. The approach is widely used, with research in a number of countries independently producing similar models [5]. The determination of the probabilities may also be a rather subjective process (Delphi technique), or may use deterministic models to examine the consequences of various starting conditions. Some deficiencies in the use of simple event trees have been recognised, and more complex “fire realm models” have evolved in an attempt to overcome these. The transition probability between events will in general not be constant, but conditional on a number of factors, and also time-dependent.

The number of interactions between the components may be large, so many conditional probabilities have to be calculated. However this is done, there is a danger that some interactions may be omitted. For example, suppose a fire growth sub-model is used to generate hot gas, which is then used as input for a smoke movement model to transport the hot gas around the building. The movement of smoke in the fire room depends partly on the movement in adjoining rooms. The level of smoke will have an effect on the fire growth. However the fire growth has been predetermined, so if the smoke level differs from that assumed by the fire growth model, this information cannot be used.

In order to include all interactions between components, we need a **Simulation Model** of the entire fire system. With all component sub-models running simultaneously, and communicating their results to one-another, no information needs to be ignored. The model is mainly deterministic, but with starting conditions and certain values drawn from suitable probability distributions. This represents not only uncertainty in the data, but also the results of stochastic processes. The overall risk is given simply by the average value of some output parameter (eg. casualties) over a suitable number of runs of the simulation.

The advantage of this technique is that the structure of the model can in principle be based directly on physical theory and experimental measurements. A subjective element remains, in terms of the degree of approximation used in the component sub-models (this is true for any method), and also "best guesses" in cases where data is unavailable. However we are not aware of any cases where simulation models have been extended beyond isolated aspects of fire engineering. Numerous deterministic models exist (in a recent survey [6], over 50 have been identified) but all only deal with certain parts of the whole system (eg. fire growth in a compartment, smoke transport between rooms, or evacuation of people). Combining these existing models into a complete simulation is not practical, so we must take the basic equations and recode them in a suitable form.

The FRS developed a model of the complete fire system (CRISP) in order to demonstrate the feasibility of this simulation approach [7]. This model has undergone extensive further development in the modelling of various physical processes. During this work it was decided to alter the structure of the model in order to use object-oriented programming techniques to ease the development, and increase the flexibility of the model.

OBJECT-ORIENTED SIMULATION.

The basic concept of object-oriented programming (OOP) is that a system may be treated as a collection of objects. Each object usually corresponds to a physical component of the real-world system. It is represented by a section of program, which defines the object's behaviour in response to stimuli (input data). The objects may interact in a number of ways, depending on the exchange of information between them. Data associated with an object may only be changed by that object's code, but may be read as input by other objects.

Although the behaviour of individual objects may be fairly simple and well understood, the behaviour of the system will be complex due to the large number of interactions occurring simultaneously. With each increment of simulation time, each object determines its behaviour depending on its input data, and modifies the value of its properties accordingly. The advantage of object-oriented programming is that, given the same stimuli, an object will respond in the same way regardless of what the rest of the system is doing. This facilitates the development of very flexible simulations, as objects can simply be added or removed as desired.

CRISP II is a zone model. This is ideally suited to the use of OOP, since each zone can be represented by one object. The object classes in CRISP II are: Items of furniture, Hot Gas Layers, Cold Air Layers (implicit at present), Vents between rooms or leading to the outside, Walls, Rooms, Smoke detectors, Fire Brigade and Occupants. The behaviours of the different types of object are described in the next section. Further object classes may be added as the model is developed in the future. OOP makes model extension easy, since the new object has only to react to data provided by other objects. The manner in which that reaction is programmed is of no concern to the other objects. (It may however be necessary to extend the programs of the other object classes, if the new object requires data that is currently not provided, likewise if the new object provides data that currently cannot be used.)

OBJECT CLASSES IN CRISP II

Burning Items

The behaviour of a burning item has three main stages: the conversion of fuel to the pyrolysed state, the conversion of pyrolysed fuel to fire products, and the transport of fire products to the hot layer together with air entrained by the plume.

For simplicity we regard all items of furniture as cylindrically-shaped stocks of fuel. Flaming combustion is assumed to start in the middle of the upper surface, and spread with a constant radial speed until the fire radius equals R_{max} (= $(A_i / \pi)^{1/2}$).For future development we will investigate making the flame spread speed a function of radiant heat flux and oxygen concentration. Once a suitable proportion of the fuel has been consumed, the area covered by the flame is considered to be "burnt out" and the pyrolysis rate per unit area decays to zero.

In a steady fire, the rate of pyrolysis from the surface covered by the flame is given by [8]

$$\dot{m}_p'' = \left(\frac{1}{2}.\lambda_r.\Phi.\dot{Q} + \dot{q}_{ext}'' \right) \Big/ L_{vap} \qquad \{1\}$$

where the symbols are defined in the Glossary. A detailed treatment of the transient behaviour of pyrolysis would need to consider the conduction of heat through the item, and the pyrolysis rate as a function of the surface temperature. The chief requirement is that the pyrolysis rate does not change the instant there is a variation in the incident heat flux. We represent this behaviour by calculating a target value of m_p" according to equation (1), and allowing the actual pyrolysis rate to approach the target by means of an exponential growth or decay.

$$\dot{m}_p''(t) = \dot{m}_p''(t - dt) + \left(dt/\tau_{grow}\right).\dot{m}_p''(t - dt) \qquad \{2\}$$

A similar equation can be witten for the decay. To reduce oscillations in m_p" we do not allow the change to overshoot the target value. The growth and decay times and the flame speed may be regarded as "free parameters" which can be adjusted until the model reproduces experimentally-derived burning rates for the item.

The pyrolysed fuel from the item is then converted into combustion products. The yields of the products depend on the ratio of oxygen to fuel [8,9]. In the latter reference this is defined by the "plume equivalence ratio", ie. the pyrolysis rate divided by the oxygen entrainment rate from the cold layer, normalised by the stoichiometric ratio for complete combustion. We specify the yields (per kg of pyrolysed fuel) of CO_2, CO, unburnt fuel, O_2 (consumption) and smoke as functions of the equivalence ratio, by means of look-up tables. The smoke yield is determined in terms of a smoke "optical mass" which is the product of the smoke mass produced in each timestep and the specific optical density of the smoke. There is some evidence that the specific optical density depends on the type of combustion [10] but for the time being we just use a constant value for the specific optical density of 7600 m^2/kg. The overall optical density of the hot layer is the total smoke optical mass divided by the layer volume.

The heat output of the flames is proportional to the rate of oxygen consumption

$$\dot{Q} = \left(\dot{m}_p''.\pi.R^2\right).Y_{O2}.\Delta H_{comb} \qquad \{3\}$$

Some fraction λ_r of the heat is lost as radiation, the remainder is convective and powers the fire plume. We take λ_r = 0.35. Only half of the radiation from the flame will be directed inwards from the flame surface (hence the factor of 0.5 in eqn (1)). The fraction of inward radiation impinging on the pyrolysing surface is given by the configuration factor Φ [11] for a flame surface (assumed conical) of radius R and height h_F. We use Heskestad's correlation [12] for the height of the flame as a function of heat output $\dot{Q}$ and fire radius R. We take Φ=1 if the flame height is calculated to be less than 0. We do not account for increased flame height in vitiated atmospheres.

The final stage concerns the plume, which entrains air from the cold layer and deposits it in the hot layer, together with combustion products and heat from the flames. The entrainment rate at some height Z [13] is proportional to $\dot{Q}^{1/3}.Z^{5/3}$ where Z is the height above the virtual origin. The height of the virtual origin above the fuel surface is a function of R and $\dot{Q}$ [14]. We calculate the mass transfer to the hot layer using $Z=h_i-z_0$, where h_i is the height of the hot:cold interface.

Burning items also produce noise depending on the type of material (eg. crackling of wood) which may alert the occupants.

Hot Gas Layers

Mass, carbon dioxide, carbon monoxide, oxygen, fuel gas, heat and smoke are transferred to the hot layer by the plume(s) of any fire(s) burning in the room. The mass transfer depends on the entrainment rate in the plume (neglecting the mass of the fuel products). Other quantities are simply transferred according to the yields calculated above. The exception is oxygen, which is entrained from the cold layer and also a (negative) yield.The negative yield may exceed the amount entrained; the deficit is made up from the hot layer (by adding a negative quantity to it).

Material may also flow between hot layers in adjoining rooms via vents. Obviously the layer must extend below the top of the vent before outflow can occur. Heat may also be lost by convection/conduction through the walls and ceiling of the room, and by radiation.

Vents

The movement of hot gases through vents is by buoyant flows [15]. The flow rates may be calculated by integrating the Bernoulli equation over the entire area of the vent. To make this easier, the heights of the hot/cold interface either side of the vent, and the heights of the top and bottom of the vent, are ranked in ascending order and used to define 3 regions of the vent.For vents leading to the outside, h_i(outside) is undefined, assumed infinite, so there are only 3 heights defining 2 regions. The flows through each region are worked out separately. No flow is possible if the region is above the top or below the bottom of the vent. The pressure in each room at the four heights is firstly calculated, then the pressure difference, and from this d(ΔP)/dh for the three regions. Each region is then checked to see whether a neutral axis exists. Finally the mass flow rate for the region is calculated

$$\dot{m} = \int_{h1}^{h2} C_d . w. \left(2\rho_{hot} . \left(\Delta P(h1) + (h - h1). \frac{d(\Delta P)}{dh} \right) \right)^{1/2} . dh \qquad \{4\}$$

The integrals are solved analytically for the various cases of whether or not a neutral axis exists, sign of d(ΔP)/dh, whether ΔP is positive, etc.If a neutral layer exists there may be inflow and outflow of hot gases; the appropriate value of ρ_{hot} must be used.

Flow rates are only calculated for hot gases. The flows of cold are assumed to be such that the total volume of gas in each room remains constant. The algorithm is designed to incorporate a non-zero ΔP_{floor}, to represent a pressurised building. However it is considered desirable to avoid modifying ΔP_{floor} at each timestep in order to get the correct inflow rates for cold air. This is not only computationally difficult and time-consuming, but contrary to OOP to make one object's behaviour dependent on data that should be inaccessible (the flows through all the other vents).

Vents have another important function, to permit the passage of people. They may present both a physical and psychological barrier. Data on the type (door or window, ground,1st or higher floor) and status (open, closed or locked) are available to the people objects to determine transit time and desirability of use. (See Occupants)

Glass in closed windows may break if the temperature difference between the two sides exceeds about 90K [16]. This breaking makes a noise which may alert the occupants of the building. The vent will open by a random fraction between 0 and 1.

Cold Air Layers

Ambient conditions are currently stored as global variables accessible to all objects. In future cold air layer objects will be developed, with similar characteristics to hot layers. Cold air layers will be particularly important for smouldering fires, where the transport of fire products will be by diffusion rather than buoyant flows.

Walls

Heat from the hot layer is simply absorbed. For now an infinite specific heat capacity is assumed, so the wall remains at ambient temperature. The wall positions determine the size of the room.

Rooms

The dimensions of the room affect the behaviour of other objects, such as the hot layer depth.

The tenability level reflects the degree of undesirability of an occupant remaining in the room. It takes integer values between 0 and 5. (These values correspond to vent degrees of difficulty - see Occupants. A tenability level of 3 is as undesirable as leaving a house by a ground-floor window rather than the front door. A level of 4 corresponds to jumping out of a first-floor window.) Factors affecting the tenability level are the radiation level, the depth of the cold layer, the temperature and obscuration of the smoke, and the difficulty in breathing (represented by the increased respiration rate as the CO_2 concentration rises). The room also calculates whether conditions would be sufficient to alert awake or sleeping occupants. The thresholds are given in Table 1 above, based on the author's subjective interpretation of values published in the literature [17]. The tenability level is reached when any one factor's threshold is exceeded, however in the case of hot layer temperature or relative respiration rate the clear depth criterion must also be satisfied.

The noise level in the room is the sum of that produced by burning items, breaking glass and activated smoke alarms.

TABLE 1. Threshold values of various factors affecting room tenability

Tenability	Heat Flux	Obscuration*	Clear Depth	Hot Layer Temp.	Resp. Rate
level 1	0.5 kW m-2	0.05	1.0 m (or less)	+ 20 K	x 1.1
level 2	1.5 kW m-2	0.15	1.5 m (or less)	+ 40 K	x 1.2
level 3	1.5 kW m-2	0.5	1.5 m (or less)	+ 80 K	x 1.5
level 4	1.5 kW m-2	1.5	1.5 m (or less)	+ 150 K	x 2.0
level 5	1.5 kW m-2	3.0	1.5 m (or less)	+ 250 K	x 2.5
alert awake	0.5 kW m-2	0.05	1.0 m (or less)	+ 20 K	x 1.1
alert sleeper	1.0 kW m-2	-	0.8 m (or less)	+ 40 K	x 2.0

* Obscuration is defined here as the product of optical density and pathlength.

Smoke Detectors

These are modelled very simply at present. If the optical density of the smoke in the hot layer of the room they are situated in exceeds a preset threshold, the detector sounds an alarm.

Fire Brigade

The activities of the fire brigade are modelled in a very abstract manner. Once summoned, they will take a variable time to arrive and set up (based on Home Office statistics [18]). After this has happened, any occupants still in the building are assumed to be rescued.

Occupants

The behaviour of the people is much more complex than any of the other objects. Firstly we consider their **physiological reaction** to the fire. The uptake of various toxic compounds is expressed in terms of their fractional effective dose (FED) [17,19]. This is defined for carbon monoxide, oxygen deficiency, carbon dioxide and convective heat. The reference also gives an expression for the FED due to hydrogen cyanide, but as the yield of this substance is not calculated by the burning item, we do not include this factor.The dose rates are integrated over each timestep, and unconsciousness occurs when one of the following three conditions are met

$$\left(FED_{CO} + FED_{HCN}\right).V_{CO_2} + FED_{O_2} > 1.0 \qquad \{5a\}$$

$$FED_{CO_2} > 1.0 \qquad \{5b\}$$

$$FED_{heat} > 1.0 \qquad \{5c\}$$

People also have a number of **sensory perceptions**. They may see smoke, hear strange noises, feel heat or notice an increase in their respiration rate (due to increased CO2 concentration). Most of these perceptions are handled by simply defining thresholds for room conditions to alert awake or sleeping occupants (see Table 1). People may also sense the tenability level of an adjoining room, if the connecting door is open. The noise heard will be the sum of the noise level in the person's current room, plus the noise levels in all the other rooms of the building, modified by attenuation factors to account for distance and whether doors are open or closed.

Once alerted to the fire, the people may undertake various **behaviours**. These are given in Table 2 below. Each action requires movement to a specified room (which may be the one currently occupied in some cases), followed by a time delay until the action is complete. In an earlier paper [20] some of the factors affecting decisions and delays are discussed.

The first stage of initiating a new action requires a destination to be decided. To make this process easier we allow the people access to information they should not have, eg. if investigating the fire, they will know which room the fire is in. If warning / rescuing others, they will know in which room(s) they are to be found. We are not interested in the actual procedure of searching the building, merely in moving the people in a manner that resembles a search, and causes them to take up the appropriate toxic doses.

The next stage requires a route to be determined. Each person carries a mental map of the building, reflecting their state of knowledge of the vents (degree of difficulty) with the vents from each room ranked in the order of utility for getting to any other room in the building, or outside. The algorithm choosing the route will pick one with the lowest possible degree of difficulty (DOD) for all vents; vents from a room with the same DOD will be chosen in the order of their utility ranking. The algorithm iterates until a route to the destination has been found, or the highest permissible DOD for the activity has been exceeded. The maximum allowed DOD's associated with various activities are given in Table 2.

In a domestic building, the vents are assigned the following DOD's depending on their type (Internal door = 1, Front door = 1, Back door = 2, Ground floor window = 3, First floor window = 4, Second or higher floor window = 5). However if a vent is known to be fully or partially blocked by smoke, this knowledge may be included by modifying the DOD of the vent. The vent DOD takes the value of the tenability level of the room beyond, if this exceeds the basic DOD of the vent. The DOD of the vent may therefore depend on the direction of transit.

The route–choosing occurs instantaneously. The person will then move following the route, at the appropriate movement speed.

At each timestep of the model, the person must decide whether to continue following the route, or choose some alternative action. A decision will be made if the destination has been reached, the room conditions become untenable, or a vent is found to have a higher DOD (due to untenable conditions in the room beyond) than believed when the route was chosen.

At present the decision process is purely deterministic, and is a function of the person's type (Leader, Led or Dependent), local room conditions, and previous action. In future a random element may be incorporated. The rules are summarised in Table 2. These rules apply to a domestic situation, and would differ for public buildings.

Each person works through the list of rules, starting from their current action until an appropriate new action is found. For example, suppose a house has two occupants (one Leader, one Led) who are both alerted by a smoke alarm. After investigating the fire and abandoning fire-fighting, the Leader's next action should be to Warn Household. This becomes the current action, but is instantly abandoned (because the Led person is also alert). The new current action becomes Rescue, but this is also instantly abandoned (because there are no Dependents to rescue). Finally, the Leader chooses to Escape, and will endeavour to leave the house. To an observer, the Leader would appear to escape immediately after abandoning fire-fighting.

TABLE 2. Rules of Human Behaviour For Domestic Dwellings

Action	Allowed DOD	Destination	Action when complete	Action when abandoned [1]
Asleep	-	current room	Wake Up	Wake Up [2]
Escape	4	outside	Safe	not applicable
Fight Fire [3]	2	fire origin	Go To Water	Warn Household
Go To Water [3]	2	kitchen / bathroom [4]	Fight Fire	Warn Household
Investigate [5]	2	fire origin	Go To Water [3]	Warn
Household				
			Warn Household [6]	
Leave Room [7]	5	nearest tenable room [8]	previous action [9]	not applicable
Rescue [5]	4	most helpless dep.	Escape	Rescue / Escape [10]
Safe	-	outside	Rescue [3] Rescue / Warn Nbr. [11] Safe [12]	not applicable
Unconscious	-	current room	not applicable	not applicable

TABLE 2. Rules of Human Behaviour For Domestic Dwellings (continued)

Action	Allowed DOD	Destination	Action when complete	Action when abandoned[1]
Waiting	4[13]	current room	Waiting	Investigate[14] Warn Household[15]
Wake Up	-	current room	Waiting	not applicable
Warn Household	4	dominant unalerted[16]	Warn H'hld/Rescue[17] Waiting[19]	Warn H'hld/Rescue[18]
Warn Neighbour[6]	-	outside	Safe	Safe[20]

Notes

1. If action abandoned because route becomes impassable, person will attempt to find an alternative route before starting a new action.
2. Abandoned when alerted.
3. Leader only.
4. Choice of room depends on location of fire origin.
5. Not Dependents.
6. Led only.
7. Over-rides current action, if tenability of current room exceeds that allowed for current action.
8. The room just left, if more tenable than current; otherwise the most tenable adjoining room (or current room, if that is best).
9. Re-attempt the action that was over-ridden.
10. Rescuing the chosen Dependent will also be abandoned if a more dominant person is going for the same target. The person will attempt to rescue someone else, or escape if no Dependents require rescue or those that do cannot be reached.
11. Led only, will warn neighbour only if no Dependents require rescue, or those that do cannot be reached
12. Dependent only.
13. Person will be alerted before tenability reaches this level.
14. Leader or Led, once alerted.
15. Dependent, once alerted, and no Leader or Led alerted.
16. If an unalerted person is encountered en-route, this person will become the chosen target.
17. Warn household unless all have been alerted, in which case start rescue.
18. Warning the chosen person will also be abandoned if a more dominant person is going for the same target. The over-ruled person will attempt to warn someone else, or rescue if no-one requires warning or those that do cannot be reached.
19. Dependent only, once a Leader or Led person is alerted.
20. Warning the neighbour will only be abandoned if the person has done this action before.

USE OF SIMULATION FOR MONTE-CARLO ESTIMATES OF RISK

As described above, the simulation is deterministic. The stochastic aspects fall into two categories - random processes, and random initial conditions (or uncertain data). Examples of the former could include flame spread and pyrolysis growth/decay, noise generation, and the delay times for various activities. Examples of the latter could include the type and location of the item first ignited, the number of occupants, their location and attributes, whether vents are open or closed, etc. In priciple any process or parameter could have random values. Whether it is worth taking the trouble to incorporate these dependes on the degree of uncertainty of the value, and the sensitivity of the overall result to the value.

It is intended that the simulation would work on buildings with a fixed floor plan. For studies concerned with general classes of buildings, we would choose a 'typical' example (or repeat the simulation for a number of examples, and average the results). The contents of the building are also fixed. For future development however, it would easily be possible to assign each item a probability of being present, or to choose an item from a list of possibilities.

In a study we are currently undertaking, we are interested only in fires in domestic houses of 2–3 storeys. The Home Office collects and stores information [18] (to which FRS has access) on all fires attended by the fire brigade. We can select fires from this database according to one or more criteria. Of all the fires in buildings of this category, we subdivide according to the season of the year(spring, summer, autumn, winter), time of day (8–hour intervals), the room of fire origin (kitchen, lounge, bedroom, hall or corridor, other/unknown) and the type of item that first ignited (cooking, furniture, structure/fittings, waste/paper, electrical, gases/liquids, other/unknown). For illustrative purposes, a small portion of this database is shown in Table 3.

For each replication of the simulation, the fire scenario is determined by the Rejection Method [21]. We randomly select the season, time, room and item. (We must make sure that each room contains at least one example of each type of item, at least for the most common fires) These 4 parameters define the scenario. We then generate a random number U(0,1) and if this is less than the relative frequency of the scenario, the scenario is accepted, otherwise it is rejected. If the scenario is rejected, we generate a new scenario and try again, until a scenario is accepted.

Families of varying composition are assisgned probabilities of occupying the building. The family type for each replication is chosen by the Rejection Method. This defines the number of and type of people present. According to each person's "occupation" ((un)employed adult, infant, child, elderly, invalid) we assign probabilities of sleeping, and being in the house if awake, as a function of the time of day. If indoors and awake, we assign the probability of being in a given room.

The strategy for locating smoke detectors (none, one, one per floor, or one in each room) is fixed for each set of replications (to facilitate comparisons between strategies). The actual location of detectors is determined by the Rejection Method from a list of probabilities for each room.

The probability of vents being open or closed depends on the season and time of day.

Once all the input conditions have been defined, the simulation then predicts how the scenario develops with time, until the fire is put out or all the occupants are dead or have escaped. The simulation is repeated many times and the distribution of the number of casualties versus frequency is built up. The mean and variance of this distribution is calculated after each run. The mean number of casulaties defines the overall (relative) risk given that a fire has started. (For absolute risk one must know the probability of ignition). Note that it is unnecessary to consider every possible fire scenario in order to derive this figure. The scenario selection by the Rejection Method automatically ensures that the most common scenarios are given greater weight (ie. simulated more often).

TABLE 3. The frequencies of various fire scenarios, from Home Office statistics. The season is Summer, and the time period 1600-2400h.

Item first ignited	Bedroom	Lounge	Kitchen	Bathroom	Hall/Stairs	Other/Unknown
Food/Fat	0	1	274	0	0	2
Furniture	60	66	4	0	1	6
Structure/Fittings	9	15	7	5	1	84
Waste/Paper	7	10	9	3	4	20
Electric Insulation	26	14	15	5	2	35
Gases/Liquids	3	7	12	2	4	18
Other/Unknown	54	32	29	3	13	57

According to the Central Limit Theorem, the standard deviation of the mean is given by

$$\sigma_{\mu} = \frac{\sigma_{pop}}{N^{1/2}} \qquad \{6\}$$

We can say that the true mean of the underlying population lies in the range $\mu \pm 2\sigma_{\mu}$ with 95% confidence.

How many replications would be needed to estimate the mean with an accuracy of (say) 20% ? To get an idea of this number, consider a slightly different problem: how many cases of 1 or more fatalities would we expect in n replications ? Each replication is independent, so if the probability of a fire being fatal is p_F, the number of fatal cases will follow a Poisson distribution, with $\mu = \sigma^2 = N.p_F$. For 20% accuracy we need $N.p_F = 100$. In 1992, about 15% of fires attended by the Fire Brigade involved injuries, and about 1% of fires involved fatalities [22]. It is estimated that brigades only attended 13% - 17% of household fires [23] (the remainder being small enough for the occupants to extinguish unaided). The probability of a fire being fatal is therefore roughly 0.15%, and 6.6×10^4 replications would be needed to estimate this with 95% confidence limits of ± 0.03%.

A number of variance reduction techniques [24] may be applied to reduce the required number of replications. By importance sampling, we may incorporate prior knowledge of the likely consequence of a particular fire scenario into the probability of selection. In practise this would involve re-generating our database (of which Table 3 is a part), selecting only those fires from the Home Office database which occurred in 2-3 storey houses and resulted in one or more casualties. Alternatively we may exploit a correlation between the number of injuries and the number of fataltiesm(if one exists). The proportion of fires resulting in injuries could be estimated to 20% accuracy in about 4000 replications (using the figures above).

For some purposes it may be sufficient to specify only an upper limit on the probability of a fire proving fatal. The probability of observing no fatal fires in n replications is $(1-p_F)^N$; the largest value of p for which there is not more than a 5% probability of this happening is given by $N.\log(1-p_F) = -1.3$. If we observe no fatal fires in 3000 replications, we may state with 95% confidence that the proportion of fatal fires is less than 10^{-3}.

SUMMARY

Simulation methods appear to be a promising approach to the assessment of fire risk. FRS is developing such a model (CRISP II) for this purpose. Object-oriented programming eases the development, and maximises the flexibility of the model. Although the model is currently deterministic from a given set of starting conditions, it would be straightforward to incorporate random aspects into the behaviour of the objects. The model will be used in Monte-Carlo studies to estimate the fire risk in a given building. It is not necessary to examine every possible fire scenario in order to do this.

ACKNOWLEDGEMENTS

The author would like to acknowledge Mr. WGB Phillips for many useful discussions during the course of this work. Mr. BB Pigott assisted with the development of the rules for human behaviour. Mr. SE Chandler extracted the data from the Home Office statistics. This work was funded by the Construction Directorate of the Department of the Environment.

REFERENCES

1. CONCAWE's Ad-Hoc Risk Assessment Group, "Methodologies for Hazard Analysis and Risk Assessment in the petroleum refining and storage industry - part I", Fire Technology, 20 : 3, 23-38, 1984
2. CONCAWE's Ad-Hoc Risk Assessment Group, "Methodologies for Hazard Analysis and Risk Assessment in the petroleum refining and storage industry - part II", Fire Technology, 20 : 4, 43-56, 1984
3. Watts JM Jnr, "Criteria for fire risk ranking", Proc. 3rd Int. Symp. on Fire Safety Science pp.457-466, Elsevier Science Publishers, 1991
4. Donegan HA, Taylor IR, Meehan,RT, "An expert system to assess fire safety in dwellings", Proc. 3rd Int. Symp. on Fire Safety Science, pp.485-494, Elsevier Science Publishers, 1991
5. Bukowski RW, "A review of international fire risk prediction methods", Interflam '93, pp.437-446, Interscience Communications Ltd, 1993
6. Friedman R, "An international survey of computer models for fire and smoke", J. of Fire Prot. Engr., 4: 3, 81-92, 1992
7. Phillips WGB, "The development of a fire risk assessment model", BRE Information paper IP 8/92, 1992
8. Tewarson A, "Generation of heat and chemical compounds in fires", in SFPE Handbook of Fire Protection Engineering, ed. PJ DiNenno, Chapter 1-13, pub. NFPA, 1st edn 1988
9. Gottuk DT, Roby RJ, Peatross MJ, Beyler CL, "Carbon Monoxide production in compartment fires", J. of Fire Prot. Engr., 4: 4, 133-150, 1992
10. Mulholland GW, "Smoke production and Properties", in SFPE Handbook of Fire Protection Engineering, ed. PJ DiNenno, Chapter 1-25, pub. NFPA, 1st edition 1988
11. Tien CL, Lee KY, Stretton AJ, "Radiation Heat Transfer", in SFPE Handbook of Fire Protection Engineering, ed. PJ DiNenno, Chapter 1-5, pub. NFPA, 1st edition 1988
12. Heskestad G, "Luminous Heights of Turbulent Diffusion Flames", Fire Safety Journal, 5: 103-108, 1983
13. Zukowski EE, "Development of a stratified ceiling layer in the early stages of a closed room fire", Fire and Matls., 2: 54-62, 1978
14. Heskestad G, "Virtual Origins of Fire Plumes", Fire Safety Journal, 5: 109-114, 1983
15. Emmons HW, "Vent flows", in SFPE Handbook of Fire Protection Engineering, ed. PJ DiNenno, Chapter 1-8, pub. NFPA, 1st edition 1988
16. Skelley MJ, "An experimental investigation of glass breakage in compartment fires", NIST Report NIST-GCR-90-578, 1990
17. Purser DA, "Modelling time to incapacitation and death from toxic and physical hazards in aircraft fires", AGARD Conference proceedings No 467 on Aircraft Fire Safety, pp. 41-1 - 41-12, NATO, 1989
18. Home Office Statistical Department, national fire data bank.
19. Purser DA, "Toxicity assessment of combustion products", in SFPE Handbook of Fire Protection Engineering, ed. PJ DiNenno, Chapter 1-13, pub. NFPA, 1st edition 1988
20. Fraser-Mitchell JN & Pigott BB, "Modelling Human Behaviour in the Fire Risk Assessment Model CRISP II", CIB W14 Conf. on Fire Safety Eng., Ulster 1993
21. Hammersley JM, Handscomb DC, Monte Carlo Methods, pp.36-37, Meuthen, 1964
22. Home Office Statistical Department, Fire Statistics United Kingdom (provisional 1992 Summary Statistics)
23. Home Office, British Crime Survey 1992
24. Hammersley JM, Handscomb DC, Monte Carlo Methods, Chapter 5, Meuthen, 1964