

Analysis of Shipboard Firefighting-Team Efficiency Using Intelligent-Agent Simulation

LCDR Sylvio F. Andrade, Brazilian Navy

55-21-38496462, sfandrade@hotmail.com, Av. Canal de Marapendi 1700/1306, Barra da Tijuca, Rio de Janeiro, RJ Brazil, CEP 22631-050

Neil C. Rowe, Naval Postgraduate School

(831) 656-2462, ncrowe@nps.navy.mil, fax (831) 656-2814, Code CS/Rp, U.S. Naval Postgraduate School, Monterey CA 93943 USA

Donald P. Gaver, Naval Postgraduate School

(831) 656-2605, dgaver@nps.navy.mil, Code OR/Gv, U.S. Naval Postgraduate School, Monterey CA 93943 USA

Patricia A. Jacobs, Naval Postgraduate School

(831) 656-2258, pajacobs@nps.navy.mil, Code OR/Jc, U.S. Naval Postgraduate School, Monterey CA 93943 USA

ABSTRACT

Analysis of the efficiency of organizational structures is important for command-and-control since an intuitively appealing structure may have unanticipated bottlenecks that occur with unexpected events and the skill levels of personnel. Good simulation can find some of these problems, but it hard to build good simulations. So we have developed tools to enable analysts to set up discrete-event multi-agent simulations in straightforward ways without having to program. We describe here our application of these tools to shipboard firefighting, which presents difficult problems for planners. Our tools uses artificial-intelligence techniques such as means-ends analysis to simulate the actions and interactions of a fire team, and uses a stochastic model for fire spread based on the kind of flammable material, its ignition and burnout rates, and the possibility of burnout and flashover effects. The duration of an action depends on the skill level of the team member in charge of the action. To assess the readiness of teams with different combinations of skills, we measured the duration of firefighting in random trials. We showed that a good scene leader is not sufficient to assure a good performance when unskilled nozzlemen and hosemen are part of the team, and we showed that skill levels did not matter much when ignition and burnout rates were high.

I. INTRODUCTION

We investigate agent-based simulation to make a quantitative assessment of the effectiveness of a team of firefighters when responding to a fire-on-board situation. An artificial-intelligence

planning model simulates human behavior. A stochastic model simulates the fire spread; it simulates the effects of fire in different compartments with different inflammable contents and anti-fire features when a team of firefighters attacks it. The simulation enables performance comparison of teams of firefighters with different skill levels [4].

An important issue in a damage-control situation involving fire is how fast that fire may spread inside a compartment with or without anti-fire devices, based on the time and skill of reaction of a team. Part of that skill is situational awareness. How quickly might a team with different skill levels put out or control a fire in a compartment? What measures of performance and effectiveness best summarize their performance (for example, the expected or median time until a fire is out, expected or median damage to compartment contents, risk/probability of total inflammation/burn up, and probability of injury or death of a team member)?

All these measures depend on factors such as time and correctness of response. What should be the composition of a team to accomplish all the tasks needed to put out a fire? How effective is a sequence of actions that comprise a standard procedure like setting boundaries for possible fire spread, use of anti-fire devices such as hose, foam, Halon, and CO₂, verification that a fire is out, testing of O₂ level, removal of smoke from the area, removal of water from the area? What happens if incorrect or improper orders are given by unskilled team members, or orders are not followed correctly, or are late? How do combinations of such mistakes affect the extinguishing of the fire?

This work developed and tested a multi-agent simulation program module of a firefighting team which integrates unique real-world elements and subjective characteristics of human beings such as knowledge, experience, communication and coordination. Since a physical experiment is very difficult to reproduce and test without any risk to human lives and expensive resources, an agent-based simulation is used to study the effectiveness of a firefighting team. The simulation models the fire and the effects of sequences of actions required to put out a fire. These actions vary with the size of compartment, how the fire spreads, the type of the fire, communications between the team members, and the skill levels of the team members.

The agents in this multi-agent simulation are models of members of the team, the command center, and the fire itself. Agents are defined as software modules with more autonomous capabilities than most software components, typically including reasoning, communication and decision-making. In this simulation, agents react to external events produced by other agents including the fire, but unlike adaptive agents, do not change their rules in response to new environments. Actions (such as approaching the fire, extinguishing the fire, testing the gases, and removing the smoke) require specific preconditions be achieved so that they can be carried out. The skills of the executor influence the duration of each action and the overall time to complete the mission. Agents use orders and reports to communicate and coordinate actions.

The Prolog programming language and a special package are used to provide artificial intelligence for the model of the situation. The artificial-intelligence means-ends algorithm is used to build action sequences from a given start state to a goal condition by a sequence of state transitions. The algorithm permits the modeler to specify the possibility of random events such

as casualties, availability of a doctor, faulty (possibly broken) lines of communication, and equipment malfunctions that can influence the outcome of the simulation.

II. BACKGROUND

A. FIRE FACTS

Once started the fire may go out if its combustibles are isolated [3, 9]. But usually the fire spreads to other material, depending on their mutual configuration. Ignition rates (rates with which combustibles catch fire) differ for solid fuels and, for example, are greater for dust or shavings than for bulky materials. While the fire is burning, the material that has been on fire gradually becomes no longer inflammable; the rate of this is called the burn-out rate (a function of the fire size and the kind of inflammable material). An anti-fire action can affect both the ignition rate and the burnout rate. As a fire is extinguished, the temperature and the ignition rate decrease. During extinguishment, the burn-out rate usually increases due to materials that become unable to burn.

“There are four classifications of fires: class A, class B, class C and class D. Class A fires involve wood and wood products, cloth, textiles and fibrous materials, and paper. Class A fires should be extinguished with water in a straight (stream) or fog pattern (spray). If the fire is deep-seated, aqueous film foam (AFFF) is more effective than seawater and can be used as a wetting agent to rapidly penetrate and extinguish the fire. Class B fires involve flammable liquids such as gasoline, diesel fuels, jet fuels, hydraulic fluid, and lubricating oil. These fires are normally extinguished with AFFF, Halon 1211, Halon 1301, or potassium bicarbonate (PKP). Class C fires are energized electrical fires that are attacked at prescribed distances using nonconductive agents such as CO₂, Halon 1211, or water spray. The most effective tactic for a class C fire is to deenergize the compartment and handle the fire as a class A fire. Class D fires involve inflammable metals such as magnesium and titanium. Application of water in quantity, using fog patterns, is the recommended tactic. When water is applied to burning class D materials, there may be small explosions; the firefighter should apply water from a safe distance, or from behind, shelter.... Most ship compartments do not have installed automatic fire-extinguishing systems such as active sprinklers or CO₂ flooding. When such compartments are on fire, mobile anti-fire apparatus operated by a team of firefighters is needed. However, special rooms such as an ammunition magazine must have active sprinklers to prevent chain reactions, while others such as the Main Machinery Room and the Pump Room must have Halon, CO₂, and AFFF installed inside or nearby.” [9].

B. FIREFIGHTING ORGANIZATION AND PROCEDURES

The most common shipboard emergency is fire [9, 12]. US Navy Ships organize their primary emergency-response team as a fire party made up of people with the different skills needed to

combat emergencies. Specific responsibilities, duties, and employment of equipment must be assigned to individuals. This information is put into a comprehensive form called a Fire Bill and is made available to all personnel addressed in the document. The purpose of the Fire Bill is to establish a firefighting organization and specify responsibilities for individuals and departments to ensure that fires and other related emergencies are effectively and quickly handled. The number of people in the fire party assigned to a fire will vary depending on the nature of fire and the number of people available. There can be a large turnover of personnel assigned to repair parties, so repair party personnel may not be familiar with the location of damage-control equipment.

The member of a team in charge of firefighting at the scene of a fire is the scene leader. He directs the attack against the fire and communicates with the repair party leader via messenger. He must be capable of making correct decisions during a changing situation based on his assessment of the current state. The scene leader is located a short distance from the fire, and his assessment must be derived from the reports he receives. Scene reports can contradict previous reports due to unexpected events such as human errors, personnel casualties, fires reflashing after being previously extinguished, and equipment malfunction. Nevertheless, a scene leader must continually reassess the needed response as a result of the reports. He must always investigate unsatisfactory or incomplete reports from the scene; for instance, if an unskilled member fails to report the end of a task or does not know what to do, the scene leader may order that action again to ensure its satisfactory completion. The scene leader relays information to the damage control assistant (at a place that we refer as the Command Center in our simulation) to keep it updated and to receive further instructions.

The fire attack team consists of an optional attack scene leader, and one or two hoses manned by nozzlemen and hosemen. The attack scene leader directs the nozzleman in employment of the hose, directs the hosemen, directs that the hose be charged or secured, selects the water pattern to be used, and directs the rotation of attack team personnel. Generally, the nozzleman works the hardest and will need to be relieved first; rotating the nozzlemen and the hosemen can extend the endurance of the attack team. The attack team can be supported by other emergency response team members such as an electrician (responsible for deenergizing, reenergizing, and desmoking the area), investigators (responsible for testing for gases and oxygen after fire is out and for monitoring the oxygen-mask control time of nozzlemen and hosemen for rotation purposes), hospital corpsman, and dewatering and desmoking teams when personnel are available.

There is a preferred sequence to completing major actions while combating a fire. For example, if a fire is not extinguished and an order to desmoke is given, the airflow from the desmoking fans will cause the fire to increase in intensity. In addition to strategic knowledge, the scene leader must understand the details of every other member's job and equipment to differentiate between satisfactory and unsatisfactory reports. For example, he must know that 16% oxygen is inadequate to sustain human breathing, whereas 21% is satisfactory indefinitely.

A sequence of actions and reports takes place during firefighting. Reports announce the class of the fire, actions taken to isolate and combat the fire, that the fire is contained, that the fire is out, that the reflash watch is set, that the compartment is ventilated, and that the compartment is tested for oxygen, flammable gases, or toxic gases. Some actions must be coordinated with

reports from the fire scene. Generally, the first step is to go to the repair locker and assemble equipment. Once at the scene of the fire, the leader appraises the state of fire and determines its precise location using men with breathing apparatus and thermal imagers for large fires involving much smoke. The leader must select the appropriate tactics to attack the fire and choose appropriate tools. Fire boundaries must be established, and the area must be deenergized to avoid potential injuries. Now team members with firefighting equipment can approach the fire. If during extinguishment the team leader perceives that the fire is not decreasing, he may order a change of tools. If the fire situation becomes worse and the fire becomes uncontrollable he may decide to isolate the compartment. However, if the fire decreases and eventually goes out, desmoking and dewatering takes place, and the area is reenergized and ventilated. Then the leader orders a watch for a reflash (means fire can start again); if nothing else happens, the team can return to the repair locker, store the equipment, and receive a debrief of the action.

C. MODELING OF FIREFIGHTING ACTIONS

Firefighting knowledge can be envisioned as a tree of tasks where nodes represent a sequence of linked actions and subactions which team members must execute to recover from a fire [10]. The tree results from a top-down stepwise refinement design. An expandable-action node corresponds to an action which can be broken down into subactions. There is a preferred order for most major firefighting actions if we are to model the way an expert scene leader would handle a fire episode. If a subaction is not completed, a negative report is received, and that subaction must be redone.

A data structure needs to represent possible states during a fire episode. Two states are special: a given start state (set of initial facts) including a fire, its size, type and location (although this is initially unknown to the fire team), and a given goal (state in which objectives have been achieved), such as "all fires are out and the ship is restored to normal operation". Searching through the possible states results in a path of states from start state to goal which can be followed by the simulation. The procedure used to find the path is means-ends analysis [5], a search method using top-down recursive decomposition of a search problem into simpler subproblems. This structure uses a "difference table" showing the recommended major action for any search problem. It is specified by assertions of the form "recommended (<difference list>, <action>)" which gives conditions for recommending an action based on the facts different between the current state and the goal. The operator will not necessarily be able to apply the recommended action immediately, but the action should be the most important one necessary to solve its search problem at any moment.

Each action has preconditions (facts that must be present in a state before we can do the action); these are specified by facts of the form "precondition (<action>, <precondition-facts>)". Once an action is done, the state is changed: some facts become true and some facts become false. For those conditions that become true, an "addpostcondition (<action>, <added-facts>)" fact is used, and for the facts that become false, a "deletepostcondition (<action>, <deleted-facts>)" fact is used. Addpostcondition and deletepostcondition facts convey both intended and side effects of an action, such as "fire is out" and "area is watery" (addpostconditions) and "it is no longer true that boundaries are set" (deletepostconditions) for the "extinguish" action. They also convey

state-dependent effects, such as "there is damage to the floor" when the user forgets to turn the power off before extinguishing the fire.

Firefighting simulations fit well the paradigm of intelligent multi-agent simulations [11]. An intelligent agent is software with autonomy, adaptation and cooperation, an enhancement of a software object [7]. Multi-agent systems are computational systems in which several agents interact to achieve their goals. In the human world, complex jobs are usually performed by several individuals because individual capabilities are limited to some extent. Analogously, agents in the multi-agent system communicate, coordinate, and negotiate with one another to achieve goals. The interaction and synchronization between agents can be accomplished by protocols for communication. Protocols govern the sending of messages between social agents and monitoring of execution. An agent has knowledge of its internal control mechanisms, or in other words, it knows how to make behaviors, problem-solving methods, and strategies work together in order to achieve its goals. Means-ends analysis accomplishes this for the firefighting simulation. Agents can also have purely reactive behaviors, a necessary ingredient for the design of real-time multi-agent systems. A fire agent is an example of reactive agent since its behavior is based on a mathematical model: it responds to other agents' actions (such as extinguishing) or events (a possible flashover).

III. PROPOSED MODEL

A. PROPOSED SIMULATION

The firefighting models proposed here are agent-based with some random behavior. More details are provided in [1]. There are six agents in our simulation: command center, team scene leader, nozzleman, hoseman, electrician and fire. The agents nozzleman and hoseman could represent several people. Each member of the firefighting team has a skill level. This parameter influences the duration of the actions "equip", "extinguish", "desmoke", "dewater" and "watch for reflashing", because these are the most important firefighting actions. Actions have durations which are assumed to be independent random variables whose mean and range are a function of the action, the agent, and the current state. The duration of the extinguishing action, however, depends on the behavior of the fire, as will be explained.

B. A MATHEMATICAL MODEL OF FIRE GROWTH

The fire agent received special attention in our simulation [8]. The fire will only be noticed if it grows above a certain limit inside a compartment. The kind of inflammable will determine the type of fire and thus the necessary tools and tactics. Then the fire will spread inside the compartment until the team of firefighters starts to extinguish it. We used a simple stochastic epidemic-like model for the fire agent [2]. The model assumes that the growth of a fire is a function of the amount of unburned inflammables and the rate with which inflammables catch fire (the ignition rate). The model includes the feature that items on fire can burn out (the combustibility of each elementary item), and once a unit is burned out, it cannot burn again. Both ignition and burn-out rates have random fluctuations during a fire and from fire to fire. The

number of unburned combustibles at any time is a function of the pattern of fire growth. When extinguishing is being carried out, the ignition rate generally decreases and the burn-out rate increases, and both depend on the skill of the team member in charge of extinguishing. If the extinguish action is successful, the fire will eventually go out (so fire size will equal 0). The fire model allows partial flashovers (when a fire grows unexpectedly but ignites only part of the intact (unburned) inflammables), and total flashovers as well. Both may happen with a certain probability. A total flashover ignites all intact combustibles, increasing the size of fire and making the number of unburned combustibles go to zero very quickly. The probability of occurrence of a flashover is a function of the size of fire, the temperature in the region, and the number of inflammables on fire.

The stochastic model for fire spread is as follows. Suppose the shipboard compartment contains a quantity C of inflammable items (in the implementation $C(0)$, the quantity at time 0, is 100). Let us assume that they are all equally inflammable for the present; also, let us ignore spatial considerations for the present. Let $X(t)$ denote the number of inflammable units on fire at time t ; this number is a continuous variable with $0 \leq X(t) \leq C(0)$. The fire starts with $X(0)=X$, and then spreads until items are on fire simultaneously. Items on fire eventually burn out (the combustibility of the item becomes exhausted).

The fire spreading described resembles a SIR (Susceptible, Infectious, and Removed) epidemic model used in epidemiology and public health; see [2]. Let $X(t)$ be the number of inflammable / combustible elements on fire at t , and let $C(t)$ be the number of inflammable / combustible units unburned at t . In epidemic theory $X(t)$ represents infectives and $C(t)$ represents susceptibles. An inflammable element burns for a while and then becomes burned out; but while it is burning it may ignite more inflammable elements. A deterministic model for the number of inflammable elements on fire or fire coverage is

$$X(t+h) = X(t) + \lambda h X(t) C(t)/C(0) - \mu h X(t) \quad (3.1)$$

Here the first term on the right side represents current fire strength, the second the new ignitions in the time interval h , and third the burnout. Since $C(t)$ inflammables are immediately reduced to $C(t+h)$ with fire spread to $C(t)$, we have

$$C(t+h) = C(t) - \lambda h X(t) C(t)/C(0) \quad (3.2)$$

Fire spread can be more realistically modeled as a random process. Suppose that the competition between fire growth and decay can be represented as a deterministic term ("drift") plus Gaussian /normal increments. Then:

$$\begin{aligned} \mathbf{X}(t+h) = & \mathbf{X}(t) + \lambda h \mathbf{X}(t) (C(t)/C(0)) - \mu h \mathbf{X}(t) \\ & + \sigma_\lambda [\mathbf{X}(t), C(t)] \mathbf{Z}_\lambda(t) - \sigma_\mu [\mathbf{X}(t), C(t)] \mathbf{Z}_\mu(t) \end{aligned} \quad (3.3)$$

where the first term on the right side represents current fire strength, the second the mean incremental fire growth, the third the mean incremental burn-out growth, the fourth the random term of the incremental random fire growth, the fifth the random term of the incremental random burn-out. $\mathbf{Z}_\lambda(t)$ and $\mathbf{Z}_\mu(t)$ are independent normally distributed random variables with mean 0 and variance 1, and $\{\mathbf{Z}_\lambda(t)\}$ and $\{\mathbf{Z}_\mu(t)\}$ are mutually independent sequences of independent and identically distributed Gaussian random variables having mean 0 and variance 1. The decline of the intact unenflamed items can be similarly modified to:

$$C(t+h) = C(t) - \lambda h \mathbf{X}(t) C(t)/C(0) - \sigma_\lambda [\mathbf{X}(t), C(t)] \mathbf{Z}_\lambda(t) \quad (3.4)$$

where $\sigma_\lambda [\mathbf{X}(t), C(t)] = (\lambda h \mathbf{X}(t) C(t)/C(0))^{0.5}$ and $\sigma_\mu [\mathbf{X}(t), C(t)] = (\mu h \mathbf{X}(t))^{0.5}$

Note if a random fluctuation in $\mathbf{X}(t)$ is positive, represented by a positive value of $\mathbf{Z}_\lambda(t)$, this must mean a negative fluctuation in the unburned material, $\mathbf{C}(t)$. Note that the last term has a negative sign although $\mathbf{Z}_\lambda(t)$ can be of either sign. It is necessary to make sure that boundary overshoots do not occur when using the difference equations; if overshoot occurs, we replace $\mathbf{X}(t)$ or $\mathbf{C}(t)$ by the boundary value 0 or C.

Equations 3.3 and 3.4 can be modified to allow for total flashover. The conditional probability of a flashover at time t given no flashover has occurred before time t is $\varphi[\mathbf{X}(t), \mathbf{C}(t)]$. If there is a flashover, then $\mathbf{X}(t+h) = \mathbf{X}(t) + \mathbf{C}(t)$ and $\mathbf{C}(t+h) = 0$. If there is no flashover then equations (3.3) and (3.4) are used. The function φ is chosen to increase monotonically with \mathbf{X} so that there is an increasing probability of flashover with increasing fire strength (representing local temperature rise). Examples of φ are $\varphi(t) = [1 - \exp(-\gamma(\mathbf{X}(t)(\mathbf{C} - \mathbf{C}(t))/\mathbf{C}))]$, where the flashover constant parameter γ is 1, and $\varphi(t) = \mathbf{X}(t)/(100\mathbf{C})$.

When the fire is being extinguished the model is assumed similar to the former spread model, with changed values of the ignition rate and the burn-out rate. The new ignition rate is a function of the old ignition rate λ_{Old} and a subjective numerical representation of the skills S of the team member in charge of extinguishing such that the greater the skills, the lower the ignition rate. We use $\lambda_{\text{New}} = (\lambda_{\text{Old}}) \exp(-10S)$. The new burn-out rate is similarly affected by the skill level: $\mu_{\text{New}} = (\mu_{\text{Old}}) \exp(S)$.

C. MODELING OF FIREFIGHTING PROCESSES

To model the human behavior in firefighting we use planning methods from artificial intelligence to determine the sequence of actions performed. A list of facts describes each planning state. The initial conditions for our model comprise the locations where fire may start (such as the engine room or ammunition magazine), the ignition and burn-out rates for the material inside the compartment, the quantity of inflammables (which was 100 for our test runs), the type of fire (e.g. type a, b, c, or d), and the initial size of fire and smoke (e.g. one and zero). Fire will not be noticed until its size is big enough (more than 1.5 in our test runs) to activate an alarm inside the compartment.

Each animate agent (fire scene leader, nozzlemen, hosemen, electrician, or command center), has a specific goal to be achieved, so they have goal-directed behavior. These agents use means-ends analysis to figure what to do when they are unoccupied. Each agent then plans a way to achieve its goals, and chooses one action from the plan to execute first. An agent can be active (doing a task), idle (if its goals are achieved), or waiting (if its goals are not achieved but it has nothing it can do). At each step in the simulation, the active agent that has been least recently updated gets a chance to plan. When an agent is idle it can be awakened by specified triggers such as orders given it or reports received by it.

Each action has a specified set of possible agent "actors"; the agent waiting the longest is assigned the action if more than one agent is available. Some complicated actions also require the presence of assistants to the actor, such as the hoseman to hold the hose while the nozzleman is extinguishing the fire; such actions require that the assistant be waiting as well as the actor for the action to proceed.

There are two different ways to choose an action. With "first doable operator" the agent executes the first action in the plan (generated as explained below) that is possible, whereas with "random doable operator" the agent chooses a random possible action in the plan. A possible action must be permitted by preconditions, recommended at the given time, and must not be in progress already. The implementation uses the "random doable operator".

The simulation of actions also adjusts the states with results of terminating concurrent actions and ensures no two actions end at the same time to prevent confusion in state reasoning. The simulation permits actions to be aborted by other actions in high-priority circumstances. For instance, a fire may reignite at random when a crewman is ventilating a compartment; the crewman then ceases ventilation and initiates immediate fire extinguishment steps. Actions are also aborted when their preconditions become false during their execution.

Means-ends analysis is responsible for the planning for each animate agent. It decomposes a complicated task into simpler subtasks until all subtasks can be accomplished by single actions (see Figure 1). A list of actions needed to achieve particular facts, not necessarily actions that can be performed immediately, is used as recommendation conditions. For example, the "ventilate" action is recommended whenever the gases or oxygen are to be made safe, and they are known to be unsafe. On the other hand, the preconditions of "ventilate" (facts that must be true beforehand) are that the fire team must be at the location of the fire, the team must be equipped, and the fire is out.

The state of the system will be affected by the "ventilate" action as some facts become false and will be deleted from the previous state ("deletepostconditions") and other facts will be true and will be added to the previous state ("addpostconditions"). The postconditions of "ventilate" are that gases and oxygen are safe, gases and oxygen are no longer tested, and smoke is removed. Random postconditions model uncertainties and accidents in the simulation. For instance, when an agent has to test gases, there is a probability (for example, of 0.3) that gases are unsafe. Random changes are applied to the state after regular postconditions are applied.

Each action has a duration which is a random variable whose mean and range can depend on other parameters, such as the skill of the assigned agent, the fire size, the kind of tool used, smoke intensity, and water magnitude. To each team member we assign a number between 0 and 1 that represents a subjective measure of their skill level, where the higher the number the better the skill. Uniform distributions are assumed for durations. For instance, order and report actions have a mean of 0.25 minutes and a range of 0.1 minutes;(the duration of an action has a uniform distribution on the interval (0.20,0.30)); and the action of watching for a reflash has a mean of 15 minutes and range of 5 minutes (12.5,17.5). The time to "dewater" has a mean of the product of 0.15 and (WaterQuantity/Skill) and a range of the product of 0.12 and (WaterQuantity/Skill), and "desmoke" has a mean of the product of 0.8 and (Smoke/Skill) and a range of the product of 0.6 and (Smoke/Skill). Other actions have a mean of 1 minute and a range of 0.5. These are just rough illustrative estimates; determination of trustworthy submodels for the entire process remains a significant practical task. There can well be substantial variability in basic parameters, such as λ , μ , and γ from fire to fire.

The fire model discussed earlier is implemented with a fire agent. A one-minute time step is used by the fire agent to update size of fire and number of intact inflammables. The fire is monitored at time steps of 5 minutes. If fire does not seem to decrease during the extinguishment process, the team realizes this and changes anti-fire tools with a probability of tool failure ($F_{size}/(F_{size}+(100*Skill))$), where F_{size} is the size of fire and $Skill$ is the skill level of the team member in charge of the extinguishment. The model does not allow the number of inflammables to be more than the initial number of inflammables or less than 0.001. This simulation has been implemented by a program written in Quintus Prolog, taking advantage of several Quintus Prolog library modules such as “math” and “random” libraries. The simulation is implemented as two files: a problem-dependent part (“fireagents”) and problem-independent machinery for running multiple agents in means-ends simulations (“meagent”). The latter is complex and includes general control of the agents' actions with code from METUTOR system [6] for procedural modeling of goal-directed behavior.

Firefighting plan tree for deterministic actions

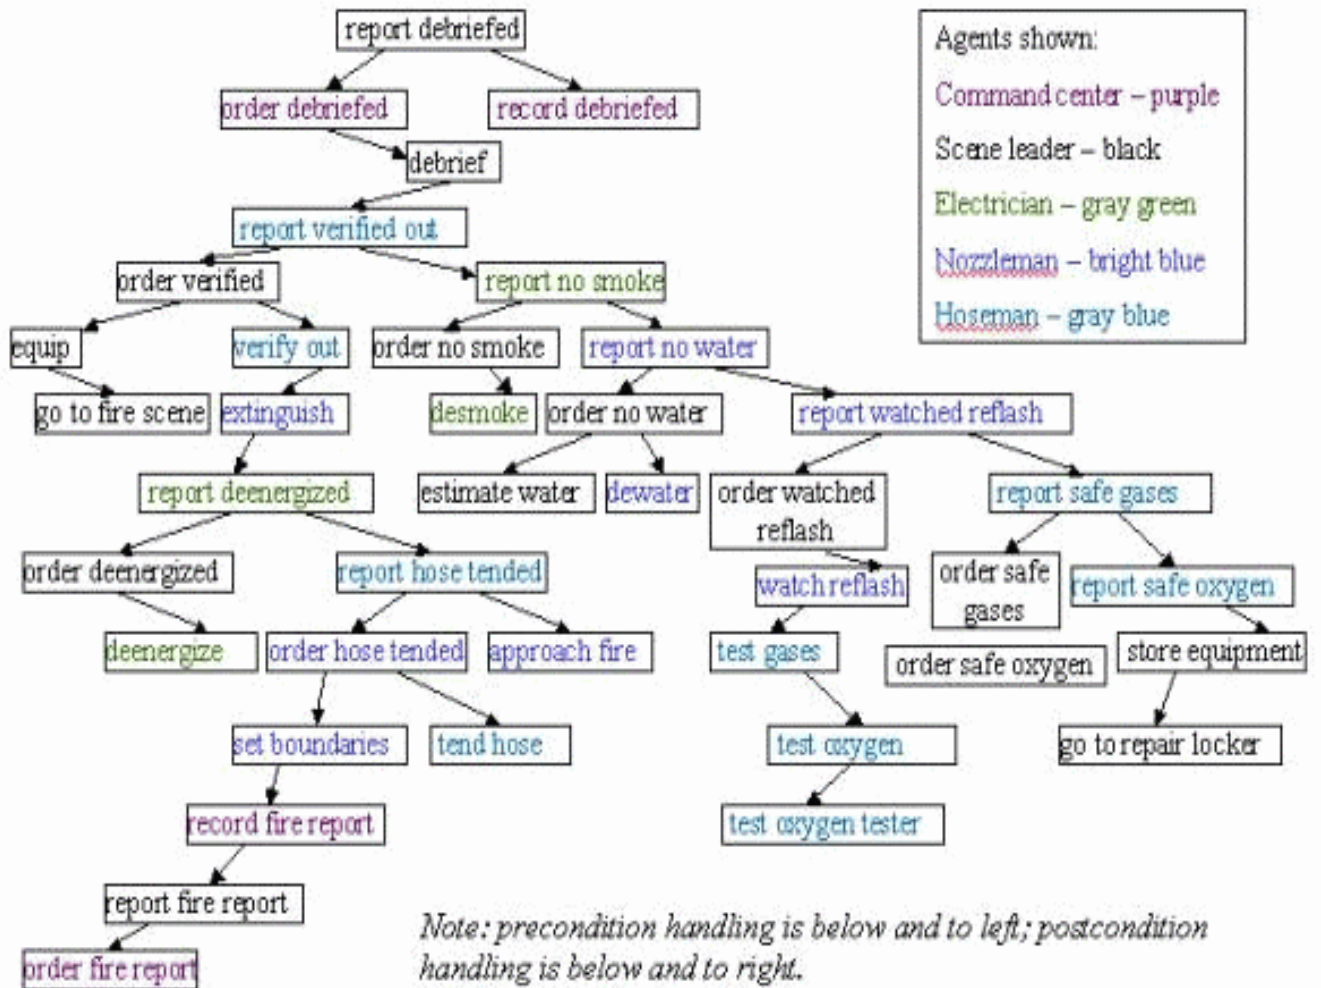


Figure 1: Plan tree for firefighting in a single compartment with no unexpected events.

D. EXAMPLE SIMULATION OUTPUT

Here is an example extract from the output file showing the tasks carried out by each agent. The parameters used for this run were an ignition rate of 0.5, a burn-out rate of 0.25, and skill level s of 0.9 for all team members. Comments are in italics.

Agent fire did

"burn([fire(5.9261,352b),inflammables(89.965,352b),smoke(1.12454,352b)],[fire(7.75211,352b), inflammables(84.7063,352b),smoke(1.25209,352b)])" from 6.0 to 7.0 minutes

The action "burn" done by the fire agent between time 6.0 and 7.0 minutes shows an increase in the fire size inside compartment 352b (from 5.9261 to 7.75211), an increase in the size of smoke inside the compartment (from 1.12454 to 1.25209), and a decrease in the percentage of intact inflammables inside the compartment (from 89.965% to 84.7063%).

Agent scene leader did "go(repairlocker,352b)" from 6.85 to 7.97 minutes.

This means that between minutes 6.85 and 7.97 the scene leader took the team from the repair locker to compartment 352b.

Agent fire did

"burn([fire(7.75211,352b),inflammables(84.7063,352b),smoke(1.25209,352b)],[fire(7.66,352b), inflammables(80.5647,352b),smoke(1.36688,352b)])" from 7.0 to 8.0 minutes

Agent command center did "order (fire(_454788,352b),command center, scene leader)" from 7.97 to 8.24 minutes.

Here the Command Center orders the scene leader to put the fire out in compartment 352b.

Agent scene leader did "order(deenergized(352b),scene leader, electrician)" from 7.97 to 8.23 minutes.

Here the electrician receives an order from the scene leader to deenergize the area of compartment 352b.

Agent fire did

"burn([fire(7.66114,352b),inflammables(80.5647,352b),smoke(1.36688,352b)],[fire(6.556,352b), inflammables(77.0206,352b),smoke(1.47019,352b)])" from 8.0 to 9.0 minutes

Agent electrician did "deenergize(352b)" from 8.23 to 9.41 minutes

Agent scene leader did "order (verified(fire_out,352b),scene leader, nozzleman)" from 8.23 to 8.45 minutes.

Agent scene leader did "report (fire(7.66114,352b),command center, scene leader)" from 8.45 to 8.67 minutes.

The scene leader reports the size of fire in compartment 352b to the Command Center.

Agent command center did "record fire(7.66114,352b)" from 9.0 to 10.02 minutes

Agent fire did

"burn([fire(6.55156,352b),inflammables(77.0206,352b),smoke(1.47019,352b)],[fire(8.276,352b), inflammables(72.1762,352b),smoke(1.56317,352b)])" from 9.0 to 10.0 minutes

Agent electrician did "report (deenergized(352b),scene leader, electrician)" from 9.41 to 9.69 minutes

Agent fire did

"burn([fire(8.27526,352b),inflammables(72.1762,352b),smoke(1.56317,352b)],[fire(11.357,352b), inflammables(65.8635,352b),smoke(1.64685,352b)])" from 10.0 to 11.0 minutes

Agent scene leader did "set (boundaries, 352b)" from 10.02 to 10.9 minutes
Agent nozzleman did "approach (fire, 352b)" from 11.0 to 11.81 minutes

IV. OUTPUT ANALYSIS AND CONCLUSIONS

A. EXPERIMENTS

Simulation trials were carried out to assess teams with different skill levels in firefighting. One hundred trials were done for each of several sets of parameters. Skill levels were illustratively chosen as follows: 0.9 for a high skill level, 0.5 for an average skill level and 0.1 for a poor skill level. There are 27 possible skill combinations for a team of four members if the nozzlemen and the hosemen have the same skill levels. Six cases of special interest were: an ideal team (all members with 0.9 for skill level); an average team (all members with 0.5 for skill level); a poorly-skilled team (all members with 0.1 for skill level); a team with highly-skilled scene leader, average-skilled hosemen and nozzlemen and poorly-skilled electrician; a team with average-skilled scene leader, poorly-skilled hosemen and nozzlemen, and highly-skilled electrician; and a team with poorly-skilled scene leader, highly-skilled hosemen and nozzlemen, and average-skilled electrician. The average skills in all of the above non-homogeneous teams are the same.

Two kinds of fire locations were used: a highly inflammable compartment with pre-extinguishing ignition and burn-out rates of 1 and 0.5 respectively; and a compartment with pre-extinguishing ignition and burn-out rates of 0.5 and 0.25 respectively (the rates are modified when the extinguishing action starts using equations 3.12 and 3.13); both locations had the same flashover probabilities. For each six combinations of skills, one hundred replicate runs were executed, and summary data were recorded.

We assume that the total time for firefighting is the time elapsed between the first appearance of the fire (time 0) and the reporting of the completion of both extinguishment and debriefing. Also, an upper limit on simulation duration of 400 minutes is enforced in cases when the fire burns out but recurs sporadically, never exceeding the alarm threshold. Otherwise, the mean time for the set of one hundred runs could be unfairly high, though the median is a useful alternative summary estimate.

Our hypothesis is that the better the scene leader, the better is the performance of the team. The following representative results are obtained (see Tables 1 and 2):

SL Skill	Nozz/ Hosemen Skill	Elec Skill	Mean Time	Standard Error	Median Time	Mean Intact Inflammables	Standard Error	Median Intact Inflammables
0.9	0.5	0.1	78.63	7.20	58.59	39.46	2.38	30.38
0.9	0.9	0.9	66.55	3.64	59.00	40.58	1.64	41.34
0.5	0.5	0.5	79.69	6.04	62.50	35.42	1.44	34.48
0.1	0.9	0.5	79.81	3.64	74.00	35.11	1.94	33.77
0.5	0.1	0.9	135.25	10.75	86.00	25.75	1.73	20.41
0.1	0.1	0.1	126.80	8.85	93.71	23.78	1.51	20.44

Table 1: Mean and median summaries for 100 runs for ignition rate = 0.5 and burn-out rate = 0.25, ordered by median total time for firefighting.

SL Skill	Nozz/ Hosemen Skill	Elec Skill	Mean Time	Standard Error	Median Time	Mean Intact Inflammables	Standard Error	Median Intact Inflammables
0.9	0.9	0.9	66.55	3.64	59.00	40.58	1.64	41.34
0.9	0.5	0.1	78.63	7.20	58.59	39.46	2.38	30.38
0.5	0.5	0.5	79.69	6.04	62.50	35.42	1.44	34.48
0.1	0.9	0.5	79.81	3.64	74.00	35.11	1.94	33.77
0.1	0.1	0.1	126.80	8.85	93.71	23.78	1.51	20.44
0.5	0.1	0.9	135.25	10.75	86.00	25.75	1.73	20.41

Table 2: Mean and median summaries for 100 runs for ignition rate = 0.5 and burn-out rate = 0.25, ordered by mean total time for firefighting.

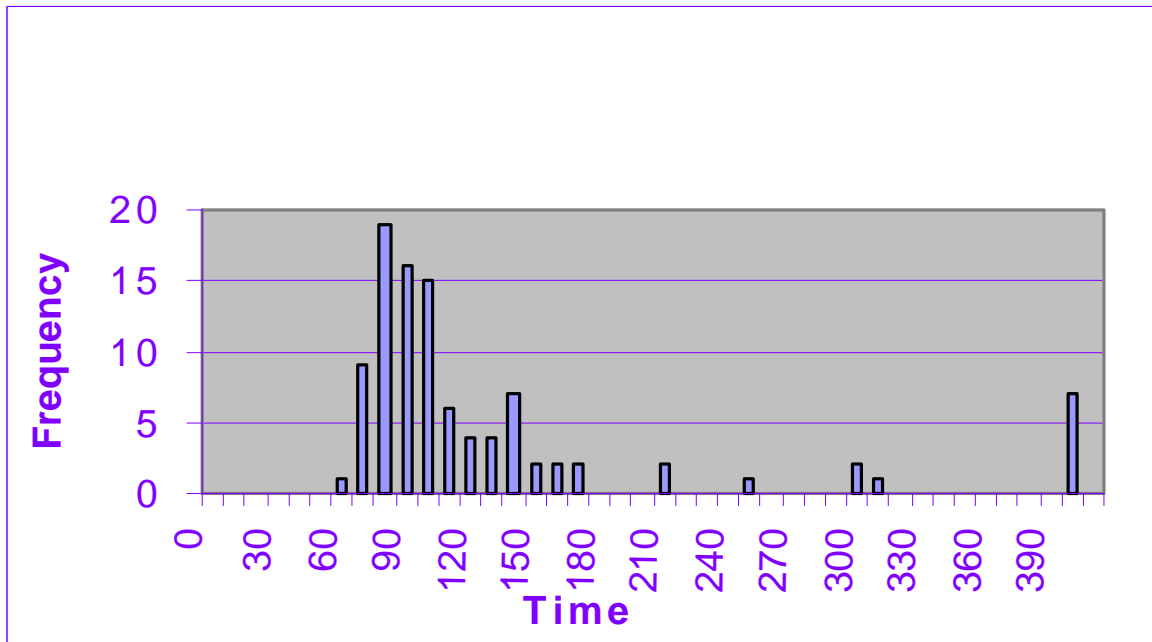


Figure 2: Histogram for total time for firefighting for ignition rate=0.5, burn-out rate=0.25, and all members with skill level 0.1.

Histograms like Figure 2 show that the 400-minute cases for total time for firefighting are more likely to occur when the skill level is 0.1. The tables also show the most important influence on total time for firefighting is low skill (0.1) of the hosemen and nozzle men. The amount of intact inflammables (unburned material) appears to correlate well with average team skill.

We next doubled the pre-extinguishing ignition rate to 1, and burn-out rate to 0.5. Results were similar, except mean and median times are smaller because the fire progresses more quickly. One hundred trials with a constant 35% probability of flashover (added to the original model, and independent of fire state), all members with skill levels 0.5, were also done. The objective of the experiment is to show that a constant probability of flashover increases the number of flashover cases, and decreases the amount of unburned inflammables. The experiment had 39 flashover cases (note that the simulation tests for flashover when the state of fire changes). When compared to a similar experiment (for same skill levels and rates) where the probability of flashover is not constant, the mean intact inflammables decreases significantly.

To analyze the performance of teams when the fire does not have a random component, runs with a near deterministic fire model were carried out for homogeneous teams (all with the same skill level). The fire model is not then completely deterministic, because there is intrinsic randomness in some human-behavior rules that cannot be changed. The results showed that the near deterministic cases have lower mean time than the non-deterministic stochastic cases. The durations of actions for the near deterministic cases have fixed means that depend on the skills, and have no variance. The final amount of unburned material is larger (about the same for case where all skill levels are 0.9 and ignition rate 1 and burn-out rate 0.5) when there is no random fire spread, flashover, and reflash. The effects of randomness on the mean and median times increases more substantially as the skill levels decrease.

B. CONCLUSIONS

The performance of a team of firefighters with the same skills seems to be different for cases with different pre-extinguishing rates in the three different experiments. The total time to complete all actions is the most sensitive to team members skill levels. When the fire spreads faster, and the amount of material that burns out is very high, the percentage of intact inflammables at the end of action is going to be low, no matter what the composition of a team. On the other hand, when fire spreads more slowly, different teams produce different results; the fire seems more difficult to extinguish. Flashover occurs in almost every run with one hundred trials at least once for all cases. The simulation shows that the skills of the electrician do not much affect the overall team performance since an electrician skill level 0.5 and the rest of the team skill level 0.9 (added to the second experiment) do not show a great difference compared to a team with all members having skill level 0.9. Two simulations also added to the second experiment (one for the scene leader skill level 0.5 and the rest of the team 0.9, and other for the hosemen/nozzle men skill level 0.5 and the rest of the team 0.9) show that when the skill levels for the hosemen/nozzle men decrease, the performance of the team also decreases. A skilled scene leader is important for the coordination of actions, and for the readiness before extinguishment; but skilled hosemen and nozzle men, who are responsible for the majority of actions, are even more important.

The methods developed in this work can be applied to the analysis of any activities accomplished by coordination of a team. Examples are anti-terrorist activities, sports teams, and office-task automation. The methods can be straightforwardly extended to explore the effects of changes in organizational structure and new methods for the tasks.

REFERENCES

- [1] Andrade, S. F., "An Intelligent Agent Simulation of Shipboard Damage Control", M. S. thesis, Operation Research Department, Naval Postgraduate School, June 2000.
- [2] Daley, P. J. and Gani, J., *Epidemic Modelling: An Introduction*. Cambridge University Press, 1999.
- [3] Federal Emergency Management Agency, United States Fire Administration, Developing Effective Standard Operating Procedures for Fire & EMS Departments, no date
- [4] Law, Averill and Kelton, David, *Simulation Modeling and Analysis, Second Edition*, McGraw-Hill, 1991.
- [5] Rowe, N., *Artificial Intelligence Through Prolog*, pp. 263-270, Prentice-Hall, Englewood Cliffs, NJ, 1988.
- [6] Rowe, N. and Galvin, T., An Authoring System for Intelligent Procedural-Skill Tutors, *IEEE Intelligent Systems*, 1998.
- [7] J. Rumbaugh, M. Blaha, W. Premerlani, F. Eddy, and W. Lorensen, *Object-Oriented Modeling and Design*. Prentice-Hall, Englewood Cliffs, NJ, 1991.
- [8] J. A. Swartz, R.F. Fahy, E.M. Connelly and D.P. Demers, Final Technical Report on Building Fire Simulation Model - Volume I, revised edition, National Fire Protection Association, Quincy MA, May 1983.
- [9] U.S. Navy, *Naval Ships' Technical Manual Chapter 555 – VOLUME 1-Surface Ship Firefighting*, pp. 1-9, 1-13, 1-22, 1-23, 1-24, 9-1, 1998.
- [10] Weingart, Stephen G., "Development of a Shipboard Damage Control Fire Team Leader Intelligent Computer Aided Instructional Tutoring System." M. S. thesis, Computer Science Department, Naval Postgraduate School, June 1986.
- [11] Zhang, Chengqi, and Lukos, Dickson, Multiagent Systems Methodologies and Applications, Lecture Notes in Artificial Intelligence, Second Australian Workshop on Distributed Artificial Intelligence, Cairns Australia, August 1996.
- [12] Web site <http://ait.nrl.navy.mil/DamageControl/Shadwell.html>.