An Agent-Based Approach for Collaboratively Modelling and Simulating Emergency-Response Organizations

Nicola Bicocchi, William Ross, Mihaela Ulieru

Abstract— **In this paper we present a modelling and simulation capability designed to enable the streamlining of joint emergency-response operations. To this end, we provide four main contributions. First, we specify a collaborative capability to model and simulate emergency-response organizations using state-of-the-art modelling and simulation frameworks. Secondly, using a simple yet relevant example, we illustrate how the proposed capability supports the flexible adaptation of the top-down policies to the crisis dynamics by accommodating the 'bottom-up' emergence of groupings of hybrid resources to respond to the unexpected events occurring 'in the field'. Finally, in order to compare our results, we propose a set of original metrics capable of capturing the effectiveness and reliability of the simulated response under various configurations.**

Index Terms— **Joint emergency-response operations; agent-based modelling and simulation; holistic security ecosystem; effective institutional policies; agile response.**

I. Introduction

Several episodes over the past ten years have exposed serious weaknesses in the emergency-response capabilities of modern countries. In general, these problems were not the result of specific conditions, but rather were the product of complex processes involving more fundamental issues. Due to their very non-linear and complex nature, these phenomena cannot be addressed solely through the application of rule sets developed through rational analysis, as they can neither be definitively described nor optimally solved. Confronting the Public Safety and Security communities reports within the past decade [9, 10, 11, 17], the major challenges that have been identified can be clustered into three major domains.

First, novelty and self-organization. Emergency responders train for predictable and routine events. This allows them to prepare in advance and to take advantage of lessons learned from previous experiences. The National Incident Management System [12] provides a flexible template for managing crisis operations that involve multiple professions, agencies, and jurisdictions. However, organizations have to learn how to adapt on-the-fly when novel situations arise as their rehearsed response plans might prove inadequate.

Secondly, scalability, surge capacity, and situational awareness. In severe incidents, where the number of endangered persons and the extent of the damage quickly surpass what the local jurisdiction is able to cope with, personnel and equipment from other jurisdictions must be brought in. The major challenge is the speed with which the necessary resources can be mobilized, as well as the cohesion of the newly created (overarching) team. Further, decision-makers must gather and assimilate key facts to make sense of a situation, even under evolving conditions with a high degree of uncertainty. They must be able to project forward the implications of the information they have gathered, so they can anticipate the likely consequences of an evolving situation [1, 8].

Finally, operational, political, and jurisdictional frictions. When difficult and even controversial tradeoffs arise, such as how to best allocate limited resources among several affected areas, effective coordination of operational commanders and political leaders who hold the authority of decision is crucial. When leadership or responsibilities need to be transferred across organizations and jurisdictions, frictions can arise which lead to unproductive behaviour including resistance to transferring full or partial responsibility to others better suited to handle the situation. Mechanisms to minimize these frictions and promote effective and spontaneous adaptation are needed.

This work concerns the development of a modelling and simulation capability to enable decision-makers to investigate alternative ways of deploying large-scale emergency

Manuscript received August 1, 2010. This work was supported by the Natural Science and Engineering Research Council of Canada Strategic Partnerships Project Grant with the Defence Research and Development Canada under Technology Investment Fund Contract W 7714-083658. The authors are grateful to Professor Valeriy Vyatkin from the University of Auckland, NZ for his valuable insights into the scenario building and to Dr. Maarten Sierhuis from NASA-AMES Research Centre for his advice with the BRAHMS simulation modelling software.

Nicola Bicocchi is with the University of New Brunswick, Computer Science Department, Fredericton, E3B5A3 Canada (phone: 506-453-3567; e-mail: nbicocch@unb.ca).

William Ross is with the University of New Brunswick, Computer Science Department, Fredericton, E3B5A3 Canada (e-mail: l5i6z@unb.ca).

Mihaela Ulieru is with the University of New Brunswick, Computer Science Department, Fredericton, E3B5A3 Canada (e-mail: ulieru@unb.ca).

operations by (i) pointing to the factors that lead to inconsistency in inter-organizational policies and (ii) allowing decision-makers to test various configurations in an effort to obtain the most effective possible course of action.

The rest of the paper is organized as follows. Section 2 describes the modelling and simulation frameworks we adopted. Section 3 demonstrates the potential of our approach on a case study scenario. Section 4 discusses our experimental results in the light of novel metrics for emergency response. Finally, Section 5 concludes the paper and highlights our directions for future work.

II. Modelling and simulation frameworks

Literature in the area of computer-aided organizational modelling and emergency-response simulation identifies several conceptual dimensions that capture the essential elements of the domain [5, 18, 19]. In this work we focus on the following five dimensions (see Figure 1):

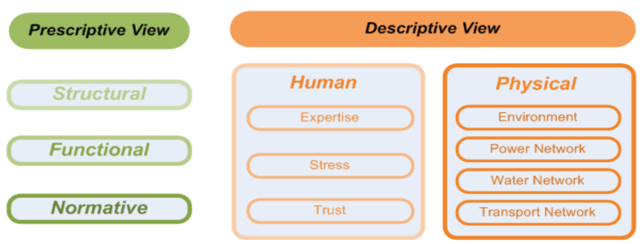

Figure 1: Major conceptual dimensions for modelling emergency-response operations.

The structural dimension refers to the relationships between organizations, departments, teams, leaders, and other individuals that comprise the emergency-response operation. To capture the dynamics of the duality autonomy-cooperation across the inter-organizational continuum we have developed a holistic approach [18] rooted in our previously developed holonic enterprise concept.

The functional dimension refers to the individual work protocol defining the role of each participant in the operation. The structural dimension arises as from the collective behaviour aggregated from the interplay between these individual roles and formal actions derived from job descriptions, where duties and tasks are explicitly outlined, altered by more ethereal, subjective factors and informal actions (e.g., social influence, gossiping) which are very difficult to capture in a simulation model.

The normative dimension refers to the expected behaviour of an organization which is based on formal norms acting as a general guideline for organizational behaviour [6, 20]. Describing what an individual should do while not forcing the individual to perform a particular action, norms may appear formally as organizational policies or informally as organizational culture.

The human dimension refers to less tangible concepts such as expertise, stress, education, risk tolerance, as well as to the norms resulting from organizational culture. This dimension influences the behaviour of individuals and, by extension, organizations. Interestingly, these influences are not apparent from studying the structural, functional, and normative dimensions alone. Real-world organizations show a similar gap between their expected behaviour (i.e., prescriptive) and their actual behaviour (i.e., descriptive). The source of this gap is strongly related to the human dimension [13, 14, 21]. As such, while it is important to represent the prescriptive view of a simulated organization within its members (i.e., software agents), it is also important to represent human-related factors which comprise the descriptive view.

The physical dimension provides a context in which the organizations can operate. It specifies features of the environment, infrastructures, natural resources, and the processes that comprise the incident modifying the environment (e.g., fires and explosions). This dimension is particularly meaningful in emergency-response where the nature of the incident and its resultant effect heavily determine the success of the involved organizations.

These five dimensions are critical for understanding organizational interoperation. The overwhelming number of variables and mutual connections require that they be addressed by using a collaborative approach which allows individual refinement using specialised knowledge while enabling their smooth interweaving into a holistic simulation model. To build on our previous holonic enterprise work, we embrace the multi-agent systems paradigm which provides a clean and disciplined approach to design, develop, and analyze complex software systems [2, 3, 22]. In the rest of this section we illustrate the modelling and simulation frameworks we have selected along with a set of requirements stemming from the above five dimensions that drove our choice.

A. *OperA Modelling Framework*

There have been several modelling frameworks proposed in the literature, yet few of them satisfy the following necessary requirements which we have identified:

1. Prescriptive view and descriptive view separation support. A clear specification of organizational structure, functions, and norms enabling agents to enforce autonomous behaviours is required.
2. Modular and incremental improvements support. The capability of modelling different organizations at various levels of scale and detail is required.
3. Adaptive structural changes support. Mechanisms to support dynamic structural updates as the number of organizations involved increases are required.

OperA [6] satisfies these three requirements in that it is able to model a wide-range of MAS, including open and closed systems, and it has been specifically developed to capture the structural, functional, and normative dimensions. In OperA, organizations are described in terms of roles, role dependencies, organizational interactions, and organizational norms. While providing an organizational template, the OperA framework does not provide any specific agent implementation. Consequently, OperA models cannot be simulated directly without a simulation framework able to describe agents. This decoupling of the abstract description of the organization from the concrete description of the individuals is consistent with the distinction between the top-down prescriptive view and the bottom-up descriptive view. OperA also supports modularity and incremental improvements through a clear separation of the structural, functional, and normative models. Changes and improvements to one model will not drastically impact another. Finally, to further facilitate collaboration, OperA provides a graphical tool, Operetta, allowing modellers to create, share, and edit organizational models. This graphical tool ensures consistency between different modules, provides a formal specification of the organizational model, and facilitates the generation of the simulation.

B. Brahms Simulation Framework

Multi-agent frameworks are promising tools to fulfill our need of transforming the top-down prescriptive view into an actual simulation. For the task of simulating emergency-response operations, we have defined the following requirements:

1. Top-down prescriptive view support. Mechanisms to represent the structural, functional, and normative dimensions expressed in OperA, while still being able to express the actual behaviour of the agents are required.
2. Human dimension support. Several facets of human cognition influence decision-making during emergencies. A robust foundation for implementing these facets is required.
3. Physical dimension support. The physical environment is highly relevant in emergency response and has to be represented.
4. Interaction among agents and physical objects. In emergency-response operations, interactions can be grouped into two categories: person-to-person (e.g., responders and civilians) and person-to-object (e.g., responders and resources). Both agents and objects need to be simulated, along with their interactions.

Brahms [16] has been specifically developed to analyze human organizations and work processes, building on the BDI paradigm with a theory of work practice and situated cognition [4], as opposed to the goal orientation embraced by most BDI frameworks. The main purpose of Brahms is to simulate human collaboration, multi-tasking, informal interactions, "off task" behaviours, and activities (which can themselves be interrupted and resumed). Brahms orchestrates these various behaviours by "activating" them once a predefined situation arises. In particular, Brahms implements a subsumption architecture enabling the simultaneous execution of general activities (e.g., extinguish fire) and specific activities (e.g., talking on the radio) inside a composite activity, allowing agents to be more reactive to changes in the environment. In addition to an excellent support for modularity, allowing functionalities defined within groups to be inherited, Brahms also provides a seamless integration of human agents and arbitrary objects (e.g., electronic communication devices, documents, and tools) within a geographic environment. It also separates agent beliefs from world facts by separating agents from the environment they are immersed in, unlike traditional BDI frameworks. This might simplify, in the long term, the integration of Brahms agents with other physical simulations. Finally, being activity-oriented, rather than goal-oriented enables Brahms simulations in conjunction with OperA modelling to encapsulate a clearer representation of both the top-down prescriptive view and bottom-up descriptive view. The top-down view is provided by the structural, functional, and normative models specified in OperA and the bottom-up view is captured in the Brahms agent code.

III. Modelling and Simulation Capability

Our modelling and simulation capability, rooted in the multi-agent system paradigm, consists of three steps:

1. Model organizations, along the structural, functional, and normative dimensions, using the OperA modelling language [6, 7].
2. Implement the OperA models using Brahms, a multi-agent simulation framework detailed in [16] so that the models can be executed and human dimensions added.
3. Incrementally improve our modelling and simulation capability for addressing the management and engineering of complex situations by (i) increasing the complexity of the incident, (ii) increasing the number of organizations involved, (iii) enriching the structural, functional, and normative models of these organizations, and (iv) improving the behavioural models of the members of these organizations.

A. Scenario

We have begun by applying our capability to an incident centred on a tanker in the harbour of a densely populated city. The tanker, filled with chemicals, enters the harbour to repair a pump needed to unload the chemicals. During the repair, something goes wrong, and this causes a fire to break out onboard the docked ship. An emergency call is dispatched to the port authority who begins coordinating the response. Four main outcomes are possible:

1. The fire can be stopped early, while it is still small, resulting in minimal damage to the ship.
2. The fire can escalate to engulf a large portion of the ship, but if more fire teams are sent at the appropriate time, the fire can be extinguished though the damage to the ship will be extensive.

3. The fire may cause the temperature onboard the ship to rise above a certain threshold. When this happens, the chemicals onboard will explode, fire will spread to the dock, and the chemicals will start spilling into the harbour. If there are sufficient fire teams at the scene, the incident can be contained so that the dock and harbour receive minimal damage.
4. If there are insufficient fire teams at the scene, the chemicals in the water will spread and catch on fire and the poisonous fumes from the fire will spread to the city, potentially resulting in widespread devastation to both civilians and marine life.

For the first iteration, we analyzed only the key organizations involved in the response: the port authority, municipal firefighters, the coast guard, and the transport authority. We also considered only the fire and explosion outcomes. In future iterations, we plan to include more organizations and enrich them with more details to address a broader range of possible outcomes.

B. Modelling Capability

There are three types of OperA models that have been used to describe the four organizations listed above:

1. The role dependency graph, to denote the relations between roles (see Figure 2);
2. The interaction structure diagram, to represent the order of important interactions within and across organizations at a high level of abstraction (see Figure 3);
3. The landmark patterns, to detail on how interactions should be achieved at the organizational level (see Figure 4).

Roles

Roles in organizations are a representation of a function or a service and provide an initial definition for both the structural and functional dimensions. Roles abstract from specific actors and instead describe groups of actors having similar functionality, rights, and capabilities. Each role inside an organization has its own set of individual objectives. For example, below we list a possible descriptive model for the Firefighter Team role. This model also specifies the objectives, rights, and norms for a Firefighter Team.

Role: Firefighter Team
Objectives: extinguishFire(F), rescuePeople(P)
Sub-objectives:
Rights: workin(jurisdiction(me))
Type: institutional
Norms:
IF DONE assigned(F) AND jurisdiction(F) isNotPartOf(jurisdiction(me)) THEN askCoordinator(F, jurisdiction(me))
IF DONE assigned(F) THEN OBLIGED move(place(F)) BEFORE TH min
IF unknownSource(F) THEN informCoordinator(F, unknownSource)
IF peoplePresent(F) THEN rescuePeople(P) BEFORE extinguishFire(F)
IF dangerHigh(F) THEN PERMITTED stopWorking(F)

The role dependency graph, showing the seven roles we identified for our scenario, is depicted in Figure 2. The direction of the dependency arrow specifies the direction of the dependency relation (i.e., an arrow pointing from role A to role B indicates that role A depends on role B for the specified objective). Each organization is managed by a single agent, called Organization Leader. The firefighters and the coast guard have additional agents within their organizations in order to provide further resolution into these organizations. Each additional agent in the firefighter organization represents a team of firefighters, while in the coast guard, each additional agent represents a fire tug team. This ability to represent different organizations at varying levels of resolution is a feature of the holonic paradigm [17].

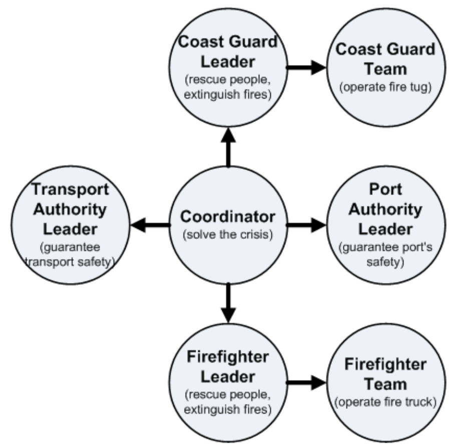

Figure 2: Role dependency graph. The Coordinator depends on all other roles to solve the incident, while the Coast Guard and Firefighter Leaders depend on their respective Teams to accomplish the objectives of their organization.

Interaction Structure

The interaction structure depicted in Figure 3 defines the most relevant scenes (represented by boxes) and transitions (represented by lines) involved in emergency-response operations and is used to further specify the functional dimension. The organization acting as the coordinator (the port authority in our scenario) begins by collecting information about the problem, followed by problem assessment. It must then determine whether it is capable of handling the problem. If it is and it has the necessary resources, the organization will proceed to attempt solving the problem before potentially concluding. Otherwise, if the available resources are not sufficient to effectively handle the emergency, other organizations become involved, additional resources are requested, and the command structure is updated. The interaction structure also defines the order of the interactions between the roles in the organization. Interactions are grouped into scenes, where a scene reflects a meaningful subset of interactions related to the achievement of a set of particular objectives.

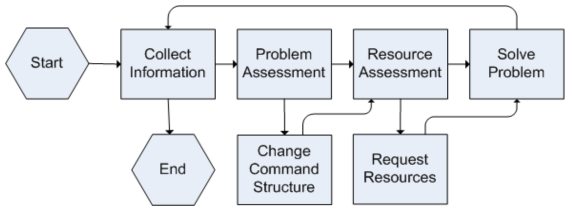

Figure 3: Interaction structure representing important high-level interactions between organizations for a general emergency-response scenario.

Landmark Patterns

Each scene in the interaction structure can be depicted in further detail by specifying landmark patterns. A landmark pattern defines an ordered pattern of important states that must be accomplished in the achievement of a scene's objectives. A specific state is considered active once all of its input states are completed (i.e., do not require further action). This allows modellers to specify the temporal relation between states, allowing different organizational behaviours to be specified from the same set of states simply by changing the temporal relations. For example, Figure 4 illustrates two alternative behaviours associated with the Firefighter Team role. Figure 4(a) directs the team to extinguish fires and rescue people simultaneously, while Figure 4(b) directs the team to first rescue people and then start extinguishing fires. In addition to specifying the states, each landmark pattern is also associated with a set of norms that specifies how the involved organization behaves prescriptively. It should be noted that the high-level interaction structure (see Figure 3) is common to all emergency-response organizations, but it is the landmark patterns that describe organization-specific objectives and interaction policies.

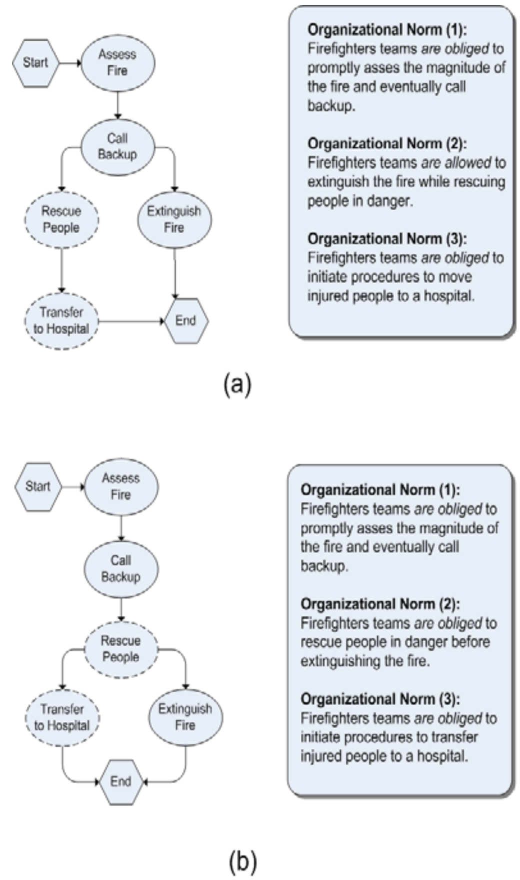

Figure 4: Landmark patterns and associated norms describing the “solve problem” scene (depicted in Figure 3) for the firefighter team role (depicted in Figure 2). In this figure, solid shapes represent portions that have currently been modelled and simulated, while dashed shapes represent those that will be in the future.

C. Simulation Capability

To implement our OperA models in Brahms, we started by applying the techniques described in [19]. Specifically, we used the following approach:

1. For the structural dimension, we implemented OperA roles as Brahms groups (see Figure 5(a)); and
2. For the functional and normative dimension (see Figure 5(b)), we implemented landmark patterns and their specific norms as Brahms workframes (see Figure 5(c)).

Workframes are a construct of the Brahms language representing situated activities. When a specific situation arises (specified by a set of preconditions in the when clause), a workframe will trigger the execution of an activity. Each workframe, representing an activity within a landmark pattern, describes the behaviour that an agent, belonging to a group, is expected to perform under a specific situation.

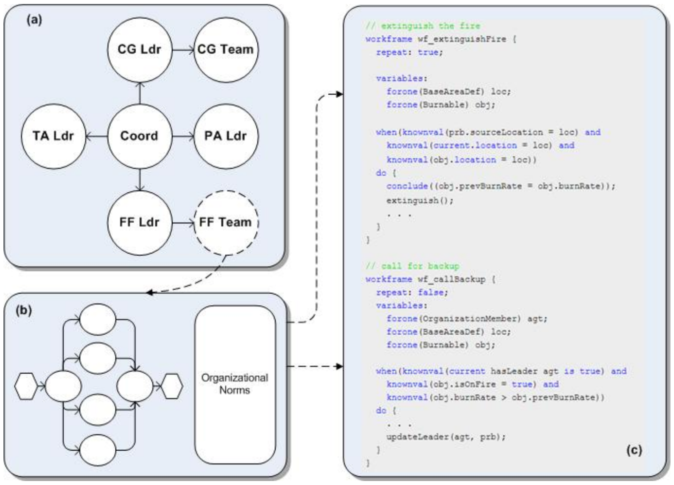

Figure 5: The implementation process from OperA to Brahms is shown for the Firefighter Team role (dashed circle). The structural (a), functional (b), and normative (b) dimensions are integrated into the corresponding Brahms code (c).

In addition to simulating the organizations, we have also simulated the physical dimension. We have defined several areas in our simulation, including those for the various organizational headquarters, for the dock, for seven districts in the city, and for four districts in the harbour. We have also specified the paths between each area. As for the incident itself, we have simulated a fire and explosion model in our physical environment. The fire is modelled as an object that can interact with the environment, affecting the temperature (T), burn rate (br), and amount of burnable material remaining (d) based on the equations in Table 1. The burn rate step (β) and temperature ratio (τ) determine, respectively, how the burn rate and temperature increase over time. Furthermore, the maximum burn rate is a function of the amount of material remaining to be burned divided by a constant (μ).Agents in the environment can interact with the fire object to reduce its intensity. However, a threshold value, which depends on the current strength of the fire, must be exceeded in order for the fire teams to begin containing the fire. If the temperature inside the ship exceeds a certain threshold, an explosion will occur.

Table 1: The equations governing a burnable object.

Equation
$d(t) = d(t-1) - br(t)$
$T(t) = br(t) * \tau \quad where\ \tau = 2$
$br(t) = \begin{cases} br(t-1) * \beta \ ; & br(t-1) < \dfrac{d(t)}{\mu} \\ \dfrac{d(t)}{\mu} \ ; & br(t-1) \geq \dfrac{d(t)}{\mu} \end{cases}$ $where\ \beta = 2,\ \mu = 10$

IV. Experimental Setup and Results

The current state of our proof-of-concept simulation is already sufficient to enable us to explore the normative (bottom-up) and structural (top-down) dimensions affecting emergency response. In general, we are able to define parameters for each dimension, set their values, run the simulation multiple times, and from the aggregated results, highlight the impact of the parameters on the system. For the normative dimension, we can specify organizational policies and agent behaviour using deontic logic, where every agent action is categorized as being obliged, prohibited, or permitted. For the structural dimension, we can modify the organizational structure by adding or removing groups and agents. To demonstrate our approach, we have defined the following three parameters:

1. Coordinator [is obliged | is not obliged] to promptly collect data about the contents of the ship once notified about the fire (normative dimension)
2. Fire Tug Team [is prohibited | is not prohibited] from entering an exclusion zone (normative dimension)
3. Fire Tug Team is owned by [Coast Guard | Firefighters] (structural dimension)

On the basis of these parameters, we have created four rule sets (see Table 2) which allow us to specify which normative and structural parameters we are investigating in the simulation, as well as the parameter values we are testing. For example, rule set 1 (RS1) specifies a simulation in which: (i) the coordinator is not obliged (O) to promptly collect data from the transport authority; (ii) the fire tug team is prohibited (P) from entering the exclusion zone; and (iii) the fire tug team is owned by the coast guard (CG). The number of variables that can be explored in our simulation is already significant. When running a simulation, all the variables not listed in a rule set are initialized to their default values.

Table 2: The four rule sets (RS) tested in our simulation.

Variable	RS1	RS2	RS3	RS4
Obliged to promptly collect data	¬O	O	¬O	¬O
Prohibited from exclusion zone	P	P	¬P	P
Fire tug team owned by	CG	CG	CG	FF

Besides normative and structural organizational dimensions, several other factors impact emergency-response operations [15]. Therefore, to improve realism and enable different results for each simulation run, we have introduced the following background noise:

1. Communication delays. Some messages take longer to convey than others; people are not always immediately reachable.
2. Team effectiveness. Teams have different levels of fatigue and experience.

3. Travel times. The time of day, traffic levels, and condition of the roads/water impact how long it takes responders to arrive at the scene.
4. Explosion threshold. The explosion should not be a strict function of the response; sometimes it will take longer (or shorter) for the explosion to occur even when the response is identical.
5. Weather conditions. Temperature, wind direction, and wind strength affect the fire.

We executed ten simulation runs for each rule set and recorded the results we observed in the simulation. These results are presented in Section 4.2. The approach of manually collecting the data was fine for our proof-of-concept. However, in the future, we plan to have the simulation automatically generate the statistics and result reports for us so that we can perform experiments using a greater number of simulation runs.

A. *Effectiveness Measures*

In order to evaluate the effectiveness of each rule set in our simulated environment, we use the six features listed in Table 3. While this list is not exhaustive, it has proven capable of comparatively assessing the results of our simulation as will be discussed in the next section.

Table 3: Effectiveness features.

Feature Description	Unit of Measure
Time needed to solve the incident	Hours
Amount of damage	Millions of dollars
Number of resources deployed	Teams
Number of resources deployed that were ineffectively used (e.g., a fire tug that is deployed but cannot enter the exclusion zone)	Teams
Overall number of organizations involved in the response	Organizations
Average logical communication distance between agents	Communication hops

B. *Results*

The bar chart portion of Figure 6 presents the results of each rule set using the features in Table 3. For each rule set, we identify the worst and best results from the simulation runs. We also display for time and amount a benchmark, which represents the result were there no response taken. As can be seen, rule sets 1 and 2 (Figures 6(a) and 6(b), respectively) have the highest overall values for each feature, suggesting the presence of conflicts in these rule sets. Specifically, these poor results are due to a previously unknown policy conflict between the fire tug team and the port authority. The port authority, who establishes an exclusion zone around the ship when the fire reaches a certain intensity level, requests that the fire tug team assist at the scene. However, if the request occurs after the exclusion zone has been established, the fire tug team, prohibited from entering an exclusion zone, is unable to approach the scene and help combat the fire. The result of this conflict is the explosion on the ship. Rule set 3, on the other hand, has lower values for most of the features, including zero for the number of resources used ineffectively (see Figure 6(c)). This points to the fact that this rule set is more effective at handling the incident (i.e., not being prohibited from entering an exclusion zone has an effect). Still, because of the noise factors presented earlier like communication and travel delays, the fire tug team is not always able to respond quickly enough to prevent the explosion. The final rule set explores the dimension involving organizational structure (see Figure 6(d)). In this rule set, a fire tug team is added to the firefighter organization, which reduces both the length of the fire and the resulting damage. By removing the coast guard from the response, valuable time is saved from not having to coordinate with another organization, and the explosion is always averted. As can be seen from the results, the firefighter organization having direct control over the fire tug team also serves to reduce the average logical communication distance, as there is one fewer organization at the scene.

In addition to the bar graphs, each rule set is also associated with a fire graph (see graph in the upper right-hand corner of Figures 6(a) - (d)). This graph shows the temperature inside the ship over time. The baseline case, when there is no response, is outlined with the solid black line, and the sudden increase in temperature is due to the explosion. As with the bar graphs, the worst case and best case results from the simulation runs are shown. For each rule set, the graph shows the effectiveness of the response in handling the fire. The difference between the best and worst cases in each graph is particularly interesting. It shows the range of the effectiveness of the rule set's responses in controlling the fire under various background noise conditions. This range can be used to determine the reliability of the response: a larger difference corresponds to a less reliable rule set, while a smaller difference corresponds to a more reliable rule set. As can be seen in Figure 6(a), rule set 1 is almost as bad as no response. Rule sets 2 and 3 (Figures 6(b) and 6(c), respectively) both have a large difference between their best and worst cases, suggesting that the effectiveness of the response is not reliable given the environmental conditions. In the worst case, both rule sets are only marginally better than no response. However, both rule sets have been able to prevent the explosion in their best case. Finally, as shown in Figure 6(d), rule set 4 is a dramatic improvement in relation to the other rule sets and the no-response case. Because we have only considered the fire and explosion outcomes, these fire-graph results agree with those of the features shown in the bar graphs.

While the organizational parameters we introduced in this section were simplistic, our simulation was able to produce a somewhat surprising result: slightly changing the structure of two organizations was significantly more important in containing the incident than changing key normative parameters of the involved organizations. While these are only preliminary results, we believe they demonstrate the potential of our approach in investigating the factors impacting organizational interoperation.

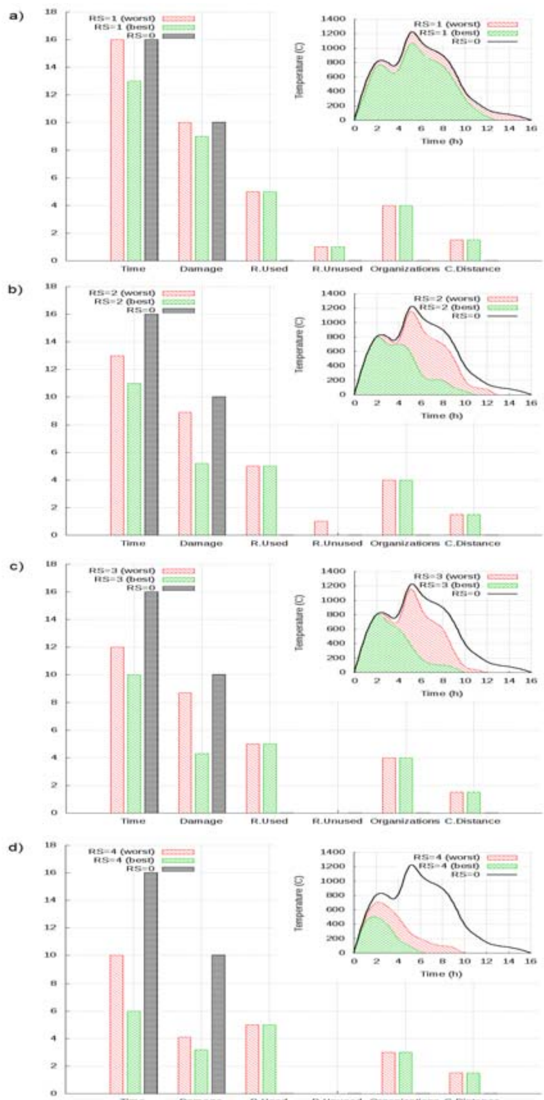

Figure 6: Aggregate effect of the four rule sets on the six features and the fire. The rule sets are as follows: (a) RS1, (b) RS2, (c) RS3, and (d) RS4.

V. Conclusion and Future Work

In this paper we outlined our collaborative capability to model and simulate emergency-response organizations using the OperA and Brahms frameworks. Furthermore, we demonstrated the benefits of our approach by exploring the impact of specific structural and normative parameters on emergency-response organizations using a proof-of-concept simulation. We also proposed a set of original metrics capable of capturing the effectiveness and reliability of the simulated response under various configurations.

An important aspect of our future efforts will consist in modelling the impact of subjective "human factors" on decision-making during emergencies. In fact, when the environment is largely unknown and it is difficult to apply predefined policies, human factors such as expertise, emotions, and professional culture often deeply influence organizational performance. Current agent-based languages founded on the BDI paradigm are a significant improvement over previous attempts in that they provide a robust representation of human agents acting organizational roles. However, although incipient work has been done in this direction [14], there is still no established modelling methodology and each endeavour requires a customized, and not necessarily reusable, solution. Our goal is to provide agents with easily reusable models of expertise, representing how the critical components of the environment function and interact with each other.

Furthermore, based on the results obtained, we aim at developing an intelligent communication backbone which enables the end-to-end management of processes running flexibly across many different organizations in many different forms. The central idea of our approach is that linking partners is on the basis of linking processes while allowing individual execution according to those processes. Particularly, we envision models for personal assistant components (e.g., PDAs) knowing about the rules of each individual organizational agent and able to suggest which ones, compiled together, may result into (sub-)optimal meta-organizational policies. Combining agents according to their protocols might simplify the "spontaneous" deployment of appropriate "ecologies" for each particular emergency situation. Our capability allows us to show, understand, and, finally, measure the impact of such new technologies on the creation and evolution of meta-organizations.

References

[1] Bobashv, G., Morris, R., Zule, W., Borshchev, A., and Hoffer L., "The Use of Agent–based Modeling in Projecting Risk Factors into the Future", Social Computing and Behavioral Modeling, Springer, 2009.
[2] Brazier, F., Jonker, C., Treur, J., "Principles of Component-Based Design of Intelligent Agents". Data and Knowledge Engineering. 41(2), 1—28, 2002.
[3] Ciancarini, P., Wooldridge M., "Agent-Oriented Software Engineering". Agent-Oriented Software Engineering, LNAI 1957, Springer, 2001.
[4] Clancey, W., "Situated cognition: on human knowledge and computer representations", Cambridge University Press, 1997.
[5] Dess, G., Beard, D., "Dimensions of Organizational Task Environments. Administrative Science Quarterly". 29(1), 52–73, 1984.
[6] Dignum, V., "A model for organizational interaction: based on agents, founded in logic". Http://igitur–archive.library.uu.nl/dissertations/2003–1218–115420/inhoud.htm, 2000.
[7] Dignum, V., V´azquez–Salceda, J., Dignum, F., "Omni: Introducing social structure, norms and ontologies into agent organizations". International Workshop on Programming Multiagent Systems Languages and tools, New York (NY), USA, 2004.
[8] Forrester J., "Counterintuitive Behavior of Social Systems", Technology Review, 1971.
[9] Howitt, A., Leonard, H., "Beyond katrina: Improving disaster response capabilities. Taubman Center Policy Briefs", 2006.
[10] Leonard, H., Herman, B., Howitt, A., "Katrina as Prelude: Preparing for and Responding to Katrina-Class Disturbances in the United States". Journal of Homeland Security and Emergency Management 3(2), 2006.
[11] Myer, R., Conte, C., Peterson S., "Human Impact Issues for Crisis Management in Organizations", 16(5) 761—770, 2007.
[12] U.S. National Management Incident System (NIMS), 2008 http://www.fema.gov/pdf/emergency/nims/NIMS_core.pdf.
[13] Pereira, D., Oliveira, E., Moreira, N., "Modelling emotional BDI agents", Workshop on Formal Approaches to Multi-Agent Systems, Riva Del Garda, Italy, 2006.

[14] Picard, R., “Affective Computing”, Mit Press, 2000.
[15] Shendarkar, A., Vasudevan, K., Lee, S., Son, Y.–J., “Crowd simulation for emergency response using BDI agents based on immersive virtual reality”. Simulation Modelling Practice and Theory 16 (9), 1415–1429, 2008.
[16] Sierhuis, M., Clancey, W., Hoof, R. V., “Brahms: A multi–agent modelling environment for simulating work processes and practices”. International Journal of Simulation and Process Modelling 3(3), 134–152, 2007.
[17] Ulieru, M., Este, R., “The Holonic Enterprise and Theory of Emergence”, International Journal Cybernetics and Human Knowing, 11(1), 79–99, 2004.
[18] Ulieru, M., “Holistic security ecosystems”. Digitial Ecosystems Technologies Conference. Istanbul, Turkey, 2009.
[19] Van Putten, B., Dignum, V., Sierhuis, M., Wolfe, S., “OperA and Brahms: A Symphony?”, Agent-Oriented Software Engineering IX, Springer, 2009.
[20] Vazquez–Salceda, J., Dignum, F., “Modelling electronic organizations”. Multi–Agent Systems and Applications, Springer–Verlag, 2003.
[21] Vicente, K., “The human factor”, Routledge, 2004.
[22] Zambonelli F., Omicini A., “Challenges and Research Directions in Agent-Oriented Software Engineering”, Journal of Autonomous Agents and Multiagent Systems, 9(3), 253–283, 2004.

Nicola Bicocchi received the M.Sc. degree in Computer Engineering from Universita’ degli Studi di Modena e Reggio Emilia, Modena, Italia, in 2004 and the Ph.D. in Computer Science from the same University in 2009. He is currently a Postdoctoral Fellow with Adaptive Risk Management Laboratory, University of New Brunswick, Fredericton, Canada. His research interests include complex systems, pervasive computing and data mining. In these areas, he has published 20 papers in international journals and conferences.

William Ross received his BCS (2005) and MCS (2007) from the University of New Brunswick. Currently he is a Ph.D. student at the same University under the supervision of Dr. Mihaela Ulieru. His research interests focus on investigating inter-organizational dynamics to improve organizational interoperation and situational understanding.

Mihaela Ulieru (M’95–SM’02) received the M.Sc. degree in Electrical Engineering and Computer Science from Politehnica University, Bucharest, Romania, in 1985 and the Dr.-Ing. from the Control Engineering and Computer Science Department, Darmstadt University of Technology, Germany, in 1995. She was a Postdoctoral Fellow with the Intelligent Robotics and Manufacturing Systems Laboratory, School of Engineering Science, Simon Fraser University, Vancouver, BC, Canada, from 1996 to 1998 and Junior NORTEL Chair in Intelligent Manufacturing at the University of Calgary (1998-2000). Since 2005 she is Canada Research Chair, Professor and the Director of Adaptive Risk Management Laboratory at the University of New Brunswick, Canada. Her research targets the design and implementation of adaptive information infrastructures fuelling tomorrow’s e-Society.