A Simulation Modelling Approach Enabling Joint Emergency Response Operations

Nicola Bicocchi, William Ross, Mihaela Ulieru
Faculty of Computer Science
University of New Brunswick
Fredericton, Canada
{nbicocch, ulieru, l5i6z}@unb.ca

***Abstract*—A novel capability for modelling and simulating intra- and inter-organizational collaboration in an emergency-response domain is presented. This capability combines the prescriptive, top-down view of organizations, which describes how they work "on paper," and the descriptive, bottom-up view, which describes how they actually work, by focusing on three components in the light of agent-based modelling and simulation tools—structural, functional, and normative. Our approach enables decision-makers to anticipate the evolution of an emerging crisis and evaluate the effectiveness of different configurations on the response. The initial results of our simulation, based on an experiment which investigates the impact of three separate parameters, are also presented and reveal the joint effectiveness of the organizations involved.**

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

I. INTRODUCTION

In an effort to investigate why emergency-response organizations do not properly orchestrate their collaboration when crises escalate, it has been synthesized from [7,8,13] that the primary causes relate to the structural, functional, and normative make-up of the involved organizations. This is especially true when organizational policies conflict or when jurisdictional issues arise. The *structural* component refers to the layout of the organization, including its command structure. The *normative* component refers to the policies and practices of the organization. Finally, the *functional* component refers to what the organization can actually do—if considered as its own entity, this would be a list of its functional capabilities (e.g., a firefighter station can put out fires).

The key purpose of our research is to determine how organizations involved in joint-response operations can better collaborate during an evolving crisis. To this end, we propose an agent-based modelling and simulation capability that will enable us to model and simulate organizations from both a top-down, prescriptive view—which describes how organizations should behave "on paper"—and the bottom-up, descriptive view—which describes how organizations behave in reality. The aim is that this capability will provide decision-makers with a better understanding of the complexities of inter-organizational collaboration and provide them with a tool to test various configurations of the important components that lie at the heart of the issue.

II. MODELLING AND SIMULATION CAPABILITY

In this section we outline our modelling and simulation capability. First, we examine the tools we selected for each task (i.e., modelling and simulation). Then, we present the scenario we developed and we discuss the key organizational models involved. To conclude, we describe how we translate our models into an actual simulation, and discuss some important aspects of our simulation.

A. Tools

There are several tools available to facilitate both the modelling and simulation of organizations that have been extensively described in [1]. For our purposes, we focused on the ones aiming to establish clean and disciplined approaches to design, develop, and analyze complex software systems based on the multi-agent systems (MAS) paradigm [15]. Further, we require that the modelling and simulation tools we select exhibit the following properties:

1) *Support a clear separation between the top-down, prescriptive view (the ideal activity of the organization) and the bottom-up, descriptive view (the activities actually performed):* In particular, we seek for a clear specification of organizational structure, functions, and norms with autonomous agents able to ignore or circumvent organizational norms;

2) *Support modularity and incremental improvements:* Enables modelling and simulation at various levels of scale and detail; and, specifically for the simulation tool,

3) *Support interaction among agents and physical objects in a geography:* 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, within a geography.

For our modelling tool, we have selected OperA [6], as it is able to fully satisfy all our modelling requirements. Specifically, it is able to model a wide-range of multi-agent systems, including open and closed systems, and has been developed to capture the structural, functional, and normative components of agent organizations. 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. Still, OperA models cannot be simulated directly

without a simulation tool 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 view and the bottom-up view.

As our simulation tool, we have selected Brahms [10, 11], which has been specifically developed to analyze human organizations and work processes, building on the Belief-Desire-Intention (BDI) paradigm with a theory of work practice and situated cognition [4]. In this paradigm, conceived by Bratman as a theory of human practical reasoning [3], *beliefs* represent the informational state of the agent—in other words its beliefs about the world including the self and other agents. *Desires* represent the motivational state of the agent—in other words the objectives the agent would like to accomplish. Finally, *intentions* represent the deliberative state of the agent—in other words what the agent has chosen to do. The success of the BDI paradigm is due to its ability to elegantly reduce the explanation for complex human behaviours to the motivational stance [5], which postulates that the causes for actions are always related to human desires.

Brahms has as main purpose to simulate human collaboration, multi-tasking, informal interactions, "off task" behaviours, and activities (which can themselves be interrupted and resumed). Brahms uses the Brooks subsumption architecture, in particular, to enable the simultaneous execution of general activities (e.g., extinguish fire) and specific ones (e.g., talking on the radio) inside a composite activity, allowing agents to be more reactive to changes in the environment. This lends it very naturally as a most suitable tool for the task of simulating emergency response operations. Moreover, Brahms provides excellent support for modularity, allowing functionality defined in groups to be inherited by all agents in that group and by subgroups. It also provides a geography in which agents can interact with other agents, as well as with objects. Finally, being activity-oriented, rather than goal-oriented, means that actions are not explicitly subordinate to goals in Brahms. This allows its simulations in conjunction with OperA modelling to encapsulate a clearer representation of both the top-down view and bottom-up view: the top-down view provided by the structural, functional, and normative models specified in OperA; and the bottom-up view captured in the Brahms agent code.

To illustrate our modelling and simulation capability, we have selected the following simplified public safety and security scenario. From it, we are able to model the organizations involved using OperA and then simulate their independent and collaborative behaviour using Brahms.

B. Scenario

Consider a chemical tanker in the harbour of a densely populated city. The tanker enters the harbour to repair the pump needed to unload its cargo. During the repair, something causes a fire to break out onboard the docked ship. An emergency call is dispatched to the port authority which begins coordinating the response. We consider three outcomes:

1) The fire can be stopped early, before it has spread, resulting in minimal damage to the ship;

2) The fire can escalate to engulf a large portion of the ship, but if the response is appropriate it 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 and fire will spread to the dock.

For the emergency response, we consider simplistically four key organizations: port authority, municipal firefighters, coast guard, and transport authority. Moreover, we also consider the emergency coordinator. In general, when a leader of an organization receives notification about a problem (e.g., a fire), it first checks whether the problem is being handled by another organization. If it is not and the leader's organization has jurisdiction in that area and is able to handle the problem, that leader will assume the role of coordinator.

In our scenario, the *port authority* leader is the first to receive an alert from the ship that a fire may have started onboard, and because it has jurisdiction and knows how to handle the problem, it assumes the role of coordinator. In this capacity, the port authority leader must determine what action to perform. For example, it may decide to contact only the firefighter organization, or to contact the transport authority to determine the contents of the ship and request resources as the situation demands. Ultimately, it is the responsibility of the coordinator to request assistance and collect information from other organizations. These requests may be generated by the coordinator or by other organizations involved in the response, so the coordinator must act as intermediary between the various organizations. Lastly, in addition to its coordinator functionalities, the port authority leader is also empowered to establish an "exclusion zone" (an area in which ships are temporarily prohibited from entering) within the harbour.

The leader of the *municipal firefighters* organization is responsible for communicating with the coordinator, managing the firefighter teams, and requesting backup through the coordinator if insufficient resources are available at its headquarters. This organization's teams are capable of travelling to and returning from an incident, extinguishing fires, and requesting backup from the leader.

Similarly, the *coast guard* leader manages a fire tug team and dispatches it upon request from the coordinator. The fire tug is a special boat that is capable of extinguishing fires at a greater rate than several teams of firefighters.

Lastly, the *transport authority* manages a manifest list of the contents of every ship in the harbour. Upon request, its leader will send any available data about the contents of a ship to the coordinator.

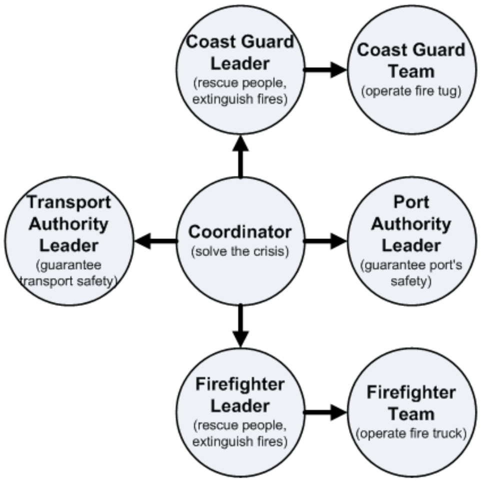

Figure 1. Role dependency graph. The Coordinator depends on all other roles to solve the incident, while the Coast Guard and Firefighter Leaders depend on their respective Teams to accomplish the objectives of their organization.

C. *Modelling Capability with OperA*

In OperA, organizations are described in terms of roles, interactions, and norms, which correspond to the structural, functional, and normative components, respectively. Roles are captured using OperA's role dependency graph (see Figure 1), while interactions are captured using its interaction structures (see Figure 2) and landmark patterns (see Figure 3). Norms, on the other hand, are associated independently with each representation—particularly with the landmark pattern—and are specified deontically as being obligatory, prohibited, or permitted.

Role Dependencies

The role dependency graph describes roles and the relationships between them. Roles provide an initial definition for both the structural and functional dimensions. They group together actors with similar capabilities, rights, and goals. Furthermore, roles may be dependent on one another for the completion of particular sub-goals. For example, the coast guard leader role, which has the goal of extinguishing fires, depends on the coast guard team role to operate the fire tug. In alignment with holonic structural modelling [12], the various organizations described in terms of roles may be enacted by a single agent or by multiple agents, depending on the scale necessary to address the particular problem. Figure 1 shows the roles we have defined for our scenario and their mutual functional dependencies.

Interaction Structure

Interaction structures provide the ability to further specify the functional dimension. Activities in an organization can be considered as the composition of multiple, distinct, and concurrent interactions, involving different actors playing different roles, and these interactions are represented in OperA through scenes. Scenes can be ordered and synchronized to represent complex activities through the use of an interaction structure.

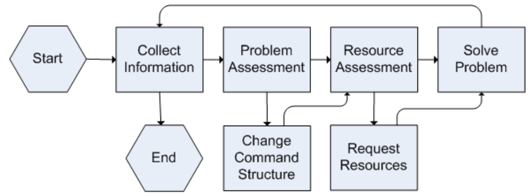

Figure 2. Interaction structure representing important high-level interactions between organizations for a general emergency-response scenario.

The interaction structure depicted in Figure 2 defines the most relevant scenes (represented by boxes) and transitions (represented by lines) involved in emergency-response operations. The organization acting as coordinator begins by collecting information about the problem, followed by problem assessment to 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. This organizational flow, able to represent a wide range of different response operations, has been extracted from an analysis of the U.S. National Incident Management System documentation [36].

Landmark Patterns

Each scene in the interaction structure can be further detailed using landmark patterns. These specify *(i)* the states an organization must reach before the scene is completed and *(ii)* the possible transitions among the states. Each state can be

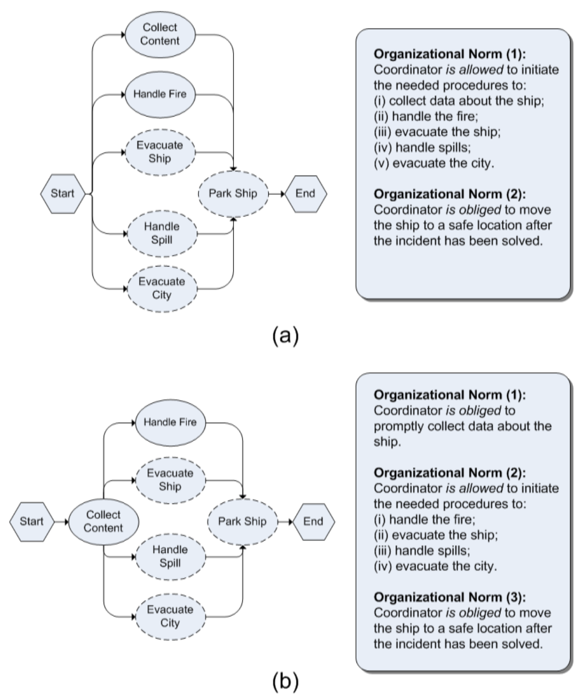

Figure 3. Landmark patterns and associated norms describing the "solve problem" scene (depicted in Figure 2) for the coordinator role (depicted in Figure 1). In this figure, solid shapes represent the portions of focus for this simple illustration, while dashed shapes complete the scenes yet are not described in this paper.

associated with one or more input and output states. A specific state is considered active once all its input states are completed

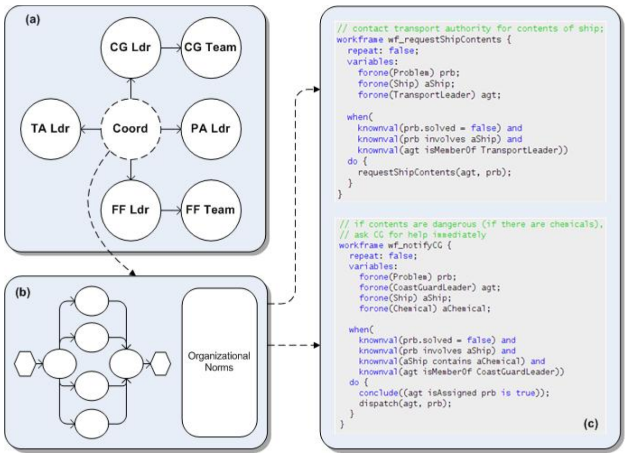

Figure 4. The implementation process from OperA to Brahms is shown. The structural (a), functional (b), and normative (b) dimensions are integrated into the corresponding Brahms code (c).

(i.e., do not require further action). This approach allows modellers to specify the temporal relation between states. Given the same set of states, different temporal relations impose different organizational behaviours. In fact, once the key roles and functionality have been defined, it is possible to investigate the impact of different organizational behaviours by proposing alternative landmark patterns and associating them to different norms.

Figure 3 illustrates this concept. The landmark pattern in Figure 3a, for example, directs the coordinator to initiate all response activities in parallel. Only once they are completed the ship can be moved to a safe location. Alternatively, the landmark pattern in Figure 3b directs the coordinator to immediately collect information about the ship before initiating the other response activities. Besides the landmark patterns, the right part of both figures shows the organizational norms defined for solving the particular problem scene.

D. Simulation Capability with Brahms

To implement our OperA models using the Brahms multi-agent software development platform, we started by applying the techniques described in [14] as follows: for the structural dimension, we implemented OperA roles as Brahms groups (see Figure 4a); and for the functional and normative dimensions (see Figure 4b), we implemented landmark patterns and their specific norms as Brahms workframes (see Figure 4c).

Workframes are a construct of the Brahms agent 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 illustrates a portion of a running simulation represented along a timeline which is created by the Brahms integrated development toolkit from a database that stores the results of each simulation run. It shows the interactions between the leaders of the port authority, transport authority, and municipal firefighters once the port authority leader receives notification about the fire. The port authority leader, having jurisdiction over docked ships, becomes the coordinator and notifies the firefighter leader about the incident (see Figure 3, which shows that the coordinator must "Handle Fire" in order to "Solve Problem" from Figure 2, and Figure 1, which shows that the coordinator depends on the firefighter leader to handle the fire). From the timeline in Figure 5, we see that the firefighter leader immediately dispatches one of its teams to the scene, and the coordinator then requests the ship's contents from the transport authority leader (see Figure 1 once again for these dependence relations). The above interaction (i.e., workframe "wf_requestShipContents") is a direct consequence of the norm illustrated in Figure 3b. The transport authority leader responds by returning the contents of the ship. The port authority leader, recognizing the potential of an explosion, requests further assistance from the coast guard leader (i.e., workframe "wf_notifyCG").

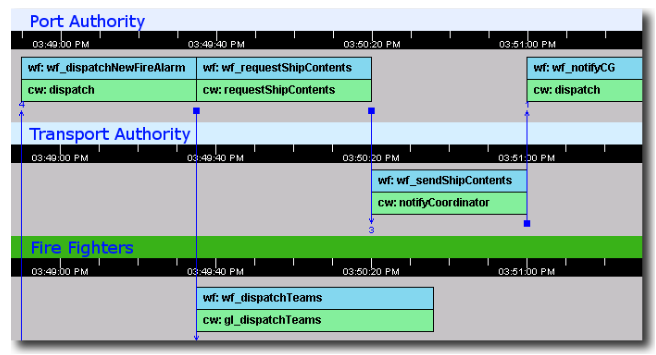

Figure 5. Brahms timeline diagram showing the Port Authority Leader (top) contacting the Transport Authority Leader (middle) and the Firefighter Leader (bottom) after receiving the distress call from the ship (arrow on the left). This shows the actual execution by our simulation framework of the workframes depicted in Figure 4c.

To simulate the incident, we implemented models for both fire and explosion events. Every burnable object within the simulation is associated with an amount of burnable material and an initial burn rate, which represents the amount of material consumed after some specified duration (currently one hour). When a fire starts, the burnable material is decreased and the burn rate is increased over time. In our simulation, organizational agents act to reduce the temperature of the fire. An explosion occurs if the fire is not contained before the internal temperature of the object exceeds the explosion threshold. The explosion results in an increase in the amount of burnable material, which in turn increases the maximum burn rate resulting in a temperature spike.

III. Experimental Results

The current state of our proof-of-concept simulation is already sufficient to enable us to explore the structural, functional, and normative components affecting emergency response. In general, we are able to define parameters for each component, set their values, run the simulation multiple times, and from the aggregated results, highlight the impact of the parameters on the system. To demonstrate our approach, we have defined the following three parameters:

1) Coordinator *[is obliged (O) | is not obliged (¬O)]* to promptly collect data about the content of the ship once notified about the fire (normative aspect);

2) Fire Tug Team *[is prohibited (P) | is not prohibited (¬P)]* from entering an exclusion zone (normative aspect); and

3) Fire Tug Team is owned by *[Coast Guard (CG) | Firefighters (FF)]* (structural aspect).

On the basis of these parameters, we have created four rule sets (see Table II) 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 I. THE FOUR RULE SETS (RS) TESTED IN OUR SIMULATION

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

B. Success Indicators

We comparatively assess the effectiveness of each rule set against the following indicators:

Time. Represents the overall amount of time, expressed in hours, needed to completely resolve the incident. Effective rule sets are expected to minimize this value.

Damage. Represents the overall economic damage produced by the incident. Every simulated object is associated with an initial economic value expressed in millions of dollars. Destructive forces decrease the value of the affected objects. Effective rule sets are expected to minimize this value.

Organizations. Represents the overall number of organizations involved in the response. Increasing the number of organizations is likely to increase the amount of inter-organizational inefficiencies and conflicts. As a consequence, effective rule sets are expected to involve the minimum possible number of organizations to solve the incident.

Resources Used. Represents the overall number of teams deployed by all the organizations. This indicator summarizes the capability of the organizations to deploy the needed resources. Effective rule sets are expected to mobilize their resources while avoiding waste.

Resources Unused. Represents the overall number of deployed teams which are not used during the response. For example, this may be because a team that is unable to handle a particular problem was deployed or because a team was prevented from participating due to conflicting organizational policies. Effective rule sets are expected to minimize this value.

Average Communication Distance. Represents the average communication distance during the response. We analyze the simulated organizations as a graph composed of nodes and edges. Every node represents one organizational agent (e.g., a team or a leader). Every team is connected to its leader. Furthermore, each team is connected to the other teams managed by the same leader. Organizational leaders join the network as soon as their organization is involved in the response. The communication distance between two nodes (i.e., agents) is defined as the number of edges along the shortest path connecting them. According to network theory, this parameter represents the radius of the communication network. Higher values mean that intra- and inter-organizational communications require, on average, a larger number of agents in order to take place. Considering that during emergencies communication processes are one of the elements most prone to errors, an effective rule set is expected to minimize this value.

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.

C. 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 6a and 6b, respectively) have the highest overall values for each feature, suggesting the presence of conflicts in these rule sets. Specifically, these 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 6c). 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 6d). 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.

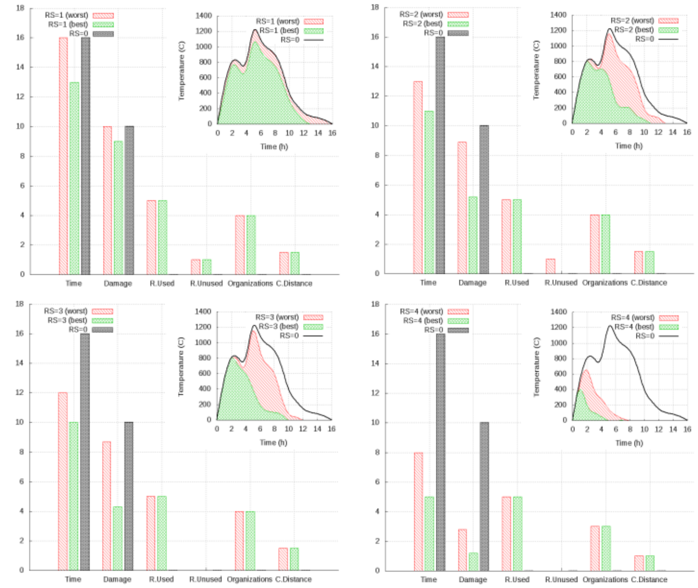

Figure 6. Comparison of the performance of four rule sets based on the features listed in Table 3. The line graph in the upper right corners represents the temperature of the fire in time.

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 6a - 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 6a, rule set 1 is almost as bad as no response. Rule sets 2 and 3 (Figures 6b and 6c, 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 6d, rule set 4 is a dramatic improvement in relation to the other rule sets and the no-response case. Not surprisingly, 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 aspects of the involved organizations.

IV. Conclusion and Future Work

Accurate modelling of human organizations is crucial in order to understand their short-term and long-term behaviours and to provide policy makers with suitable models for testing and evaluating new policies [2]. As presented in this paper, one of the most successful techniques to simulate organizations is agent-based simulation. Our capability, which couples OperA and Brahms, allows for the development of realistic organizational models immersed in a simulated environment. These models, and their performance, can be also tested and evaluated under various configurations to improve inter- and intra- organization synergies.

Our future work concerns the use of our modelling and simulation capability to analyze the important trade-offs that must be decided upon when choosing to transition from single organization operation to collaborative endeavours, as well as how to capture the coordination logic over a joint alliance using an agent-based approach to implement an overarching operational layer that enables optimal synergy from the interactions of hybrid individual participants.

V. References

[1] Bicocchi, N., Ross, W., Ulieru, M., "Addressing the Challenges of Complex Public Safety and Security Operations through Simulation Modeling". Transactions of the IEEE Systems, Man and Cybernetics, Part A (Systems and Humans), submitted May 7, 2010.

[2] 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 US, 2009.

[3] Bratman, M., "Intention, Plans, and Practical Reason". Harvard University Press, 1999.

[4] Clancey, W., "Situated cognition: on human knowledge and computer representations", Cambridge University Press, 1997.

[5] D. Dennett. The Intentional Stance. Bradford Books, 1987.

[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] Howitt, A., Leonard, H., "Beyond katrina: Improving disaster response capabilities. Taubman Center Policy Briefs", 2006.

[8] 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.

[9] U.S. National Management Incident System (NIMS), 2008 http://www.fema.gov/pdf/emergency/nims/NIMS_core.pdf.

[10] 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.

[11] Sierhuis, M., Jonker, C., Riemsdijk, B. v., & Hindriks, K., "Towards Organization Aware Agent–based Simulation". International Journal of Intelligent Control and Systems, 14(1), 62—76, 2009.

[12] Ulieru, M., Este, R., "The Holonic Enterprise and Theory of Emergence", International Journal Cybernetics and Human Knowing, 11(1), 79–99, 2004.

[13] Ulieru, M., "Holistic security ecosystems". Digitial Ecosystems Technologies Conference. Istanbul, Turkey, 2009.

[14] Van Putten, B., Dignum, V., Sierhuis, M., Wolfe, S., "OperA and Brahms: A Symphony?", Agent-Oriented Software Engineering IX, Springer, 2009.

[15] 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.