Modelling a Chain of Command in the Incident Command System using Sequential CFGs

Tabajara Krausburg[1 [0000−0002−8252−4099]], Jürgen Dix[1 [0000−0002−8528−1440]], and Rafael H. Bordini[2 [0000−0001−8688−9901]]

[1] Department of Informatics, TUC, Clausthal-Zellerfeld, Germany
{tkr19,dix}@tu-clausthal.de
[2] School of Technology, PUCRS, Porto Alegre, Brazil
rafael.bordini@pucrs.br

Abstract. Disaster response is a major challenge given the social and economic impact on the communities affected by disaster incidents. We investigate how *coalition formation* can be used for the problem of forming a hierarchy of resources (e.g., personnel responding to the incident). As a case study, we consider the roaring river flood scenario and model the Incident Command System (ICS) framework—providing guidelines on cooperatively responding to disaster incidents. Our main tool is sequential characteristic-function games induced by size-based valuation structures. We show that our approach can deliver a hierarchy required by the Operations Section in the ICS and provides a promising way to analyse the computed solution. Future work will focus on collecting feedback from experts on the topic regarding the proposed modelling.

Keywords: Sequential CFGs · Coalition structure generation · Incident command system · Real-world applications.

1 Introduction and Related work

Coalition formation has long attracted attention from the scientific community due to its theoretical challenges and practical applications. The aim is to group agents together to accomplish a particular goal, e.g. ride-sharing where agents represent the users of the system [3]. Another interesting application is disaster response given its social and economical impact on the affected communities. Usually, a team of responders is required to perform several tasks, for instance, save victims [17, 2]. So the set of available resources (e.g., responders or equipment) has to be distributed into teams; each possessing expertise and abilities to comply with the designated objectives.

This can be modelled with Characteristic-Function Games (CFG) [18] in which one uses a valuating function to evaluate how well the agents work together. The outcome is then a Coalition Structure (CS) which is a partition of the set of agents into coalitions. In this paper, we put forward the work on modelling disaster response applications under the perspective of the Coalition

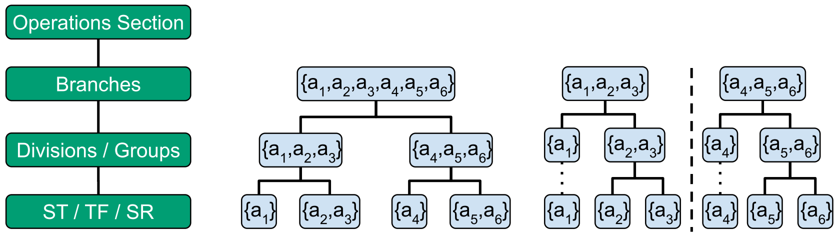

Fig. 1: The Operation Sections hierarchical organisation on the left-hand side. On the right-hand side, two sequences of length three of CSs.

Structure Generation (CSG) problem [13]. In particular applications that require the use of the Incident Command System (ICS) [7, 5]: a framework largely put into practice to guide the response to disaster incidents. In general, a disaster response application can be modelled by a single CFG. However, our main goal is to model a chain of command imposed by the ICS making the usage of a single CFG no longer possible. We suggest to use an extension of Sequential Characteristic-Function Games (SCFG) [10] to model additional constraints on the coalitions allowed to form; this sort of game is named SEQSVS [8].

After introducing the ICS and the SEQSVS frameworks in the sections below, we discuss in Section 1.3 related work for modelling similar disaster response problems. In Section 2, we introduce the problem we aim to solve. Section 3 introduces our experiments with the proposed model. Finally, in Section 4, we discuss the main findings and achievements in our approach.

1.1 The Incident Command System Framework

The Incident Command System (ICS) [7] is designed to deal with multiple organisations that respond to a disaster incident. One of the key guidelines the ICS establishes is the span of control which determines the ratio of command between superiors and subordinate units: $\frac{1}{\lambda}$. That means, one superior is responsible for at most λ subordinate units. It constraints how one builds the hierarchical structure of resources put into action in a disaster response operation.

A resource in this context can be either a personnel, a team, a facility, an equipment, or a supply. This characterisation is called *kind*. Apart from its kind, the resources will differ from each other based on [5]: (i) its core *capabilities* for which it is the most useful (e.g., critical transportation); (ii) its *category* which is the function for which a resource would be most useful (e.g., health and medical); (iii) a *type* describing a resource's level of minimum capability to perform its function (e.g., Type 2). This whole process is known as the identification of a resource. One may use the Resource Typing Library Tool (RTLT[3]) to retrieve standard specifications commonly used in practice. To illustrate the core idea, consider an ambulance ground team [6, page 193] in the example below.

[3] https://rtlt.preptoolkit.fema.gov

Example 1 (Ambulance Ground Team). As *capability*, an ambulance ground team provides basic life support and transports 2 non-ambulatory patients. It is tagged with *category* emergency medical services. It has resource *kind* team. Finally, the resource is tagged as Type 3, which means a crew of 2 personnel.

Out of the five main sections defined in the ICS, we focus on the Operations Section responsible for achieving command objectives, tactical operations, etc. We depict in the left-hand side of Figure 1 the operations section hierarchical organisation. Branches, at the upper level, are divided into divisions/groups. Divisions are allocated to geographical areas of the incident (e.g., because of different jurisdictions). Groups represent a functional aspect of the operation (e.g., rescue group). The number of branches, divisions, and groups depends on the span-of-control ratio. At the bottom level, resources are organised into Task Force (TF), Strike Team (ST), or even Single Resource (SR). Any combination of resources of different kinds and/or types constitute a TF, whilst of same kind and type form an ST. The number of organisational levels follows a *bottom-up approach*. Once a superior is overwhelmed a new upper level is created to make it fit into the span of control.

This mechanism establishes a *chain of command*, in which each resource *must* have a supervisor. The authority system within a section is as follows: a chief, a director, a supervisor, and a leader are responsible for a section, a branch, a division/group, and a TF/ST respectively. We shall use the term *superior* to capture the notion of leadership in a unity. The disaster response operation is guided by Incident Objectives (IO), and the superiors should identify the requirements for accomplishing them and then gather the relevant resources.

1.2 Sequential CFGs Induced by Size-Based Valuation Structures

A Sequential CFGs Induced by Size-Based Valuation Structures (SEQSVS) [8] is a framework that extends the valuation structure components in a SEQVS [9]. As introduced in [9], a valuation structure is a tuple $\langle G, S\rangle$, where G is an interaction graph (A, E) on the set of agents A and S is a set of pivotal agents $S \subseteq A$. A Size-Based Valuation Structure (SVS) π adds to a valuation structure a set of permitted sizes that a feasible coalition may have.

Definition 1 (Size-based Valuation Structure π). *A size-based Valuation Structure is a tuple π $\langle G, S, Z\rangle$ where G is an interaction graph, S is a set of pivotal agents, and $Z \subseteq \mathbb{N}$. A coalition C belongs to the set of feasible coalitions $C \in \boldsymbol{\mathcal{C}}^\pi$ (i.e., is allowed to form) if:*

(i) the induced sub-graph of C over G is connected;
(ii) $|C \cap S| \leq 1$; and
(iii) $|C| \in Z$.

An SVS π is then induced over a Characteristic Function Game (CFG) $\Gamma = \langle A, v\rangle$, where $v : 2^A \rightarrow \mathbb{R}$, to constrain the coalitions allowed to form in that game (denoted Γ^π). A Coalition Structure (CS) *CS* is formally defined as a set of coalitions such that for all $C, C' \in CS$, $C \neq C'$, $C \cap C' = \emptyset$, and $\bigcup_{C \in CS} = A$. Then, given the set of all CSs $\boldsymbol{\mathcal{CS}}^A$, one can construct the set of allowed CSs $\boldsymbol{\mathcal{CS}}^\pi = \{CS \in \boldsymbol{\mathcal{CS}}^A \mid CS \subseteq \boldsymbol{\mathcal{C}}^\pi\}$. Then, we can formally define an SEQSVS.

Definition 2 (Sequential CFGs induced by SVSs). *A sequential* CFG *induced by a sequence of* SVSs *is a tuple* $\mathcal{G} = \langle A, \mathcal{H}, \Pi, \mathcal{R}\rangle$*, where:*

- A *is a set of agents;*
- $\mathcal{H}$ *is a totally ordered set* $\Gamma_1 = \langle A, v_1\rangle, \ldots, \Gamma_h = \langle A, v_h\rangle$ *of* CFGs*;*
- Π *is a totally ordered set consisting of* $\pi_i = \langle G_i, S_i, Z_i\rangle$ *with the* SVS *parameters described in Definition 1, one for each game* $i = 1, \ldots, h$*;*
- $\mathcal{R}$ *is a binary relation on* $\boldsymbol{\mathcal{CS}}^A \cup \{\emptyset\}$.

This tuple determines the sequence $\boldsymbol{\Gamma} = \langle \Gamma_1^{\pi_1}, \ldots, \Gamma_h^{\pi_h}\rangle$ *of* CFGs *induced by* SVSs*, with* $\Gamma_i = \langle A, v_i\rangle$ *and* $\pi_i = \langle G_i, S_i, Z_i\rangle$*, where* v_i *are characteristic functions and* $\pi_i \in \Pi$*, for* $i = 1, \ldots, h$.

We now introduce the solution concept for such games.

Definition 3 (SEQSVS Optimisation Problem). *Given a sequence* $\boldsymbol{CS} = \langle CS_1, \ldots, CS_h\rangle$ *of coalition structures from* $\boldsymbol{\mathcal{CS}}^A$*,* $\boldsymbol{CS}$ *is a potential solution for* $\langle \Gamma_1^{\pi_1}, \ldots, \Gamma_h^{\pi_h}\rangle$ *if:*

(i) each CS_i *is a solution of* $\Gamma_i^{\pi_i}$*; and*
(ii) it respects the relation $\mathcal{R}$*:* $CS_i \; \mathcal{R} \; CS_{i+1}$*,* $0 \leq i < h$ *(we set* $CS_0 = \emptyset$*).*

*Such a sequence is called a Feasible Coalition-Structure Sequence (*FCSS*). A function* $\mathcal{V}$ *determines a value* $\mathcal{V}(\boldsymbol{CS})$ *for any* FCSS $\boldsymbol{CS}$*, which is given by* $\sum_{i=1}^h V_i(CS_i)$*, where* $V_i(CS_i) = \sum_{C \in CS_i} v_i(C)$.

A solution for a SEQVS *game instance is an optimal* FCSS

$$\boldsymbol{CS}^* = \arg\max_{\boldsymbol{CS}} \mathcal{V}(\boldsymbol{CS}).$$

To solve SEQSVS problems, an algorithm based on Monte Carlo tree search, called UCT-Seq, was proposed [8]. It iteratively constructs a tree in which each node corresponds to a CS and the root node corresponds to $\emptyset$. Whenever UCT-Seq reaches level h in the tree, it will have found a corresponding FCSS. Three parameters guide the tree search, they are:

Degree $\bar{b}$ is used during a simulation (i.e., a roll out) to determine the number of attempts to find a compatible coalition structure at the current simulated level.

Depth $\bar{d}$ is the maximum relative depth a simulation can reach before updating the visited nodes' reward throughout the path back to the root.

Exploration factor γ determines when it is desirable to expand the frontier of a given node.

We refer to [8] for additional information on how this algorithm works and how to set it up to solve the ICS problem.

1.3 Related Work

Usually, disaster response operations inspire problems in which one assigns agents to accomplish a set of tasks announced by the command centre. In [1], the authors model a disaster response incident in which tasks have different priorities, deadlines and are composed of sub-tasks. To carry out a task, an agent needs a number of resources, and therefore the agents responding to the event form coalitions to meet a task requirement. The authors in [11] focus specifically on robots. Each robot has two vectors of capabilities, one for sensing (e.g., cameras) and other for acting (e.g., arms). Similarly, to complete a task, the required sensing and acting capabilities must be met. To deal with tasks announced during a flood-inspired disaster event (e.g., to collect samples of water), the authors in [2] investigate composed tasks; each demanding particular roles to be carried out. Roles are mapped onto capabilities which are possessed by the robots in the system. Tasks affected by spatial and temporal constraints were also studied [14]. This sort of game is motivated by the RoboCup Rescue competition [17] which is inspired by an earthquake incident in Japan.

The coalition structure generation problem has already been used to model a response to a disaster incident. The authors in [18] model a scenario in which a satellite, powered by radioactive fuel, has crashed in a sub-urban area (this scenario was first introduced in [15] as an agent-based planning problem). The emergency services that respond to this incident are composed of medics, soldiers, transporters and fire-fighters (the roles). The scenario is modelled as a grid where responders must accomplish a set of rescue tasks by dropping off the *targets* at specific locations in the grid. A target can either be a victim, animal, fuel, or other resource; therefore, four targets are considered. Each task demands a set of roles in order to be carried out (a mapping targets onto roles). Each agent has a different level of capabilities while playing each of the four roles. This information may represent its training for that role, past experience, etc. Doing so, one can estimate the performance p_k^C of a coalition C to rescue a specific target k by summing up the capabilities of each individual responder in C for each role required by target k. To determine a partition of the resources, the function $v(C) = \sum_{k=1}^{4} p_k^C$ is used. This way, coalitions that can deal with as many targets as possible are formed.

Although the work above have considered disaster response applications, none of them addressed the ICS problem specifically. To properly form a hierarchy of resources, the ICS requires a chain of command among all participants in the response operation which is then subjected to a span of control. Without carefully observing this feature, superiors in the chain of command can get easily overwhelmed. This leads us to tackle the overall problem under a hierarchical perspective in which the coalitions are influenced and constrained by upper hierarchical levels. Therefore, we model the problem using many games that are interdependent instead of designing a single game that models the entire disaster response operation in contrast to previous approaches proposed in the literature.

Table 1: The incident objectives for the roaring river flood scenario and their corresponding groups and required roles.

General Incident Task	Group	Role Ids
euthanize suffering animals	Euthanasia	3
begin the disposal operation	Disposal	2, 4, 6
identify relocation sites and relocate animals	Relocation	4, 5, 7
control the movement of host material	Regulatory	1
eradicate the fruit flies	Control	2
survey for fruit fly locations	Survey/ID	8
identify fruit flies	Survey/ID	9

2 A Disaster Response Application

Our main goal is to model the distribution of resources throughout a hierarchy. We do not consider only the last hierarchical level for twofold reasons: (i) a coalition in the last hierarchical level (i.e., a TF, ST, or SR) is influenced not only by its corresponding leader, but also by the superiors of upper levels; (ii) the need to enforce the span of control between subsequent levels. The CSG problem for this application (see right-hand side of Figure 1) states that a coalition in the upper organisational level may be split into at most λ coalitions in the subsequent level (i.e., the subordinate units). This structure can be modelled by the SCFG framework [10]. We now model and solve a realistic problem.

2.1 The Roaring River Flood

We aim to model a fraction of the hierarchy provided in the the Roaring River Flood (RRF) scenario [16]. It is a part of an ICS training course of the United States department of agriculture. In this scenario, due to heavy rains, a severe flooding is now inundating the roaring river valley. Of special concern are: (i) the contamination of food processing plants; (ii) the heavy damage to a federal fruit fly research facility which determined the release of thousands of fruit flies; and (ii) the widespread livestock losses. We depict in Figure 2 the entire Operations Section for this particular scenario. We shall focus on modelling only the Vet Services and PPQ branches (the red-dashed square). These two branches offer a good balance between complexity and easy explainability.

The incident objectives follow from the ICS training course [16] and are introduced in Table 1.

Although the RRF scenario provides a good explanation of the disaster response operation, it does not provide all the necessary information to model and solve the problem of distributing resources to act upon the damaged zone. For instance, it does not provide which resources have been identified and are available for that particular operation (see Section 1.1). Therefore, we looked up roles in the RTLT that are expected to provide capabilities required to solve the incident objectives above. We introduce the resulting set of roles in Table 2. In

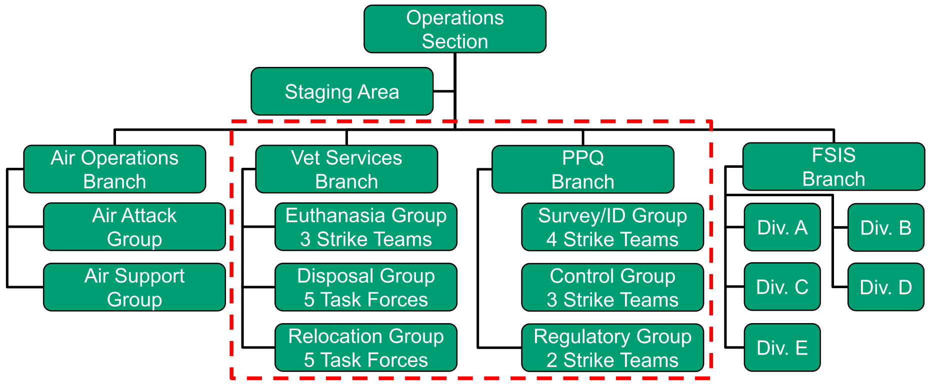

Fig. 2: The Operations Section for the Roaring River Valley scenario [16].

Table 2: The roles for the roaring river flood scenario collected from the resource typing library tool.

Role	Id	Acronym
Animal Control/Humane Officer	(1)	ACHO
Animal Depopulation Specialist	(2)	ADS
Animal Depopulation Team	(3)	ADT
Animal Evacuation, Transport, and Re-Entry Team	(4)	AETRT
Animal Search and Rescue Team	(5)	ASRT
Animal Search and Rescue Technician	(6)	ASRTec
Veterinary Medical Team	(7)	VMT
Animal and Agriculture Damage Assessment Team	(8)	AADAT
Animal Behavior Specialist	(9)	ABS

this work we consider only personnel-related *kinds* (i.e., *personnel* and *team*), even though the ICS framework specifies also *equipment*, *facilities* and *supplies* as kinds [5].

2.2 A Single CFG Problem for Disaster Response

In our particular application, we consider a CFG $\Gamma = \langle A, v \rangle$, where $A = \{a_1, \ldots, a_n\}$ is the set of resources assigned to the Operations Section. Moreover, let R be a set of roles and $adopt(a_i) \subseteq R$ the roles that agent a_i may adopt. The Operations Section must achieve a set of incident objectives $IO = \{o_1, \ldots, o_7\}$. Each of them demands a set of roles, $demand : IO \to 2^R$ (see Table 1).

A CFG Γ is induced by a sized-based valuation structure $\pi = \langle G, S, Z \rangle$ (see Definition 1). Moreover, for each special agent in a game $a_i \in S$ (representing a superior), we assume he/she is responsible for accomplishing a set of incident objectives, $resp : S \to 2^{IO}$. Now, for a superior $a_i \in S$, we can construct the corresponding set of roles a_i expects to have at its disposal to achieve its assigned IOs, $expect(a_i) = \{r \in R \mid o \in resp(a_i), r \in demand(o)\}$.

In this modelling we assume that a disaster response management system is available in which the information above can be queried. This system will map resources, who possess capabilities, onto a set of roles R. This assumption is not disregarded from real-world operations as agencies that undertake them use systems to help managing resources. For instance, consider the interagency resource ordering capability[4] system used by agencies in the USA to manage resources as well as the incident itself. Unfortunately, anonymised information about the resources was not disclosed.

Assigning Roles to the Agents In our modelling, each pivotal agent adopts a random *personnel* role. We select $|R|$ ordinary agents (i.e., resources who are not superiors) and assign to each of them a single different role in R. The remaining agents adopt a role based on a probability associated with each role in R. This is done based on a bound on the expected quantity of roles. Assume a leader $s \in S_1$ in game Γ_1. Given a span of control λ, let e^r be the number of expected agents adopting role r computed as follows:

$$e^r = \sum_{s \in S_1 : r \in expect(s)} \left\lceil \frac{\lambda}{|expect(s)|} \right\rceil .$$

The main idea is that we look up in the leaders (of both TF and ST) to figure out which roles they need. Then, we can estimate the number of expected agents that should assume each role. Consider the example below.

Example 2. Assume $S_1 = \{s_1, s_2\}$ and the span of control $\lambda = 2$. Leader s_1 expects roles $\{r_1, r_2\}$, and leader s_2 expects roles $\{r_3\}$. Then, in the overall operation we expect $e^{r_1} = 1$, $e^{r_2} = 1$ and $e^{r_3} = 2$ (leader s_2 expects a single role and the span of control is set to 2).

The probability of agent $a \in A$ adopting role r is then computed using Equation 1. Note that we distinguish between roles of kind personnel and team. A resource has 50% chance of being one of the two kinds.

$$P(a, r) = \frac{e^r}{\sum_{r' \in R} e^{r'}} \tag{1}$$

Example 3. Consider Example 2. The probabilities of agent a adopting a given role (assuming all roles to be of the same kind) are as follows $P(a, r_1) = 0.25$, $P(a, r_2) = 0.25$ and $P(a, r_3) = 0.5$.

Based on the procedure above, we assign roles to all resources in any given instance.

[4] https://famit.nwcg.gov/applications/IROC

Valuation Functions The valuation functions have two important components. The first one is the relationship among resources. We assume a function that evaluates the relationship between any two agents, $relationship : A \times A \rightarrow [0, 1]$. Similarly, the relationship within any coalition, $relationship : 2^A \rightarrow [0, 1]$. The second aspect addresses the roles available in a coalition. We define a function $disturbance : 2^A \rightarrow [0, 1]$ to calculate how unbalanced are the roles within any coalition. We combine these two terms in Equation 2.

$$v_i(C) = |C| \times (relationship_i(C) - disturbance_i(C)) \quad (2)$$

We show the exact instantiation of both components in Section 3.

The rationale behind the characteristic function in Equation 2 is: the more resources in a coalition the better, given the condition that the coalition has a balanced number of roles. Usually in the literature the size of a coalition is given as a sub-additive function, i.e., the more agents in it the greater the loss in value [3]. However, in our case, the span of control set for the disaster response operation already takes this matter into account. Therefore, we are only considering coalitions that are manageable by its corresponding superiors.

2.3 A Sequence of CFG Problems for Disaster Response

To connect each level of the hierarchy we use the SEQSVS model $\mathcal{G} = \langle A, \mathcal{H}, \Pi, \mathcal{R} \rangle$. The Operations Section contains three hierarchical levels: branch, group/division, and TF/ST levels. We follow bottom-up approach to model the problem. That means, our first CFG corresponds to the bottom level of the hierarchy (i.e., TF/ST level). Moreover, we use an additional CFG to keep track of the span of control of the operations section chief. Therefore, $\boldsymbol{\Gamma} = \langle \Gamma_1^{\pi_1}, \Gamma_2^{\pi_2}, \Gamma_3^{\pi_3}, \Gamma_4^{\pi_4} \rangle$. The first three games in the sequence $\Gamma_i \in \mathcal{H}$ use the valuation defined in Equation 2. In the last game, $v_4 : C \mapsto 0$ to not modify the overall value.

Regarding the size-based valuation structures (Definition 1). The set of pivotal agents in each game are the corresponding superiors at each hierarchical level. Let L be the set of all superiors; that is, $L = \bigcup_i^4 S_i$. Also, we set $Z_2 = Z_3 = \{1, \ldots, n\}$. In the first game, a coalition can have at most $\lambda + 1$ members (a superior and λ resources). Therefore, $Z_1 = \{1, \ldots, \lambda + 1\}$. In the last game, only the grand coalition is allowed, $Z_4 = \{n\}$. We introduce the interaction graphs in Section 3 as they are specific to each problem at hand.

We set the span of control to $\lambda = 5$ (1 superior responsible for 5 subordinates units) following the maximum number of task forces in the TF/ST level (see Figure 2). To connect the outcomes of the CFGS, we design $\mathcal{R}$ as follows.

Definition 4 ($\mathcal{R}^{\textbf{basis}}$). *Given $CS, CS' \in \boldsymbol{\mathcal{CS}}^A$, the pair $(CS, CS') \in \mathcal{R}^{\text{BASIS}}$ iff:*

1. *$|CS| > |CS'|$; and*
2. *for all $C' \in CS'$ there exist at most $\lambda + 1$ coalitions $C \in CS$ such that $\bigcup_j C = C' : 1 \leq j \leq \lambda + 1$;*

Moreover, $(\emptyset, CS) \in \mathcal{R}^{\text{BASIS}}$ for all $CS \in \boldsymbol{\mathcal{CS}}^A$.

Note that, without the fourth game (i.e., the section level) a coalition structure of the branch level CS_3 could have any number of coalitions. In fact, it could easily overwhelm the operations section chief. By adding the fourth game, in which the CS CS_4 contains only the grand coalition, we enforce that $|CS_3| \leq \lambda + 1$.

2.4 Roaring River Flood Instances

We assume instances of the RRF problem containing a set of 101 and 141 resources. A full hierarchy is formed with 141 resources: 31 superiors (including the operations section chief) and 22×5 ordinary agents at the TF/ST level. We also assume that *personnel* resources may adopt at most 3 roles. This can be the case, for instance, when a resource is of the highest type (i.e., Type 1) and therefore is skilled enough to assume any of the remaining types. We assume that resources of kind *team* must stick with a single-role-adoption rule. We understand that the corresponding team roles (in Table 2) are too dissimilar in the RRF problem. However, it is completely feasible to consider a team that can play different types. We followed the procedure described above to assign a role to each agent. The resulting distributions of roles in both instances are depicted in Figure 3.

The green bars stand for the number of roles required by the leaders at the TF/ST level. That is, summing up the roles expected by each leader $s \in S_1$. The first two bars (left to right) show the number of superiors (in all levels) in both RRF instances. The blue bars depict the number of agents that may adopt each role in an instance containing 101 resources and the the purple bars in an instance containing 141 resources. Finally, the yellow bars stand for a bound on the quantity of roles required to solve the problem. This is calculated based on the requirement of each leader (either of a task force or strike team) and the span of control. For instance, given the span of control $\lambda = 5$ and 2 leaders who require 5 resources of the same role r_1 each, then the bound for role r_1 will be equal to 10.

In the role distribution in Figure 3 we note that all team-related roles match at least the quantity of leaders requiring them. Regarding personnel-related roles, one can see that the demand is exceeded for all of them. Recall that a resource can work only for a single coalition at the TF/ST level. Once it has decided to adopt a role, the remaining two roles are no longer available to be assigned to other coalitions. Moreover, we assume that for each branch, group, ST, and TF a superior has already been assigned by the operations section chief. We arrange the IOs in such a way that they form subsets as we move down in the hierarchical structure. For instance, the director of PPQ branch is responsible as well for achieving the goals of the Survey/ID, Control, and Regulatory groups.

3 Evaluation

In this section we experiment with the instances introduced above. We run UCT-Seq [8] using the following parameters (see Section 1.2): (i) $\bar{b} = 5$; (ii) $\bar{d} = 3$; (ii)

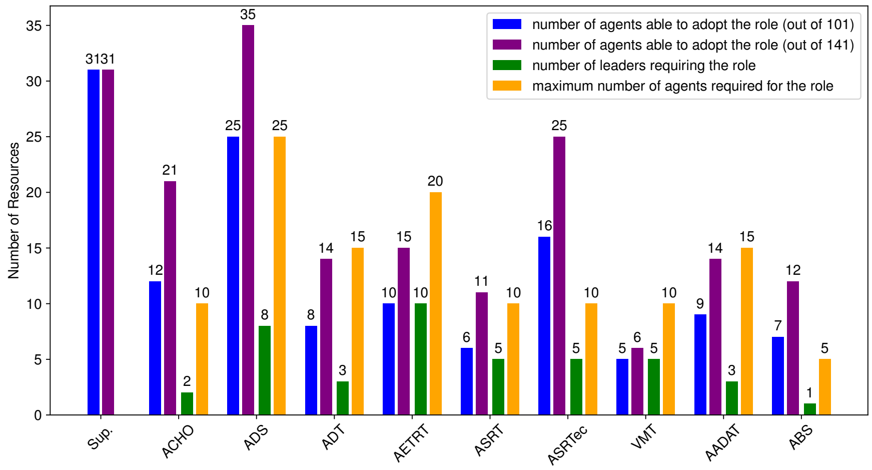

Fig. 3: Distributions of roles in the RRF instances. Role descriptions in Table 2.

$\gamma = 0.7$. These are determined empirically. All experiments were conducted in a machine with 32 GB of RAM and a CPU with four single cores of 3400 MHz each. The algorithm is implemented in Python 3.8.

3.1 A Hierarchy of Resources for the RRF Problem

We aim to form a hierarchy of resources in which each ordinary resource is either accompanied by a superior of the corresponding level or isolated in a singleton coalition. We refer to this problem as a fixed-hierarchy.

Preliminaries In this problem, ordinary agents are supposed to be subordinated to a superior. Therefore, we evaluate the relationship between subordinate agents and the superiors of the corresponding level. To do so, we use Equation 3, where C is a coalition, and S_i is the set of superiors of the corresponding level i.

$$relationship_i(C) = \frac{1}{|C \setminus S_i|} \times \sum_{a \in C \setminus S_i} relationship(a, C \cap S_i) \tag{3}$$

In case C is a singleton coalition, then the relationship is set to 0.

Regarding the roles, we do not assume that each superior is requiring a precise number of agents adopting a given role at the last level of the hierarchy (nonetheless, for the upper hierarchical levels this can be computed). To figure out which roles are required in a coalition C, we look at the pivotal agents of the TF/ST level that are in C. Let $require(C)$ be a *multiset* that counts the role demands (which role and which quantity) required by a coalition C. That is, $require(C) = \{r \in expect(a) \mid a \in C \cap S_1\}$. Note that the quantity per role in $expect(a)$ is always a single unity. Let $m_{require(C)}(.)$ count the multiplicity of an element in a multiset $require(C)$.

Example 4. Assume a leader $l_1 \in S_1$ and a coalition $C : C \cap S_1 = \emptyset$ formed for the third game Γ_3 (i.e., the branch level). Moreover, let $expect(l_1) = \{r_1, r_2\}$. Thus, $require(C \cup \{l_1\}) = \{r_1 : 1, r_2 : 1\}$. Now, assume another leader $l_2 \in S_1$ such that $expect(l_2) = \{r_2\}$. Thus, $require(C \cup \{l_1\} \cup \{l_2\}) = \{r_1 : 1, r_2 : 2\}$.

Similarly, let $available(C)$ be a multiset of roles in C adopted by its ordinary agents. That is, $available(C) = \{r \in adopt(a) \mid a \in C \setminus L\}$.

To calculate the role disturbance we use the relative entropy (Kullback-Leible distance) [4]. Let $P(r)$ be the proportion of agents in C that adopt role r. This is computed using Equation 4.

$$P(r) = \frac{m_{available(C)}(r)}{\sum_{r' \in available(C)}(m_{available(C)}(r'))} \tag{4}$$

Similarly, let $Q(r)$ be the proportion of agents that must adopt role r computed using Equation 5.

$$Q(r) = \frac{m_{require(C)}(r)}{\sum_{r' \in require(C)}(m_{require(C)}(r'))} \tag{5}$$

Then, we can compute the relative entropy using Equation 6.

$$D(P \parallel Q) = \sum_{r \in require(C)} P(r) \ln \frac{P(r)}{Q(r)} \tag{6}$$

To compute the role disturbance within a coalition in the fixed-hierarchy problem we use Equation 7.

$$disturbance_i(C) = \frac{D(P \parallel Q)}{\ln |require(C)|} \tag{7}$$

Consider the example below.

Example 5. Assume a coalition C such that $require(C) = \{r_1 : 1, r_2 : 2\}$ and $available(C) = \{r_1 : 1, r_2 : 2\}$. Then, $disturbance(C) = 0$. In case, $available(C) = \{r_2 : 2\}$, then $disturbance(C) \approx 0.58$. If the quantity of available roles move away from the baseline requirement, then the disturbance increases as well. For instance, if $available(C) = \{r_1 : 3, r_2 : 2\}$, then $disturbance(C) \approx 0.21$. If $available(C) = \{r_1 : 4, r_2 : 2\}$, then $disturbance(C) \approx 0.33$.

However, it is not always the case that the disturbance function above can/should compute a valid result. We consider some special cases:

- if either $available(C) = \emptyset$ or $require(C) = \emptyset$, then $disturbance(C) = 1$;
- if $|require(C)| = 1$, then $disturbance(C) = 0$.

All that remains to be done is for us to model the constraints on the interaction graphs. To do so, given an interaction graph $G_i : 1 \leq i \leq h$, an edge $(u, r) \in E_i$ connects a superior $u \in S_i$ to any resource if either:

- $r \in L \setminus S_i$ and $resp(r) \subseteq resp(u)$; or
- $|adopt(r) \cap expect(u)| \geq 1$.

Experiments We set a timeout of 30 minutes and record the best hierarchy computed within the time budget. In our first experiment, we aim to form a hierarchy of 101 resources. We use the relation $\mathcal{R}^{\text{BASIS}}$ (Definition 4) to determine the feasibility in an FCSS. We depict the resulting hierarchy (without the section level) in a bar chart representation in Figure 4a. Note that we depict the hierarchy in its natural representation, that is, first the branch level, then the group level and finally, the TF and ST level. One can link the coalitions of different levels by following the coalition Ids. We refer to them using the letter 'C' followed by a numerical identifier, but we omit it in the charts for readability. For instance, in Figure 4a, at the uppermost level there exists coalitions C1, C2, and C3. The coalitions subordinated to C1 are the coalitions C11, C12, C13 and C14 in the second level. Similarly, the coalitions subordinated to C11 are C111 through C116. Underneath each coalition Id, we depict the corresponding coalition size. In case a singleton coalition corresponds to a superior of an upper level, we set both relationship and disturbance metrics to 0 (e.g., C14).

The resulting hierarchy in Figure 4a let no ordinary resource or superior isolated; all coalitions have a corresponding superior and a subset of subordinate resources. One can note a greater role disturbance in coalition C125. The leader responsible for that coalition required 3 different team roles. However, only two were available. The missing role AETRT is available in the instance in the exact number as there are leaders requiring it (see Figure 3). Nevertheless, coalition C112 contains 2 resources adopting that role. The leader of C112 requires the role AETRT in addition to 2 other personnel-related roles, which are of expertise of a particular resource present in that coalition. This makes the overall role disturbance value of 0 combined with a great relationship value of 94%. Moreover, the average relationship value is 49% in the branch level, followed by 55% and $\approx$ 79% in the group and TF/ST level respectively.

Next, we investigate the hierarchy formed by 141 resources. This is the precise number of resources required to form a full hierarchy. That is, one has enough resources in the RRF problem to assign to every leader at the TF/ST level λ resources. We depict the resulting hierarchy in Figure 4b. In this setting we can see that three ordinary resources were left apart in the TF/ST level (C225, C226, and C233) and three coalitions contain exactly 5 members (C125, C133, and C214). All those leader required team-related roles. However, the resources available were of kind personnel. The spare resources could have been removed from the hierarchy, but they are still manageable by the supervisors of the group level, and therefore, kept in the hierarchy. When it comes to the relationship value in the hierarchy, in the branch level we have $\approx$ 46%, followed by $\approx$ 49% and $\approx$ 80% in the group and TF/ST level respectively.

3.2 A Hybrid Hierarchy of Resources

In the previous section, the resources had no choice other than to form coalitions with predefined superiors or stay in singleton coalitions. We relax the fact that all superiors are given as input; some may be given and others will be decided depending on the result of the computed FCSS, hence the name *hybrid*.

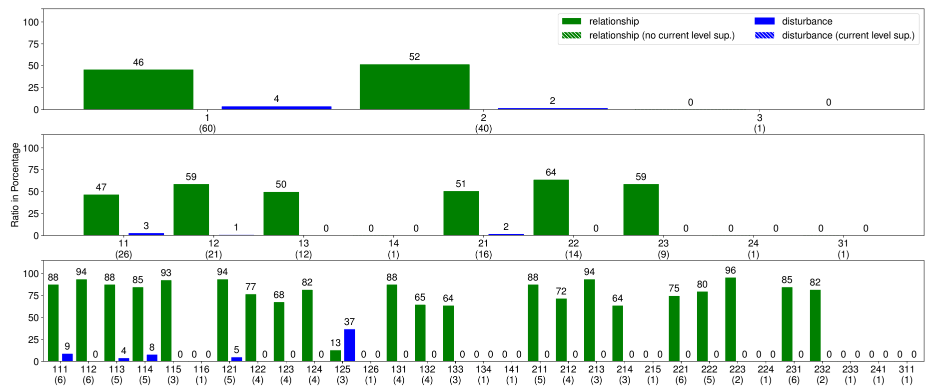

(a) 101 resources.

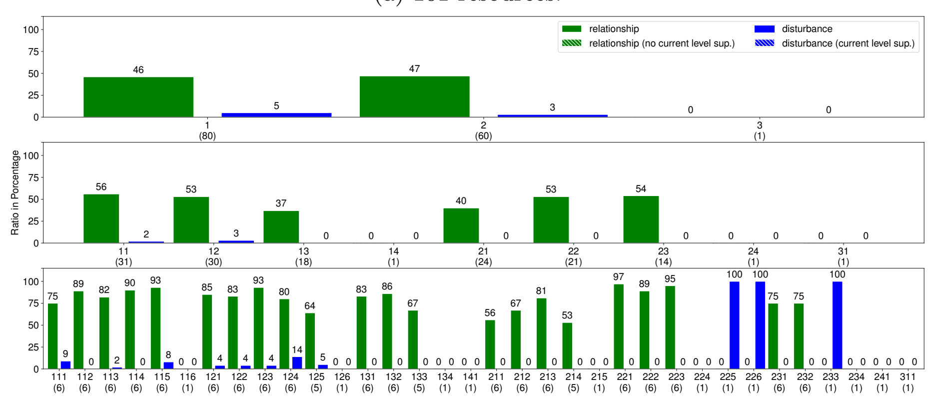

(b) 141 resources.

Fig. 4: Fixed hierarchies of resources computed by UCT-Seq for two RRF instances. The third chart (in both figures) corresponds to the TF and ST level.

Preliminary Given the fact that ordinary resources do not need to be in the same coalition as the pivotal agents, we modify the relationship component to take into account any two agents in a coalition. This is given in Equation 8.

$$relationship_i(C) = \frac{1}{\binom{|C|}{2}} \times \sum_{a,a' \in C : a \neq a'} relationship(a, a') \tag{8}$$

In case C is a singleton coalition, then the relationship value is set to 0.

Regarding the role disturbance, we just need to adjust how we compute the proportion of roles to ignore roles required by superiors. Recall that $available(C)$ is a multiset counting the quantity of roles adopted by the ordinary agents in a coalition C. This time we use an entropy value [4] computed using Equation 9.

Recall that $P(r)$ is the proportion of agents in coalition C adopting role r.

$$H(P) = -\sum_{r \in available(C)} P(r) \ln P(r) \tag{9}$$

Then we compute the role disturbance using Equation 10.

$$disturbance_i(C) = \frac{H(P)}{\ln |available(C)|} \tag{10}$$

We also consider some special cases in this problem:

- if $available(C) = \emptyset$, then $disturbance(C) = 1$; and
- if $|available(C)| = 1$, then $disturbance(C) = 0$.

To allow resources to form coalitions that do not contain superiors, we add more edges in the interaction graphs. Let us introduce an auxiliary notation. Let R^a be a set of roles that are compatible with the roles that resource a may adopt. That is, $R^a = \{r \in demand(o) \mid o \in IO, adopt(a) \cap demand(o) \neq \emptyset\}$. Then, given an interaction graph $G_i : 1 \leq i \leq h$, an edge $(u, r) \in E_i$ such that $r \notin L$ connects any two resources if either:

- $u \in L$ and $adopt(r) \cap expect(u) \neq \emptyset$; or
- $u \notin L$ and $adopt(r) \cap R^u \neq \emptyset$.

We slightly modify the binary relation $\mathcal{R}^{\text{BASIS}}$ (Definition 4) to achieve a hierarchy with no mandatory superiors.

Definition 5 ($\mathcal{R}^{\mathbf{hybrid}}$). *Given $CS, CS' \in \boldsymbol{\mathcal{CS}}^A$, $(CS, CS') \in \mathcal{R}^{\text{HYBRID}}$ iff:*

1. $|CS| > |CS'|$;
2. *for all $C' \in CS'$ there exist at most $\lambda + 1$ coalitions $C \in CS$ such that $\bigcup_j C = C' : 1 \leq j \leq \lambda + 1$;*
3. *if $j = \lambda + 1$, then $\exists \widehat{C} \in CS : \widehat{C} \subseteq C', |\widehat{C}| = 1, \widehat{C} \cap L \neq \emptyset$.*

Moreover, $(\emptyset, CS) \in \mathcal{R}^{\text{HYBRID}}$ for all $CS \in \boldsymbol{\mathcal{CS}}^A$.

We introduced Rule 3 to make sure a superior is assigned to a set of coalitions whenever we reach the span of control plus one. If that is not the case, the algorithm forms at most λ coalitions and no superior *must* be assigned to them.

Experiments For this set of experiments we set a timeout of 1 hour to compute a solution for the problem. We experiment first with 101 resources and depict the results in Figure 5a. In this setting 3 new coalitions are formed in the hierarchy if compared to the fixed one: 2 in the group level and 3 in the TF/ST level. Interestingly, a leader is left in a singleton coalition (C135). This leader was responsible for the coalition of greatest role disturbance in the fixed hierarchy (C125 in Figure 4a). When it comes to the relationship metric, the hierarchy reached 51%, $\approx 60\%$, and $\approx 80\%$ in the branch level down to the TF/ST level

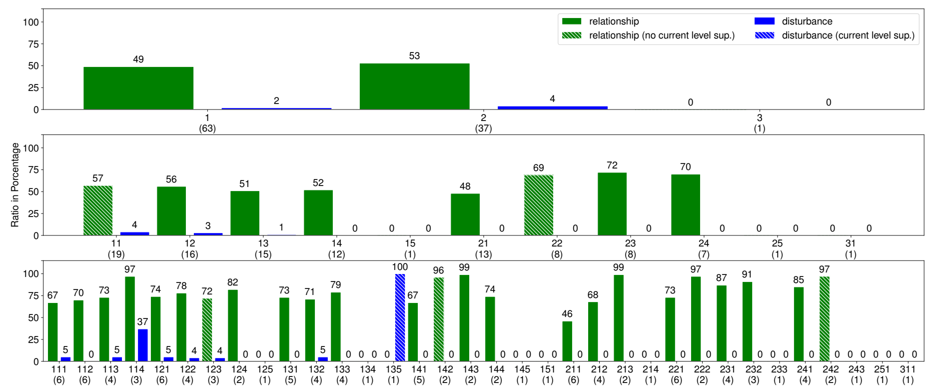

(a) 101 resources.

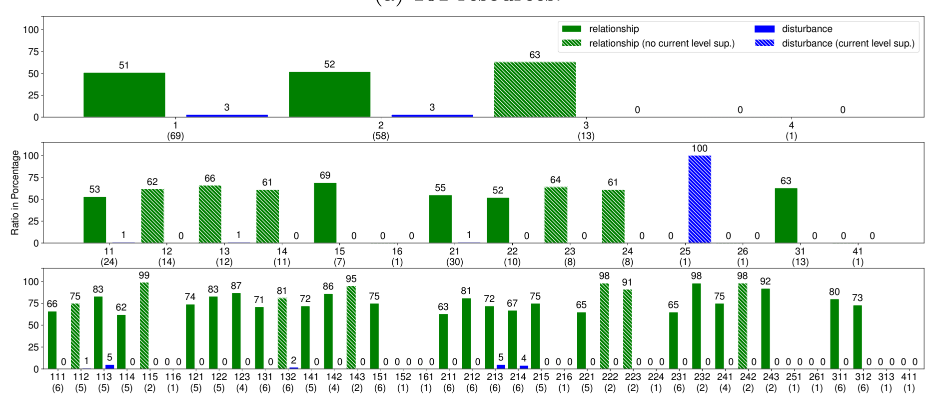

(b) 141 resources.

Fig. 5: Hybrid hierarchies of resources computed by UCT-Seq for two RRF instances. The third chart (in both figures) corresponds to the TF and ST level.

respectively. That means the hybrid hierarchy surpassed the fixed one in terms of relationship in all levels.

Next, we increase the number of resources to 141 and introduce the resulting hierarchy in Figure 5b. In total, 13 coalitions with no predefined superior are formed: 1 in the branch level, 5 in the group level, and 7 in the TF/ST level. The new coalition at the branch level remains the same in the group level with an assigned supervisor which is then split into 2 coalitions in the immediate lower level. The three ordinary resources who did not participate in either a TF or a ST in the fixed hierarchy are now members of coalitions of the corresponding level. Moreover, a supervisor of a group is not used in the hierarchy, instead two new groups are formed with no predefined supervisors. Regarding the relationship metric, if compared to fixed hierarchy, the branch and group levels increase significantly to $\approx 55\%$ and $\approx 61\%$ respectively. At the TF/ST level we note a small decrease of the value from $\approx 80\%$ to $\approx 79\%$.

4 Discussion and Future Work

Our overall goal was to model a real-world application, in particular, the Roaring River Flooding (RRF) scenario [16]. The problem of interest for us is how to form a hierarchy of resources taking into account the designated superiors throughout the chain of command (i.e., a hierarchy) as well as the span of control which defines a manageable ratio between superiors and subordinate units. Our approach is to interpret each level of the hierarchy as a different characteristic function game. The overall hierarchy is then modelled using the SEQSVS framework.

We showed how to model task forces and strike teams in our framework. Although not required in the RRF problem, we believe single resources can also be modelled easily in our proposed formalism. Regarding modelling divisions, one approach could connect agents over the interaction graph that are in the same jurisdiction; that is, edges connecting agents of different jurisdictions are removed. Doing so, one can compute all main elements in the Operations Section. Our modelling might be extended to take into account the entire ICS hierarchy as well. However, the problem then turns out to be how to design valuation functions that are general enough to describe well coalitions aimed for quite different purposes. For instance, coalitions of the Planning and Finance Sections [5]. An important characteristic of multi-agent systems and yet to be addressed in the ICS is autonomous software entities. In the light of the advances in disaster robotics [12], in which an agent might control a robot, many challenges arise, e.g., how to design a valuation function that brings together resources of kind personnel and team as well as those autonomous entities. Those are challenges for any coalition formation framework, not restricted only to SEQSVS.

When it comes to modelling the ICS problem using SEQSVS, the interaction graphs are the main tools to model different sorts of constraints. For instance, in our modelling we used them to link the roles adopted by the resources to the ones required by the superiors. Although those constraints can be very narrow (i.e., a few edges in an interaction graph), for the problems introduced in this work many FCSSs can be formed. We depict in Figure 6 the quality improvement of the solutions computed by UCT-Seq over time (in all experiments). We mark in the chart the precise moment an FCSS of greater value was found. Moreover, for each FCSS computed by UCT-Seq, we mark with a dot the precise point in time it was found and the current value of the best FCSS found so far; that explains the discontinuous lines in Figure 6. Recall that the time budget in our experiments was 30 minutes for instances requiring a fixed hierarchy and 1 hour for the hybrid ones.

One can see that the number of improvements on solutions found by UCT-Seq is greater for hybrid hierarchies, regardless of the time budget. In fact, in those instances, after 30 minutes only a single update on an FCSS were carried out. That suggests the heuristic applied by UCT-Seq works well as we can see that many more FCSSs were computed. For a fixed hierarchy of 101 resources 9.459 FCSSs were computed, whilst 11.751 FCSSs for a hybrid one. For 141 resources, UCT-Seq computed 5.745 and 7.532 FCSSs for a fixed and a hybrid

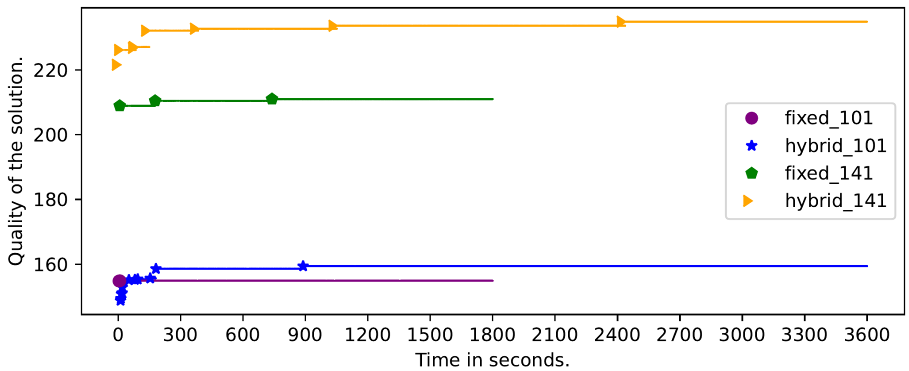

Fig. 6: UCT-Seq running time to progressively improve on the solution quality.

hierarchy respectively. Note that in all cases the algorithm found FCSSs until the time budget is finished.

As a matter of fact, one should bear in mind that the ICS application modelled here lacks feedback from experts on disaster response operations. In particular, incident commanders and section chiefs. They might modify (narrow/relax) the constraints as they see fit as well as the valuation functions. Our main motivation was to show that modelling a real-world problem is feasible using SEQSVS. Further effort is required to validate the proposed modelling of the ICS with experts on the topic, although we followed the ICS, a well-known approach in that area, exactly as reported in the literature. Also, we aim to gather realistic data regarding disaster response operations to feed into our proposed model. Unfortunately, for the time being, this sort of information is not publicly available, even for research purposes. Nevertheless, we believe the produced material is mature enough to be brought into discussion in the disaster response community.

References

1. Ayari, E., Hadouaj, S., Ghedira, K.: A dynamic decentralised coalition formation approach for task allocation under tasks priority constraints. In: Proceedings of the 18th International Conference on Advanced Robotics. pp. 250–255 (2017)
2. Basegio, T.L., Bordini, R.H.: Allocating structured tasks in heterogeneous agent teams. Computational Intelligence **35**, 124–155 (Nov 2019)
3. Bistaffa, F., Farinelli, A., Chalkiadakis, G., Ramchurn, S.D.: A cooperative game-theoretic approach to the social ridesharing problem. Artificial Intelligence **246**, 86–117 (May 2017)
4. Dragomir, S., Scholz, M., Sunde, J.: Some upper bounds for relative entropy and applications. Computers & Mathematics with Applications **39**, 91–100 (May 2000)
5. FEMA: National Incident Management System. Independently Published, 3 edn. (Oct 2017)
6. FEMA: IS-0200.c Basic Incident Command System for Initial Response. Independently Published, 1 edn. (Sep 2020)

7. Irwin, R.L.: The incident command system (ICS). In: Auf der Heide, E. (ed.) Disaster response: Principles of preparation and coordination. C.V. Mosby, St. Louis, MO (1989)
8. Krausburg, T.: Sequences of Coalition Structures in Multi-Agent Systems Applied to Disaster Response. Ph.D. thesis, Pontifical Catholic University of Rio Grande do Sul (Forthcoming)
9. Krausburg, T., Dix, J., Bordini, R.H.: Computing sequences of coalition structures. In: IEEE Symposium Series on Computational Intelligence, SSCI 2021, Orlando, FL, USA, December 5-7, 2021. pp. 1–7. IEEE (2021)
10. Krausburg, T., Dix, J., Bordini, R.H.: Feasible coalition sequences. In: Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems. p. 719–727. AAMAS '21, IFAAMAS, Richland, SC (2021)
11. Mouradian, C., Sahoo, J., Glitho, R.H., Morrow, M.J., Polakos, P.A.: A coalition formation algorithm for multi-robot task allocation in large-scale natural disasters. Computing Research Repository **abs/1704.05905**, 1909–1914 (Apr 2017)
12. Murphy, R.R.: Disaster Robotics. The MIT Press (2014)
13. Rahwan, T., Michalak, T.P., Wooldridge, M., Jennings, N.R.: Coalition structure generation: a survey. Artificial Intelligence **229**, 139–174 (Dec 2015)
14. Ramchurn, S.D., Farinelli, A., Macarthur, K.S., Jennings, N.R.: Decentralized Coordination in RoboCup Rescue. The Computer Journal **53**, 1447–1461 (03 2010)
15. Ramchurn, S.D., Wu, F., Jiang, W., Fischer, J.E., Reece, S., Roberts, S., Jennings, N.R., Rodden, T., Greenhalgh, C.: Human-agent collaboration for disaster response. JAAMAS **30**, 82–111 (Jan 2016)
16. U.S. Department of Agriculture: ICS 200 – incident command system. https://www.usda.gov/sites/default/files/documents/ICS200.pdf (2022)
17. Visser, A., Ito, N., Kleiner, A.: Robocup rescue simulation innovation strategy. In: RoboCup 2014: Robot World Cup XVIII. pp. 661–672 (2014)
18. Wu, F., Ramchurn, S.D.: Monte-carlo tree search for scalable coalition formation. In: Proceedings of the 29th International Joint Conference on Artificial Intelligence. pp. 407–413. IJCAI Organization (Jul 2020)