

Coordination in Adaptive Organisations: Extending Shared Plans with Knowledge Cultivation

Kathleen Keogh^{1,2}, Liz Sonenberg², and Wally Smith²

¹ School of Information Technology and Mathematical Sciences
University of Ballarat, Mt Helen, VIC 3353, Australia
k.keogh@ballarat.edu.au

² Department of Information Systems, University of Melbourne, Parkville, VIC 3052

Abstract. Agent-based simulation can be used to investigate behavioural requirements, capabilities and strategies that might be helpful in complex, dynamic and adaptive situations, and can be used in training scenarios. In this paper, we study the requirements of coordination in complex unfolding scenarios in which agents may come and go and where there is a dynamic organisational structure. This is a step on the way to developing a simulation framework that can be part of a training system in the domain of emergency management. We argue the need for an extension to the SharedPlans formalism required to support the sharing of knowledge about a dynamically unfolding situation, specifically: who is in the team? and who holds relevant knowledge? Our rationale for such an extension is presented based on a prior case study of a railway accident and a further analysis of the coordination and communication activities amongst the disaster management team during its recovery. We conclude that in addition to the obligations imposed by the standard SharedPlans framework, agents in complex unfolding scenarios also need *knowledge cultivation* processes to reason about the dynamic organisational structure and the changing world state. We briefly express the requirements of knowledge cultivation as obligations that could be imposed on agents. We argue that in order to facilitate appropriate knowledge cultivation, agents need access to explicit models of organisational knowledge. This knowledge encapsulates the relational structure of the team, along with shared beliefs, goals and plans. We briefly present a formal representation of this model in order to clearly identify the rich information needed in an adaptive organisation.

1 Introduction

We are motivated to develop believable artificial agents to work with humans in simulations of complex situations such as Disaster Management (DM). Such agents could replace human team members and enable research into team coordination using simulations. Using synthetic agents to provide expert feedback and guided practice in training has also been shown to be helpful [7,22] but the skills

required to coordinate an expert team need to be developed in a realistic and suitably complex simulation environment [2,22]. Teaming agents with humans requires a model for collaboration and coordination.

Team coordination in dynamic situations can be very complex as dependencies between tasks must be managed and priorities can change as the situation changes. Additional complexities are introduced if the team organisational structure is allowed to evolve during the collaboration. So, when people and agents engage in joint activity, agents need to be sensitive to the needs and knowledge of the people. There needs to be an interpredictability of behaviour so that each team player can predict the behaviour of others in order to coordinate their own plans [11].

In the real world domain of DM, the team is distributed across different locations and communication may be limited by radio bandwidth and accessibility. It is usual for multiple service agencies to be involved. The organisational structure is dynamic, people may leave and join the recovery team during the management of the disaster and roles may change. Due to the emergent nature of the situation and the uniqueness of each emergency, team members mutually adjust their behaviour to work with and support others in the team, to achieve a goal. There are shared goals, default plans and some agreed protocols (e.g., the disaster site must be secured by the Fire agency before the Ambulance officers can enter), but the overall coordination is adaptive and flexible.

This coordination is not explicitly controlled by fixed norms. There are significant subtleties evident in an analysis of real-world behaviour that point to a richness of awareness and communication. It could be dangerous to attempt to simplify such to a mere list of requirements, however in an attempt to move closer to a form that might be computationally achievable, we analyse these requirements based on a real world case study in section 2.

The SharedPlans framework [8,9] offers a form of adaptive dynamic team planning needed in DM. Team Situation Awareness and Shared Mental Models of the task and team are not made explicit in SharedPlans, although they are important in enabling adaption in human teams [2]. The SharedPlans formalism focuses on knowledge related to decision making, and although intentions-that in SharedPlans could be employed to address a group's commitment to maintaining accurate and up-to-date shared beliefs about the environment, this has not been explicitly addressed to date. In this paper, we seek to describe the essential requirements and hence capabilities to be enabled in artificial teams for adaptive strategic and coordinated collaboration. We extend SharedPlans to accommodate commitments concerning beliefs about the unfolding situation - including the state of the environment which is only partly observable by any individual, and information about the organisation of individuals involved in resolving the problem.

DM involves the establishment and maintenance of a management system as well as directing and controlling operational tasks. The management system is the dynamic group of agents and resources engaged with sharing the goal of resolving the disaster. Following Smith and Dowell [15], we define DM coordination

as “the resolution of interdependencies between the activities of the disparate resources of the incident organization”. An organisation is a set of actors, with a social order and working toward a common goal. The organisational structure defines roles and enables coordination[3]. Coordination requires adaptivity that occurs within ‘coordination loops’ or collections of groups that form to achieve a goal. Such groups require synchronisation and mutual adjustment of tasks[20]. Within the DM system, sub-teams or groups form to plan and act together to resolve subgoals (e.g., secure site, attend to injured). These sub-teams may involve people from different agencies working together, who form a team to collaborate on a shared goal. When such a (sub)-team emerges during DM, without a need for explicit negotiation, there must be assumed obligations and responsibilities taken on regarding the need to communicate and share related knowledge - knowledge of the situation and regarding planning for action, as well as knowledge regarding who is involved in the team [2]. In our abstract architecture, we attempt to explicitly describe the processes and artefacts needed to enact these obligations. The next step is to formalise these proposed extensions.

When teams of people are working to resolve a situation, as the following case analysis will highlight, communication to share vital information regarding the accident situation can be crucial to a coordinated response. When artificial agents become part of teams collaborating in such domains, we suggest it is important to provide a mechanism for awareness and knowledge sharing in agents to avoid problems derived from decision making based on incorrect knowledge. Teams need to establish common ground - sharing pertinent knowledge, beliefs and assumptions [11]. For effective team coordination and collaboration, beliefs need to be shared and to some degree beliefs need to be held as mutual. In order to facilitate appropriate knowledge sharing, there needs to be awareness of who is ‘in the team’, ‘who knows what’ and ‘who might need to know’.

Agent technology has been shown to be useful in aiding human teams in response to emergency situations (e.g. [1,4,13]). Tate and colleagues relied on human initiative and intuition to control information availability to other humans in the team [1], others have used adjustable autonomy allowing agents to defer to humans for complex decisions [13]. Using an information system to deliver filtered, context based situational knowledge to appropriate people in the DM team could help human performance[19]. Our architecture is aimed at naming explicit obligations that the artificial agents need to have to compel the appropriate automated sharing of knowledge. Having shared dynamic artefacts available to enable human collaboration has been shown to be effective [1]. Sharing mental models and agents having a representation to enable the identification of relevant knowledge to share has been demonstrated to improve human behaviour in human-agent teams [4]. Sharing mental models including shared situation awareness is important to the functioning of adaptive human teams [2,12,19]. Our interest is in the design of realistic (believable), autonomous, adaptive, artificial team agents.

In the remainder of this paper, we outline requirements of adaptive coordinating teams in DM by analysis of an existing real world case study. We then suggest

the capabilities needed for adaptive team agents. We formalise the requirements in terms of obligations on each team member and suggest an extension to the SharedPlans framework to address these issues. We have chosen SharedPlans as it provides an intention driven formalism based on adaptive human behaviour. We hope it would result in believable and predictable agent behaviour. We present a model of the components that comprise an adaptive organisation and express these using an extension to an existing formalism. This formal organisational model provides a language for describing the adaptive organisations in our case study.

2 Case Study – Ais Gill Train Accident

In this section, we describe the scenario of a train accident based on previous work by Smith and Dowell [15]. Our analysis is focusing on identifying beliefs and knowledge emerging about the situation and the associated communication to share this. We highlight the need for appropriate sharing of beliefs across a distributed team in order to successfully coordinate a response. This case study motivates our proposal in section 4.

2.1 Overview of Accident Scenario

At 18:49, 31st January, 1995, UK Emergency Services were notified that a train had become derailed somewhere between Kirkby Steven and Blea Moor in the county of Cumbria. Six minutes later, a second train crashed into the derailed train resulting in escalation of the incident. A train conductor was killed, 6 passengers and a train driver were seriously injured. A significant period of time elapsed (about an hour) during which the exact location of the train crash was not clear. Initially, the emergency services were unaware that the second train had crashed into the derailed train and the number of injured was thought only to be 2. A number of agencies were involved in the response: Ambulance, Fire, Police, Volunteer Mountain Rescue, and Railtrack. The accident site was not easily accessed by road and the first to arrive on the scene were fire crews 30 minutes after the accident. When ambulance crew from Brough arrived (55 minutes after the initial accident) it became known that there were closer to 30 injured people and one deceased. It took a further hour after the location and number of injured was known, before decisions regarding how to transport the injured to hospital were finally settled.

The seriously injured driver was carried along the train track to meet an ambulance at a road bridge. It was raining and this proved hazardous. Earlier, a request was made inquiring about the possibility of sending a rescue train to transport the injured via rail to Carlisle train station, to be met by ambulances. This emerged into a reality, and so other injured people were kept dry in the train waiting for the rescue train. Communication difficulties occurred with the inappropriate dissemination of this decision: ambulances were not informed in a timely way and continued en-route to the disaster site rather than being

redirected to Carlisle railway station. There was also confusion regarding the eventual destination of the injured, so one hospital remained on standby longer than necessary.

2.2 Development of Situational Knowledge Regarding the Incident

As with many disaster situations, the knowledge regarding details of the incident was not clear initially and were established with difficulty over an extended period. In a DM team, the people involved need to share obligations to pass on relevant information to others in the team. As the actual situation changed at Ais Gill, beliefs were shared and revised. This is ‘fill in the blanks’ type of coordination - establishing uncertain facts such as location and number of injured. During this first phase, the team focus is information gathering and situation awareness[2]. Information is shared according to protocols and obligations to keep others in the disaster recovery system informed and up to date. Some incorrect conjectures were made whilst details were uncertain. It was assumed that there wouldn’t have been many passengers on the train, so the number of injured passengers was (incorrectly) assumed to be small.

Communication was significantly constrained. Firstly, in the hilly terrain there was disruption to radio. Second, ambiguity and uncertainty about the situation made it difficult to know what to communicate. Third, while many agencies are in the process of building up their response, each individual agent knows little about who to communicate with. The difficulties and errors in communication highlight these challenges in such a dynamic situation. But, passing on relevant information to others in the system is crucial. The fire crew, arriving at 19:25 passed on accurate location information ”1 mile north of Ais Gill” to the fire control, who passed it to railtrack, who passed it to ambulance control centre in Carlisle, who then passed it to the Brough ambulance crew with instructions to change course en route to the site. Examining this communication with the eyes of a potential designer of a response system including artificial agents, we conclude that agents engaged in such a scenario need obligations to follow similar protocols to humans: to pass on new or revised situational information such as the accident location and number of injured.

Observation 1. The agents in an adaptive team need obligations defined to pass on new or revised information about the disaster situation to relevant others in the system. This implies that agents need to have a mechanism to judge the relevance of information to pass on. This is complicated by the continuous nature of change in the task and that relevance is relative to the recipients’ context - knowledge and experience that the sender doesn’t necessarily possess.

2.3 Coordinating Interdependent Dynamic Goals

It could be said that the people involved were all motivated by the high level goal to resolve the disaster, enacting a sub-goal to mobilize resources to the site. Control centres off site have responsibility for strategic and planning goals.

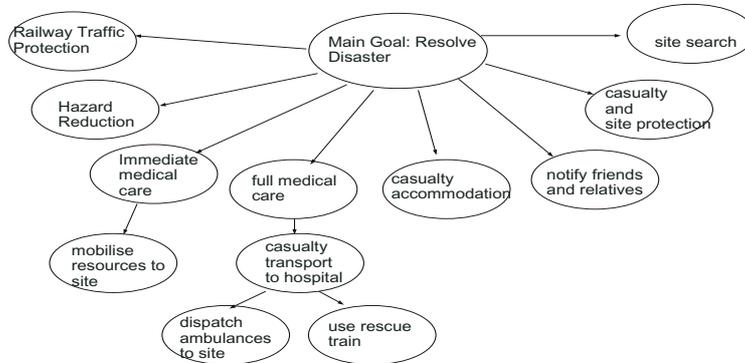


Fig. 1. Initial goal decomposition

Recipe for **GOAL**: Save Lives and Relieve Suffering

SUBGOALS: (listed in order of priority)

subgoal	group allocated to enact this goal
1.Railway Traffic Protection	Railtrack
2.Hazard Reduction: Fire, instability	Fire and Regional Railway Company
3.Immediate Medical Care	Fire and Ambulance
4.Casualty Transport (to hospital)	Fire, Ambulance, Volunteers, RailTrack, Regional Railway Company
5.Full medical care	Hospital H
6.Casualty Accommodation	Regional Railway Company
7.Notification of friends and relatives	Police
8.Casualty and Site Protection	Fire and Police
9.Site search	Police

Fig. 2. Generic High level Incident Plan

Unknown parameters such as exact location and number of injured need to be established. Based on an analysis of this incident [15], we can suggest a potential partial hierarchy of goals (Figure 1) that represent the system early in the response . This is based on the high level incident plan - a default allocation of roles and responsibilities in any incident (Figure 2). Each agency would likely operate following generic recipes that are initially partial and then elaborated and verified as the situation unfolds. In this case of DM at least, there was the need to maintain multiple possible goals and to make (and enact) multiple, possible alternative (partial) plans based on available information. It is not practical to wait and establish facts in isolation before any action is taken, so it is better to enact multiple options until it becomes clear which will be fully enacted. Plans are revised or dropped as more details are established as will be discussed in the next section.

The high level incident plan for the disaster response in Figure 2 shows the broad allocation of responsibilities to each agency [16]. Sometimes these goals

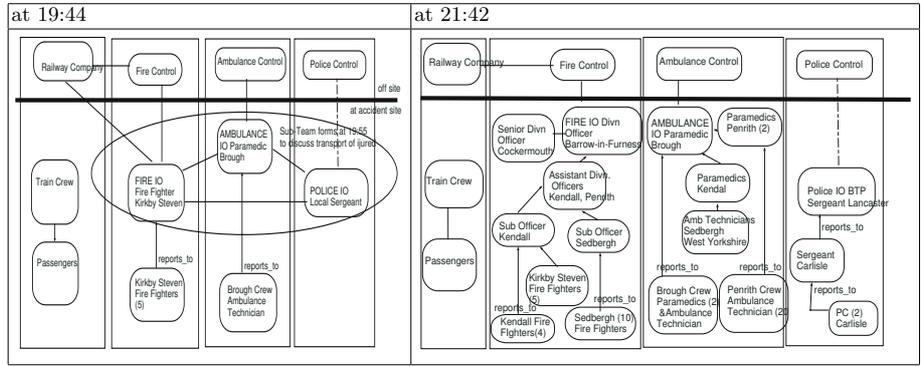


Fig. 3. Organisation Structures at Ais Gill

might be enacted autonomously, though in most cases, interdependencies require collaboration and communication between agencies for coordination.

At 21:42 when the rescue train arrived at the accident site, for example, collaboration and coordination of interdependencies occurred between many agencies: fire, police, ambulance, railtrack and the railway company.

2.4 Evolution of the Organisational Structure

The system which manages a disaster - agents, resources and technologies - is not only distributed over space but its structure evolves over time.

We can picture the disaster recovery system that includes each of the response agencies: police, fire, ambulance, plus organizations such as railtrack, volunteers and emergency services. A simplified organisational network focusing on the 3 main agencies: Fire, Ambulance and Police depicting how system might look early in the response and then about 2 hours later, is shown in Figure 3.

The command structure develops organically. Roles of individuals are adjusted flexibly responding to the disaster with resources available [16]. Organic role adjustment occurs with the Fire Incident Officer (IO) role reallocated to 4 different people, driven by new arrivals at the site. At 19:20, the leading Fire Fighter from Kirkby Steven assumes the IO role. As a more senior ranking officer arrives at the scene, that person may take over as IO. At 19:35 a sub-officer from Sedbergh became IO. At 20:01 two Assistant Divisional Officers (ADOs) were present and one of these (Penrith) took over as IO, with assistance from the other (Kendal). At 20:11 the Divisional Officer from Barrow-in-Furness arrived, worked collaboratively with the 2 ADOs, then gradually took over the IO role. In some cases, the role of IO is not handed over - for example, when the Senior Divisional Officer from Cockermouth arrived at 21:00, he made it explicit that he would not be taking command from the existing IO, but would remain and observe.

As the organisational structure changes, the team needs to be aware of current team structures to enable appropriate adjustment and sharing of planning and situational knowledge. At Ais Gill, there were difficulties in maintaining a shared

mental model of the organisation [15] and as will be described in the next section, this hampered communication of a change in the plan regarding transport of injured to hospital. This led to further coordination problems. Team shared mental models of resources, tasks and team members are essential for team processes [2]. A shared representation of the organisation is needed to enable knowledge transfer and task level shared mental models[15].

Observation 2. In order for appropriate sharing of relevant information, adaptive team agents/people need to have some representation of the current DM system's organisational structure.

*Observation 3. Adaptive Team Agents need defined obligations to explicitly **maintain** an appropriate shared mental model of the DM system's organisational structure.*

Observation 4. As the organisation is dynamically adapting structurally and members may join/leave, there is a complexity beyond a simple predetermined hierarchy or pattern. This subtlety requires an awareness of relevancy linking knowledge to actors in the organisation, to enable appropriate knowledge sharing.

2.5 Development (or Revision) of Plans in Response to the Incident

During an incident, decision making occurs to assign resources to operational or management tasks. These decisions are informed by the current incident plan and knowledge of the current situation. Both of these are changing with time as more details are known. The incident plan is initially based on generic predefined plans described at an organizational level; it is then developed or revised in response to the actual situation. In DM, there are typically two types of coordination: Filling in the unknown details into existing predefined default plans (e.g., location of incident, type of incident, number and type of casualties); and revising a chosen plan recipe in response to the situation, (e.g., the decision regarding the transport of injured to hospital).

Knowledge regarding the Ais Gill incident was distributed between people and the decision to use a rescue train was shared. A sub-team including the paramedic, police sergeant and leading fire fighter were partially responsible for the revised decision to use a rescue train, although it remains unclear when this proposition became an actual decision. They queried if a rescue train might be possible at 20:00. At 20:10 the reply, without further communication or collaboration was 'ETA train 60 minutes and also coach from Robinson of Appleby is mobile'.

This revised plan was not shared with ambulances off site, resulting in them continuing to travel toward the roadside near the accident where they could not be useful. The ambulance controller in Carlisle didn't know about the rescue train until 20:55. During the 45 minutes prior to 20:55, ambulances and PTS were dispatched toward the site where they weren't needed. This illustrates the importance of keeping relevant others informed of new/revised plans. The communication regarding the decision and adopting of a revised plan to transport

- | |
|--|
| 1 Establish location of accident (LA)
2a CHOICE 1: transport casualties to Hospital, H. via ROAD using ambulances with pick up at LA access point. [DEFAULT option assumed in initial plan]
2b CHOICE 2: transport casualties to Hospital, H. using a RESCUE TRAIN to railway station RS
2c CHOICE 3: transport casualties to Hospital, H. on foot along train track then via ambulance with pick up at LA - nearby road-bridge
unknown parameters: LA, RS, H(either Carlisle or Lancaster) |
|--|

Fig. 4. Potential recipe for transporting injured to hospital

the injured to hospital using a rescue train involved a plan revision. This revised decision was not well communicated and led to coordination problems later (for example, not enough ambulances were available in the correct location at the train station when the rescue train arrived; staff at Lancaster hospital that could not have been used due to the limitations on the direction of the train were kept on standby for longer than necessary; and adequate police resources were not at Carlisle station to provide protection to the injured passengers from media).

2.6 Multiple Uncertain, Partial and Adaptive Plans

Typical of disaster recovery, at Ais Gill, multiple plans were concurrently considered and partially enacted, due to uncertainty and incomplete knowledge availability. A disaster team cannot afford to do nothing until all information is known. Planning and action occur in parallel. As in this case, decisions may not be made explicitly, but may emerge as the best option and the actual selection of the final decision may not be clearly distinguished. Our architecture for knowledge sharing in a human-agent team needs to account for **multiple plans** and may well need to **provide a mechanism for recognising the certainty related to knowledge** (for example: uncertain, possible, probable).

Assuming predefined recipes outlining actions toward goals, existed at least implicitly for the people involved as Ais Gill, we might imagine that there was a recipe for the goal: Transport Injured to Hospital. It could look something like the recipe shown in Figure 4. Initially, all the participants would have adopted the default plan assuming that the injured would be transported to hospital via road using ambulances. Due to particular constraints of this disaster, the plan was revised. At Ais Gill, all three choices were partially enacted, before choice 2 became the final decision. We would presume the poor communication regarding the change of plan was not due to the people concerned being unaware that others needed to know, but it could well be the case that the obligation to tell was not clearly defined. Who was responsible for telling those away from the disaster scene? It could have been presumed by the people at the scene that others away from the scene had already been told by the railway authority. During phase 2 of the team response - plan formulation [2], this example serves to indicate that when creating artificial agents to engage with people in coordination in such a domain, **obligations to share knowledge about (new or revised) plans**

need to be clearly stated. This is in part addressed by the SharedPlans formalisation within the decision making team, though needs extension to include explicit sharing with others in the system who may need to know. This will be discussed further in section 4.

Observation 5. Adaptive Agents need obligations to share knowledge about changed elements of plans to relevant others in the organisational system.

This particular incident exemplifies some general features of DM that are important to note. The changing demands of a disaster situation result in instability of the incident organisation [16]. Decision making can be highly reactive due to the novelty and instability of most disasters. Decisions can be made in part by following predefined rules and protocol but also a level of flexible adjustment of these plans is needed to deal with the developing nature of an incident.

During phase 3 - plan enactment [2], further team adaption and flexibility is needed to respond to changing plans and awareness is needed understanding uncertainty relating to multiple plans being partially enacted concurrently. At Ais Gill, the ambulance controller recalled the ambulance crews en route to site at 20:55 when he learnt of the alternative plan to use the rescue train.

3 Defining Components of an Adaptive Coordination Team

In order to be coordinated whilst being distributed across multiple locations and between various agencies, as our case study shows, we need to be able to share up to date, mutual knowledge regarding the situation. This mutual knowledge includes details of the situation, current and future plans, the current team organization and ‘who knows what’. In order to know ‘who else might be interested’ in new or revised knowledge, the current system organization needs to be available to ‘look up’ and attached to the people/agents in the system, we need an indicator of relevancy: who is interested in particular beliefs.

In this section, we describe in abstract terms features of a team capable of the coordination described thus far. In DM, there are many teams that work as part of one overall system response team. Each sub-team that forms is linked by either sharing plans, or sharing knowledge or an interest in certain knowledge. Overall, the team is distributed and it is unlikely that all members have access to all knowledge. Such a team comprises of actors or agents in multiple organisational networks based on agents’ relationships with others in terms of shared knowledge interests, hierarchy of organisational structure, and shared goals. These organisational groups are similar to the coordination loops described for adaptive and complex situations[20]. Coordination loops are defined in three classes: horizontal, vertical and projective. These different loops indicate multiple ways that individuals can be connected and inter-related. Such loops define reciprocity obligations and enable anticipatory action and describe the complexity of interdependencies between tasks. An awareness of the coordination loops/organisational networks provides knowledge about relevancy and connectedness.

The SharedPlans theory [9,8] accounts for group obligations regarding the cultivation of shared group intentions and dynamic plans, but doesn't explicitly address sharing of knowledge - regarding the team/organisational structure or the situation. In order to establish common knowledge, other architectures for coordinating agent-human teams have successfully employed mental models. RCAST [4], a system for sharing of situation awareness in teams, incorporates an explicit Shared Mental Model (SMM) of team structure and task knowledge which we would also propose in an adaptive team. The RCAST architecture is based on a recognition primed decision making process rather than the SharedPlans decision making process, however using explicit shared mental models complemented by intention based obligations to maintain these suits the adaptive nature of our case study. In RCAST, SMM components include team processes, team structure, shared domain knowledge and information-needs graphs [5]. Other generic agent team models for collaboration exist (e.g. [18]).

Situation assessment and representation of situation (e.g., mental representation, status boards, maps, reports) inform the decision making in a disaster response [16]. SMMs have been used in various multi-agent team experiments, as explicit separate artefacts and implicit in shared ontologies. Yen categorises 3 types of SMM: Blackboard based, joint mental attitudes and rule based systems [6]. The model in the RCAST architecture is more than a set of mutual beliefs, all agents share a commitment toward maintaining shared awareness and proactive sharing of relevant information[21]. This model is consistent with our proposed extension to SharedPlans. Rather than use a shared artifact to represent shared knowledge, an alternative is to have individual beliefs that are shared and kept consistent by agreed update processes. When agents work together in a coalition, it is possible to broadcast interests to other agents in the group to establish information needs and enable mutual support[14]. When the team structure is dynamic, the appropriate sharing of knowledge relies upon an ability to recognise relevancy of knowledge associated with others.

Maintaining mutual beliefs between individuals requires that each individual has beliefs about the beliefs of others in the team. In DM, due to the distributed nature of the system, one central shared artifact is not feasible, so we seek to describe *processes* to ensure sharing and maintaining of knowledge about the situation. This includes beliefs such as location of the incident, number of injured, current resources deployed etc. Knowledge cultivation as described in section 4 begins to describe the details of how we might achieve sharing of relevant knowledge in adaptive teams. Such a shared model of knowledge and team organization, is comprised of an organizational network, a belief network and a relevancy network; plus the obligations to cultivate this shared knowledge.

4 Knowledge Cultivation – Extending Shared Plans

Grosz and Kraus have defined obligations that must be adhered to for a group of individuals to maintain shared (individual) plans around a shared goal [9]. These *intention cultivation* obligations enable dynamic, real time planning and

decision making in a team without relying on shared artefacts, but controlling the mutual knowledge by obligations regarding reasoning about intentions. Grosz and Hunsberger, [8] argue that group decision making in planning and intention updating are crucial to establishing collective intentionality for coordinating dynamically. Coordinated updates of group-related intentions are the key to modelling collective intentions.

The SharedPlans formalism satisfies two requirements in our DM system team: adaptive planning - enabling a group of agents to share decision making dynamically; and distributed team work, such that a problem can be decomposed into sub-goals that can be allocated to independent sub-teams to autonomously enact. In DM the team involved is not fixed at the start of the problem – team structures emerge as the event unfolds.

We introduce *knowledge cultivation obligations*; and *explicit models of the organizational network* - linking agents/people and knowledge.

In SharedPlans, intentions are categorised into four specific types - *intentions-to perform an action*, *intentions-that a proposition be held true*, and *potential intentions* (identified, but not yet adopted) [9]. Recipes for achieving a group activity, are defined in terms of sub-acts and parameters that can be filled dynamically. Collective group intentions are used to constrain that all participants have uniform intentions, and that any agent updates their intentions only in accordance with the group [8].

“Agents have a Shared Plan to do α if and only if they hold the following beliefs and intentions: 1. individual intentions that the group perform α ; 2. mutual belief of a (partial) recipe for α ; 3. individual or group plans for the sub-acts in the (partial) recipe; 4. intentions that the selected agents or subgroups succeed in doing their sub-acts (for all sub-acts that have been assigned to some agent or group); and 5. (in the case of a partial Shared Plan) subsidiary commitments to group decision-making processes aimed at completing the group’s partial plan.” [8]

Our extension includes three components:

- TSMM - a task level SMM including situational knowledge of the problem;
- RSMM - a Reflexive SMM representing the structural organization of the dynamic team [15]; and
- KC - knowledge cultivation obligations: intentions to establish and maintain common shared content amongst the team.

In the remainder of this section, we define Knowledge Cultivation obligations. In section 5, we describe a representation of the structural organization that can be used to formalise RSMM and TSMM.

We adopt the formalisms used in SharedPlans formalism, where possible. We have used the operators Mutual Belief ($MB(GR, \phi)$) interpreted as group GR mutually believes proposition ϕ ; $Bel(G, \phi)$ meaning Agent G believes proposition ϕ ; Intention.To ($Int.To(G, A)$) Agent G intends to do action A; and Intention.That ($Int.Th(G, \phi)$) representing agent G intends that proposition ϕ holds [9,10]. We introduce additional predicates to informally represent elements needed in the descriptions, described in table 1.

We define SGR, the System GRoup, to represent all the agents/people who are part of the DM system at any stage during the DM response. For now, we represent any organisational network (coordination loop) as RSMM and the state of the world at some time as TSMM (these will be more formally elaborated upon in section 5). The following obligations describe the processes that we impose on the SGR.

- SGR mutually believes that all members of SGR are committed to success of the goal to resolve the disaster;
- SGR mutually believes that all members of RSMM are committed to updating RSMM so that, at any time, the union of all RSMM (multiple organisational subteams) reflects the current management system structure;
- SGR mutually believes the need to update and maintain relevant knowledge about the state of the world (TSMM).

These obligations are imposed on any sub-team that forms. The SGR is a union of all RSMM that occur during the team disaster response.

We explain these obligations based on examples from the case study introduced in section 2.

1. Awareness of system group goals. The system group represents all involved - the injured passengers, the ambulance, fire and police agencies, British Rail and volunteers. The high level shared goal is to resolve the disaster. Each agency does not work in isolation, but is aware that other agencies are involved. There are dynamic sub-teams representing each agency, Fire, Ambulance, Police, Railway Company, similar to figure 3. There are also sub-teams that form between agencies such as the Incident Officers responsible for decision making considering options for transport of the injured. Even though, undoubtedly there was an awareness of others off site, there were problems with communication of their revised plans to ambulance control center in a timely fashion. To avoid this kind of error occurring with artificial agent teams, we need to be explicit in creating an awareness of the system group (distributed over time and space).

This group SGR has a mutual belief that all members have a goal to achieve α - the goal to resolve the incident.

$$MB(SGR, \forall g \in GR_i : g \in SGR \wedge Int.th(g, \alpha))$$

Table 1. Predicate definitions

RM(gr,t)	Reflexive Model for group gr at time t. This represents the organisational structure of group gr and identifies relationships between agents.
A Rel.to B	Agent A is Related to agent B. A and B are both linked in an existing organisational hierarchy or coordination loop - a reflexive model(RM).
Add(x,gr)	Add agent x to group gr.
Update(x,RM)	Update RM to include agent x.
Relevant(b, BS)	belief b is relevant to a set of beliefs BS.
Focus(b,GR)	Belief b is in the relevant focus of interest to an agent in group GR. This may be established by a goal that is related to this belief or a role allocated to an agent in GR that is associated with this belief

2. Obligation to update RSMM as needed. As figure 3 indicates, the organisational network, thus RSMM representing this is very dynamic. The communications during initial phase of establishing the details of the situation relied on appropriate notification to others in the DM system. This could only occur because there was knowledge of who else was involved. At Ais Gill, the decision-making team on the ground needed to know to share their decision regarding how to transport the injured with the ambulance controller to enable plan revision and strategic response to reallocate resources (the hospital on standby and the ambulances en route). We capture this as:

The group mutually believes that all members are committed to updating the RSMM to include new members as they join the system, where the current RSMM at time, t is represented by $(RM(SGR,t))$.

$$MB(SGR, (\forall g_j \in SGR, Int.th(\exists g_k, \exists g_i : \neg g_k \in SGR \wedge g_i \in SGR : g_k Rel.to(g_i))) \\ \Rightarrow Int.to(g_j, add(g_k, SGR) \wedge update(g_k, RM(SGR,t))))$$

3. Knowledge cultivation - obligations to share relevant knowledge. The group has a mutual belief regarding need to update shared situational task knowledge TSMM with new relevant beliefs that are not already shared. Relevant beliefs need to be identified as such by a relationship to existing beliefs or tags to show relevancy in the belief knowledge set, BS. Updates can also be motivated by recognising that some agents have an existing focus (interest) in this knowledge.

$$MB(SGR, \exists g \in GR : Bel(g, b) \wedge \neg b \in BS \wedge \\ (Relevant(b, BS) \vee Focus(b, GR)) \Rightarrow Int.to(g, add(b, BS)))$$

As in R-CAST[5], when information is distributed, and proactive sharing required, there is a need to represent information relevancy so that it is possible to ascertain if a particular piece of information is relevant to another agent. Creating an associated obligation and commitment to agreed knowledge updates would enable agent responsibility regarding communication of knowledge to other agents to be formally represented. To recognise relevance, an agent needs a representation of their own and others' focus [17]. In R-CAST, information needs graphs are used to associate information with context, enabling matching of knowledge to previous experiences and providing a mechanism for identifying when information might be of future relevance to another team-member [5]. In the DM case, shared focus may be identified by relational links between agents within a sub-team or organisation, or identifying the coordination loops[20] that are present. In the next section, we begin to represent the structure of knowledge representing the disaster management organisation more explicitly.

5 Representing the Organisational Knowledge at Ais Gill

In this section, we formally represent the shared mental models in the organisations involved in our case study. We identify elements that define an organisation

and give some examples of the RSMM organisations that exist based on the Ais Gill case study. This analysis takes a step closer to expressing an adaptive organisational system in a computational tractable way. Knowledge cultivation obligations define intentions to maintain these models within an organisation.

To represent each organisation involved over time during the disaster scenario, we use a formal organisational model[3]: Given W , an organization O is defined in a world $w \in W$ as: $O^w = \{A_o^w, \leq_o^w, D_o^w, S_o^w\}$ where $A_o^w = \{a_1, \dots, a_n\}$ is the set of agents, $\{A_o^w, \leq_o^w\}$ is a partial order relation on A reflecting the structure of the organization, $D_o^w \subseteq \Phi$ is a set of objectives (states to achieve), and $S_o^w \subseteq \Phi$ is the set of current states relevant to this organisation holding at a given moment. We extend this model to include two additional components in an organisation: B_o^w is the set of current mutual beliefs of the organisation and SP_o^w is the set of current shared plans held by the organisation.

At Ais Gill, there were multiple organisations, with some overlapping objectives and some agents belong to more than one organisation. Each organisation, at any time, can be described using the above formalism. The SGR includes all agent sets involved in each organisation during the scenario, so SGR in this case is a union of the fire, ambulance, police and incident officer teams over time. The RSMM representing each organisation is the shared model of structure (A_o^w, \leq_o^w) . The TSMM is the union of beliefs, objectives (goals) and shared plans: B_o^w, D_o^w and SP_o^w . The set of objectives (D_o^w) can be broken into two subsets - action based objectives and information needs objectives. In the case of the DM system, as discussed, it is necessary for the organisation to access networks of relevancy. Expressing our organisational model using this formalism imposes a structure through which relevance can be identified.

Agents in our DM system include Fire Fighters (ff), Ambulance Paramedic (ap), Control (fc), Fire Incident Officer (fio), Ambulance Incident Officer (aio), Police Sergeant (ps), Objectives include: Transport to site (tS), Hazard Reduction (hR), Immediate Medical Care (mC), Casualty Transport (cT), Site Protection (sPr), Assess situation (aS), Share location details (shLoc). Information needs objectives include: Establish Location of Incident (getLoc), Establish Number of Injured (noInj). Each organisation that forms has a shared objective - either strategic or practical.

The Fireteam and Ambteam organisations are hierarchically based and share generic goals as shown in Figure 1. At Time 1, 19:20, the organisational model for the fireteam could be described as follows.

$$O_{oft}^1 = \{A_{fireteam}^1, \leq_{oft}^1, D_{oft}^1, B_{oft}^1, SP_{oft}^1, S_{oft}^1\} \text{ where}$$

$$A_{fireteam}^1 = \{fc, fio, ff1, ff2, ff3, ff4, ff5\}.$$

$\{A_{fireteam}^1, \leq_{oft}^1\}$ is a hierarchical structure:

$$fc \leq_{oft}^1 fio, fio \leq_{oft}^1 ff1, fio \leq_{oft}^1 ff2, fio \leq_{oft}^1 ff3, fio \leq_{oft}^1 ff4, fio \leq_{oft}^1 ff5.$$

$$D_{oft}^1 = (hR, mC, cT, sPr).$$

$$B_{oft}^1 = (traincrash, loc : 1mileNthAisGill, 2injured).$$

$$SP_{oft}^1 = (shLoc, aS).$$

$$S_{oft}^1 = (traincrash, loc : 1mileNthAisGill, 30injured, 1DOA).$$

Information was shared from the fire crew from Kirkby Steven who arrived at 19:25 to fire control to give accurate details on the location (originally thought to be Birkett Tunnel). Fire control then passed this information to Railtrack who passed it to Ambulance control and the Brough ambulance crew en route to Birkett Tunnel changed route to take B6259 to the accident site. The presence of the RSMM $\{A_{fireteam}^1, \leq_{oft}^1\}$ provides the awareness that enables this information sharing within the fireteam organisation.

At time 2, 19:55, the ambulance crew from Brough had arrived on scene. The more senior ambulance paramedic assumed the role of Ambulance IO. The ambteam organisation hadn't yet changed in structure, but TSMM begins to develop new beliefs and SharedPlans. These would be represented in new values for B and SP respectively in the model of the ambteam organisation.

The IOteam forms on site at 19:55. This comprises the Ambulance IO, Fire IO and Police Sergeant, shown circled in Figure 3. Members of the IOteam are in multiple organisations - at least one for their agency (Fire/Ambulance) and one for the IOteam. The IOteam generates SharedPlans for how to transport the injured to hospital. Multiple plans were enacted with ultimately one being chosen eventually.

Knowledge Cultivation obligations would motivate the IOteam to share this eventual plan choice with *relevant* others in SGR. In this case, each IO would be obliged to share the plan with their relevant control officer who is in their other organisation RSMM - Fireteam (oft) or Ambteam (oat). The IOteam structure is not a hierarchy, so \leq_{oio}^2 does not contain any ordering relations. Their SharedPlan SP_{oio}^2 would reflect Figure 4. The IOteam is represented as:

$$O_{oio}^2 = \{A_{ioteam}^2, \leq_{oio}^2, D_{oio}^2, SP_{oio}^2, S_{oio}^2\}. A_{ioteam}^2 = \{fio, aio, ps\}. D_{oio}^2 = (PlanCT).$$

As the above examples show, appropriate transfer of knowledge relies upon a structure for representing relevance of knowledge. Associating knowledge with relational organisational structures provides a network of linking knowledge with groups of agents and enables agents in an organisation to determine which other agents may have an interest in knowledge. As agents may belong to multiple organisations, each with a different focus, this limits the scope for broadcasting information, but enables information to be passed on to appropriate agents.

6 Conclusion and Future Work

We have described the complex demands of disaster management and proposed an extension to SharedPlans that offers a step toward meeting these demands. SharedPlans provides a framework for sharing and maintaining a team mental model of decision-making and intentions but does not account for changing team structures. We have used an existing formalism to represent the additional requirements for an adaptive organisational system.

The contribution in this paper is to introduce knowledge cultivation processes to complement the intention cultivation processes already in SharedPlans. Our underlying motivation is similar to that of Tambe and colleagues, e.g. [18], who

argued for a generic teamwork model as a way to reduce the burden on the agent designer in the development of complex coordinative scenarios. Here we present the motivation for key new processes that must be present, and highlight new data requirements for achieving effective coordination in situations where not only the external environment changes but the organisation itself is organic. Further work is required to develop these ideas to implementable models. Our ambition is to work toward creating believable team agents that can work with humans and other artificial agents in simulations of such complex domains.

References

1. Allsopp, D., Beutement, P., Kirton, M., Tate, A., Bradshaw, J., Suri, N., Burstein, M.: The coalition agents experiment: Network-enabled coalition operations. *Journal of Defence Science, Special Issue on Network-enabled Capabilities* 8(3) (September 2003)
2. Burke, C., Stagl, K., Salas, E., Pierce, L., Kendall, D.: Understanding team adaptation: A conceptual analysis and model. *Journal of Applied Psychology* 91(6), 1189–1207 (2006)
3. Dignum, V., Tick, C.: Agent-based analysis of organizations: Formalization and simulation. In: *Proceedings of IAT 2007: The 2007 IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, pp. 244–247. IEEE, Los Alamitos (2007)
4. Fan, X., Sun, B., Sun, S.: RPD-enabled agents teaming with humans for multi-context decision making. In: *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multi Agent Systems, Hakodate, Japan (May 2006)*
5. Fan, X., Sun, S., Sun, B., Airy, G., McNeese, M., Yen, J.: Collaborative RPD-enabled agents assisting the three-block challenge in command and control in complex and urban terrain. In: *Proceedings of 2005 BRIMS Conference Behavior Representation in Modeling and Simulation, May 2005*, pp. 113–123. Universal City, CA (2005)
6. Fan, X., Yen, J.: Modeling and simulating human teamwork behaviors using intelligent agents. *Journal of Physics of Life Reviews* 1(3), 173–201 (2004)
7. Gonzalez, C.: Decision support for real-time dynamic decision making tasks. *Organizational Behavior and Human Decision Processes* 96, 142–154 (2005)
8. Grosz, B., Hunsberger, L.: The dynamics of intention in collaborative activity. *Cognitive Systems Research* 7(2-3), 259–272 (2006)
9. Grosz, B., Kraus, S.: The evolution of shared plans. In: *Foundations and Theories of Rational Agency*, pp. 227–262 (1999)
10. Hunsberger, L.: Making sharedplans more concise and easier to reason about. In: Rao, A.S., Singh, M.P., Müller, J.P. (eds.) *ATAL 1998. LNCS*, vol. 1555, pp. 81–98. Springer, Heidelberg (1999)
11. Klein, G., Feltovich, P., Bradshaw, J., Woods, D.D.: Common ground and coordination in joint activity. In: Rouse, W., Boff, K. (eds.) *Organizational simulation*. New York, Wiley, Chichester (2005)
12. Lipshitz, R., Omodei, M., McLennan, J., Wearing, A.: What’s burning? the RAWFS heuristic on the fire ground. In: *Expertise Out of Context*, pp. 97–112 (2007)

13. Schurr, N., Pratik, P., Pighin, F., Tambe, M.: Lessons learned from disaster management. In: Proceedings of First International Workshop on Agent Technology for Disaster Management (2006)
14. Siebra, C., Tate, A.: Integrating collaboration and activity-oriented planning for coalition operations support. In: Bredenfled, A., Jacoff, A., Noda, I., Takahashi, Y. (eds.) RoboCup 2005. LNCS, vol. 4020, pp. 561–568. Springer, Heidelberg (2006)
15. Smith, W., Dowell, J.: A case study of co-ordinative decision-making in disaster management. *Ergonomics* 43(8), 1153–1166 (2000)
16. Smith, W., Dowell, J.: Inter-agency coordination in disaster management: A cognitive engineering framework
17. So, R., Sonenberg, L.: Situation awareness as a form of meta-level control. In: Proceedings of the First International Workshop on Metareasoning in Agent-Based Systems at the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems, Honolulu, Hawaii (2007)
18. Tambe, M.: Towards flexible teamwork. *Journal of Artificial Intelligence Research* 7, 83–124 (1997)
19. van der Lee, M.D.E., van Vugt, M.: Imi - an information system for effective multidisciplinary incident management. In: Proceedings of ISCRAM 2004, Brussels (May 2004)
20. Voshell, M., Woods, D., Prue, B., Fern, L.: Coordination Loops: A New Unit of Analysis for Distributed Work. Erlbaum, Mahwah (2007)
21. Yen, J., Fan, X., Sun, S., Hanratty, T., Dumer, J.: Agents with shared mental models for enhancing team decision-makings. *Journal of Decision Support Systems: Special issue on Intelligence and Security Informatics* 41(2) (2005)
22. Zachary, W., Weiland, W., Sclaro, D., Sclaro, J., Santarelli, T.: Instructorless team training using synthetic teammates and instructors. In: Proceedings of the Human Factors & Ergonomics Society 46th Annual Meeting, pp. 2035–2038. Human Factors & Ergonomics Society (2002)