

Safety in the Context of Coordination via Adjustable Autonomy

Paul Scerri*, Katia Sycara* and Milind Tambe⁺

* Carnegie Mellon University

⁺ University of Southern California

Abstract. When large heterogeneous robot and agent teams operate in the real-world, it is essential that a human operator has overall control to ensure safety. However, giving an operator the required control is difficult due to the complexity of the activities such teams engage in and the infeasibility of simply stopping the team whenever human input is required. Our approach to interaction in such a context has three key components which allow us to leverage human expertise by giving them responsibility for key coordination decisions, without risks to the coordination due to slow responses. First, to deal with the dynamic nature of the situation, we use pre-planned sequences of transfer of control actions called *transfer-of-control strategies*. Second, to allow identification of key coordination issues in a distributed way, individual coordination tasks are explicitly represented as coordination roles, rather than being implicitly represented within a monolithic protocol. Such a representation allows *meta-reasoning* about those roles to determine when human input may be useful. Third, the meta-reasoning and transfer-of-control strategies are encapsulated in a mobile agent that moves around the group to either get human input or autonomously make a decision. In this paper, we describe this approach and present initial results from interaction between a large number of UAVs and a small number of humans.

1 Introduction

Recent developments in a variety of technologies are opening up the possibility of deploying large teams of robots or agents to achieve complex goals in complex domains[1, 2]. In domains such as space[3], military[4], disaster response[5] and hospitals[6], hundreds or thousands of intelligent entities might be coordinated to achieve goals more cheaply, efficiently and safely than they can currently be performed. Several interacting, distributed coordination algorithms are required to flexibly, robustly and efficiently allow the team to achieve their goals in complex, dynamic and sometimes hostile environments. The key algorithms may leverage a range of approaches, from logic[7] to decision-theoretic reasoning[8] to constraint satisfaction[9] to achieve their requirements. Since these teams are acting in real-world environments they are capable of causing harm and hence safety is a key requirement of deployment. However, traditional notions of, and approaches to, safety may not always be enough. For example, in some situations, in domains like rescue response, there may be no way of not causing some

harm, but the harm (or risk) must simply be minimized. Moreover, approaches such as stopping when there is risk are typically not applicable because stopping the team may also cause significant harm.

Our hypothesis is that safety in large teams should be achieved by empowering humans to make critical decisions when incorrect decisions can cause physical, financial or psychological harm. These critical decisions include not only the decisions that, for moral or political reasons, humans must be allowed to make but also coordination decisions that humans are better at making because of their particular cognitive skills. The key assumption is that giving responsibility for such decisions to a human, even in systems with the best designed and most reliable software, improves the overall safety. In some cases, human insight can result in better decisions (and hence better behavior) than would be achieved by following, typically sub-optimal, coordination algorithms. For example, allocating agents to tasks, taking into account potential future failures is extremely complex[10], however humans may have experience that allows them to rapidly make reasonable decisions that take into account future failures. In other cases, human experience or simply preference should be imposed on the way the team performs its tasks. Human decision making may be of a higher quality because of access to tool support or information that is not available to agents in the team. It is not necessarily the case that a human would make a better decision with the same resources as the agent. If better coordination results in lower costs, damage or harm (or risk thereof) then developers are obliged to give responsibility for decisions which humans can make better to those humans. However, that human decision-making is a valuable and expensive resource and its use is not without cost. Much work has looked at how and when to transfer control of a critical decision from a single agent to a single person to utilize human expertise in critical situations[11, 12]. While some decisions to be made in the context of coordination are effectively isolated from other activities of the team, and hence the previous work is applicable, there are an additional class of decisions due to coordination, that are not covered by the previous work. In some cases, multiple entities are involved and coordination continues while the decision is being made, introducing the possibility of mis-coordination.

To allow humans to make these critical decisions in the context of coordination is an interesting challenge. First, it is infeasible for humans to monitor ongoing coordination and intervene quickly enough to make the critical decisions unless the team explicitly transfers control of the decision to the human. Since the coordination is distributed, the complete state of the team will typically not be completely known by any individual member of the team. Moreover, in a sufficiently large team or sufficiently dynamic domain, it will be infeasible to continually present an accurate picture of the team to any observer. Thus, the team must proactively transfer control of key decisions to human experts. Second, even when decision-making control is explicitly given, in a dynamic domain, with many possible decisions to make, human decision-making will not always be available or cannot be made quickly enough to allow the team to continue to operate correctly. It should not be the case that delays in making decisions intended

to improve coordination end up causing miscoordination. Thirdly, there may be multiple human experts who can make decisions and the decisions should be sent to the expert in the best position to make them in a timely manner. These three problems have not been adequately addressed by a complete solution in previous work. An effective solution must identify decisions where human input is necessary or useful in a distributed way then transfer control of those decisions to humans capable and available to make those decisions without compromising ongoing coordination with decision-making delays.

In this paper, we present an approach embodying three key ideas. To allow the team to identify critical decisions to be made by humans, we use *coordination meta-reasoning* which uses heuristics to find coordination phenomena that may indicate problems. For example, when there are two high risk alternative courses of action that the team cannot autonomously distinguish, humans may draw on additional experience to choose between. We explicitly represent coordination tasks, such as initiating a team plan or allocating a resource, explicitly via *coordination roles* allowing meta-reasoning to simply identify cases where role performance is poor. Critically, the meta-reasoning is performed “out in the team”, based on local information of individual team members and hence does not rely on an aggregation of coordination information at some central point. However, distributed identification of decisions for potential human input is a double edged sword: on the one hand it removes the need to generate and maintain a centralized state, but on the other it means that identification must be performed with only local knowledge, resulting in less accurate identification of key decisions.

The second part of our approach is that when a decision is to be made by a human, a *transfer-of-control* strategy is used to ensure that lack of a timely response does not negatively impact the performance of the team[13]. A transfer-of-control strategy is a pre-planned sequence of actions that are designed to balance the benefits of getting human input against the costs of that input not coming in a timely manner. Each action in a transfer-of-control strategy either transfers decision-making control to some entity, human or agent, or takes an action to buy more time for the decision to be made. Previously, a mathematical model of transfer-of-control strategies was presented and operationalized via Markov Decision Processes[14]. In that work, although the mathematical model supported the possibility of having multiple humans available to give input, experimental results used only one human expert. In this work, we make multiple human experts available to the agent team and allow the transfer-of-control strategies to reason about transferring control to each. People are modelled by the types of meta-reasoning they can perform and the agents maintain models of what tasks each person is currently performing, in order to create appropriate transfer-of-control strategies.

We are embedding the approach to human interaction with teams into coordination software called Machinetta[14]. Using Machinetta, each team member is given a *proxy* which encapsulates the coordination reasoning and works with other proxies to achieve the coordination. Coordination roles are encapsulated

within simple mobile agents. Thus, each Machinetta proxy acts as essentially an intelligent mobile agent platform and the mobile agents move around the network to achieve the coordination. For example, there will be a mobile agent for finding a team member to perform a particular task and another to pass on a piece of potentially useful information. The third aspect of the approach to human interaction with teams is to create an additional type of mobile agent, called an *adjustable autonomy agent*, that encapsulate pieces of the interaction with humans. An adjustable autonomy agent is created when a proxy identifies some situation that may require human input, then creates and executes a transfer-of-control strategy to get the required input while minimizing costs. The adjustable autonomy agents can also encapsulate the intelligence required to fix the problem, if the transfer-of-control strategy decides an autonomous action is the best way forward.

The approach was evaluated in a simulation of 80 WASMs and two human decision-makers. These preliminary experiments did not involve real humans, but were designed to understand the working of the underlying approach. Many decisions were identified for meta-reasoning and adjustable autonomy agents effectively chose between autonomous and human decision-making to ensure timely decisions. However, some additional issues related were identified, including the need to prioritize decisions that might be presented to the human and the need to tune heuristics so to limit the number of meta-reasoning decisions.

2 Wide Area Search Munitions

Our current domain of interest is coordination of large groups of Wide Area Search Munitions (WASMs). WASMs are a cross between an unmanned aerial vehicle and a standard munition. The WASM has fuel for about 30 minutes of flight, after being launched from an aircraft. The WASM cannot land, hence it will either end up hitting a target or self destructing. The sensors on the WASM are focused on the ground and include video with automatic target recognition, ladar and GPS. It is not currently envisioned that WASMs will have an ability to sense other objects in the air. WASMs will have reliable high bandwidth communication with other WASMs and with manned aircraft in the environment. These communication channels will be required to transmit data, including video streams, to human controllers, as well as for the WASM coordination.

The concept of operations for WASMs are still under development, however, a wide range of potential missions are emerging as interesting[15,16]. A driving example for our work is for teams of WASMs to be launched from AC-130 aircraft supporting special operations forces on the ground. The AC-130 is a large, lumbering aircraft, vulnerable to attack from the ground. While it has an impressive array of sensors, those sensors are focused directly on the small area of ground where the special operations forces are operating making it vulnerable to attack. The WASMs will be launched as the AC-130s enter the battlespace. The WASMs will protect the flight path of the manned aircraft into the area of operations of the special forces, destroying ground based threats as required.

Once an AC-130 enters a circling pattern around the special forces operation, the WASMs will set up a perimeter defense, destroying targets of opportunity both to protect the AC-130 and to support the soldiers on the ground. Even under ideal conditions there will be only one human operator on board each AC-130 responsible for monitoring and controlling the WASMs. Hence, high levels of autonomous operation and coordination are required of the WASMs themselves. However, because the complexity of the battlefield environment and the severe consequences of incorrect decisions, it is expected that human experience and reasoning will be extremely useful in assisting the team in effectively and safely achieving their goals.

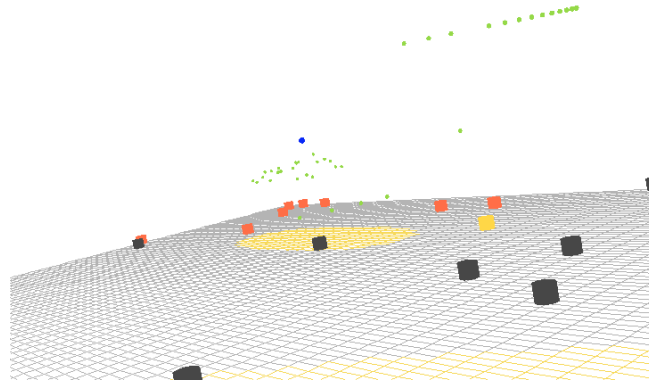


Fig. 1. A screenshot of the WASM coordination simulation environment. A large group of WASMs (small spheres) are flying in protection of a single aircraft (large sphere). Various SAM sites (cylinders) are scattered around the environment. Terrain type is indicated by the color of the ground.

Many other operations are possible for WASMs, if issues related to coordinating large groups can be adequately resolved. Given their relatively low cost compared to Surface-to-Air Missiles (SAMs), WASMs can be used simply as decoys, finding SAMs and drawing fire. WASMs can be used as communication relays for forward operations, forming an adhoc network to provide robust, high bandwidth communications for ground forces in a battle zone. Since a WASM is “expendible”, it can be used for reconnaissance in dangerous areas, providing real-time video for forward operating forces. While our domain of interest is teams of WASMs, the issues that need to be addressed have close analogies in a variety of other domains. For example, coordinating resources for disaster response involves many of the same issues[5], as does intelligent manufacturing[17] and business processes.

3 Coordination Meta-Reasoning

In a large scale team, it will typically be infeasible for a human (or humans) to monitor ongoing coordination and pre-emptively take actions that improve the overall coordination. This is especially the case because the complete state of the coordination will not be known at any central point that can be monitored. Each team member will have some knowledge of the overall system and human users can certainly be given more information, but the communication required to have the complete coordination state continuously known by any team member are unreasonable. In fact, even if complete state were made available to a human, we believe it would be too complex and too dynamic for the human to reasonably make sense of. Hence, the team must, in a distributed way, identify situations where human input may improve coordination and explicitly transfer responsibility for making that decision to a human.

Due to the computational complexity of optimal decision-making, coordination of large teams is typically governed by a set of heuristics. The heuristics may be developed in very principled ways and be provably near optimal in a very high percentage of situations that will be encountered. However, by their nature, heuristics will sometimes perform poorly. If these situations can be detected and referred to a human then overall performance can be improved. However, this is somewhat paradoxical, since if the situations can be reliably and accurately identified then, often, additional heuristics can be found to perform well in that situation. Moreover, performance based techniques for identifying coordination problems, i.e., reasoning that there must be a coordination problem when the team is not achieving their goal, are inadequate, because in some cases, optimal coordination will not allow a team to achieve its goal. Hence meta-reasoning based on domain performance will say more about the domain than it will about the coordination.

Typical approaches to multiagent coordination do not explicitly represent the individual tasks required to coordinate a group. Instead, the tasks are implicitly captured via some protocol that the group executes to achieve coordination. This implicit representation makes meta-reasoning difficult because specific issues can be difficult to isolate (and even more difficult to rectify.) By explicitly representing individual coordination activities as *roles*, it is more straightforward to reason about the performance of those tasks[14]. Individual coordination tasks, such as allocating a resource or initiating a team plan, represented as explicit roles, can be monitored, assessed and changed automatically, since they are decoupled from other coordination tasks.

In practice, autonomously identifying coordination problems that might be brought to the attention of a human expert is imprecise. Rather than reliably finding poor coordination, the meta-reasoning must find *potentially* poor coordination and let the humans determine the actually poor coordination. (In the next section, we describe the techniques that are used to ensure that this does not lead to the humans being overloaded.) Notice that while we allow humans to attempt to rectify problems with agent coordination, it is currently an open question whether humans can actually make better coordination decisions than the

agents. To date, we have identified three phenomena that may be symptomatic of poor coordination and bring these to the attention of the human:

- *Unfilled task allocations.* In previous work, meta-reasoning identified unallocated tasks as a symptom of potentially poor coordination[14]. When there are more tasks than team members able to perform tasks, some will necessarily be unallocated. However, due to the sub-optimality of task allocation algorithms¹, better overall performance might be achieved if a different task was left unallocated.

When a role cannot be allocated, three things can be done. First, the role allocation process can be allowed to continue, using communication resources but getting the role allocated as soon as possible. Second, the role allocation process for the role can be suspended for some time, allowing the situation to change, e.g., other roles completed. This option uses less communication resources but potentially delays execution of the role. Thirdly, the role and its associated plan can be cancelled. Choosing between these three options is an ideal decision for a human since it requires some estimate of how the situation will change in the future, something for which human experience is far superior to an agents. If a human is not available to make a decision, the agent will autonomously decide on suspending the allocation of the role for some time, letting the role allocation continue.

- *Untasked team members.* When there are more team members than tasks, some team members will be untasked. Untasked physical team members might be moved or reconfigured to be best positioned for likely failures or future tasks. Potentially, the team can make use of these team members to achieve goals in different ways, e.g., with different plans using more team members, or preempt future problems by assigning untasked team members to preventative tasks. Thus, untasked team members may be symptomatic of the team not effectively positioning resources to achieve current and future objectives.

There are currently two things that can be done when a team member does not have a task for an extended period: do nothing or move the agent to some other physical location. Doing nothing minimizes use of the agent's resources, while moving it around the environment can get it in better position for future tasks. Again, this decision is ideally suited for human decision-making because it requires estimates of future activities for which they can draw upon their experience. If a human is not available to make a decision, the agent will autonomously decide to do nothing.

- *Unusual Plan Performance Characteristics.* Team plans and sub-plans, executed by team members to achieve goals and sub-goals will typically have logical conditions indicating that the plan has become unachievable or irrelevant[]. However, in some cases, due to factors unknown to or not understood

¹ In simple domains, it may be possible to use optimal, typically centralized, task allocation algorithms, but decentralized task allocation algorithms involving non-trivial tasks and large numbers of agents will be sub-optimal.

by the team, plans may be unachievable or irrelevant without the logical conditions for their termination becoming true. In such cases, meta-reasoning about the plan's performance can bring it to the attention of a human for assessment. Specifically, we currently allow a plan designer to specify an expected length of time that a plan will usually take and bring to the attention of the human plans that exceed this expected time. We envision allowing specification of other conditions in the future, for example, limits on expected resource use or number of failure recoveries.

When a plan execution does not meet normal performance metrics, there are two things that can be done: cancel the plan or allow it to continue. Cancelling the plan conserves resources that are being used on the plan, but robs the team of any value for successfully completing the plan (if it were possible.) If the agents knew the plan was unachievable they would autonomously cancel it, so this meta-reasoning will only be invoked when there are circumstances outside the agents sensing or comprehension that are causing plan failure. Thus, in this case the humans are uniquely placed to cancel the plan if it is indeed unachievable.² Since the agent will have no evidence that the plan is unachievable, if required to act autonomously, it will allow the plan to continue.

Notice that when meta-reasoning finds some coordination phenomena that may indicate sub-standard coordination the team does not stop execution but allows the human to take corrective actions while its activities continue. This philosophy enables us to be liberal in detecting potential problems thus ensuring that most genuine problems are subsumed by the identified problems. However, from the perspective of the humans, this can lead to being inundated with spurious reports of potential problems. In the next section we present an approach to dealing with this potential overload.

4 Transfer-of-Control Strategies

Coordination meta-reasoning decisions must be made in a timely manner or the value of the decision is lessened (or lost all together). In fact, we observe that in some cases, timely decisions of a lower quality can have more positive impact on the team than untimely high quality decisions. For example, if resources are being wasted on an unachievable plan, time taken to cancel the plan incurs a cost, but to cancel a plan that is achievable results in the team losing the value of that plan. In this case, *quality* refers loosely to the likelihood an entity will make an optimal or near optimal decision. To leverage the advantage of rapid decision-making we have developed simple, autonomous meta-reasoning. Decision-theoretic reasoning is then used to decide whether to use slower, presumably higher quality, human reasoning to make a meta-reasoning decision or whether to use the simple, though fast agent reasoning. Additionally, the system

² However, humans will not infallibly detect plan unachievability, either.

has the option to use a *delaying* action to reduce the costs of waiting for a human decision, if there is value in getting human input. A pre-planned sequence of actions either transferring control of a meta-reasoning decision to some entity or taking an action to buy time is called a *transfer-of-control strategy*. Transfer-of-control strategies were first introduced and mathematically modeled in [13]. An optimal transfer-of-control strategy optimally balances the risks of not getting a high quality decision against the risk of costs incurred due to a delay in getting that decision. Thus, the computation to find an optimal transfer-of-control strategy takes as input a model of the expected “quality” of each entity’s decision-making ability and a model of the expected costs incurred per unit time until the decision is made. Additionally, the impact of any delaying action needs to be modeled. In previous work, these models were used to define a Markov Decision Process that found the optimal strategy.

4.1 Transfer of Control Strategies

In this section, we briefly review mathematical model of transfer-of-control strategies presented in [13]. A decision, d , needs to be made. There are n entities, $e_1 \dots e_n$, who can potentially make the decision. These entities can be human users or other agents. The expected quality of decisions made by each of the entities, $\mathbf{EQ} = \{EQ_{e_i}^d(t) : \mathcal{R} \rightarrow \mathcal{R}\}_{i=1}^n$, is known, though perhaps not exactly. $\mathbf{P} = \{P_\tau(t) : \mathcal{R} \rightarrow \mathcal{R}\}$ represent continuous probability distributions over the time that the entity in control will respond with a decision of quality $EQ_e^d(t)$.

We model the cost of delaying a decision until time t as $\{\mathcal{W} : t \rightarrow R\}$. The set of possible wait-cost functions is \mathbf{W} . We assume $\mathcal{W}(t)$ is non-decreasing and that there is some point in time, Γ , when the costs of waiting stop accumulating (i.e., $\forall t \geq \Gamma, \forall \mathcal{W} \in \mathbf{W}, \mathcal{W}(t) = \mathcal{W}(\Gamma)$).

Finally, there is an additional action, with cost \mathcal{D}_{cost} , with the result of reducing the rate at which wait costs accumulate. We call such an action a *deadline delaying action* and denote it \mathcal{D} . For example, a \mathcal{D} action might be as simple as informing the party waiting for the decision that there has been a delay, or more complex, such as reordering tasks. We model the value of the \mathcal{D} by letting \mathcal{W} be a function of $t - \mathcal{D}_{value}$ (rather than t) after the \mathcal{D} action.

We define the set \mathbf{S} to be all possible transfer-of-control strategies. The problem can then be defined as:

Definition 1. For a decision d , select $s \in \mathbf{S}$ such that $\forall s' \in \mathbf{S}, s' \neq s, EU_s^d t \geq EU_{s'}^d t$

We define a simple shorthand for referring to particular transfer-of-control strategies by simply writing the order that entities receive control or, alternatively, \mathcal{D} s are executed. For example,

$a_{fred}a_{barney}\mathcal{D}a_{barney}$ is shorthand for a strategy where control is given to the agent *fred*, then given to the agent *barney*, then does a *Dis* performedq, and finally given indefinitely to *barney*. Notice that the shorthand does not record the timing of the transfers of control. In the following discussion we assume that there is some agent A that can always make the decision instantly.

To calculate the EU of an arbitrary strategy, we multiply the probability of response at each instant of time with the expected utility of receiving a response at that instant, and then sum the products. Hence, for an arbitrary continuous probability distribution:

$$EU = \int_0^{\infty} P_{\top}(t) EU_{e_c}^d(t) .dt \quad (1)$$

where e_c represents the entity currently in decision-making control.

Since we are primarily interested in the effects of delayed response, we can decompose the expected utility of a decision at a certain instant, $EU_{e_c}^d(t)$, into two terms. The first term captures the quality of the decision, independent of delay costs, and the second captures the costs of delay, i.e.,: $EU_{e_c}^d t = EQ_e^d(t) - \mathcal{W}(t)$. A \mathcal{D} action affects the future cost of waiting. For example, the wait cost after performing a \mathcal{D} at $t = \Delta$ at cost \mathcal{D}_{cost} is : $\mathcal{W}(t|\mathcal{D}) = \mathcal{W}(\Delta) - \mathcal{W}(\Delta - \mathcal{D}_{value}) + \mathcal{W}(t' - \mathcal{D}_{value}) + \mathcal{D}_{cost}$.

To calculate the EU of a strategy, we need to ensure that the probability of response function and the wait-cost calculation reflect the control situation at that point in the strategy. For example, if the user has control at time t , $P_{\top}(t)$ should reflect the user's probability of responding at t . To do this simply, we can break the integral from Equation 1 into separate terms, with each term representing one segment of the strategy, e.g., for a strategy eA there would be one term for when e has control and another for when A has control.

Using this basic technique, we can now write down the equations for some general transfer-of-control strategies. Equations 2-6 are the general EU equations for the AA strategies A , e , eA and $eDeA$ respectively. We create the equations by writing down the integral for each of the segments of the strategy, as described above. T is the time when the agent takes control from the user, and Δ is the time at which the \mathcal{D} occurs. One can write down the equations for more complex strategies in the same way.

In the case of the large team of WASMs, we currently use simple models of $EQ_e^d(t)$, $\mathcal{W}(t)$ and $P_{\top}(t)$. We assume that each human expert is equally capable of making any meta-reasoning decision. In each case, wait costs accrue linearly, i.e., the wait costs accrued in the second minute are the same as those accrued in the first minute. We assume a Markovian response probability from the human, though this is a model that future work could dramatically improve. For two of the three meta-reasoning decisions, we model the quality of the agent's autonomous reasoning as improving over time. This is not because the reasoning changes, but because the default response is more likely correct the more time that passes. Specifically, over time it makes more sense to let a role allocation continue, since there has more likely been changes in the team that allow it to complete. Likewise, terminating a long running plan is more reasonable as time passes, since it becomes more likely that something is actually preventing completion of the plan. However, notice that the rate at which the quality of the autonomous reasoning for the long running plan increases is much slower than for the role allocation. $EQ_{human}^d(t)$ is highest for the long running plan, since

$$EU_A^d = EQ_A^d(0) - \mathcal{W}(0) \quad (2)$$

$$EU_e^d = \int_0^T P_\tau(t) \times (EQ_e^d(t) - \mathcal{W}(t)).dt + \int_T^\infty P_\tau(t) \times (EQ_e^d(t) - \mathcal{W}(D)).dt \quad (3)$$

$$EU_{eA}^d = \int_0^T P_\tau(t) \times (EQ_e^d(t) - \mathcal{W}(t)).dt + \int_T^\infty P_\tau(t).dt \times (EQ_a^d(T) - \mathcal{W}(T)) \quad (4)$$

$$EU_{U\mathcal{D}eA}^d = \int_0^\Delta P_\tau(t)(EQ_e^d(t) - \mathcal{W}(t)).dt + \int_\Delta^T P_\tau(t)(EQ_e^d(t) - \mathcal{W}(\Delta) + \mathcal{W}(\Delta - \mathcal{D}_{value}) - \mathcal{W}(t - \mathcal{D}_{value}) - \mathcal{D}_{cost}).dt + \int_T^\infty P_\tau(t)(EQ_A^d(t) - \mathcal{W}(\Delta) + \mathcal{W}(\Delta - \mathcal{D}_{value}) - \mathcal{W}(T - \mathcal{D}_{value}) - \mathcal{D}_{cost}).dt \quad (5)$$

Table 1. General AA EU equations for simple transfer of control strategies.

it is a relatively decoupled decision that requires the expertise of the human, whereas the other decisions are more tightly coupled with other coordination and involve more coordination expertise which the team can have. Over time, the autonomous reasoning to let an untasked agent stay where it is does not change, hence the autonomous model of the quality is a constant. Currently, there are no deadline delaying actions available to the team.

5 Agentifying Adjustable Autonomy

The coordination algorithms are encapsulated in a domain independent *proxy*[18–20, 14]. Such proxies are the emerging standard for implementing teamwork between heterogeneous team members. There is one proxy for each WASM and one for each human expert. The basic architecture is shown in Figure 2. The proxy communicates via a high level, domain specific protocol with either the WASM or human to collect information and convey requests from the team. The proxies communicate between themselves to facilitate the coordination. Most of the proxy code is domain independent and can be readily used in new domains requiring distributed control. The current version of the code, known as *Macchinetta*, is a substantially extended and updated version of the *TEAMCORE*

proxy code[19]. TEAMCORE proxies implement *teamwork* as described by the STEAM algorithms[8], which are in turn based on the theory of *joint intentions*[21, 22].

In a dynamic, distributed system, protocols for performing coordination need to be extremely robust, since the large numbers of agents ensures that anything that can go wrong typically does. When the size of a team is scaled to hundreds of agents, this does not simply imply the need to write “bug-free” code. Instead abstractions and designs that promote robustness are needed. Towards this end, we are encapsulating “chunks” of coordination in *coordination agents*. Each coordination agent manages one specific piece of the overall coordination. When control over that piece of coordination moves from one proxy to another proxy, the coordination agent moves from proxy to proxy, taking with it any relevant state information. There are coordination agents for each plan or subplan (PlanAgents), each role (RoleAgents) and each piece of information that needs to be shared (InformationAgents). For example, a RoleAgent looks after everything to do with a specific role. This encapsulation makes it far easier to build robust coordination. Since, the coordination agents actually implement the coordination, the proxy can be viewed simply as a mobile agent platform that facilitates the functioning of the coordination agents. However, the proxies play the additional important role of providing and storing local information.

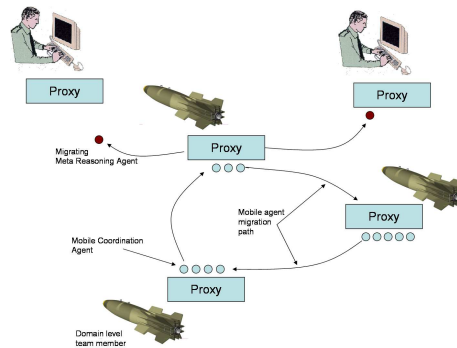


Fig. 2. Overall system architecture.

To integrate the meta-reasoning into a Machinetta controlled team, an additional type of mobile agent is introduced, called an *adjustable autonomy mobile agent*. Each adjustable autonomy mobile agent has responsibility for a single piece of interaction with the humans. In the case of coordination meta-reasoning, an adjustable autonomy mobile agent is created when a proxy detects that a situation has occurred that requires meta-reasoning. To get the decision made, the adjustable autonomy mobile agent creates and executes a transfer-of-control

strategy. The adjustable autonomy mobile agent exists until a decision has been made, either autonomously or by some human.

Information is collected by the proxy as mobile agents of various types pass through to implement the coordination. Critically, information agents move information about the current workload of each human expert to the proxy of each other human expert. These models currently contain a simple numeric value for the human’s current workload, though more detailed models can be envisioned. When an adjustable autonomy agent is created it can first move to the proxy of any human expert to get the models it requires to create an optimal transfer-of-control strategy. Note that these models will not always be complete, but provide estimates of workload the agent can use to construct a strategy. In some cases, if the decision can easily be made autonomously, the adjustable autonomy agent will not move to the proxy of the human expert, because their specific availability is not important.

Importantly, as described above, simple reasoning for making the meta-reasoning decisions are encapsulated by the adjustable autonomy mobile agent so that a decision can be made autonomously if human expertise is not readily available or the decision is not sufficiently important. For example, the reasoning on unallocated roles is to let the role “continue” in an attempt to allocate it.

Figure 3 shows the interface for presenting decisions to human experts. Notice, that while the specific information about the information is presented, contextual information is not. This is a serious shortcoming of this work to date and a key area for for future work. Presenting enough context to allow the human to make effective decisions will involve both getting the information to the expert and, more challengingly, presenting it in an appropriate manner.

6 Preliminary Experiments

Objective evaluation of the approach described above is difficult for two reasons. First, improvements in the performance of the team will rely heavily on the interfaces that present the decisions and related information to the human. Such interfaces are not a focus of this work. More pragmatically, a simulation run with a reasonable number of WASMs takes around 30 minutes to run, hence to run a meaningful set of experiments would take an large amount of human time. Moreover, domain experts required to make such decisions are hard to come by. While such experiments will eventually have to be run to validate the effectiveness of this approach, our initial experiments do not rely on human subjects. Instead, the object of these experiments is to evaluate how the underlying algorithms work in finding potential team problems and deal with the possibility that a human is not available to make these decisions when they arise. These preliminary experiments form a basis for future experiments, involving actual human subjects.

To remove the need for many hours of human input, the interfaces were augmented with code that made decisions as if they were made by the human. These “human” decisions were made between five seconds and two minutes after



Fig. 3. A view of the interface the human uses to make decisions. Three decisions are shown. At the top is a decision whether to cancel a long running plan. Below are two unallocated roles from the same plan, for destroying a particular target.

control was transferred to the human. The experiments involved a team of 80 WASMs operating in a large environment. The primary task of the team was to protect a manned aircraft by finding and destroying surface-to-air missile sites spread around the environment. Half the team spread out across the environment searching for targets while the other half stayed near the manned aircraft destroying surface-to-air sites as they were found near the aircraft. Plans were simple, requiring a single WASM to hit each found target. If a target was not hit within three minutes of being found, this was considered abnormal plan execution and meta-reasoning would be invoked. Meta-reasoning was also invoked when a WASM was not allocated to hit any target for five minutes. These times are low, but reasonable since the simulation ran at approximately four times real-time. Finally, meta-reasoning was invoked when no WASM was available to hit a found target. Two human commanders were available to make meta-reasoning decisions (although, as discussed above there were not “real” human commanders).

Six different scenarios were used, each differing the number of surface-to-air missile sites. Each configuration was run ten times, thus the results below represent around 30 hours of simulation time (120 hours of real-time). As the number of missile sites increases, the team will have more to do with the same number of WASMs, thus we can expect more meta-reasoning decisions.

Figure 4 shows that the total number of meta-reasoning decisions does increase with the number of targets. Over the course of a simulation, there are around 100 meta-reasoning decisions or about one per agent. However, as Figure 5 shows, only about 20% of these get transferred to a human. The large number of decisions that are made autonomously is primarily because humans




Fig. 4. The number of meta-reasoning decisions to be made as the number of targets in the environment increases.

are not available to make those decisions. This suggests work may need to be done to prioritize decisions for a user, to prevent high priority decisions being left to an agent, while the user is busy with low priority decisions. However, an appropriate solution is not obvious, since new decisions arrive asynchronously and it will likely not be appropriate to continually change the list of decisions the human is working on. Finally, notice in Figure 6 that a large percentage of the meta-decisions are to potentially cancel long running plans. The large number of such decisions illustrates a need to carefully tune the meta-reasoning heuristics in order to avoid overloading the system with superfluous decisions. However, in this specific case, the problem of deciding whether to cancel a long running plan was the most appropriate for the human, hence the large percentage of such decisions for the human is reasonable.

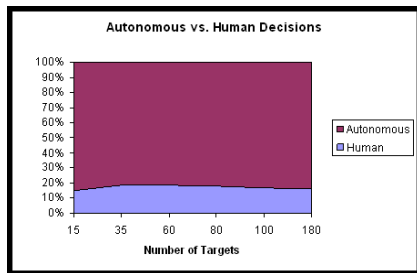


Fig. 5. The percentage of decisions transferred to humans versus the percentage made autonomously.

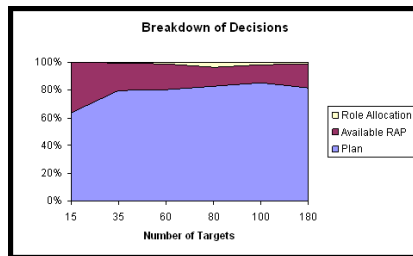


Fig. 6. Ratios of different types of meta-reasoning decisions presented to the user.

7 Conclusions and Future Directions

This chapter presented an integrated approach to leveraging human expertise to improve the safety of a large team. Via the use of coordination meta-reasoning, key decisions could be brought to the attention of human experts allowing the

team to leverage their expertise. Transfer-of-control strategies ensured that miscoordination was not caused by delays waiting for human input. The approach was encapsulated in adjustable autonomy agents that are part of the Machinetta proxy approach to coordination. Initial experiments showed that the approach was able to balance the need for human input against the potential for overloading them. However, further experiments are required to determine whether humans actually make a system safer. In the future it is likely that special functionality will be required in the operator interface to ensure mistakes are not made that actually cause harm when the system would not have. However, imposing such constraints on the operator's interface will lead to new challenges.

References

1. Horling, B., Mailler, R., Sims, M., and Lesser, V., "Using and Maintaining Organization in a Large-Scale Distributed Sensor Network," *In Proceedings of the Workshop on Autonomy, Delegation, and Control (AAMAS03)*, 2003.
2. *Proceedings of AAMAS'04 Workshop on Challenges in the Coordination of Large Scale MultiAgent Systems*, 2004.
3. Goldberg, D., Cicirello, V., Dias, M. B., Simmons, R., Smith, S., and Stentz, A. T., "Market-Based Multi-Robot Planning in a Distributed Layered Architecture," *Multi-Robot Systems: From Swarms to Intelligent Automata: Proceedings from the 2003 International Workshop on Multi-Robot Systems*, Vol. 2, Kluwer Academic Publishers, 2003, pp. 27–38.
4. Vick, A., Moore, R. M., Pirnie, B. R., and Stillion, J., *Aerospace Operations Against Elusive Ground Targets*, RAND Documents, 2001.
5. Kitano, H., Tadokoro, S., Noda, I., Matsubara, H., Takahashi, T., Shinjoh, A., and Shimada, S., "RoboCup Rescue: Search and Rescue in Large-Scale Disasters as a Domain for Autonomous Agents Research," *Proc. 1999 IEEE Intl. Conf. on Systems, Man and Cybernetics*, Vol. VI, Tokyo, October 1999, pp. 739–743.
6. on Visionary Manufacturing Challenges, C., "Visionary Manufacturing Challenges for 2020," National Research Council.
7. Pollack, J. L., "The logical foundations of goal-regression planning in autonomous agents," *Artificial Intelligence*, Vol. 106, 1998, pp. 267–334.
8. Tambe, M., "Agent Architectures for Flexible, Practical Teamwork," *National Conference on AI (AAAI97)*, 1997, pp. 22–28.
9. Modi, P. J., Shen, W.-M., Tambe, M., and Yokoo, M., "An Asynchronous Complete Method for Distributed Constraint Optimization," *Proceedings of Autonomous Agents and Multi-Agent Systems*, 2003.
10. Nair, R., Tambe, M., and Marsella, S., "Role allocation and reallocation in multiagent teams: Towards a practical analysis," *Proceedings of the second International Joint conference on agents and multiagent systems (AAMAS)*, 2003.
11. Ferguson, G., Allen, J., and Miller, B., "TRAINS-95 : Towards a mixed-initiative planning assistant," *Proceedings of the Third Conference on Artificial Intelligence Planning Systems*, May 1996, pp. 70–77.
12. Veloso, M., Mulvehill, A., and Cox, M., "Rationale-supported mixed-initiative case-based planning," *Proceedings of the fourteenth national conference on artificial intelligence and ninth innovative applications of artificial intelligence conference*, 1997, pp. 1072–1077.

13. Scerri, P., Pynadath, D., and Tambe, M., "Towards Adjustable Autonomy for the Real World," *Journal of Artificial Intelligence Research*, Vol. 17, 2002, pp. 171–228".
14. Scerri, P., Pynadath, D. V., Johnson, L., Rosenbloom, P., Schurr, N., Si, M., and Tambe, M., "A Prototype Infrastructure for Distributed Robot-Agent-Person Teams," *The Second International Joint Conference on Autonomous Agents and Multiagent Systems*, 2003.
15. Clark, R., *Uninhabited Combat Air Vehicles: Airpower by the people, for the people but not with the people*, Air University Press, 2000.
16. Defense Science Board, "Defense Science Board Study on Unmanned Aerial Vehicles and Uninhabited Combat Aerial Vehicles," Tech. rep., Office of the Under Secretary of Defense for Acquisition, Technology and Logistics, 2004.
17. Ranky, P., *An Introduction to Flexible Automation, Manufacturing and Assembly Cells and Systems in CIM (Computer Integrated Manufacturing), Methods, Tools and Case Studies*, CIMware, 1997.
18. Jennings, N., "The ARCHON Systems and its Applications," Project Report, 1995.
19. Tambe, M., Shen, W.-M., Mataric, M., Pynadath, D., Goldberg, D., Modi, P. J., Qiu, Z., and Salemi, B., "Teamwork in Cyberspace: using TEAMCORE to make agents team-ready," *AAAI Spring Symposium on agents in cyberspace*, 1999.
20. Pynadath, D. and Tambe, M., "Multiagent Teamwork: Analyzing the Optimality and Complexity of Key Theories and Models," *First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'02)*, 2002.
21. Jennings, N. R., "Specification and Implementation of a Belief-Desire-Joint-Intention Architecture for Collaborative Problem Solving," *Intl. Journal of Intelligent and Cooperative Information Systems*, Vol. 2, No. 3, 1993, pp. 289–318.
22. Cohen, P. R. and Levesque, H. J., "Teamwork," *Nous*, Vol. 25, No. 4, 1991, pp. 487–512.
23. Burstein, M. H. and Diller, D. E., "A Framework for Dynamic Information Flow in Mixed-Initiative Human/Agent Organizations," *Applied Intelligence on Agents and Process Management*, 2004, Forthcoming.
24. Hartrum, T. and DeLoach, S., "Design Issues for Mixed-Initiative Agent Systems," *Proceedings of AAAI workshop on mixed-initiative intelligence*, 1999.
25. T. Fong, C. T. and Baur, C., "Advanced Interfaces for Vehicle Teleoperation: collaborative control, sensor fusion displays, and Web-based Tools," *Vehicle Teleoperation Interfaces Workshop, IEEE International Conference on Robotics and Automation*, San Francisco, CA, April 2000.
26. Malin, J., Thronesbery, C., and Schreckenghost, D., "Progress in human-centered automation: Communicating situation information," 1996".
27. Kortenkamp, D., Schreckenghost, D., and Martin, C., "User Interaction with Multi-Robot Systems," *Proceedings of Workshop on Multi-Robot Systems*, 2002.
28. Goodrich, M., Olsen, D., Crandall, J., and Palmer, T., "Experiments in Adjustable Autonomy," *Proceedings of IJCAI Workshop on Autonomy, Delegation and Control: Interacting with Intelligent Agents*, edited by H. Hexmoor, C. Castelfranchi, R. Falcone, and M. Cox, 2001.