

Mutual State-Based Capabilities for Role Assignment in Heterogeneous Teams

Somchaya Liemhetcharat
The Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213, USA
som@ri.cmu.edu

Manuela Veloso
Computer Science Department
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213, USA
veloso@cs.cmu.edu

ABSTRACT

In a heterogeneous team, agents have different capabilities with regards to the actions relevant to the task. Roles are typically assigned to individual agents in such a team, where each role is responsible for a certain aspect of the joint team goal. In this paper, we focus on role assignment in a heterogeneous team, where an agent's capability depends on its teammate and their mutual state, i.e., the agent's state and its teammate's state. The capabilities of an agent are represented by a mean and variance, to capture the uncertainty in the agent's actions as well as the uncertainty in the world. We present a formal framework for representing this problem, and illustrate our framework using a robot soccer example. We formally describe how to compute the value of a role assignment policy, as well as the computation of the optimal role assignment policy, using a notion of risk. Further, we show that finding the optimal role assignment can be difficult, and describe approximation algorithms that can be used to solve this problem. We provide an analysis of these algorithms in our model and empirically show that they perform well in general problems of this domain, compared to market-based techniques.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent systems

General Terms

Algorithms

Keywords

Capability, Role Assignment, Heterogeneous Teams, Multi-Agent

1. INTRODUCTION

Role assignment in heterogeneous teams, i.e., teams with members of different capabilities, has been an extensive topic

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PCAR '10, May 10, 2010, Toronto, Canada

of research. Different capabilities among the team make certain team members better suited for specific roles, and role assignment algorithms plan based on these capabilities. Common approaches to this problem, which we further discuss in the related work section, include market-based techniques, where agents bid over the roles, and the capabilities and states of the agents are used in generating the bids. Most current approaches that explicitly model capabilities assume that the agents have fixed capabilities for tasks, which may not be an accurate reflection in dynamic environments, since the state of the world affects an agent's ability to perform an action.

We are interested in explicitly modelling the capabilities of the agents in dynamic environments, where state is an important factor that agents should be aware of. We address the case where an agent's ability in performing an action depends on the teammate that is affected by the action, as well as their mutual state, i.e., the agent's state and its teammate's state. Modelling capabilities in this way captures information about both the innate abilities of the agent in performing the action, as well as how effective a particular pairing of agents would be when the action is taken in their mutual states. For example, an agent may perform an action well with a particular teammate in some mutual states, but not in other mutual states, or with any other agent. Alternatively, any two agents may work well together at one particular mutual state. We contribute a model that captures these situations.

We are also interested in a pick-up team scenario, where agents who have had no prior contact are put together to form a team, and must learn to coordinate well. The agents do not have prior knowledge of the capabilities of their teammates, nor how well they work together. Thus, we assume that information about capabilities is incrementally obtained through interaction and observation, which implies that agents are aware of their actions and learn about the actions' effects. In this paper, we formally model the capabilities of the agents, and do not discuss the techniques used to learn the capabilities from observations.

Since the capabilities of agents are incrementally obtained, we define the capability of an agent performing an action (with a particular teammate in a mutual state) as the mean and variance of the utility achieved by the action. Representing a capability with a mean and variance captures both

the uncertainty in the agent’s action and the uncertainty in the world. In addition, we introduce the concept of risk in role assignments, by evaluating the mean and variance of role assignment policies. Having a parameter for risk allows a tradeoff between the mean and variance in the optimal role assignment policy, so conservative role assignment policies are chosen over riskier ones, i.e., policies with higher mean utilities but correspondingly higher variances, if so desired. For example, by varying the amount of risk over time, a pick-up team of robots can quickly learn to work together and discover the best role-assignment of the entire team.

We contribute the Mutual State Capability-Based Role Assignment (MuSCRA) model, which includes the states of the world, the actions that can be taken, the roles in the team, and the capabilities of the agents. Roles have association strengths with states of the world, i.e., the states that are important to the role, and different roles have different emphasis on actions, i.e., which actions are important to the role. We introduce the concept of an optimal role assignment in our model, and discuss approximation algorithms that can be used to find effective role assignments. We compare the performance of the approximation algorithms with a market-based approach.

There are several real scenarios in dynamic environments where MuSCRA can be applied. In robot soccer, robots may have different capabilities in kicking the ball accurately and passing the ball to one another. Role assignment in this case would involve choosing the “best” players for the team (if there were more robots than actual players allowed on the field), and giving them the right assignments such as attacker or defender. The risk parameter would determine whether an aggressive role assignment is preferred over a conservative one. Similarly, in urban search and rescue, different robots have different capabilities, for example the ability to cross rough terrain, the speed of movement, and the ability to detect trapped humans. Furthermore, such abilities could depend on the composition and state of a team, e.g., a robot’s ability to detect a human accurately could depend on the sensor readings of another robot. Roles could be diggers that clear the path for other robots, and searchers that travel along the cleared path to search for humans.

The rest of the paper is structured as follows: in Section 2, we discuss related work. In Section 3, we formalize the MuSCRA model. We use robot soccer as a detailed example for our model, and instantiate the model for other application domains. In Section 4, we describe how to solve the MuSCRA model for an optimal role assignment, and discuss approximation algorithms that can be used to find effective role assignments. We discuss our experiments and results in Section 5, and we conclude with a discussion of MuSCRA and its contributions in Section 6.

2. RELATED WORK

Task allocation in multi-agent systems is very similar to role assignment, where the tasks in the domain can be likened to the actions in our approach, and roles represent responsibilities over certain groups of tasks. Gerkey and Mataric provide a detailed taxonomy of task allocation in multi-robot systems [5]. The authors also use the concept of utility, which they posit “is carried out somewhere in every au-

tonomous task allocation system”. Our approach uses utility, but instead of a single value, we use a mean and variance to represent the utility’s distribution instead.

Market-based approaches are frequently used in task allocation domains, and Dias et al. provide a comprehensive survey [2]. In addition, market-based approaches allow agents to perform role assignment [3]. While market-based approaches typically can determine a role assignment in $O(n)$ time or better (where n is the number of agents), the complexity of the domain is delegated to the method of generating the bid. Instead of encapsulating this complexity and embedding the capabilities, our approach explicitly models capabilities, and allows the utility of actions to be learnt over time through observation.

Besides market-based approaches, formal models such as Multi-Agent Markov Decision Processes have recently been used to solve the multi-agent role assignment problem [11]. However, they do not explicitly model the capabilities of individual agents. He and Ioerger provide a model of capabilities in multi-agent systems [8], but they do not model the capabilities of agents as a function of the team composition. In our approach, agents are able to update their capabilities as they acquire new observations. These capabilities are dependent on the configuration of the world, as well as the agents occupying the roles of the team, and their mutual states.

Guttmann models uncertainty in the capabilities of teams in performing tasks, using means and standard deviations [7]. We model capabilities probabilistically as a function of the agent, a teammate, and their mutual state. Kok and Vlassis use states to model teammates to coordinate actions [9]; Gmytrasiewicz and Doshi model agents to select optimal actions [6]; Garland and Alterman study conventions to coordinate actions [4]. We model capabilities to select a role assignment to form a team with a high utility, calculating the performance of the role assignment by using the statistics of the capabilities.

Role assignment can be used in a variety of domains, such as robot soccer [13], formation control [1], and assembly [12]. McMillen and Veloso use the concept of Skills, Tactics, and Plays, where a play determines the available roles for the robots on the team [10], while Vail and Veloso use potential fields to determine the role assignment of the robots [14]. However, in both cases, the robots are assumed to be homogeneous and equally capable of performing all tasks. Our approach instead focuses on the case when robots are heterogeneous and have differing capabilities.

3. MODELLING AGENT CAPABILITIES

A heterogeneous team of agents consists of agents with different capabilities, and the goal is to find an assignment of roles for the agents such that the best team configuration is achieved, in terms of the utility attained. The capability of an agent to successfully perform an action with a teammate depends on their mutual state, i.e., the agent’s state as well as the teammate’s. These capabilities are assumed to be obtained incrementally from observation, and are represented by the mean and variance of the utility of performing the action. Role assignment of a team of robots incorporates a




Figure 1: Robot soccer scenario with 2 teams (α and β), showing the defensive and offensive regions, and goals.

risk factor, which represents the tradeoff between the mean and variance of the utility of a role assignment policy.

3.1 The MuSCRA Model

DEFINITION 1. A *Mutual State Capability-Based Role Assignment (MuSCRA)* is a tuple $\{\mathcal{X}, \mathcal{A}, \mathbf{a}, \mathcal{R}, S, E, C, \rho\}$

- \mathcal{X} is the set of states
- \mathcal{A} is the set of actions
- \mathbf{a} is the set of agents
- \mathcal{R} is the set of roles
- $S : \mathcal{R} \times \mathcal{X} \rightarrow \mathbb{R}$ is the association between roles and states
- $E : \mathcal{R} \times \mathcal{A} \rightarrow \mathbb{R}$ is the emphasis of actions in roles
- $C : \mathbf{a} \times \mathcal{X} \times \mathcal{A} \times \mathbf{a} \times \mathcal{X} \rightarrow (\mu_C, \sigma_C^2)$ is the function of capabilities, where

$$C(a_1, x_1, A, a_2, x_2)$$

returns the mean and variance of the utility obtained when agent a_1 in state x_1 performs action A while agent a_2 is in state x_2

- $\rho \in (0, 1)$ is the amount of risk to take in assigning roles

To aid in the explanation of the MuSCRA model, we use a single scenario extensively. Consider a robot soccer scenario where a team of 2 robots plays against another team of 2 robots. The goal of each team is to score as many points into the opponent's goal as possible, while minimizing the number of points scored against themselves. Fig. 1 shows the field and the positions of the 2 teams α and β . The teams α and β score into goals G_β and G_α respectively. Although the examples involve 2 agents in a robot soccer scenario, the MuSCRA model is applicable to scenarios with many agents, and in other domains, which we discuss later.

States

The set of states \mathcal{X} is the set of all possible states of the agents. \mathcal{X} is not the joint state-space of the team — each $x \in \mathcal{X}$ represents a state that a single agent can be in.

In robot soccer, the state of each robot has 2 features: its physical position on the field (defensive/offensive half), and whether it is blocked by an opponent. However, for ease




Figure 2: Actions of soccer robots. Large circles indicate the robots, and small black circles indicate the ball. The thick bold line represents the goal area. Solid and dotted arrows indicate the path of the ball and robot respectively.

of explanation, we only use the first feature; we will elaborate on the 4-states case (using both features) later in this section. We thus define the set of states:

$$\mathcal{X} = \{x_d, x_o\}$$

where the x_d means that the robot is in the defensive half, and x_o means that the robot is in offensive half. In Fig. 1, robot α_1 is in state x_d and robot α_2 is in state x_o .

Actions

The set of actions \mathcal{A} represents actions that the agents can take. Similar to \mathcal{X} , \mathcal{A} is not the joint action-space of the team, but instead individual actions that each agent can take. These actions may take a teammate as a parameter, if the action involves a teammate.

In the robot soccer scenario, each robot can perform 3 actions, with varying capabilities. These actions (and their corresponding symbols) are: dribble (A_d), pass (A_p), and score (A_s). Thus, the set of actions is:

$$\mathcal{A} = \{A_d, A_p, A_s\}$$

Dribbling involves the robot moving the ball as along the field (typically around an opponent), while passing involves kicking a ball to a teammate. Scoring involves shooting the ball directly at the opponent's goal. Fig. 2 shows an example of these actions in the field. The solid arrows indicate the direction of travel of the ball, while the dotted arrows represent the path of the robot.

In most situations, the utility obtained by an agent performing an action is affected by its current state, as well as its teammate and its teammate's state. For example, it is more likely for a robot in an open position to successfully pass to a teammate who is also in an open position, while it is more unlikely to succeed if either or both of the robots are blocked. Also, a particular robot may be better at receiving passes than other robots in the team.

Agents

The set of agents, \mathbf{a} , represent the team of cooperative agents whose roles are being assigned. Only the team of agents are considered — adversarial agents (or agents that cannot be controlled) are not part of this set.

Thus, for the robot soccer example, we only control the team of shaded robots, and the team of agents is defined as:

$$\mathbf{a} = \{\alpha_1, \alpha_2\}$$

Roles

The roles of the team, \mathcal{R} , represent associations with certain states of the world, as well as an emphasis in certain actions. Each role is assigned to a single agent, and so roles can be viewed as the smallest element of a team. In particular, the number of roles should be no larger than the number of agents, i.e., $|\mathcal{R}| \leq |\mathbf{a}|$, which ensures that every role can be fulfilled by an agent.

In robot soccer, we can define the set of roles as:

$$\mathcal{R} = \{R_d, R_a\}$$

where R_d is a defender and R_a is an attacker.

The defender’s responsibility is primarily to prevent a goal from being scored against the team, and secondarily to pass the ball upfield to the attacker. The attacker’s responsibility is to score goals for the team.

Association between Roles and States

Roles are associated with states of the world, and this is represented by the function $S : \mathcal{R} \times \mathcal{X} \rightarrow \mathbb{R}$, which indicates how strongly associated a state and role are, where a higher value indicates a stronger association. S is such that:

$$\forall R \in \mathcal{R}, x \in \mathcal{X}, 0 \leq S(R, x) \leq 1 \quad (1)$$

$$\forall R \in \mathcal{R}, \sum_{x \in \mathcal{X}} S(R, x) = 1 \quad (2)$$

Eqn. 1 states that all associations of roles and states are between 0 and 1. Eqn. 2 states that the sum of all associations of any given role is 1, which provides normalization so that the weighting of every role is equal. A state x that is unassociated with any role, i.e., $\forall R \in \mathcal{R}, S(R, x) = 0$, means that the state is unimportant to all roles in the team.

Table 1 displays the values of the function S in the robot soccer example. The state x_d (i.e., the robot is physically in the defensive half) is associated with the defender role R_d , and the state x_o (i.e., the robot is physically in the offensive half) is associated with the attacker role.

$S(R, x)$	x_d	x_o
R_d	1	0
R_a	0	1

Table 1: Regions defined in robot soccer

Emphasis of Actions in Roles

Different roles may emphasize different actions, to indicate how important the action is to a role. This emphasis is given by the function $E : \mathcal{R} \times \mathcal{A} \rightarrow \mathbb{R}$, such that:

$$\forall R \in \mathcal{R}, A \in \mathcal{A}, 0 \leq E(R, A) \leq 1 \quad (3)$$

$$\forall R \in \mathcal{R}, \sum_{A \in \mathcal{A}} E(R, A) = 1 \quad (4)$$

Eqn. 3 states that all emphases of actions in roles are between 0 and 1 (inclusive), and Eqn. 4 states that the sum of emphases for any role is 1, which ensures normalization.

In our robot soccer example, the function E is shown in Table 2. The defender role R_d places a high emphasis on the dribbling action A_d and the passing action A_p , with a little emphasis on the scoring action A_s . This reflects the defender’s role in preventing a goal from being scored and pushing the ball upfield. The small emphasis on scoring reflects that the defender may take a direct shot on goal if an opportunity arises. The attacker role R_a places a very high emphasis on scoring (which reflects its main responsibility), and dribbling (which it uses to go around opponents). There is no emphasis on passing, since there is no other player upfield worth passing the ball to.

$E(R, a)$	A_d (Dribble)	A_p (Pass)	A_s (Score)
R_d	0.5	0.4	0.1
R_a	0.3	0	0.7

Table 2: Emphasis of actions in roles in robot soccer

Mutual State Capabilities

In a heterogeneous team, agents have different capabilities in the actions and skills relevant to the task. The capability function, $C : \mathbf{a} \times \mathcal{X} \times \mathcal{A} \times \mathbf{a} \times \mathcal{X} \rightarrow (\mu_C, \sigma_C^2)$, represents these capabilities of the agents.

The first two parameters of C represent the agent and its state. The third parameter, \mathcal{A} , represents the action being performed, and the last two parameters represent the teammate and the teammate’s state. Such a parametrization of capabilities takes into account that an agent’s ability to perform an action and achieve the desired outcome depends on the teammate, and their mutual state, i.e., the agent’s state and the teammate’s state. If an action does not depend on a teammate, e.g., the scoring action A_s in the robot soccer example, then the values returned by C is equal for all values of teammates and teammate states.

C returns a mean and variance, which represent the distribution of the utility obtained by performing the action. The variance in the utility is affected by the uncertainty in performing the action (how likely the robot is able to carry out the action) as well as the uncertainty in the world (how likely the desired outcome of the action is obtained).

In addition, the means and variances returned by C are intended to be obtained through a data-driven approach, e.g., learning or modelling. Thus, by including a variance term, the model captures an idea of how much data has been collected and how consistent a robot is in its actions. For example, if only the mean was captured, then the value of a robot that is consistently observed to be average would be equal to that of another robot that has only been observed twice, once really well and once really poorly.

We henceforth adopt the notation C_μ to represent the function derived from C that returns only the mean (μ_C), and

C_{σ_2} to represent the function derived from C that returns only the variance (σ_C^2).

Table 3 shows the different capabilities of the 2 robots in the robot soccer example. Robot α_2 is slightly better (in terms of the mean) than α_1 in passing and scoring (A_p and A_s respectively), as shown in italic for α_1 and bold for α_2 in Table 3. However, the variance of α_2 's utility in these actions are also higher than that of α_1 . Depending on the amount of risk taken (ρ) in the role-assignment, the role assignments of α_1 and α_2 may change, as described later.

Risk

The ρ term in the MuSCRA model represents how much risk to take while assigning roles to the team. The utility U_r of a role assignment is normally distributed with a mean and variance, and given a certain value of $\rho \in (0, 1)$, u is a value such that $U_r \leq u$, i.e.,

$$P(U_r \leq u) = \rho \quad (5)$$

Thus, as ρ increases, the probability that the role assignment value takes on a lower value than u increases, and so ρ is a measure of risk taken in the role assignment. Further details on using ρ in calculate the value of a policy is shown later.

3.2 Applications of the MuSCRA Model

The robot soccer scenario described in detail earlier in this section is a good application of the MuSCRA model. As mentioned before, the states of each robot have two main features, its position on the field, and whether the robot is blocked by an opponent robot. The states of the robots would then be {(defensive, clear), (defensive, blocked), (offensive, clear), (offensive, blocked)}. Such a formulation allows the encapsulation of information about the world beyond positional information. For example, the capability of a robot passing a ball to a teammate, when both of them are blocked by opponents, would take into account the robots' abilities to manoeuvre around the opponents, as well as the opponent's abilities to get in the way. Thus, even though the opponent robots are not explicitly modelled, their strengths and weaknesses affect the values in the team's capabilities.

Table 4 lists a number of domains that the MuSCRA model can be applied to. Besides robot soccer, the MuSCRA model works well in other domains that involve role assignment to heterogeneous teams, e.g., urban search and rescue (USAR). The success of agents' actions in USAR depends on their teammates and their mutual states, e.g., the ability for a robot to move through a blocked area depends on which teammate is carrying rubble to help clear the path. Roles can be split into searchers, whose main job is to search for humans, and diggers, who clear rubble from the area and allow the searchers to travel quickly.

The MuSCRA model can also be applied to task allocation domains, such as in preparing a presentation. Actions include analyzing data for the presentation, creating slides, and actually making the presentation. These actions depend on teammates and mutual states as well. For example, an agent's capability in drawing graphics depends on whether it has data available and which teammate planned the layout of the slides.

4. ROLE-ASSIGNMENT IN MUSCRA

In order to find the optimal assignment of roles in the MuSCRA model, we define a role assignment policy:

DEFINITION 2. A role assignment policy $\pi : \mathcal{R} \rightarrow \mathbf{a}$ is an assignment of roles to agents such that every agent has at most 1 role, i.e., $\pi(R) = \pi(R') \Rightarrow R = R'$.

Def. 2 states that role assignments are unique; no agent has more than 1 role. When there are more agents than there are roles, it is possible for some agents to have no role. Role assignment policies are similar to the single-task robots, single-robot tasks, instantaneous assignment (ST-SR-IA) in Gerkey and Mataric's task allocation taxonomy [5].

4.1 Finding the Optimal Policy

Given a role assignment policy π , we determine the utility of the team thus assigned, taking into account the capabilities of each agent and its assigned role. We define the utility of a role assignment policy:

DEFINITION 3. The utility of a role assignment policy is determined via the function:

$$U : \pi \rightarrow (\mu_\pi, \sigma_\pi^2)$$

where μ_π and σ_π^2 represent the mean and variance of the policy's utility.

We denote U_μ as the function derived from U that returns the mean, and U_{σ^2} as the function derived from U that returns the variance. The functions are computed as follows:

$$U_\mu(\pi) = \sum_{\substack{R \in \mathcal{R} \\ x \in \mathcal{X} \\ A \in \mathcal{A} \\ R' \in \mathcal{R}: R \neq R' \\ y \in \mathcal{X}}} \phi(\cdot) C_\mu(\pi(R), x, A, \pi(R'), y)$$

$$U_{\sigma^2}(\pi) = \sum_{\substack{R \in \mathcal{R} \\ x \in \mathcal{X} \\ A \in \mathcal{A} \\ R' \in \mathcal{R}: R \neq R' \\ y \in \mathcal{X}}} \phi(\cdot) C_{\sigma^2}(\pi(R), x, A, \pi(R'), y)$$

where ϕ is a weight function:

$$\phi(R, x, A, R', y) = E(R, A)S(R, x)S(R', y)$$

Using the action emphasis function E and role-state association function S , ϕ determines how much weight to place on the utility of an action taken by a role. Thus, actions with more emphasis in the role will reflect a higher weight in ϕ . Similarly, highly associated states of the agent and its teammate will have higher weights during U 's calculation.

It may seem that the calculation of a policy's utility involves a massive amount of computation. However, the states in which the roles are valid (i.e., $S(R, x) > 0$ or $S(R', y) > 0$) are typically much smaller than the entire state space. Thus,

Agent	Agent's State	Action	Teammate	Teammate's State	Mean (μ_C)	Variance (σ_C^2)
α_1	x_d	A_d	α_2	x_o	2	1
α_1	x_d	A_p	α_2	x_o	8	2
α_1	x_d	A_s	α_2	x_o	3	1
α_1	x_o	A_d	α_2	x_d	5	2
α_1	x_o	A_p	α_2	x_d	-3	2
α_1	x_o	A_s	α_2	x_d	10	3
α_2	x_d	A_d	α_1	x_o	2	1
α_2	x_d	A_p	α_1	x_o	9	3
α_2	x_d	A_s	α_1	x_o	4	3
α_2	x_o	A_d	α_1	x_d	5	2
α_2	x_o	A_p	α_1	x_d	-2	3
α_2	x_o	A_s	α_1	x_d	12	7

Table 3: Capabilities in robot soccer example

Domain	State Features	Actions	Roles
Robot Soccer	Position on field Clear / Blocked by opponent	Dribble Pass Score	Defender Attacker
Urban Search & Rescue	Near human Blocked / Open area Carrying rubble	Detect humans Dig through rubble Carry rubble Move	Searcher Digger
Assembly	Holding object Near object Object in sight	Lift object Rotate object Drill Bolt	Coarse Manipulator Fine Manipulator
Presentation	Data available Graphics created Layout planned	Analyze data Draw graphics Create slides Make presentation	Slide creator Graphic Designer Presenter

Table 4: Examples of the MuSCRA model in different domains

the computation of U can be optimized. Furthermore, calculation of a policy's mean utility and variance is only a constant factor increase as compared to calculation of the policy's mean utility alone.

In the robot soccer example, 2 possible role assignment policies exist. The policy $\pi_1 = (R_d \rightarrow \alpha_1, R_a \rightarrow \alpha_2)$ has a mean utility of 14.4 with variance 6.9, while the policy $\pi_2 = (R_d \rightarrow \alpha_2, R_a \rightarrow \alpha_1)$ has a mean utility of 13.5 with variance 4.7.

Incorporating Risk into Utility

The utility of a policy π is normally distributed with a mean and variance, as computed by U_μ and U_{σ^2} respectively. Given the risk parameter ρ , we define the value of a policy as follows:

DEFINITION 4. *The value of a policy is given by the function $V : \pi \rightarrow \mathbb{R}$, where:*

$$V(\pi) = U_\mu(\pi) + \sqrt{U_{\sigma^2}(\pi)}\Phi^{-1}(\rho)$$

where Φ^{-1} is the inverse of the standard normal cumulative distribution function.

Thus, ρ is the probability that the value of a role assignment policy π is lower than $V(\pi)$, and so ρ represents the risk taken in the role assignment policy.

Optimal Role-Assignment Policy

With a value function V as defined above, we define the optimal (or best) policy:

DEFINITION 5. *The optimal policy π^* is the policy with the highest value:*

$$\forall \pi, V(\pi^*) \geq V(\pi)$$

The optimal policy π^* has the highest value among all possible policies. Thus, we can find π^* with the following:

$$\pi^* = \underset{\pi}{\operatorname{argmax}} V(\pi)$$

Referring back to robot soccer, when $\rho = 0.2$, the policies π_1 and π_2 have values 8.6 and 9.5 respectively, so π_2 is the optimal policy. However, when $\rho = 0.8$, the values of π_1 and π_2 are 20.2 and 17.4 respectively, and π_1 becomes the optimal policy. Thus, varying the value of ρ can result in different optimal policies, e.g., π_2 , the policy with a lower

mean and variance, has a higher value than π_1 when ρ is low, i.e., a conservative policy is preferred.

4.2 Approximation Algorithms

Finding the optimal policy π^* using a brute-force method involves searching the entire space of policies, which is factorially large. As such, we consider approximation algorithms such as hill-climbing.

Hill-Climbing

To perform hill-climbing or other approximation techniques, we define neighbors of policies:

DEFINITION 6. *Role assignment policies π_1 and π_2 are neighbors if:*

- π_1 and π_2 have a swapped value, i.e., $\exists x_i, x_j$ s.t.
 $(\pi_1(x_i) = \pi_2(x_j)) \wedge (\pi_1(x_j) = \pi_2(x_i)) \wedge$
 $(x_i \neq x_j) \wedge (\forall x \neq x_i, x \neq x_j, \pi_1(x) = \pi_2(x))$

OR

- π_1 and π_2 differ by one value, i.e., $\exists x$ s.t.
 $(\pi_1(x) \neq \pi_2(x)) \wedge (\forall x' \neq x, \pi_1(x') = \pi_2(x'))$

Hill-climbing is performed by selecting an initial policy π and evaluating its neighbors and choosing the best neighbor to iterate on. It continues until all neighbors of the policy have equal or lower value.

Hill-Climbing with Random Restarts

The performance of the hill-climbing algorithm is dependent on the initial starting policy. In order to circumvent this issue, we can perform hill-climbing with random restarts. The algorithm calls the hill-climbing function repeatedly, starting with random initial policies, and returns the best policy from all the hill climbs.

5. EXPERIMENTS AND RESULTS

In order to test our model and solving algorithms, we first generated a robot soccer scenario, with different types of agents, and used the approximate algorithms as a proof-of-concept of our model. Next, we generated random instances of the MuSCRA model to experiment on, and evaluated the performance of the approximate algorithms versus a market-based bidding technique. While the capability function is intended to be created via a data-driven (or learning) method, the focus of this paper is on the MuSCRA model and not the learning technique, and as such, we generated the capabilities based on normal distributions, as described below for each domain.

5.1 Robot Soccer Domain

To simulate robot soccer, we defined 3 states in the world: a defensive zone, mid-field, and an offensive zone. We defined 4 possible actions: passing, clearing, dribbling and scoring. Next, we created 3 **types** of roles for the agents: defenders, mid-fielders, and attackers. Given the number of agents n as input, we created $\lfloor n/3 \rfloor$ copies of the mid-fielder role,

$\lfloor \frac{n - \lfloor n/3 \rfloor}{2} \rfloor$ defender roles, and the rest as attacker roles. As such, there were n roles to be assigned to n agents.

Table 5 shows the role-state associations for the defender, mid-fielder, and attacker roles, and Table 6 shows the emphasis of actions for the roles. There were **multiple copies** of each role type (defender, mid-fielder, attacker), and they each took on the relevant values shown in Tables 5 and 6.

	Defensive	Mid-Field	Offensive
Defenders	0.8	0.2	0
Mid-Fielders	0.2	0.6	0.2
Attackers	0	0.2	0.8

Table 5: Role-State associations for robot soccers

	Pass	Clear	Dribble	Score
Defenders	0.2	0.5	0.3	0
Mid-Fielders	0.4	0.2	0.3	0.1
Attackers	0.2	0	0.2	0.6

Table 6: Essential actions in roles for robot soccer

We generated the capabilities of the n agents by doing the following: given the number of defender, mid-fielder and attacker roles, we first pre-assigned each agent to a role. Next, we generated capabilities such that agents pre-assigned as defenders had higher mean utilities for the clearing action, mid-fielders had higher mean utilities for passing, and attackers had higher mean utilities for scoring.

The goal of this setup was to confirm that the approximation algorithms would be able to find the “right” role assignments, given the capabilities that we generated. We ran the hill-climbing algorithms, varying the number of robots from 3 to 7. In all cases, we found the optimal role assignment policy, which matched the pre-assigned roles of the agents.

5.2 Random Domains

We created a simulator in Java that, given as input the number of states $|\mathcal{X}|$, the number of actions $|\mathcal{A}|$, the number of agents $|\mathbf{a}|$ and the number of roles $|R|$, generates the functions S, E, and C randomly, as described below.

To generate the associations S between roles and states, we did the following: for each role and state, we generated a random number uniformly distributed between 0 and 1. We then normalized the values for each role, so that the sum of its associations was 1, i.e., $\forall R \in \mathcal{R}, \sum_{x \in \mathcal{X}} S(R, x) = 1$.

We created the emphasis function E similarly, by generating a random number uniformly distributed between 0 and 1 for each role-action pair. We then normalized these values as well to ensure that $\forall R \in \mathcal{R}, \sum_{A \in \mathcal{A}} E(R, A) = 1$.

Lastly, we created the capabilities of robots by generating a mean between -1 and 1 (using a standard Normal distribution, with tail ends removed), and a variance between 0 and 1, taking the absolute value of a standard Normal distribution, with tail ends removed.

5.3 Evaluating the Algorithms

To compare the approximate algorithms against a market-based approach, a bidding technique had to be devised. The agents would bid for each role sequentially, and the agent with the highest bid is assigned the role. To generate the bid, each agent a calculated its individual capability for a role R , as shown below:

$$c(a, R) = \sum_{\substack{A \in \mathcal{A} \\ x \in \mathcal{X}}} E(R, A) S(R, x) \frac{\sum_{\substack{a' \in \mathcal{A} \\ x' \in \mathcal{X}}} C(a, x, A, a', x')}{\sum_{\substack{a' \in \mathcal{A} \\ x' \in \mathcal{X}}} 1}$$

Using the mean and variance of its individual capability, each agent would then form the bid by incorporating the risk factor ρ . Essentially, each robot calculated its capability based on the average of its teammates and their states, since at the time of bidding, the agents do not have any knowledge about what roles their teammates will take.

We then implemented the hill-climbing and hill-climbing with random restarts algorithms and compared them with the bidding technique. To create a random policy, we chose a policy at random from the entire space of policies.

In the experiments for the random domains, we varied n to be between 5 and 7, and set $|\mathcal{X}| = |\mathcal{A}| = |\mathbf{a}| = |\mathcal{R}| = n$. We generated 300 MuSCRA models and ran the algorithms. We fixed the risk factor ρ to be 0.5 for all the experiments. The number of random restarts was set to be 5% of the number of possible policies. We limited the value of n to 7 since a brute-force search of large policy spaces (to obtain the optimal policy) would take extremely long to compute when $n > 7$.

We defined the effectiveness of the algorithm, by comparing the policy found π against the optimal policy π^* and the worst policy π_{min} , i.e., the policy with the minimum value:

$$effectiveness = \frac{V(\pi) - V(\pi_{min})}{V(\pi^*) - V(\pi_{min})}$$

Table 7 displays the value of the best policy found by the algorithms, as a percentage compared with the brute-force method (which evaluates all policies). Hill-climbing and random restarted performed very well, compared to the market-based technique, without exploring much of the policy space (2.0% and 8.9% respectively when $n = 7$).

Effectiveness of policy in random domains	n		
	5	6	7
Hill-climbing	98.1%	97.5%	97.2%
Random restarts	98.3%	100%	100%
Market-based	68.1%	69.6%	70.2%

Table 7: Effectiveness of policies found

To ensure that the risk factor ρ did not have an effect on the algorithms, we repeated the experiments on hill-climbing described above, fixing the value of n to 5, and varying ρ from 0 to 1, and running 1000 trials each. We found that the performance of the algorithms were not affected by the value of ρ , because the algorithms are blind to the value function, and the definition of neighbors in policy space allowed

an effective exploration of the space. However, for a given random domain, different values of ρ will result in different optimal policies found.

6. CONCLUSION

We formally defined the Mutual State Capability-Based Role Assignment (MuSCRA) model, and described each of its components in detail, using robot soccer as an example of an instantiation of the model. We briefly described other domains where MuSCRA can be applied, giving examples of possible states, actions and roles in those domains. Capabilities of agents in MuSCRA are defined not only as a pairing between an agent and action, but also incorporates the teammate, and their mutual state, i.e., the state of the agent and its teammate. This allows a generalization of capabilities to include the fact that the success of an action in a team depends on the composition of the team, as well as the state of the world. Information about the world can be embedded as part of the agent’s (or its teammate’s) state, such as whether a robot is blocked by an opponent in robot soccer. In addition, capability is represented by a mean and variance, to signify the uncertainty in the actions and world, as well as the reliability of the data collected from observations. Since the values of capability are assumed to be obtained through observation, the variance modelled in the capability function provides a measure of how much data has been collected from an agent, as well as how consistent the agent’s performance is.

We defined a role assignment policy, as well as how to determine the utility of such a policy, represented by a mean and variance. We then described how to incorporate the risk factor to retrieve the value of a policy. The risk factor adjusts the mean-to-variance trade-off in the optimal role assignment, and is an important contribution of our model. We then discussed approximation algorithms such as hill-climbing and hill-climbing with random restarts, after defining the concept of a neighbor in a role assignment policy. We ran extensive experiments on the approximation algorithms, and showed that hill-climbing (both the simple version and with random restarts) found effective policies (as compared to the optimal), performing better than market-based techniques. Different values of risk affected the optimal policy found, but did not affect the performance of the algorithms.

A possible application of MuSCRA is in a pick-up team scenario, where a group of cooperative robots collectively learn about their abilities in a team and discover an optimal role assignment for the team. However, there are certain drawbacks to our approach. Firstly, the calculation of a policy’s utility can take extremely long in a worst-case scenario. Also, it may take a long time for the team to have sufficient observations to fill out the capability function before a good role assignment is obtained. An agent that is aware of its own skills may not be able to readily transfer this knowledge to the rest of the team, since the capabilities are defined as a function of the teammates as well. Thus, robots will need to experiment with their team to learn this information. However, learning techniques can be applied to our model, in order to generate the capability function used in role assignment.

We are currently working on methods to generate the capa-

bility function from observational data, as well as methods to determine what actions robots should take in order to explore the capability space more effectively. In addition, we hope to be able to use MuSCRA in adversarial conditions like robot soccer, so that robots can learn their capabilities with regards to the opponents, and develop an effective counter-strategy.

Acknowledgments

This work was partially supported by the Lockheed Martin, Inc. under subcontract 8100001629/1041062, and the Agency for Science, Technology, and Research (A*STAR), Singapore. The views and conclusions contained in this document are those of the authors only.

7. REFERENCES

- [1] Y. Chen and Y. Wang. Obstacle avoidance and role assignment algorithms for robot formation control. In *Int. Conf. Intelligent Robots and Systems*, pages 4201–4206, 2007.
- [2] M. B. Dias, R. Zlot, N. Kalra, and A. Stentz. Market-based multirobot coordination: A survey and analysis. *Proc. IEEE*, 94(1):1257–1270, 2006.
- [3] V. Frias-Martinez, E. Sklar, and S. Parsons. Exploring auction mechanisms for role assignment in teams of autonomous robots. In *Proc. RoboCup Symp.*, pages 532–539, 2004.
- [4] A. Garland and R. Alterman. Autonomous agents that learn to better coordinate. In *Proc. 3rd Int. Conf. Autonomous Agents and Multiagent Systems*, pages 267–301, 2004.
- [5] B. P. Gerkey and M. J. Mataric. A formal analysis and taxonomy of task allocation in multi-robot systems. *Int. J. Robotics Research*, 23(9):939–954, 2004.
- [6] P. J. Gmytrasiewicz and P. Doshi. A framework for sequential planning in multi-agent settings. *J. Artificial Intelligence Research*, 24:24–49, 2004.
- [7] C. Guttman. Making allocations collectively: Iterative group decision making under uncertainty. In *Proc. 6th German Conf. Multiagent System Technologies*, pages 73–85, 2008.
- [8] L. He and T. R. Ioerger. A quantitative model of capabilities in multi-agent systems. In *Proc. Int. Conf. Artificial Intelligence*, pages 730–736, 2003.
- [9] J. Kok and N. Vlassis. Mutual modeling of teammate behavior. Technical Report IAS-UVA-02-04, Computer Science Institute, University of Amsterdam, The Netherlands, 2004.
- [10] C. McMillen and M. Veloso. Distributed, Play-Based Role Assignment for Robot Teams in Dynamic Environments. In *Proc. Int. Symp. Distributed Autonomous Robotic Systems*, pages 145–154, 2006.
- [11] S. Proper and P. Tadepalli. Solving multiagent assignment markov decision processes. In *Proc. 8th Int. Conf. Autonomous Agents and Multiagent Systems*, pages 681–688, 2009.
- [12] R. Simmons, S. Singh, D. Hershberger, J. Ramos, and T. Smith. First results in the coordination of heterogeneous robots for large-scale assembly. In *Proc. 7th Int. Symp. Experimental Robotics*, pages 323–332, 2000.
- [13] P. Stone and M. Veloso. Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. *J. Artificial Intelligence*, 110(2):241–273, 1999.
- [14] D. Vail and M. Veloso. Dynamic multi-robot coordination. In *Multi-Robot Systems: From Swarms to Intelligent Automata, Volume II*, pages 87–100, 2003.