Coordination without Negotiation in Teams of Heterogeneous Robots

Michael Isik[1], Freek Stulp[1], Gerd Mayer[2], and Hans Utz[2]

[1] Technische Universität München, Boltzmannstr. 3, D-85747 München, Germany
[2] University of Ulm, James-Franck-Ring, D-89069 Ulm, Germany
{isik,stulp}@cs.tum.de, {gerd.mayer,hans.utz}@informatik.uni-ulm.de

Abstract. A key feature of human cooperation is that we can coordinate well without communication or negotiation. We achieve this by anticipating the intentions and actions of others, and adapting our own actions to them accordingly. In contrast, most multi-robot systems rely on extensive communication to exchange their intentions.
This paper describes the joint approach of our two research groups to enable a heterogeneous team of robots to coordinate *implicitly*, without negotiation. We apply implicit coordination to a typical coordination task from robotic soccer: regaining ball possession. We discuss the benefits and drawbacks of implicit coordination, and evaluate it by conducting several experiments with our robotic soccer teams.

1 Introduction

Coordination of actions is essential to solve multi-agent tasks effectively. A striking aspect of human coordination is that we can achieve it with little or no communication. Consider two people assembling a bookcase. With apparent ease, actions are *anticipated* and coordinated: if I see you grab a screwdriver, I will take one of the shelves, and hold it in place, and so forth. Instead of communicating, humans achieve this by inferring the intentions of others. Once the beliefs and desires of the cooperating party are known, we simply imagine what we would do in that situation. This is called the Intentional Stance [1].

In previous research, we have achieved negotiation-free coordination, also called *implicit* coordination, within the Agilo RoboCuppers group [2]. Here, we extend and integrate this with another line of research, which is the formation of a mixed team [3]. Due to scientific as well as pragmatic reasons, there is a growing interest in the robotics field to join the efforts of different labs to form mixed teams of autonomous mobile robots. For many tasks, a group of heterogeneous robots with diverse capabilities and strengths is likely to perform better than one system that tries to encapsulate them all. Also, for many groups, the increasing cost of acquiring and maintaining autonomous mobile robots keeps them from forming a mixed team themselves. Furthermore, to allow all mid-size teams to participate at RoboCup 2006, many of them are required to form a mixed team. Our two groups, who have individually taken part in the RoboCup mid-size league since 1998, have formed a mixed team with robots from the different labs. As almost all robots in this league, our robots are custom built research platforms with unique sensors, actuators, and software architectures. Therefore, forming a heterogeneous cooperative team presents an exciting challenge. One of these challenges is achieving robust coordination.

The standard solution in robotic teams is not to anticipate the actions of others, as humans do, but instead to extensively communicate utilities or intentions in a negotiation scheme. Previous work on robot coordination seems to have focussed almost exclusively on explicit coordination, as an overview paper on the key architectures shows [4].

In this paper, we discuss the benefits of implicit coordination, and implement it for our heterogeneous team. We apply implicit coordination to a typical coordination task from robotic soccer: regaining ball possession. Acquiring ball possession is a goal for the team as a whole, but only one of the field players is needed to achieve it. Of course, the robots must agree upon which robot will approach the ball. The intuitive underlying rule is that only the robot who is quickest to the ball should approach it. To infer the intentions of others, the agents first learn utility prediction models from observed experience. For the ball approach task, the utility measure is time, so the robots learn to predict how long it will take to approach the ball. During task execution, the robots locally predict the utilities for all robots, and globally coordinate accordingly.

The main contributions of this paper are: 1) learning temporal prediction models that take technical differences between the robot platforms into account 2) using these models to enable implicit coordination within a heterogeneous team of robots 3) demonstrating that coordination based on belief states is more robust than explicit coordination.

The rest of this paper is organized as follows. In the next section we describe how implicit coordination was implemented in our teams. Experimental results are presented in Section 3. In this section, we also discuss the benefits, as well as some drawbacks, of implicit coordination. Related work is presented in Section 4, and we conclude with Section 5.

2 Applying Implicit Coordination

In [2], we introduced a computational model for implicit coordination, that specifies three components are necessary for implicit coordination: 1) utility prediction models 2) knowledge of the states of others 3) the robots should have a shared performance model for joint actions. In our scenario, the latter component is a locker-room agreement [5] that only the quickest robot should approach the ball. Here, we apply this computational model to a team of heterogeneous robots. After presenting the two teams, we discuss the first two components of the computational model in more detail.

The Ulm Sparrows [6] are custom built robots, with infrared based near range finders and a directed camera. The available actuators are a differential drive, a pneumatic kicking device and a pan unit to rotate the camera horizontally (270^o). Each robot acts upon an egocentric belief state, using the camera as its main sensor. The AGILO RoboCuppers [7] are customized Pioneer I robots, with differential drive and a fixed forward facing color CCD camera. They act upon an allocentric belief state, which is acquired by cooperative state estimation. For the experiments we will present later, it is important that the variables are controllable and reproducible. Therefore, we have used our ground truth system to determine the positions of the robots with even more accuracy. This system uses three ceiling cameras to detect colored markers on top of the robots.

Utility Prediction Models. All robots must be able to predict their own ball approach time, as well as that of others. Therefore, they learn temporal prediction models from observed experience. Examples are gathered by navigating to random targets on the field, thereby measuring the time it took to approach the target. These measurements were acquired through the ground truth system. A model tree is then trained with these examples. Model trees are functions that map continuous or nominal features to a continuous value. They recursively partition the data, and fit linear models to the data in each partition. In previous research [2], we have shown that executing 300 navigation tasks yields a sufficient amount of training examples to learn an accurate prediction model. This takes about half an hour, so this is well within the continuous operational range of the robots. The mean absolute error of the models on a separate test set was 0.25s for the Ulm Sparrows, and 0.18s for the AGILO RoboCuppers. For more information about model trees, and how they can be used to learn action models of navigation tasks, we refer to [2].

Knowledge of the states of others. Predicting utilities for others, called perspective taking, can only be done if the robots have estimates of the other's states, which can be difficult if the robots only have local sensors. For instance, due to the limited view of our cameras, it is often not possible to see all the teammates. Therefore, our robots communicate their belief states to each other to achieve more coherent and complete beliefs about the world [3], which they use to determine their (joint) actions. This might seem contrary to the paradigm that we want to achieve coordination without communication. However, there are some important differences between communicating *intentions* and communicating *beliefs*, as we shall discuss in Section 3.2.

3 Experimental Evaluation

To evaluate if the learned prediction models and shared representations are sufficiently accurate for implicit coordination, we have conducted three experiments, one in a dynamic, one in a static environment, and one in simulation. For each experiment, we used one Sparrow and one AGILO robot. Each robot has a temporal prediction model for both robot types.

Dynamic environment experiment. In this experiment, the robots continuously navigated to random targets on the field, for about half an hour. The paths were generated such that interference between the robots was excluded. At 10Hz, each robot records its own position and orientation, as well as that of its teammate and the ball. Each robot also logs the predicted approach time for both robots, and based on these times, which robot should approach the ball, in their view. Note that the robots never actually approach the ball.

Static environment experiment. In the previous experiment, it is impossible to measure if the temporal predictions were actually correct, and if potential inaccuracies caused the robots' estimate of who is quickest to be incorrect. Therefore a second experiment was conducted. The experimental set-up was as follows: Both robots navigate to different random positions and wait there. During the experiment, the target to approach is fixed and the same for both robots. Then, the robots are requested to record their own state, as well as that of their team mate. The robots compute the predicted approach times, and add them to the log-file. Then, one after the other, the robots are requested to drive to the goal position, and the actual approach duration is recorded. The log-files so acquired are almost identical to the ones in the dynamic experiment. The only difference is that they also contain the actual observed time for the robot. This static environment is less realistic, but allows us to compare the predicted time with the actually measured time for each robot.

Simulated experiment. Here, the experimental set-up is identical to the dynamic experiment. The simulator allows us to vary two variables that most strongly influence the success of implicit coordination. The first is communication quality. At random times, and for random durations, communication is switched off in both directions. By controlling the length of the intervals, we can vary between perfect (100%) and no (0%) communication. The second is the field of view of the robot. We can set the view angle of the robot's forward facing camera between 0 (blind) and 360 (omni-directional vision) degrees. The other robot and the ball are only perceived when in the field of view. Gaussian noise with a standard deviation of 9, 25 and 22 cm is added to the robot's estimates of the position of itself, the teammate and the ball respectively. These correspond to the errors we have observed on the real robots.

3.1 Results

Do the robots agree upon who should approach the ball? To answer this question, we simply determined how often the two robots agreed on which robot should approach the ball in the dynamic experiment, which was 96%.

Do the robots choose the quickest one? We would also like to know if the chosen robot is actually the quickest one to approach the ball. Of course, this could only be determined in the static experiment, in which the actual times it took each robot to approach the ball are known. A robot's decision to coordinate is deemed correct, if the robot that was the quickest was indeed predicted to be the quickest. The robots' choice was correct 92% of the time.

Are temporal prediction models necessary, or would a more simple value such as distance suffice? Using only distance as a rough estimate of the approach time, would save us the trouble of learning models. Although time is certainly strongly correlated with distance, using distance alone leads to significantly more incorrect coordinations. Agreement is still very good (95%), but the robot that is really the quickest is chosen only 68% of the time. So, when using distance, the robots are still very sure about who should approach it, but they are also wrong about it much more often.

When does implicit coordination fail? In the dynamic experiment, coordination succeeds 96% of the time. In the log-file, we labeled all examples in which exactly one robot decided to approach the ball with 'Success', and others with 'Fail'. A decision tree was then trained to predict this value. The learned tree is represented graphically in Figure 1. The main rule is that if the difference in predicted times between two robots is small, coordination is likely to fail, and if it is large, it is likely to succeed. This is intuitive, because if the difference between the times is large, it is less likely that estimation errors will invert which time is the smallest. Note that in between these two limits, there is a 'gray' area, in which some other rules were learned. They only accounted for a small number of example, so for clarity, we will not discuss them here.

In sports like soccer or volleyball, it is sometimes not completely clear who should go for the ball. Humans solve this problem by communicating their intention through an exclamation: "Mine!", or "Leave it!". The decision tree essentially provides the robots with similar awareness, as they predict when implicit coordination failure is likely. So, they could be used for instance to determine when robots should resort to explicit coordination.

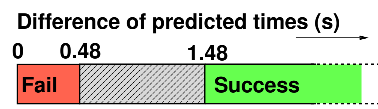

Fig. 1. Visualization of the decision tree that predicts coordination failure

How do communication quality and state estimation accuracy influence coordination? The results of the simulation experiment, which show how the performance of different coordination strategies depends on the quality of communication and the field of view, are depicted in Figure 2. Communication quality is the percentage of packets that arrive, and field of view is in degrees. The z-axis depicts coordination success, which is the percentage that only one robot intended to approach the ball.

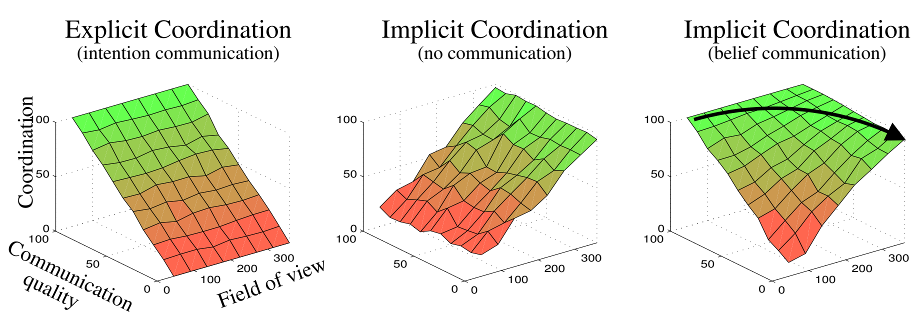

Fig. 2. Results of the simulation experiment, which show how the performance of coordination strategies depends on the quality of communication and the field of view.

Since explicit coordination is based completely on communication, it is not surprising that it perfectly correlates with the quality of the communication, but is independent of the size of the field of view. No communications means no coordination, and perfect communication means perfect coordination. For implicit coordination without communication, the relation is converse. If a robot is able to estimate the states of others better, it is able to coordinate better. The third graph shows implicit coordination with belief state exchange (as used on our real robots). If the robot has another in its field of view, it determines the other's state through state estimation, otherwise it uses communication (if possible) to exchange beliefs. These states are then used to predict the utilities of others, independent if they were perceived or communicated. The graph clearly shows that this approach combines the benefits of both.

3.2 Discussion

There are several important benefits that implicit coordination without communication has over explicit coordination. First of all, protocols and arbitration mechanisms must be adopted between communicating entities to enable intention communication, which adds complexities and can degrade the system. It is generally argued that communication can add unacceptable delays in information gathering and should be kept minimal [8]. Furthermore, rescue robotics and autonomous vehicles operating in traffic are examples of domains in which robust communication is not guaranteed, but where correct coordination and action anticipation is a matter of life and death. Finally, human-robot interaction, a current research focus in for instance space exploration or rescue robotics, it cannot be expected of humans to continuously communicate their intentions. Instead, the robot must be able to anticipate a human's intentions, based on predictive models of human behavior. We consider implicit coordination to be essential for natural interaction between robots and humans, so adhering to explicit coordination will prevent robots from making a break-through into these application domains.

The most difficult aspect of implicit coordination is estimating the states of others. Especially for robots with a limited field of view, such as ours, this is problematic. Therefore, we resorted to the communication of beliefs to acquire a shared representation. This might seem contrary to our communication-free paradigm, but there is an important difference between communicating intentions and beliefs. We believe that improvements in sensor technology and state estimation methods will allow robots to autonomously acquire a increasingly complete and accurate estimation of the states of others. In RoboCup for instance, almost all mid-size teams have resorted to omni-directional vision to achieve exactly that. So, beliefs needed to infer the intentions of others are becoming more complete and accurate, independent of communication. The arrow in the third graph in Figure 2 depicts this trend. More accurate state estimation can essentially replace communication. This is certainly not the case for explicit coordination, which will always fully rely on communication.

Furthermore, the third graph in Figure 2 clearly shows that implicit coordination with belief exchange achieves better performance with communication loss than explicit coordination alone. Instead of complete coordination failure in case of communication loss, there is a graceful decay, because a second system based on state estimation can still be used to estimate the intentions of others.

Summarizing, improvements in sensor and state estimation will allow implicit coordination to depend less and less on belief communication. This is necessary to simplify communication schemes, increase coordination robustness, and enable human-robot cooperation. This work proposes a step in this direction.

4 Related Work

The idea of cross team cooperation has some tradition within the RoboCup leagues. The most similar mixed team cooperation effort was the Azzurra Robot Team, a mid-size team from various Italian universities. Their focus was on explicit role assignment and communication-based coordination strategies among the field players [9].

Previous research on cooperation has focussed almost exclusively on explicit coordination [4]. On the other hand, work on implicit coordination usually assumes that all agents have access to a central and global representation of the world, which is enabled by simulation, as in [10], or global perception, as in the RoboCup small-size league [8,11]. In all these papers, teammates are not reasoned about explicitly, but are considered to be mere environment entities, that influence behavior in similar ways to obstacles or opponents.

In [5] the issue of low band-width communication in the simulation league is dealt with by *locker-room agreements*, in which players agree on assigning identification labels to certain formations. During the game, only these labels, instead of complete formations, must be communicated.

Most similar to our work is [12], in which robots in the legged-league also coordinate through implicit coordination which is based on representations which are completed through the communication of belief states. Communication is essential, and assumed to be flawless. It is not investigated how communication loss influences coordination. The utility measure is a sum of heuristic functions, which are represented as potential fields. Whereas our utility models are grounded in observed experience, and have a well-defined meaning (e.g. execution duration in seconds), these heuristic functions have no clear semantics. Therefore, customizing these functions to individual robots is difficult, as the semantics of and interactions between them are not fully understood. However, this customization is essential for achieving efficient coordination in a heterogeneous team with robots with different dynamics and capabilities.

5 Conclusion

In this paper, we have discussed the necessity for implicit coordination in domains in which communication is unreliable or impossible. Relying on intention communication will prevent multi-robot systems from being applied in these domains. We have presented a system that achieves implicit coordination by predicting the utility of itself and others, and adapt its actions to the predicted intentions of others. Knowing the states of others is essential, so belief states are communicated. We have shown that this approach is more robust than communicating intentions, and have argued that improvements in sensors and state estimation will allow implicit coordination to become increasingly independent of communication. We have applied the system to a ball approach task from robotic soccer, and demonstrated its performance in several experiments.

Our current work aims at learning temporal models that take opponent robots into account. Because the state space of this problem is much larger, more training examples are needed. We will also learn more complex models that take into account the player's roles, as well as strategic considerations.

Acknowledgments This work was partially funded by the Deutsche Forschungsgemeinschaft (German Research Foundation), in the SPP-1125, "Cooperating Teams of Mobile Robots in Dynamic Environments".

References

1. Dennett, D.: The Intentional Stance. MIT-Press (1987)
2. Stulp, F., Isik, M., Beetz, M.: Implicit coordination in robotic teams using learned prediction models. Accepted for the IEEE Intl. Conf. on Robotics and Automation (ICRA). (2006)
3. Utz, H., Stulp, F., Mühlenfeld, A.: Sharing belief in teams of heterogeneous robots. In: Proceedings of the International RoboCup Symposium (2004)
4. Gerkey, B.P., Matarić, M.J.: Multi-robot task allocation: Analyzing the complexity and optimality of key architectures. In: Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA). (2003)
5. Stone, P., Veloso, M.: Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. Artificial Intelligence **110** (1999) 241–273
6. Kraetzschmar, Mayer, Utz, et al: The Ulm Sparrows 2003. In: Proc. of the Intl. RoboCup Symposium (2004)
7. Beetz, M., Schmitt, T., Hanek, R., Buck, S., Stulp, F., Schröter, D., Radig, B.: The AGILO robot soccer team experience-based learning and probabilistic reasoning in autonomous robot control. Autonomous Robots **17** (2004) 55–77
8. Tews, A., Wyeth, G.: Thinking as one: Coordination of multiple mobile robots by shared representations. In: Intl. Conf. on Robotics and Systems (IROS). (2000)
9. Castelpietra, C. et al.: Coordination among heterogenous robotic soccer players. In: Proc. of the Intl. Conference on Intelligent Robots and Systems (IROS). (2000)
10. Sen, S., Sekaran, M., Hale, J.: Learning to coordinate without sharing information. In: Proceedings of the Twelfth National Conference on Artificial Intelligence. (1994)
11. Veloso, M., Stone, P., Bowlin, M.: Anticipation as a key for collaboration in a team of agents: A case study in robotic soccer. In: Proceedings of SPIE Sensor Fusion and Decentralized Control in Robotic Systems II. (1999)
12. Vail, D., Veloso, M.: Multi-robot dynamic role assignment and coordination through shared potential fields. In: Multi-Robot Systems. Kluwer (2003)