Beyond Distributed AI, Agent Teamwork in Ubiquitous Computing

Harry Chen
Department of CSEE
University of Maryland Baltimore County
1000 Hilltop Circle
Baltimore, MD 21250
hchen4@cs.umbc.edu

Tim Finin
Department of CSEE
University of Maryland Baltimore County
1000 Hilltop Circle
Baltimore, MD 21250
finin@cs.umbc.edu

ABSTRACT

Agent teamwork has been widely studied in the fields of Distributed AI. Much success has been achieved in defining theory to explain how agents should act together as a team and in developing programming frameworks to simulate team coordinations in a software environment. With the emergence of ubiquitous computing, agents are facing a new set of challeneges, such as perception limitations, planning limiation and device mobility, when they engage in teamwork activities.

1. INTRODUCTION

Teamwork has been widely studied in the fields of Distributed AI [3, 10, 4, 5, 7, 15]. Much success has been achieved in defining theoretical foundations for guiding agent cooperation and coordination in course of teamwork activities [15, 3, 9, 5] and in developing pragmatic framework for programming teamwork agents [11, 12]. Numerous teamwork domains that often involve highly complex group activities have been explored in software simulation environment environment, such as RoboCup97 soccer games [10, 2, 11, 13], teamwork in military helicopter flying simulation [11], cooking [5] and distributed industrial applications [6]. With the advent of ubiquitous computing technology, in particular in the fields of wireless communication and mobile computing, agent teamwork research is pushed to its limit, facing a new set of challenges that arise in the new computing paradigm.

Many of the existing teamwork research have assumed agents will conduct their group activites in a controllable environment, for example, in which agents can have complete knowledge and unlimited computational resources. As the near future ubiquitous computing environment is open and dynamic, it is questionable whether those assumptiosn will still be guaranteed to hold for agent teamwork in the new environment.

Because teamwork is of great importance to any group activities that demand flexible and robust coordination and cooperation, understanding the kinds of real-world teamwork challenges is crucial to the realization of the ubiquitous computing vision [14]. The objective of this paper is to raise some of the research issues that need to be addressed in order to enable ubiquitous agents to engage in cooperative teamwork activities.

In next section we will give a brief overview of the state of art teamwork research in Distributed AI. Although many different approaches have been proposed for agent teamwork, because our main concern is to discover issues that span across various aspects of agent group activities, we will concentrate our discussion on a teamwork theory that deals with teamwork from the beginning to the end – the Wooldridge-Jennings CPS model [15]. In Sec. 3 we will discuss some of the rising teamwork challenges: perception limitations, planning limitations, and device mobility. Remarks and future works are given at end of this paper.

2. TEAMWORK IN COOPERATIVE PROBLEM SOLVING

Teamwork is a concept in the human world. Often when a group of people work together to solve a problem or to achieve an objective (e.g. building a house), not only they can increase productivity and efficiency, but also they can solve a problem that otherwise cannot be solved by an individual alone (e.g. moving a heavy object). Motivated by the great benefits of working in teams, the Cooperative Problem Solving (CPS) research within the domain of Distributed AI is to define theories that would explain how a group of agents can be made to conduct joint actions, and to develop methods that would guide the coordination and cooperation processes of the team.

Before we begin our discussion on the Wooldridge-Jenning CPS model, in the next subsection we will introduce the Joint Intention theory, which sets the background knowledge for understanding the model that we will discuss next.

2.1 Joint Intentions

The key mental states that control agent behavior are intentions and joint intentions – the former define local asocial behavior, the latter control social behavior [1]. Based on these two concepts, Cohen and Levesque [3] developed the

Joint Intention model that specifies how a group of agents can jointly act together as an aggregated agent by sharing certain mental beliefs about the cooperative actions. At the core of their model is the notion of commitment. While individual agents can have beliefs about the world and the goals that they want to achieve, they can also form commitments for actions – a commitment to an action is defined as a persistent goal.

A persistent goal has certain implications on the beliefs of an agents. Once a persistent goal is adapted, the agent believes that the goal is current false, and it wants to the goal to be true. The agent holds these two beliefs until either the goal becomes true, or it will never be true, or the goal becomes irrelevant with respect to some other higher level goals.

The importance of a persistent goal are threefold [3]: 1) once the goal is adapted, an agent cannot drop them freely; 2) the agent must keep the goal at least until certain conditions arise; 3) the agent will try again and again to achieve the goal should the initial attempts fail. Another important property of the persistent goal is that an agent can form commitment to the act of other agents. For example, when agent i is committed to the goal that agent j should perform action φ, although agent i may not intend to perform the action φ itself, it could monitor and instruct the behaviors of agent j if necessary.

While a persistent goal, or an individual commitment, controls how an individual agent should behavior, a joint commitment guides how a group of agents should act together as an aggregated single agent. When a group of agents act together to achieve a common goal, there is a possibility that some agent may diverse its belief and drop its commitment to achieve the goal. Consequently, this could lead to the breakdown of the team effort. For example, when two people decide to drive together in a convoy, one person leads to show the other the way to go home. In the half way, without any communications, the person who leads the way can suddenly speed ahead because he wrongly concludes that the other person indeed knows her way home [3]. In order to prevent teamwork breakdowns that may cause by the belief divergence in individual agents, joint commitments extends the individual commitment to include an additional concept, namely the weak achievement goal, to ensure a mutual belief about the team goal is believed by all team members.

The definition of the weak achievement goal is the same as the persistent goal, except that the weak achievement goal has an additional condition. This condition specifies that when a member agent privately becomes to believe the state of a team goal is changed, and other members do not share a common belief about this change, then this member agent will adapt a new goal to inform the team members about the change. This new condition introduces a kind of obligation that individual agents should follow when their beliefs about the team goal have changes.

This obligation is key to the specification of joint persistent goal, which defines how a team of agents should act together to achieve a common goal. When a team of agents individually adapts a joint persistent goal, they mutually believe that the goal is current false; they mutually know that they all want the goal to be true eventually; it is true (and mutual knowledge) that until they come to mutually believe either that the goal is achieved or can never be achieved, or that it becomes irrelevant, they will mutually believe that they each have the goal as a weak achievement goal.

2.2 The Wooldridge-Jennings CPS Model

Based on the Joint Intention theory, Wooldridge and Jennings developed a formal model of cooperative problem solving, the Wooldridge-Jennings CPS model. A four-stages teamwork process is described in this model: 1) recognition, 2) team formation, 3) plan formation, and 4) team action [15]. Without giving a complete mathematical description of this model, in this section we will briefly describe each of the four stages and some of the assumptions that have been made in this model in order to lay the ground for the teamwork discussion[1].

Recognition

Teamwork in CPS begins when an agent in a multi-agent community has a goal, and recognizes the potential for cooperative action with respect to that goal [15]. In particular, the model defines two possible conditions which can trigger an agent to recognize the need for a teamwork effort. These two conditions are defined as the following:

> **Definition: (Potential for cooperation)** With respect to Agent i's goal φ, there is potential for cooperation if and only if there is some group of g such that i believes that g can jointly achieve φ; and either (1) i cannot achieve φ in isolation; or (2) i believes that for every action α that it cold perform which achieves φ, it has a goal of not performing α.

The above two conditions implies two possible alternatives. First, an agent in a group discover that it has a goal which itself cannot achieve in isolation, perhaps due to lack of resource, but it believes that cooperative action can achieve it. Second, an agent may have the resource to achieve the goal, but it does not want to use them, for example, it may believe that a cooperative effort can improve efficiency.

Team Formation

After an agent has recognized the need for cooperative actions, in this stage, the agent solicits assistance from other agents in the community. If the agent is successful in doing so, the result will be a group of agents having a joint commitment to collective action. Because an agent cannot always guarantee that a team can be successfully formed ever after it has tried its best effort possible, a notion of *attempt* is defined for describing an agent's attempt to bring about a team. The model of attempt is adapted from the work developed by Cohen-Levesque [15].

Based on the notion of attempt, an agent hopes to bring about an goal, the second stage team formation is defined as the following:

[1] The definition of each stage is adapted from [15]

Definition: (Team formation) An agent i, who believes that there is potential for cooperative action with respect to its goal φ, will eventually attempt to bring about in some group g, (that it believes can jointly achieve φ), a state wherein: (1) it is mutually believed in g that g can jointly achieve φ, and g are jointly committed to team action with respect to i's goal φ; or, failing that, to at least cause in g (2) the mutual belief that i has a goal of φ and the mutual belief that i believes g can jointly achieve φ.

Plan Formation

Plan formation is the third stage in the Wooldridge-Jennings CPS model. In this stage a team of agents, which has already established a common objective that they have committed to achieve, need to come to an agreement on the course of action that they will follow to achieve the team goal. Plan negotiation is the solution that is proposed in the model for agents to come into agreements.

Because negotiation in general is an extremely complex process [8], the Wooldridge-Jennings CPS model propose a simple assumption about the behavior of agents during plan negotiation. The assumption is that the agents will attempt to bring about their preferences. For example, if an agent has an objection to a particular plan, then the agent will try to prevent the plan from to be adapted by the team; if it has a preference for some plan, it will attempt to convince the team members to adapt the plan.

Team Action

If a team of agents is successful in coming to an agreement on a plan that they all follow, then the team will be coming into the forth stage of the model: team action. The model simply requires that the team to be jointly intend to some appropriate action by following the specification that is defined in the Joint Intention model. In addition, the agents are also required to adapt certain social conventions [7] for monitoring their teamwork progress. The definition for team action is defined as the following:

Definition: (Team Action) A group g are considered a team with respect to i's goal φ if and only if there is some action α, such that (1) α achieves φ; and (2) g have a joint intention of α, relative to i having a goal of φ.

3. SOME CHALLENGES IN UBIQUITOUS COMPUTING

However interesting the existing approach may be, we believe that with these theories alone are not sufficient to guarantee successful teamwork in a ubiquitous computing environment. Ubiquitous computing environment is an open environment. In such environment fundamental assumptions of the existing theory cannot always be guaranteed to persist through out the course of agent collaborations. In this section we will describe three of the problems that we have explored so far in our preliminary work, namely perception limitations, planning limitations, and device mobility.

3.1 Perception Limitations

The ability to perceive is of importance for agents to recognize a potential teamwork problem. For a ubiquitous agent to recognize a potential problem that require group attention is a complex task. First the agent needs have a comprehensive understanding of the problem domain. Second it needs to have adequate sensors that allow the agent to construct an accurate view of the problem domain. These two abilities can often be simplified in the defining theoretical foundations for teamwork. However, such simplifications cannot be safely assumed in a ubiquitous computing environment.

Ubiquitous agents often have limited perceptions to acquire knowledge about the world in which they live in because of computing resource limitations or physical obstacles in the world. For example, a coffeemaker agent in a meeting room may have limited perceptions to know if any participants are still interested to have more coffee; a camera agent in the room may be unable capture the movement of a person because a bookshelf is in the way. Perception limitations can greatly reduce an agent's ability to participate in teamwork. Unable to perceive the world could lead to the inability to recognize problems that indeed require group effects.

3.2 Planning Limitations

Planning is an intricate part of any teamwork. In order for team agents to conduct the kind of planning or re-planning that is described in the Wooldridge-Jennings CPS model, first the agents need to be capable of constantly acquiring comprehensive knowledge about the state of the world, which includes the state of the team goal, the commitments of individual team members, and the plans that the individual agents intend to follow.

Unfortunately, these requirements cannot always be met in a ubiquitous computing environment. For example,

- We cannot safely assume that all ubiquitous agents that participate in teamwork are capable of planning, perhaps due to lack of computing resources.
- We cannot safely assume that plans can be freely exchanged or negotiated between agents. Some agents may not be willing to share plans with other agents, or some agents may not be able to describe their action behaviors in term of expressive planning languages.
- We cannot safely assume all agents share a common plan representation or a common ontology for describing their plans and goals.

3.3 Device Mobility

Device mobility is another challenge that ubiquitous agents will face when teamwork behaviors are demanded. Device mobility refers to situations where computing devices are constantly moving in and out of a network environment. As we expect ubiquitous agents to be embedded in our everyday objects, such as cell phones, key chains, Smart Cards, and wearable-computers, their expectancy to participate in teamwork in any given network environment is completely depended on the duration of their presence in such environment. For example, after a projector agent is committed to engage in teamwork with other Audio/Video equipment agents to provide presentation services, because some person removes the projector from the room, the projector agent would no longer be able to perform the tasks that it has previously committed.

Mobility may also reduce the number of assumptions that agents can make about other prospective team members, particularly when they are in the stage of team formation and team action. In the case of team formation, for example, when an agent enters a community of which it has no knowledge about, the agent would not be able to initiate team formation protocols even though it has recognized a potential problem that requires group efforts. On the other hand, in the case of team action, some team members may be force to leave the network, that is to leave the team, which could lead to teamwork failures if the absent agent do not signal others about its departure.

4. REMARKS AND FUTURE WORKS

The openness of the near future ubiquitous computing presents a set of challenges for agents to engage in teamwork. Although we have described some of those challenging issues, by no means they make a complete list of all possible problems. We anticipate additional issues will arise as more mature technologies for constructing ubiquitous computing become available and more comprehensive real-world ubiquitous computing systems are prototyped.

As parts of our future works, we plan to further investigate how perception limitation, planning limitation and mobility will affect agent teamwork and to propose some pragmatic infrastructures to overcome some of those issues. In addition, we are interested to explore teamwork issues that are related to human social constraints, context awareness, security and privacy.

5. REFERENCES

[1] M. E. Bratman. *Intention, Plans, and Practical Reason*. Harvard University Press, 1987.

[2] S. Ch'ng and L. Padgham. From roles to teamwork: a framework and architecture. *Applied Artificial Intelligence Journal*, 1997.

[3] P. R. Cohen and H. J. Levesque. Teamwork. *Handbook of MultiAgent Systems*, 25(4):487–512, 1991.

[4] F. Dignum, B. Dunin-Keplicz, and R. Verbrugge. Agent theory for team formation by dialogue. In *Agent Theories, Architectures, and Languages*, pages 150–166, 2000.

[5] B. J. Grosz and S. Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269–357, 1996.

[6] N. R. Jennings. On being responsible. In E. Werner and Y. Demazeau, editors, *Decentralized A. I. 3*, pages 93–102. North-Holland, 1992.

[7] N. R. Jennings. Commitments and conventions: The foundation of coordination in multi-agent systems. *The Knowledge Engineering Review*, 8(3):223–250, 1993.

[8] S. Kraus, K. P. Sycara, and A. Evenchik. Reaching agreements through argumentation: A logical model and implementation. *Artificial Intelligence*, 104(1-2):1–69, 1998.

[9] H. J. Levesque and P. R. Cohen. On acting together. In *Proceedings of AAAI-90*, 1990.

[10] D. V. Pynadath, M. Tambe, N. Chauvat, and L. Cavedon. Toward team-oriented programming. In *Agent Theories, Architectures, and Languages*, pages 233–247, 1999.

[11] M. Tambe. Agent architectures for flexible, practical teamwork. In *National Conference on Artificial Intelligence (AAAI-97)*, 1997.

[12] M. Tambe. Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7:83–124, 1997.

[13] M. Veloso, P. Stone, and M. Bowling. Anticipation as a key for collaboration in a team of agents: A case study in robotic soccer, 1999.

[14] M. Weiser. The computer for the 21 st century. *Scientific America*, 265(3):94–104, 1991.

[15] M. Wooldridge and N. R. Jennings. Towards a theory of cooperative problem solving. In *Proc. Modelling Autonomous Agents in a Multi-Agent World (MAAMAW-94)*, pages 15–26, Odense, Denmark, 1994.