

Toward the Advancement of the Social Skills of Agents in Human-Agent Coordination

Matthew Johnson¹, Catholijn Jonker², Paul J. Feltovich¹, and Birna van Riemsdijk²

¹ Florida Institute for Human and Machine Cognition,
40 South Alcaniz, Pensacola, Florida, USA

²EEMCS, Delft University of Technology, Delft, The Netherlands
{mjohanson@ihmc.us, c.m.jonker@tudelft.nl, pfeltovich@ihmc.us, and
m.b.vanriemsdijk@tudelft.nl}

Abstract. Interdependent activity is the hallmark of social life (human and machine). In treatments of this kind of activity, such concepts as cooperation, collaboration, teamwork and coordination have been used indiscriminately, often applied as almost interchangeable. The problem of coordination is common to many fields, and its importance has been noted in each. Similarly, the problem of terminology has also been noted. This paper addresses this problem for one of the concepts, coordination, which is argued to have special status in the successful conduct of social activity. Traditional treatments are analyzed and critiqued, and an effort is made to systematize the application of this concept. Additionally, we suggest that increasing autonomy necessitates an increase in coordination. Moreover, it requires different and potentially more advanced types of coordination than currently employed. Hence, the second goal is to identify the potential requirements for new capabilities that must be addressed as autonomous systems take more responsibility, particularly in systems that include humans.

Keywords: coordination, cooperation, collaboration, interdependence

1 Introduction

Agent technology is ready for new levels of capability, especially increased flexibility and functionality in coordination. The current state of the art is geared towards planning/scheduling on the basis of complete information (which tasks, and when, requiring which type of resources), standard protocols (no explicit choice in which protocols to use in the current situation), and reasoning capabilities regarding the coordination process itself that are assumed standard (i.e., either hierarchic or in which all agents can take the lead and know what tasks are required). Interdependent activity is the hallmark of social life, whether human or extrapolated to human-agents "societies" [1]. In treatments of this kind of activity, such concepts as cooperation, collaboration, teamwork and coordination have been used indiscriminately, often

applied as almost interchangeable (but see, [2], for a systematic approach based on the loci of interdependency). This inconsistency of terminology has been noted by many [3-6]. The diversity of fields in which these concepts are relevant, as shown in Malone's [3] interdisciplinary study, is also a contributing factor to the non-standard terminology. In this paper, we address this problem for one of the concepts, coordination, which is argued to have special status in the successful conduct of social activity. This paper will pull together previous research to try to systematize the application of this concept in a way that can lead to a technical formalization, ideally adhered to across diverse applications. Traditional treatments are analyzed and critiqued, and an effort is made to systematize application.

We choose coordination as a starting point because it seems to be foundational to most of the other terms. Additionally, coordination connotes something more mechanical than terms like cooperation, collaboration or teamwork. Instead of providing an additional definition to the vast array of existing ones, we will adopt one of the more common definitions, one that has demonstrated some consistency across domains and is fairly inclusive. The definition is:

Coordination is managing dependencies between activities[3]

We argue that this definition captures the key feature of coordination, as it reflects a common theme across many current definitions from a variety of domains. However, it still leaves much room for specification.

This paper is concerned with coordination amongst heterogeneous agents of which some may be humans. Today's systems continue to increase in complexity, and researchers continue to push the boundaries of agent autonomy. March [7] sights decentralization as the factor that necessitates coordination. This is consistent with much of the previous work in fields such as Distributed Computing, Parallel Computing, Distributed AI, and Multi-Agent Systems where coordination has received significant attention. We suggest that increasing autonomy similarly necessitates an increase in coordination. Moreover, it requires *different and potentially more advanced* types of coordination. Hence, the second goal is to identify the potential requirements for new capabilities that must be addressed as autonomous systems take more responsibility, particularly in systems that include humans.

2 Characterizing Types of Dependency

Coordination is required largely because of interdependencies among activities [8]. Understanding the nature of the interdependencies involved in the coordination is an important part of determining the capability requirements of agents and designing a solution. There are many types of dependencies, and Malone [3] categorized some of the most common; use of shared resources, producer/consumer relationships, simultaneity of processes, and task/subtask roles. In Table 1 we regroup and extend this original listing of types of dependencies. While this listing is still non-exhaustive, we provide it as an embellishment of the original. In that regard, the

correlation of our terms to Malone's original terms is shown in italics. Some modifications are also made, as now discussed.

Dependency for resources has received much attention in the literature and is the same as Malone's "shared resource." We combine his producer/consumer and simultaneity constraints into temporal dependency, to capture all ordering dependencies, as well as the more constrained synchronization dependencies. Compositional dependency covers Malone's task/subtask category, and we provide additional examples of relations for this category. However, the main additions to the list are *information dependency* and *monitoring dependency*.

Types of Dependency	
Resource <i>(Shared Resource)</i>	Processes A and B must share the same limited resource Processes A and B must share the same limited space Processes A and B must be tasked to complete an objective <i>(Task Allocation)</i>
Mutual exclusion <i>(Simultaneity constraints)</i>	Process A cannot be performed if Process B is being performed
Mutual dependence <i>(Simultaneity constraints)</i>	Process A can only be performed at the same time as process B
Temporal	
None	Process A and B can be completed independently
Serial <i>(Producer/Consumer)</i>	Process A is not authorized to start until Process B is complete Process A is dependant on the results of Process B Process A uses the results of Process B as input
Parallel <i>(Simultaneity constraints)</i>	Processes A and B must start at the same time Processes A and B must end at the same time Processes A and B must be performed concurrently Processes A and B may be performed concurrently
Synchronization <i>(Synchronization)</i>	Process A must end at precisely the time Process B starts Process A must temporary intersect Process B at some precise time Process A must be going on for the entire period of process B Process A must start and finish while process B is going on
Compositional <i>(Task/subtask)</i>	Process A is a subpart of Process B Process A is composed of several processes including Process B Process A and B are part of a larger Process
Information	Process A must transport/communicate results to Process B <i>(Transfer)</i> Process A must output results compatible with Process B requirements <i>(Usability)</i> Process A is affected by information from Process B Process A may have information that could affect Process B
Monitoring	Process A must monitor the status of Process B

Table 1 Types of Dependencies

"Information dependency" can be a hard or soft constraint. Hard information dependency is a form of resource dependency; Process A must have information (e.g. output) from Process B in order to proceed. Soft information dependency is the more interesting concept in that, while not required, if provided it could potentially alter the behavior of the recipient. Some examples would be progress appraisals [9] ("I'm running late"), warnings ("Watch your step") and unexpected events ("It has started to rain"). This type of dependency can lead to much richer and more interesting types of coordination than have typically been considered.

"Monitoring dependency" is similar to information dependency, except the information is not communicated, but instead is just observed in the situation. Monitoring dependency requires both observability and interpretation of another, ongoing interdependent activity. Monitoring is a basic requirement for more advanced types of coordination such as error notification and supporting behavior.

The existence of different kinds of interdependency (including others we may have missed), some abstract and complex, will call for new types of coordination and impose new requirements.

3 Types of Coordination

With such a variety of dependency types, it is clear that there will be quite a variety in the types of coordination necessary to address them. The planning and scheduling community has done a lot to address resource, temporal and compositional dependencies. However, there has been less emphasis on the information and monitoring dependencies. These tend to shift the problem from coordinating the task work to considering the coordination process itself more directly. In other words, there is coordination of both content (what they intend to do) and process (physical and mental systems to carry out the former)[10]. The research areas that have focused on similar problems are Linguistics (e.g. [10]) and Natural Language (e.g. [11]). This change in the type of communication was noted by Allen:

"The only type of interactions supported by a typical state-of-the-art planning system (namely, adding a new course of action) handles less than 25% of the interactions. Much of the interaction was concerned with maintaining the communication (summarizing and clarifying, for example) or managing the collaboration (discussing the problem solving strategy)."[11]

In Table 2 we extend Malone's [3] table to include how different disciplines have analyzed coordination, adding Linguistics and Natural Language to broaden the picture of coordination as a process.

Type of Process	Techniques and Mechanisms of Coordination			
	Computer Science/ AI/ MAS	Economics	Organization / Teamwork Theory	Linguistics / Natural Language
Managing Resource Dependency	processor scheduling	analyses of markets, optimization techniques, utility and game theory	analyses of organizational structure, resource dependence	Turn taking
Managing Temporal Dependency	data flow modeling, Petri net analysis, synchronization techniques	PERT charts, critical path methods	meeting scheduling, process modeling	signaling
Managing Compositional Dependency	modularization techniques, Planning in AI, coalition formation techniques	economies of scale and scope	strategic planning, team formation techniques	dialog management
Managing Information Dependency	communication protocols, interface specification, process algebras		standardization, team training, competence identification	speech act, nonverbal communication techniques, joint action ladder
Managing Monitoring Dependency	listener model, callback model		feedback, supporting behavior	attention management

Table 2 How Different Disciplines Have Analyzed Coordination Processes

4 Coordination as a Process

Current planning solutions generally provide a recipe, which if followed and if everything works flawlessly, will be successful. This almost never holds, given any kind of realistic problem or work environment. Resource, temporal and composition dependencies lend themselves to this recipe style of coordination, where coordination takes on the appearance of a simple decision in one moment in time. While a single, point event may sometimes serve to coordinate a particular dependency, coordination in general is a process. The information and monitoring dependencies don't fit nicely into the recipe model and take on a form more common in Linguistics, Natural Language and Joint Intention Theory. Much treatment of coordination in the literature involves preplanned coordination, without any responsibility for such things as acknowledgement, monitoring, feedback or the maintenance of common ground (adequate common understanding). While not essential to handle all dependencies, these concepts play a vital role in the more sophisticated types of coordination such as progress appraisal and supporting behavior. Klein's [4] discussion of Clark's

construct of Common Ground [12] differentiates several activities to support common ground; preparations, sustaining activities, updating, monitoring, detecting anomalies and repairing the loss of common ground. These highlight the *dynamic* nature of coordination and point to several of the different types of coordination activities required to support high level and effective coordination. Sustaining, updating and repairing are all types of information dependency. Inspecting and detecting anomalies are clearly monitoring dependencies.

One of the conceptual strengths of the work in Linguistics and Natural Language work is that it is inherently bi-direction. These fields consider the various aspects of sending and receiving information and the need to close the communication loop. In fact, Clark has suggested that unilateral adaptation is not coordination at all [5]. As coordination moves beyond simply signaling of resource availability and task assignment, bi-directional communication will need to play a greater role.

6 Challenges in Coordination

There are many challenges in coordination, and some that have been cited in the literature (e.g. Klein et al.'s 10 Challenges [13]). We will focus on the ones directly related to the coordination issues associated with information dependency and monitoring dependency. The main challenges are:

- 1) Determining when coordination is needed
- 2) Establishing of coordination protocols
- 3) Recognizing relevance
- 4) Determining coordination costs and benefits

Determining when coordination is needed is itself a challenge. For the resource, temporal and composition interdependency the coordination requirements might be clear, but how can an agent determine when some information it has might be appropriate for another agent? How can an agent query generally for information on topics that it is interested in? When is it appropriate to monitor an activity and at what level does the monitoring need to take place?

Currently coordination protocols tend to be implied. They are built into the language, built into middleware, or built into agent intelligence. "In closed situations, the participants know their roles, rights, duties and potential joint purposes. All they need is to establish is the joint purpose for the occasion. That they can do with a routine procedure. Many open situations must establish their roles, rights, and obligations [10]." Everything is about protocols, but how one comes to a particular protocol is implicit. Setting up the cooperation requires a different type of coordination. No matter what, in the end there has to be a common ground where the agents that intend to cooperate understand each other. That common ground should be such that at least some intention recognition is possible, either through the observation of each other's actions, or through communication, probably through a combination of both. From that common ground a number of actions and communicative acts are undertaken, leading to increased understanding and progress in the coordination. The

processes identified at the start of this section are typically intertwined. For example, based on the increased understanding among the parties, the costs of coordination become clearer, whereas the expected costs of the coordination form an incentive to either continue or disrupt the explorative phase of getting to know and understand each other.

Recognition and interpretation are two areas still struggling in the machine world, while we can only marginally recognize known items and interpret only the most basic signal in the environment. Determining general relevance is still a far off dream. Even conveying what needs to be recognized is difficult, though it may be possible to leverage work with the semantic web to formally specify information needs.

All coordination comes at a cost, though not all provides a benefit. Determining the cost of coordination is still an open problem across many fields. The agent community has the potential, given its strong computational and formalization background, to lead the way towards measuring coordination cost.

7 Analysis of Requirements for Coordination

Coordination requirements can impose capabilities requirements on the participants. For example, the monitoring dependency requires the agent to be able to observe the dependant activity and interpret its status. Communication requirements also change drastically based on the dependencies, mechanisms and type of coordination. Rasmussen [14] suggests that a change in the level of abstraction involves a shift in concepts and structure for representation, as well as a change in suitable information. We propose increasing autonomy as the next level higher, and thus requiring another leap in information requirement. March points out that when communications are primitive – relative to needs – so will be the system of coordination [7]. We now turn to addressing directly the requirements and capabilities necessary to support three specific exemplar of types of coordination associated with information and monitoring dependence; progress appraisal, supporting behavior, and common ground.

7.1 Progress Appraisal

Progress appraisal requires the ability to understand how well one's current activity is proceeding toward successful completion. It also requires understanding of with whom to share this information. Progress appraisal introduces four main challenges: determining when progress appraisals are appropriate, establishing the expectation of who will be appraised of what and when, the capability to recognize ones progress, and considering the cost and benefit of the update on the greater, overall activity.

7.2 Supporting Behavior

Supporting behavior is even more complex, involving both the monitoring requirements, the judgment to know when to support another activity, and the capability to perform the other agent's activity on its behalf. There are various levels of support that include merely alerting the other party about something, to providing advice, to assisting in the other parties activity, to performing the complete activity for the other party. Depending on the type of support, it may actually require a capabilities overlap between the parties involved. Many questions still remain, such as how can a supporting behavior be requested? Is supporting behavior inherently implied in all teams? What coordination requirements apply during execution of a supporting behavior to avoid confusion and conflict? Poorly coordinated automated support has been the cause of accidents in the aviation community.

7.3 Common Ground

Clark suggests that Common Ground is a prerequisite for coordination (in joint actions) [5]. This issue associated with Common Ground is covered thoroughly in [4]. In short, some degree of common ambitions, interpretations, presuppositions, and knowledge are requisite for smoothly working together. Still the questions remain, how to determine the “ground” that should be common? How is it established? How can we recognize breakdowns or even better divergence to preempt the breakdowns?

8 Discussion

While it is interesting to ponder what an ideal coordination design might be, it may be more directly relevant to consider what coordination options are *available*, given the limited capabilities of today’s computational systems to observe, interpret and associate. Larson [15] suggests there are three broad team objectives; problem solving, creative and tactical. Problem solving focuses on issues and requires sound judgment to grow out of a careful, investigative, and mutually-informing process. Creativeness focuses on exploring possibilities and alternatives and requires innovation. Tactical focuses on execution and requires directive, highly focused tasks, role clarity, well-defined operational standards, and accuracy. Activities can exhibit all three objectives but tend to have one dominant (loaded) high level objective. Humans can play in all three areas reasonably well, but automation has limitations. Automation has demonstrated very little potential for creativity. Although automation can support problem solving, it lacks many of the abilities to lead. Tactical applications seem most relevant in conditions where automation is most likely to be beneficial, thus the importance of coordination. Figure 1 shows a hypothetical visualization of the problem space. We are not suggesting that automation cannot play a role in the other areas, but when it encounters a new or unexpected situation requiring problem solving or creativeness, it is likely time to make sure that humans are in the lead.

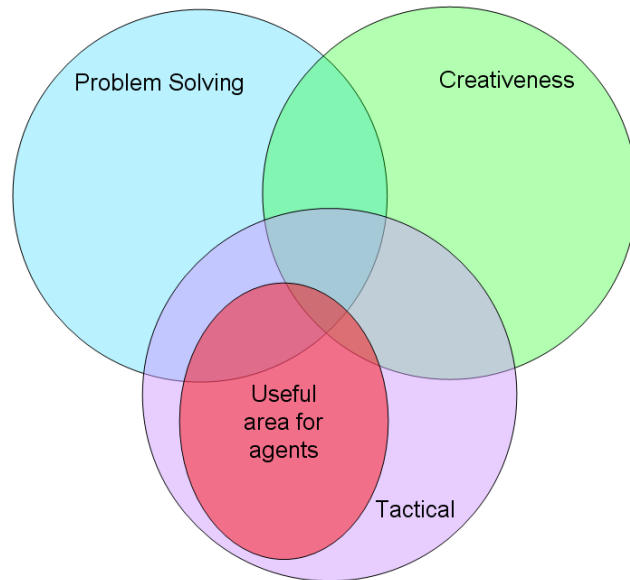


Figure 1 Suggested Useful Area for Agent Applications

As the agent community moves toward greater and greater autonomy of the agents, several researchers have expressed concerns. Norman states that “the danger [of intelligent agents] comes when agents start wresting away control, doing things behind your back, making decisions on your behalf, taking actions and, in general, taking over.”[16] One solution to this concern may be coordination. The human’s need for trust can be thought of as an information dependency, requiring the system to provide sufficient information to assuage the human’s fears.

Additionally, as agents systems perform in increasing new and difficult tasks, coordination will be as important as task competence. Christoffersen states that “the ability of a joint system to perform effectively in the face of difficult problems depends intimately on the ability of the human and machine agents to coordinate and capitalize upon the unique abilities and information to which each agent has access [17].” To do this, participants will need to be much more flexible, and coordination is one tool that can potentially help. Agents have not become integral members of teams partly because current software agents lack the dynamism and adaptiveness [18]. A precise of the issue can be obtained by applying Salas’ statement on team members to agents systems:

“It is not sufficient that members be technical experts – they must also be experts in the social interactions that lead to adaptive coordination action (i.e. teamwork)” [19]

9 Summary

We have attempted to present previous research from diverse fields to try to systematize the application of the concept of coordination in a way that can lead to a technical formalization, ideally adhered to across diverse applications. Starting with a candidate definition of coordination (managing dependencies between activities), we have extended Malone's [3] listing of dependencies to include *information dependency* and *monitoring dependency*, which will require different and potentially more advanced types of coordination. We propose these two dependencies emphasize that coordination is a process. They also bring to the forefront four main challenges; determining when coordination is needed, establishing of coordination protocols, recognizing relevance, and determining coordination costs and benefits.

As today's systems increase in complexity and we push the boundaries of autonomy, the role coordination will likewise grow in complexity and importance. These requirements for new coordination capabilities must be addressed as autonomous systems take more responsibility, particularly in systems that include humans.

References

1. Fessler, D.M.T., From appeasement to conformity: Evolutionary and cultural perspectives on shame, competition, and cooperation, in *The self-conscious emotions: Theory and research*, J.L. Tracy, R.W. Robins, and J.P. Tangney, Editors. 2007 New York: Guilford Press. p. 174-193
2. Feltovich, P.J., et al. Toward an Ontology of Regulation: Socially-Based Support for Coordination in Human and Machine Joint Activity. in *Engineering Societies in the Agents World VII, 7th International Workshop, ESAW 2006*. 2006. Dublin, Ireland.
3. Malone, T.W. and K. Crowston, The interdisciplinary study of coordination. *ACM Comput. Surv.*, 1994. 26(1): p. 87-119.
4. Klein, G., et al., Common Ground and Coordination in Joint Activity, in *Organizational Simulation*, K.R.B. William B. Rouse, Editor. 2005. p. 139-184.
5. McIntyre, R.M. and E. Salas, Measuring and managing for team performance: Emerging principles from complex environments., in *Team Effectiveness and Decision Making in Organizations*, R. Guzzo and E. Salas, Editors. 1995, Jossey-Bass.: San Francisco, CA. p. 149-203.
6. Cannon-Bowers, J.A., et al., Defining competencies and establishing team training requirements, R. Guzzo and E. Salas, Editors. 1995 Jossey Bass: San Francisco. p. 33 - 380.
7. March, J.G., H.A. Simon, and H.S. Guetzkow, *Organizations*. 2nd ed. 1993, Cambridge, Mass., USA: Blackwell. ix, 287 p.
8. Feltovich, P.J., et al. Toward an Ontology of Regulation: Socially-Based Support for Coordination in Human and Machine Joint Activity. in *Pre-Proceedings of the Engineering Societies in the Agent's World 06 (ESAW06)*. 2007. Athens, Greece.
9. Feltovich, P.J., et al. Progress Appraisal as a Challenging Element of Coordination in Human and Machine Joint Activity. in *Engineering Societies in the Agents World VIII, 8th International Workshop, ESAW 2007*. 2007. Athens, Greece.
10. Clark, H.H., *Using language*. 1996, Cambridge [England] ; New York: Cambridge University Press. xi, 432 p.

11. Allen, J.E., C.I. Guinn, and E. Horvitz, Mixed-Initiative Interaction. IEEE Intelligent Systems, 1999. 14(5): p. 14-23.
12. Clark, H.H. and S. Brennan, Grounding in communication. Perspectives on Socially Shared Cognition, 1991: p. 127-149.
13. Klein, G., et al., Ten Challenges for Making Automation a "Team Player" in Joint Human-Agent Activity. IEEE Intelligent Systems, 2004. 19(6): p. 91-95.
14. Rasmussen, J., Information Processing and Human-Machine Interaction: An Approach to Cognitive Engineering. 1986: Elsevier Science Inc. 228.
15. Larson, C.E. and F.M.J. LaFasto, Teamwork : what must go right, what can go wrong. Sage series in interpersonal communication. 1989, Newbury Park, Calif.: SAGE Publications. 150 p.
16. Norman, D.A., The invisible computer : why good products can fail, the personal computer is so complex, and information appliances are the solution. 1998, Cambridge, Mass.: MIT Press. xii, 302 p.
17. Christoffersen, K., et al. How to Make Automated Systems Team Players. 2002.
18. Sycara, K. and G. Sukthankar, Literature Review of Teamwork Models. 2006, Robotics Institute, Carnegie Mellon University: Pittsburgh, PA.
19. Salas, E., et al., The making of a dream team: when expert teams do best, in The Cambridge Handbook of Expertise, K.A. Ericsson, et al., Editors. 2006