10th INTERNATIONAL COMMAND AND CONTROL RESEARCH AND TECHNOLOGY SYMPOSIUM THE FUTURE OF COMMAND AND CONTROL

Title of Paper : A Simplified Taxonomy of Command and Control Structures for Robot Teams

Topic : Decision Making and Cognitive Analysis

Name of Author(s) : Dr Robert Grabowski, Dr Alan Christiansen

Point of Contact : Dr Robert Grabowski

Name of Organization : The MITRE Corporation

Complete Address : 7515 Colshire Drive, McLean, VA 22102-7508

Telephone /Fax : 703-883-3490 / 703-883-6857

E-mail Address : rgrabowski@mitre.org

A Simplified Taxonomy of Command and Control Structures for Robot Teams

Dr. Robert Grabowski **Dr. Alan Christiansen**

The MITRE Corporation
7515 Colshire Drive
McLean, Virginia 22102
{rgrabowski, adc}@mitre.org

Abstract —***For a team to be effective, it must coordinate and cooperate in some fashion. Often this ability is a direct function of the way the team is put together. Selecting the right architecture is driven by many factors including the skill of the individual, ability to communicate, availability of resources, and the size of the team. In this paper, we examine the issue of command and control from the perspective of coordinating a team of robots. We look at the existing field of robotics and select several representative teams that cover the spectrum from teleoperation to peer to peer interaction. We identify and examine the mechanisms that facilitate coordination and define a taxonomy that describes the coordination complexity. Finally we look at the role of the human as he interacts with the team and how this interaction influences the coordination between members of the team.***

Keywords-component; robot taxonomy, robot teams, command and control

I. INTRODUCTION

Congress has established goals that "by 2010, one-third of the operational deep strike aircraft of the Armed Forces are unmanned; and by 2015, one-third of the operational ground combat vehicles of the Armed Forces are unmanned."[1] The coming age of unmanned military systems raises many questions about how robots will work together as teams, and how humans and robots will work together within teams.

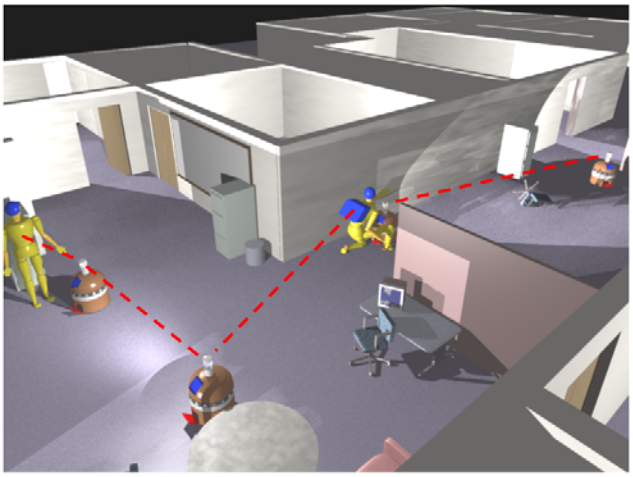

Figure 1. Cooperative Human Robot Teams – the future vision of technology has humans and robots working closely together. Above is a mission we are developing at MITRE where robots interact directly with humans in a building takedown scenario.

Robots of today are diverse and span many different domains. Multiple robots can coordinate action and sensing to extend a collection of individual entities to a single, cohesive group. A coordinated team can fuse information from a variety of different platforms to build global maps, exploit the proximity of neighbors to localize, assist one another to manipulate objects and coordinate sensing and action to monitor dynamic events such as monitoring and tracking of intruders.

In this paper we analyze the underlying mechanisms that robot teams employ to coordinate action and intent. We compare these mechanisms to one another to propose a simple taxonomy of coordination. Development of a robot coordination taxonomy serves two useful purposes. First, it exposes the current state and capabilities of existing robotic systems. Second, it forces us to more closely analyze the fundamental elements of coordination. While robots are modeled after biological systems, they are still machines. However, they often share similar mechanisms of coordination. By analyzing the mechanisms exploited in the robotic domain, we gain some insight into the command and control of multiple agents with disparate resources, skills, and degrees of autonomy.

To establish this taxonomy, we will closely examine several robot team candidates and determine the mechanisms that they employ to facilitate coordination. Some forms of coordination are implicit and occur simply by working in close proximity of one another or by exchanging information through a third party. Conversely, other forms are more explicit and involve direct peer-to-peer communications. We will examine the properties of a chosen architecture with respect to its intended application, and then analyze their strengths and weaknesses.

[1] Senate Armed Services Committee Bill S.2549, National Defense Authorization Act for FY 2001

To date, we have identified eight representative robot teams that characterize the spectrum of coordination complexity. Within this spectrum we have identified three distinct mechanisms that separate these architectures.

As part of the analysis, we also examine the role of the human in a robot team. For the most part, the role of the human has not been considered when talking about collections of robots. We have identified several classes of human involvement and discuss the impact on the taxonomy.

In Section II, we define the fundamental concept of team with respect to multiple robots. We then discuss the different factors that facilitate coordination and how they impact complexity. In Section III we develop a simple, single spectrum of coordination that organizes teams by the complexity of the mechanisms that facilitate coordination. Across this spectrum, we enumerate several representative robot teams. Finally in Section IV we classify the different roles of the human in coordination and control of robot teams.

II. Elements of Coordination

At its most fundamental level, a team can be described as a collection of *n* agents with individual capabilities, connected by a communications structure, and moderated by a communications protocol that interact to achieve a common goal. The agents may be either human or machines. The machines can either be mobile or immobile. While most researchers only consider the robotic teams as a collection of mobile machines, we will consider the role of the human in the loop as well as any immobile resources (e.g., central broker or graphical user interface) that facilitate coordination. Agent capabilities include sensing, action, memory, and reasoning. Actions modify the physical environment through mobility and manipulation. Mobility primarily defines the domains in which a robot can operate but those domains influence the degree of coordination essential to the team. Reasoning refers to the individual's ability to relate sensing and action to an understanding of the world and other agents.

A communications structure is defined by a collection of links between agents. Each link is characterized by its bandwidth and reliability, both of which can vary over time and distance. A communications structure also defines the underlying architecture of these links and falls across two camps. Centralized communications architectures route information and control through a central authority while a distributed architecture relies on peer-to-peer interaction.

A coordination protocol defines the language of how robots interact. Interaction generally takes the form of messages and includes commands, requests and reports of state. At any instance in time, a team can communicate its state at one of three levels. At the lowest level, state describes properties of an agent such as its location or its current sensor readings. At the middle level, state describes the status of an individual such as its current task, the status of that task, or estimated time to completion. At the highest level, state describes the current goal of the group or its progress toward completing the team's mission.

Finally, any discussion of team coordination must also consider the composition of the team—including the number of agents and the types of robots. Homogenous teams are composed of a collection of identical robots, enabling robots who know their own capabilities to likewise know the capabilities of their teammates. On the other hand, a heterogeneous team is composed of robots of different types. Coordination must not only account for the availability of a particular individual but the capabilities of that robot as well.

Each of these factors impacts the ability and requirements of a team with respect to coordination and impacts the available choices of architecture. For example, the coordination mechanisms employed by a team of five robots may not easily support a team of five hundred robots. Similarly, a team composed of a set of identical copies may have a different set of advantages over a team of specialized agents. It is easy to see that the space of all possible combinations and factors that must be considered when dealing with multiple robots is quite large.

A few researchers have attempted to encapsulate the properties of multi-robot teams in a systematic fashion. Dudek defines seven major axes that characterize multirobot systems that include, team size, communications range, topology and bandwidth, reconfigurability, processing ability and composition [2]. Each axis is discretized according to its degree into a subset of classes. For example, team size is divided into

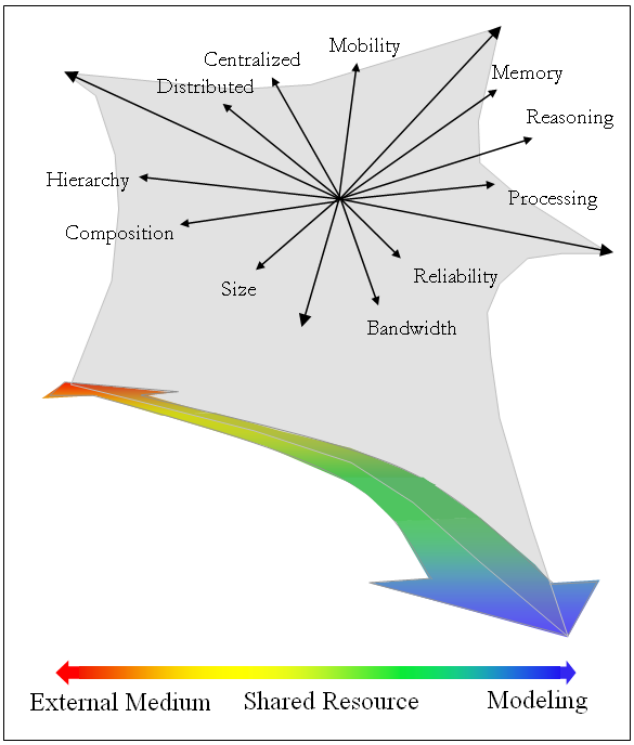

Figure 2. Simplified Spectrum – The factors that determine the complexity of coordination for robot teams make up a large, multi-dimensional space. To allow us to compare and evaluate the mechanisms of cooperation, we collapse these axes to make a simplified spectrum that represents the degree of coordination complexity.

the following classes: one, two, many, and infinite. However, while discretization provides an implicit degree of ranking and allows robots to be more easily clustered, the sparseness of existing robot teams and the large number of combinations still make it hard to compare and evaluate robot teams that fall across different axes.

Iocchi takes a more specific approach and looks at the classification of multirobot systems with respect to coordination [8]. His main focus is to investigate the division between reactive and deliberative control. He too defines several axes that include composition, awareness, reactivity, etc. However, he is left with a similarly large combination of choices that are difficult to directly compare to one another.

III. A Simple Spectrum

Realistically, any true analysis of robot coordination will require a consideration of many discrete and continuous variables. The resulting analysis would leave us with a complex structure that makes it difficult to compare different robot teams with different purposes operating in different domains. Instead, we take a different approach. We look at the multitude of existing robot teams and organize them by the complexity of their coordination mechanism. Instead of generating multi-dimensional axes, we fit these robot teams onto a single axis according to the degree of cooperation they exhibit (Figure 2). In a sense, we have collapsed the entire multi-dimensional space onto a more manageable, single-axis spectrum. Even through this simple exercise, we see a rough taxonomy begin to emerge. We note that teams naturally divide into more or less three classes. At the lower end of the spectrum are teams that rely on an external medium, such as properties of the environment to provide a natural means of coordination. In the middle of the spectrum are teams that rely on an explicit central entity to arbitrate conflict and coordinate action. At the higher end of the spectrum are robots that internally model the environment, mission, or fellow teammates to aid in coordination.

We have chosen eight representative teams to investigate this taxonomy. Each represents an increasingly more complex system with respect to coordination. That does not imply that the robots at the higher end of the spectrum are more complex or that the individual capabilities are more advanced. A team of simple robots can have a fairly sophisticated degree of coordination whereas a team of highly competent individual robots may choose a simpler scheme. Instead we look at complexity in terms of the ability of one robot to share information, coordinate action or convey intent to its teammates.

A. *Lower Spectrum*

At the lowest end of the coordination spectrum are teams that rely on an external force or medium to coordinate action and manage conflict. In these systems, little or no explicit, mechanism exists for communication between robots. Essentially, each robot acts as if it is alone in the world. However, even when communications are not explicit, a team can act as a somewhat cohesive unit.

The greatest utility of robot teams at this end of the spectrum is derived from multiplicity. For the most part, these loosely-coupled teams are essentially composed of multiple copies of the same individual. Multiple robots, as opposed to a single robot, can perform actions concurrently—often completing a given task in a shorter period of time. Multiplicity can also provide a measure of resiliency. Loss of an individual does not necessarily mean failure of the entire mission. On the logistics side (programming and design), systems that do not require coordination are easier to construct and field. Literally,

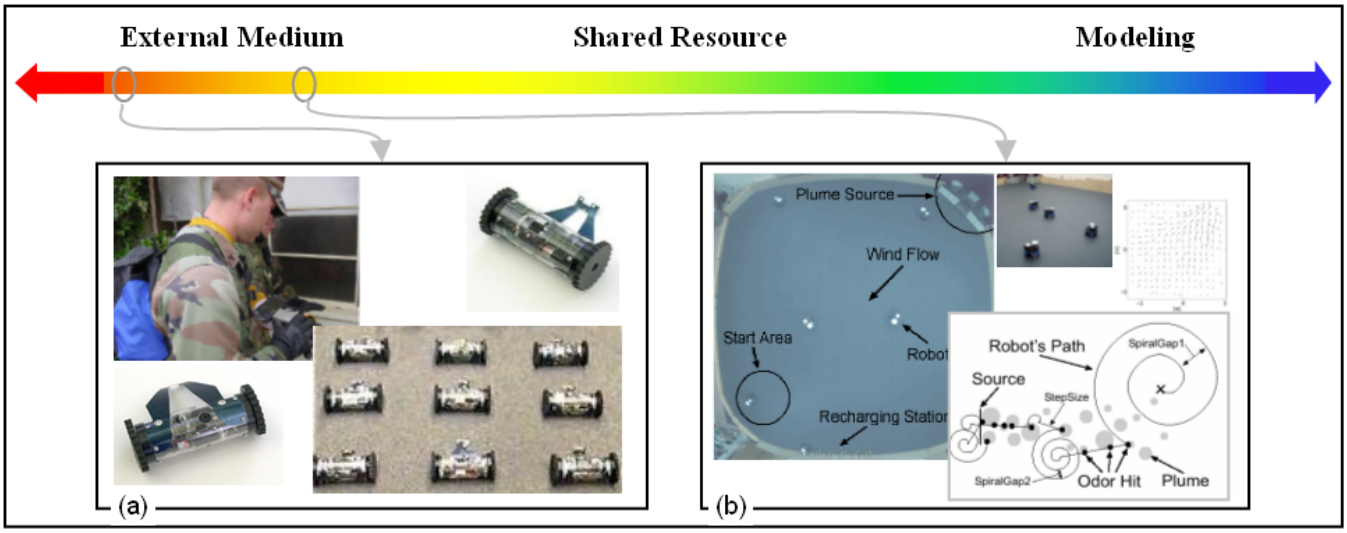

Figure 3. Lower portion of the spectrum – teams that coordinate through an external medium. **a)** The Individual Perspective - Each robot only knows about itself and does not communicate with other team members. Minnesota Scouts are small, two wheeled robots that have a small video transmitter for video teleoperation. Each robot provides a video link to a single user. Interference is mitigated by deploying multiple robots in different areas. **b)** Environmentally directed teams - The nature of the environment provides an implicit mechanism for coordination. Plume Robots use internal chemical, heat or humidity sensors to detect the gradient of a plume. Each robot moves along the gradient towards the source. The instantaneous position of the group as a whole captures the information about the plume.

this sort of team is simply a sum of its component parts.

However, the total lack of coordination can be a big disadvantage as well. When multiple robots work close to one another there is an increased risk that they might interfere with one another. Two robots may compete for the same physical space or attempt to access the same resources. Without an external coordination mechanism, unintended interactions can rapidly degrade the performance of a team effectively wiping out any advantages gained from duplicity.

One example of a team that falls on this end of the coordination spectrum is given by the Minnesota Scouts [8] (Figure 3a). Minnesota Scouts are small robot platforms that were developed to provide distributed visual sensing to the war fighter. Each robot has the form factor of a soda can placed on its side (Figure 3a). Wheels on either end provide mobility with a pliable fin extending from the center to prevent the robot from rolling while driving. A small video camera with a wireless transmitter is embedded in the center of the body to provide remote vision sensing. Real-time video images are relayed to a mobile operator providing a valuable remote view of that space. Through this feedback, the operator can direct the robot where to move to gain a better vantage. The main task of these robots is to provide remote video to the operator in the field so coordination between robots is not essential. To effectively operate several of these robots in the same space would require some coordination at the operator level. A more manageable approach is to deploy the robots in different areas.

A direct communications link between robots is not always necessary to achieve a coordinated effort. In some applications, the properties of the environment can be exploited to provide a natural mechanism for coordinating movement and arbitrating conflict. For example, robots can be equipped with internal chemical, heat or humidity sensors that allow them to follow the gradient of a plume [5] (Figure 3b). Each robot is guided by its own sensors. However, the entire team moves collectively towards the greater concentration. Robots of this organization do not share explicit information with one another and except for direct teleoperation they do not share information with a central entity. The instantaneous position of the group as a whole captures the information about the plume.

B. Mid Spectrum

In the middle of the spectrum are robots that begin to coordinate indirectly through a third party or shared resource (Figure 4). A common organization technique is for the individuals to pass sensor information to a central interface where that information can be fused into a common representation such as an occupancy map or topological graph. The fusion process can be automated, but is often performed by the human operator, producing a global view of the status of the team. The central authority also has the primary responsibility of coordinating the information to minimize interference and maximize cooperation. In turn, the central authority evaluates that information to either direct or guide the individual. While information is being collected from multiple robots, the robots themselves do not have explicit knowledge of teammates nor do they communicate directly with one another. Instead, they coordinate by acting on the combined knowledge of team.

This form of coordination exploits the multiple, disparate views that a team of robots can generate without dealing with the difficulties of coordinating communications between them. Robots need not be aware of or build models of their teammates. They act as if they are alone in the world and just happen to gain information derived from other sources.

The disadvantage of centralized coordination is that the central coordinator becomes the weakest point in the system. If something happens to the central point, either through physical loss or degradation of the communications channel, the entire notion of the team dissolves. Moreover, a centralized coordination scheme does not scale well. Eventually, a single central process simply cannot handle the quantity of information being passed to it. Imagine the processing load of hundreds of robots passing sensor information to a single central map builder. As more robots are added to the team, the resources can become saturated and delays can become significant. However, for small systems that can handle the load, the reduction in complexity as a whole outweighs the sensitivity to using a central coordinator.

An example application that benefits from a central arbitrator is search and rescue. In search and rescue, time is the most valuable commodity and naturally benefits from multiple entities working in concert. However, to be effective, the information often needs to be centralized to maximize coordination of resources. Moreover, robots can also exploit this centralized information to provide a level of autonomy relieving the burden placed on the operators. To this end, MITRE has developed a team of robots that perform a coordinated search and rescue task [3] (Figure 4a). Each robot is equipped with a suite of sensors including a laser range finder, a pan-tilt camera and a pyro detector. The laser rangefinder gives the robot the ability to build detailed occupancy maps while the pyro detector can detect the heat from warm bodies. Cameras allow an optional direct link back to the operator when a potential victim has been identified. While these robots support direct teleoperation, each robot is also able to exhibit a degree of autonomy including self-localization, obstacle avoidance and exploration. The robot utilizes its own sensor information to build a local map to aid in obstacle avoidance and near range navigation. However, the robot also passes sensor information back to a central authority which fuses the information with data from other robots into a larger global map. In turn, the map is processed as a whole and used to guide the robot to new areas of exploration—sometimes beyond what the individual can sense. The same map also provides needed information to a human manager who can redirect the team towards possible victims. However, while the information itself is being coordinated, these robots are unaware of the presence or existence of teammates. They each work to maximize their utility based on the collective information stored in the team map.

Higher degrees of coordination are possible when robots start reacting to one another as individuals. The most fundamental understanding one robot can have about another is its position in space. With this simple knowledge, robots can

begin to coordinate both action and sensing. Knowledge of teammates can be sensed, explicitly communicated or inferred through a central entity. However, while robots have the inherent ability to detect the presence of others through sensing, some form of communications is still required to disambiguate a fellow robot from background clutter.

Howard shows how simple knowledge of teammates can be used to direct the motions of individuals but ultimately influence the entire team [6] (Figure 4b). The underlying mission of these robots is exploration. However, as previously discussed, multiple robots in close proximity can hamper the movements of one another. An exploring robot attempting to reach an open space often finds another robot in its path forcing it to find a different route. However, since the other robots are reacting the same way, the new route may quickly become blocked. Without some arbitration, the entire team may grid lock. To mitigate this reaction, Howard describes a simple coordination mechanism that reduces this effect. In addition to sharing a central map, each robot exploits its awareness of all the other robots around it. Robots at the edges start to move immediately towards open space. Robots in the middle, calculate a heading vector that would best maximizes the distance to all of its neighbors. The overall effect is that the robots begin to spread out. The result is a coordinated action without explicitly conveying intent.

However, while this method keeps the robots from interfering with one another in a physical sense, it does not prevent two robots from competing for the same open space eventually leading to future conflict and lower efficiency. A more advanced variant of this type of coordination allows the robots to arbitrate conflict through management of the external resource—in this case, open space. For example, Simmons describes an exploration task that is performed simultaneously by multiple robots [9] (Figure 4c). To be effective, these robots initially seek to explore spaces apart from each other to avoid potential interference. This naïve approach would have the robots spread out to minimize potential conflict. However, such a policy can lead to excessive movement of teammates as they move to seek separation. Moreover, as robots move farther away, they lose advantages gained by proximity such as the ability to assist each other in localization. Instead, Simmons describes how robots can arbitrate potential conflict by evaluating the cost of competing for a given space. This analysis not only accounts for the cost of movement but also takes into account for whether that space is also being claimed by another. If that space is claimed, the cost is devalued and the robot choices another more appealing option.

With this coordination mechanism, robots are able to explore space in an efficient manner and still maintain close proximity to one another. Even though the robots are effectively interacting with one another, they do not explicitly communicate nor are they aware of the abilities or roles of their neighbors.

C. Upper Spectrum

At the higher end of the spectrum are teams that begin to model their environment and neighbors to predict the effects of interaction (Figure 5).

The most decoupled method of modeling is to assign specialized roles to each robot. That is, the role a robot plays in a task is an inherent property of that robot type. Specialization naturally partitions complexity as each robot understands its role in the task implicitly. In some domains, this partitioning also maps to a spatial element. That is, robots position themselves in a particular location or area as part of their role. One popular example is robotic soccer teams where two robot

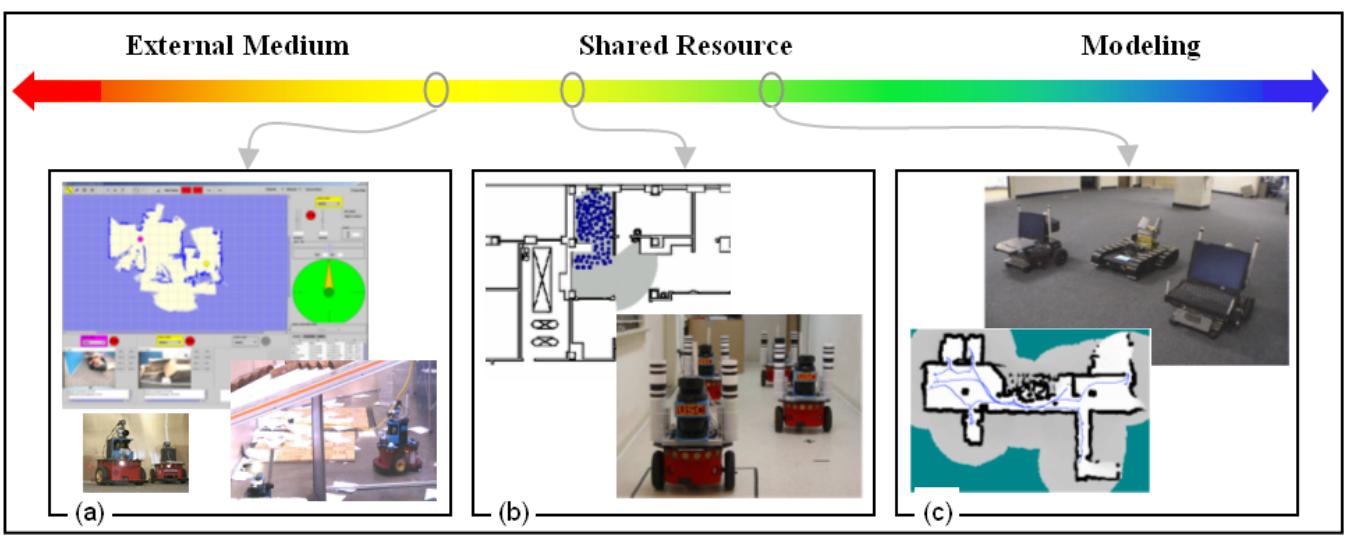

Figure 4. Middle portion of spectrum - Centrally Shared Resources – Robots share information with a central entity but are generally unaware of others in the group. **c)** MITRE Search and Rescue. Multiple robots are deployed to find victims in a search and rescue scenario. Robots use a global map to determine where to explore next but are generally unaware of their teammates. **a)** Robots at USC – robots are aware of the position of teammates. Each converts this range information into a force vector that works to separate the robots. Map information is also exploited to add a guiding force that moves the robot towards open space. The result is behavior that is very similar to an ideal gas. **b)** Robots at CMU – Multiple robots explore a space by sharing information to build a global map. By calculating the cost of moving to a new space and devaluing spaces already claimed, these robots are able to efficiently explore a space even when in close proximity to one another.

teams compete to score the most goals. Veloso describes a team of robots that model their specific roles to facilitate coordination [11] (Figure 5a). Goalies know to stay at one end of the field while other members are loosely responsible for a specific area of the field. The role of each robot and how it coordinates with one another is preprogrammed into the robots.

Higher on the spectrum of coordination are robots that begin to model parameters of the mission through the mechanisms of social interaction. A good example of this is robots that adhere to a "free market model" of interaction. In these systems robots bid for access to a common resource—such as accessibility to a given space or promotion to leader status. The coordination mechanism itself is centralized, in this case a broker that arbitrates between multiple bids for the same service. However, the decision process and resulting interaction between robots is distributed. Each robot acts as its own agent and seeks to maximize its own profits. Robots receive 'payment' for completing tasks but have to 'pay' for access to that service. This payment includes the cost of that access including the energy they would have to expend as well as the time to accomplish the task. When two or more robots bid for the same resource, the one that is willing and able to pay the most for access is awarded the bid.

Dias describes a team of robots called Traderbots that use the free market architecture to accomplish common tasks such as item retrieval and exploration [1] (Figure 5b). Robots bid for rights to a particular object or access to potentially open space. The cost for reaching that target is calculated based on the particular abilities of that robot as well as environmental factors such as the distance from the target and the time to complete the task. Potential reward is calculated by the amount of items a robot can pick up or the coverage it can provide with its suite of sensors.

This social approach to coordination provides a natural way to manage mission complexity while encouraging cooperation. Robots decide on which actions to pursue based on increasing their individual profit, not on the cost of achieving a goal. Conversely, the market sets the revenue and the rewards for each task, dictated by the needs of the mission. In this fashion, tasks that are more critical can be given higher reward and are pursued more aggressively. However, less important tasks are also incorporated when their margins with respect to the individual are more appealing. The difficulty with this social approach is how to model commodities and currency. For example, it is difficult to determine how to relate sensor coverage to distance when determining whether to bring a more equipped robot from farther away. However, while robots are machines and differ in subtle but significant ways, we can leverage much from what we know about how societies interact.

An even higher form of cooperation in our taxonomy involves robots that coordinate by modeling the actions and potential interactions of one robot with another. Stroupe describes a team of robots that attempt to localize a set of obstacles by viewing them from multiple vantage points [10] (Figure 5c). In this example, robots estimate the combined utility of moving and sensing in a new position assuming that a nearby robot will also be following a similar policy. When two robots are in agreement, the combined action is greater that the sum of the individual actions. In this case, they cooperate to localize a given obstacle better than could be achieved with two independent measurements.

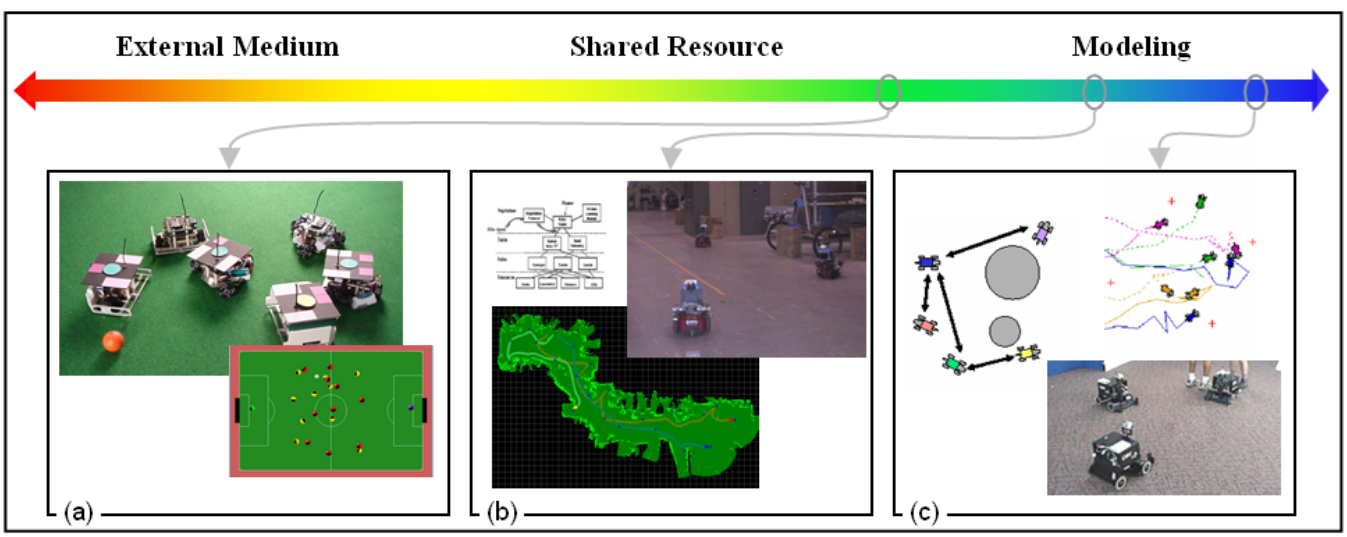

Figure 5. Upper portion of spectrum – teams that model the task or other members in the team to increase coordination - **a)** Robots at CMU model their specific role in a task to facilitate coordination. Specialization naturally partitions complexity as each robot understands its role in the task implicitly. In the case of soccer robots, this partitioning also maps to a spatial element. Goalies know to stay at one end of the field while other members are loosely responsible for a specific area of the field. **b)** Robots at CMU model the interaction of competing agents in a free market architecture. Robots compete for a resource using bids through a central broker. The robot with the greatest bid is given that particular resource. In this example robots bid for a region of space to explore. **c)** Robots at CMU exhibit a higher degree of coordination by robots modeling the interaction with neighbors. In this example, robots estimate the combined utility of moving to a space assuming that a nearby robot will also be following a similar policy. When two robots are in agreement, the combined action is greater that the sum of the individual actions.

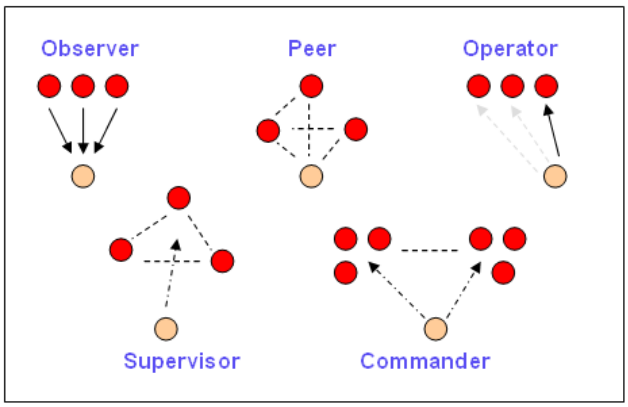

Figure 6. Role of the Human – the role of the human can be quite diverse. a) Observer – the human essentially watches the interaction of the team but does not participate during operation. b) Peer – the human is an operator in the group. His perspective is most likely limited like his robot counterpart. c) Operator – the human takes control of the actions of the robot and directs movement based on direct feedback either through direct visual contact or by directly viewing the sensor feedback from the robot. d) Supervisor – the human directs the actions of the team as a whole either by changing the desired goal or task set. e) Commander – the human interacts with the team by changing its intentions or goals but not how the team achieves those goals.

IV. Role of the Human

So far, our discussion has focused on the cooperation among machines, but most future systems will require humans and machines to cooperate. Humans can bring a vast wealth of experience and domain knowledge to a robot team. The impact on complexity of adding a human to the equation depends on the coupling between human and robot perceptions. That is, to what degree is the human able to impart this experience and to what degree are the robots able to provide situational awareness back to the user. This ability to couple perception is, in large part, a function of the role the human plays in the team. We now consider the diverse roles that a human can assume in the execution of a mission. Figure 6 shows a graphical representation of five roles that a human may perform within a human-robot team.

It may be that the human has little involvement in the coordination of the team and is a mere observer of events. As an observer, the human has no impact on the way the team coordinates. Information flows from the robot to the human but no feedback path is established.

A second role for a human is as a peer to the robots within the team. In this case, the human and robot may exchange information, but the human does not control the robot in any way.

For most current robot systems, the human plays a more direct role in the operation of the individuals in a team. In this case, the human is called the operator. Usually the human acts as the surrogate intelligence for the robot. In this role, the human subsumes the thinking role for a specific robot and directly dictates its actions. Since this dictation is almost always done remotely, it is called *teleoperation*. The human either has a direct visual link to the robot or the robot feeds its sensors directly back to the human. While teleoperation gives the operator greater control over the robot, it also places a heavy burden on the attention of that operator. Consequently, it is typically difficult for a human to operate more than one robot at a time.

A fourth possible role for a human is as a supervisor of multiple robots. In supervisory control [9], the robot possesses certain minimal competencies to perform some tasks, and the human supervisor monitors task progress and only occasionally needs to give guidance or direction.

A fifth role for a human is as a commander of a team that contains robots. In this case, the human does not directly task robots, but rather defines mission objectives. In this case, the human needs to understand the capabilities and limitations of the robots, but intimate knowledge of robot control is not required.

With these roles defined, we would describe a command and control system in terms of the available communication channels and the messages that can pass from one member of the team to another. These messages may allow one agent to grant authority to another agent, or to order the execution of a plan known to the parties. Messages can also be used to report status and outcomes of tasks. From our previous discussion of reported approaches to robot teams, one can see that some command and control approaches may require quite sophisticated messages.

V. Conclusion

In this paper we have proposed a simple spectrum for representing the complexity of cooperation in robot teams. Given a desired mission, we believe that the design of an effective human-robot team starts with two fundamental choices: First, one must choose the degree of coordination desired of the robots. Greater coordination requires mechanisms for sharing data, task status, and mission models. Second, one must choose the role(s) to be played by the humans within the team. Given these two choices, one then proceeds to build infrastructure to support the communication required by these roles in interacting with the rest of the team.

References

[1] M., Dias and A., Stentz, "A Free Market Architecture for Distributed Control of a Multirobot System," Proceedings of the 6th International Conference on Intelligent Autonomous Systems (IAS-6), 2000.

[2] G. Dudek, M. Jenkin, E. Milios, and D. Wilkes. A taxonomy for multi-agent robotics. Autonomous Robots 3, (4):375--397, 1996.

[3] J., Drury, L., Riek, A., Christiansen, Z., Eyler-Walker, A. Maggi, and D. Smith "Evaluating Human-Robot Interaction in a Search-and-Rescue Context," presented at Robotics and Automation, 2004.

[4] L., Parker, B., Kannan, X., Fu, and Y., Tang, "Heterogeneous Mobile Sensor Net Deployment Using Robot Herding and Line-of-Sight Formations," ," Proceedings of the 2003 IEEE/RSJ Intl. Conference on Intelligent Robots and Systems Las Vegas, Nevada · October 2003.

[5] A., Hayes, A., Martinoli, and R., Goodman. Swarm Robotic Odor Localization. Proc. of the 2001 IEEE/RSJ Int. Conference on Intelligent Robots and Systems IROS-2001, Wailea, Hawaii, October 2001, pp. 1073-1078.

[6] A. Howard, M. Mataric´ and G. Sukhatme, "An Incremental Self-Deployment Algorithm for Mobile Sensor Networks", Autonomous Robots, Special Issue on Intelligent Embedded Systems, Vol 13 No 2, pages 113--126, 2002.

[7] L., Iocchi, D., Nardi, M., Salerno, "Reactivity and Deliberation: a survey on Multi-Robot Systems," In Balancing Reactivity and Social Deliberation in Multi-Agent Systems (LNAI 2103) M. Hannebauer, J. Wendler, E. Pagello Eds. Springer, 2001.

[8] P., Rybski, I., Burt, T., Dahlin, M., Gini, D., Hougen, D., Krantz, F., Nageotte, N., Papanikolopoulos, S., Stoeter, "System Architecture for Versatile Autonomous and Teleoperated Control of Multiple Miniature Robots," Proceedings of the 2001 IEEE International Conference on Robotics and Automation, Seoul, Korea, May 2001.

[9] R. Simmons, D. Apfelbaum, D. Fox, R. Goldman, K. Zita Haigh, D. Musliner, M. Pelican, and S. Thrun, "Coordinated Deployment of Multiple, Heterogeneous Robots," In Proceedings of the Conference on Intelligent Robots and Systems (IROS), Takamatsu Japan, October 2000.

[10] A. Stroupe, "Collaborative Execution of Exploration and Tracking using Move Value Estimation for Robot Teams (MVERT)," Thesis dissertation, Carnegie Mellon University, 2003.

[11] M. Veloso, P. Stone, K. Han, and S. Achim, "The CMUnited-97 Small Robot Team." In Proceedings of RoboCup-97: The First Robot World Cup Soccer Games and Conferences, Kitano, H. (ed.). Springer Verlag, Berlin. 1998.