

# Principled Approaches to the Design of Multi-Robot Systems

Chris Jones, Dylan Shell, Maja J Matarić  
Computer Science Department  
University of Southern California  
941 West 37th Place  
Los Angeles, CA 90089-0781 USA  
Email: {cvjones|dshell|maja@robotics.usc.edu}

Brian Gerkey  
Artificial Intelligence Lab  
Stanford University  
Stanford, CA 94305 USA  
Email: gerkey@stanford.edu

## Abstract

Coordinated multi-robot systems (MRS) have been successfully demonstrated in a variety of task domains. However, the design of these systems is typically *ad hoc* and centered around resource-intensive trial-and-error design processes. To fully exploit the power and potentials of MRS, the community is in need of additional principled design methodologies. In this paper, we present an overview of our ongoing work on three different principled methodologies, each of which addresses a different type of MRS, based on the scale of the system or mechanisms by which coordination is achieved. First, we present a formal study and analysis of task allocation in MRS, including a domain-independent taxonomy of multi-robot task allocation (MRTA) problems. Second, we present a principled controller design methodology for MRS executing sequential tasks. This design methodology provides a formal framework which serves as the foundation for an integrated set of systematic MRS controller synthesis and analysis methods. Third, we discuss the application of a formal analytical approach to the study of large-scale MRS motivated by work from the physics community. Through the use of principled methods such as these, we hope to provide a formal footing for the design of MRS.

## I. INTRODUCTION

The study of multi-robot systems (MRS) has grown significantly in recent years. This is not surprising as continually improving technology has made the deployment of MRS consisting of increasingly larger numbers of robots possible. Task domains well-suited to multi-robot solutions are expanding. With the growing interest in MRS comes the expectation that, at least in some important respects, multiple robots will be superior to a single robot in achieving a given task.

Potential advantages of MRS over a single robot system (SRS) are frequently discussed in the literature. For example, total system cost, it is frequently claimed, may be reduced in many domains by utilizing multiple simple and cheap robots as opposed to a single complex and expensive robot. The inherent complexity of some task environments may *require* the use of multiple robots as the necessary capabilities are too substantial to be met by a single robot. Furthermore, multiple robots are often claimed to increase system efficiency, robustness, and flexibility by taking advantage of inherent parallelism and redundancy.

Our work focuses on *distributed* MRS, a class of MRS in which each robot can potentially communicate with other robots, but each robot ultimately operates independently under local sensing and control, with coordinated system-level behavior arising from local interactions among the robots and between the robots and the task environment. Distributed MRS comprise an important subset of possible MRS designs and are frequently studied in the literature. Distributed solutions are typically more likely to scale to large numbers of robots and provide more robust performance as compared to more centralized designs, which are more likely to have communication and computational bottlenecks. In this paper we address the design of distributed MRS.

Utilization of distributed MRS poses potential disadvantages and additional challenges that must be addressed. A poorly designed MRS, with individual robots working toward opposing goals, can be less effective than an appropriately designed SRS. In most domains just taking a suitable SRS design and scaling it up to multiple robots is not adequate. Designing a MRS is in many important respects a different paradigm than designing a SRS and, as such, requires new types of design tools and methodologies. A paramount challenge that must be directly addressed in the design of effective MRS is managing the complexity introduced by multiple, interacting robots.

### A. Coordination in Multi-Robot Systems

In order to maximize the effectiveness of a MRS, the robots' actions must be spatio-temporally coordinated and directed towards the achievement of a given system-level task. Just having robots interact is generally not sufficient in itself to produce interesting or practical system-level coordinated behavior. The design of MRS can be quite challenging because unexpected system-level behaviors may emerge due to unanticipated ramifications of the robots' local interactions. In order for the interacting robots to produce coherent task-directed behavior there must be some overarching coordination mechanism that spatio-temporally organizes the interactions in a manner appropriate to the task.

The design of such coordination mechanisms can be difficult; nonetheless, many elegant hand-crafted distributed MRS have been demonstrated, both in simulation and on physical robots [6, 13]. The methods by which these systems have

achieved task-directed coordination are diverse and the possibilities are seemingly limited only by the ingenuity of the designer. From a few robots performing a manipulation task [4, 18], to tens of robots exploring a large indoor area [24, 31], to potentially thousands of ecosystem monitoring nano-robots [39, 43], as the number of robots in the system increases, so does the necessity and importance of coordination.

An often used and important spectrum to classify coordination is that of explicit vs. implicit coordination. An *explicit* coordination mechanism is one in which the robots' interactions involve explicit task-directed communications or negotiations about global resource usage or task assignments in order to achieve coordinated behavior. One such mechanism, for example, is market-based coordination in which auctions are used to continuously assign tasks to robots [9, 18]. In this approach, individual robots competitively bid for tasks which they must either complete or report as a broken contract.

An *implicit* coordination mechanism is one in which each robot operates under local sensing and control and coordinated system-level behavior arises from local robot-robot and robot-world interactions *without* an explicit notion of task-directed communication or negotiation. For example, one such mechanism for achieving coordination through local interactions is stigmergy [3, 23]. In this approach, communication is mediated through task-related changes in the environment rather than through explicit communication [33].

Admittedly, the distinction between explicit and implicit coordination is not crisp. It represents a continuous spectrum and rarely does a given coordination mechanism(s) fall cleanly at one end of the spectrum or the other. Nonetheless, it is a useful spectrum by which to compare and contrast different approaches to coordination.

At this time, it is still an open and energetically debated issue as to the relative merit of explicit and implicit coordination. MRS coordinated by explicit mechanisms are typically able to achieve closer to optimal task performance and are more amenable to formal analysis. In many instances, however, they are not flexible or robust to individual robot failure and, due to computational and communication complexity, do not scale well to large group sizes. On the other hand, implicit mechanisms are generally more robust and are often times most useful in MRS composed of a large number of robots; however, implicitly coordinated MRS are typically more difficult to formally analyze and the task performance can be less optimal than an explicitly coordinated alternative.

Unfortunately, the design of distributed MRS remains more of an art than a science. A MRS designer is equipped with few formal and general tools or methodologies to systematically and objectively evaluate the relative merits of different design options in order to make principled design decisions. Current design methodologies are frequently driven by informal and undocumented expert knowledge and, in the worst cases, they are driven by resource-intensive trial-and-error processes. Available design methodologies are typically task-specific and are infrequently accompanied by formal analysis of the expected system performance or domains of applicability. Furthermore, formal explanations are rarely provided to justify the superiority of the resulting system design relative to possible alternatives.

## B. Formalizing Multi-Robot Systems

The design of coordination mechanisms for MRS has proved to be a difficult problem. In the last decade the design of a variety of such mechanisms over a wide range of task domains has been studied [12, 6]. Although the literature highlights some elegant solutions, they are generally domain-specific and provide only indirect insight into important questions, such as how appropriate a given coordination mechanism is for a particular domain, what performance characteristics should one expect from it, how it is related to other coordination mechanisms, and how can one modify it to improve system performance. These questions must be answered in a formal way before one can quickly and efficiently produce an effective MRS for a new task domain.

To more fully utilize the power and potentials of MRS and to move the design process closer to a science, the community must have at its disposal principled design tools and methodologies. Such tools and methodologies will provide a solid foundation about which to construct increasingly capable, robust, and efficient MRS. While it is unrealistic to expect a completely general methodology, we present a summary of our research directed toward the development of a collection of compatible and principled MRS design methodologies. Three methodologies are presented, each of which focuses on MRS at a different location along the implicit vs. explicit coordination spectrum, as shown in Figure 1.

First, in Section II, we discuss work on a formal study of multi-robot task allocation. This work includes a domain-independent taxonomy of multi-robot task allocation (MRTA) problems and shows how many such problems can be viewed as instances of other, well-studied, optimization problems. Such a study provides formal knowledge on which to influence design decisions in multi-robot task allocation domains, including insight as to communication and computational complexity of different approaches. This approach to coordination in MRS falls toward the explicit end of the coordination spectrum as robots are intentionally performing task-related negotiations to allocate robots to available tasks.

Next, in Section III, we present a principled controller design methodology for distributed homogeneous MRS performing sequential tasks. The design methodology provides a formal framework which serves as the foundation for an integrated set of systematic MRS controller synthesis and analysis methods. The integrated nature of the synthesis and analysis enables controller optimization based on a desired performance metric. We place this design methodology to the implicit side of the center of the coordination spectrum. The methodology allows the use of robots which maintain




Fig. 1. Placement of the 3 design methodologies along the explicit vs. implicit coordination spectrum, according to the types of MRS on which they focus.

internal state and communication; however, the communication is quite simple and is does not involve task-directed negotiation.

In Section IV we discuss the problems associated with very large numbers of robots and why addressing these challenges may lead to better understanding of coordination. We outline ongoing work in this direction that makes use of a class of simple interaction dynamics that permits the application of analytical approaches originally developed in the physics community. The discussion aims to show that although there are currently no operational systems of this sort, they will likely require unique formal methods. We place this methodology toward the implicit end of the coordination spectrum as the group size is very large and all interactions are local.

## II. FORMAL ANALYSIS OF MULTI-ROBOT TASK ALLOCATION

In this section we present a formal approach to the study of multi-robot task allocation (MRTA). As MRS are now becoming considered as feasible, and even desirable, solutions in many problem domains, MRTA has become a key area of study. An implicit questions that is addressed by all MRS coordination mechanisms, in one way or another, is ‘which robot should execute which task?’. This question is at the core of the MRTA problem. Explicit coordination mechanisms usually address this question directly through intentional communication or negotiation between the robots. Implicit coordination mechanisms usually address this question in an indirect manner, with the assignment of robots to tasks occurring due to local interactions and emergent effects. To date, there exist empirically validated methods that are primarily *ad hoc*, but there is relatively little work in the development of formal approaches to MRTA. Some MRTA architectures have been demonstrated, but there are only limited formal attempts to evaluate or compare these architectures, either analytically or empirically.

Next, we present a particular taxonomy for studying MRTA, based on organizational theory from several fields, including operations research, economics, scheduling, network flows, and combinatorial optimization. We demonstrate how this taxonomy can be used to analyze and classify MRTA problems, and evaluate and compare proposed solutions. In this paper, we just give a brief overview of this work. Additional details on the methods and results can be found in [17, 19, 20].

### A. MRTA Taxonomy

We propose a taxonomy of MRTA problems based on the three axes listed below. In Gerkey and Mataric [19], we used this taxonomy of MRTA problems to analyze and classify important MRTA problems and evaluate and compare proposed solutions.

- **single-task robots (ST) vs. multi-task robots (MT)**

ST means that each robot is capable of executing at most one task at a time. MT means that some robots can execute multiple tasks simultaneously.

- **single-robot tasks (SR) vs. multi-robot tasks (MR)**

SR means that each task requires exactly one robot to achieve it. MR means that some tasks can require multiple robots.

- **instantaneous assignment (IA) vs. time-extended assignment (TA)**

IA means that the available information concerning the robots, the tasks, and the environment permits only an instantaneous allocation of tasks to robots, with no planning for future allocations. TA means that more information is available, such as the set of all tasks that will need to be assigned, or a model of how tasks are expected to arrive over time.

We denote a particular MRTA problem by a triple of two-letter abbreviations drawn from this list. For example, a problem in which multi-robot tasks must be allocated once to single-task robots is designated ST-MR-IA.

Name	Computation / iteration	Communication / iteration	Solution quality
ALLIANCE [37]	$O(mn)$	$O(m)$	at least 2-competitive
BLE [42]	$O(mn)$	$O(mn)$	2-competitive
M+ [5]	$O(mn)$	$O(mn)$	2-competitive

TABLE I

Summary of selected iterated assignment architectures for MRTA. Shown here for each architecture are the computational and communication requirements, as well as solution quality.

Name	Computation / task	Communication / task	Solution quality
MURDOCH [18]	$O(1)$ / bidder $O(m)$ / auctioneer	$O(m)$	3-competitive
First-price auctions [8]	$O(1)$ / bidder $O(m)$ / auctioneer	$O(m)$	at least 3-competitive
Dynamic role assignment [7]	$O(1)$ / bidder $O(m)$ / auctioneer	$O(m)$	at least 3-competitive

TABLE II

Summary of selected online assignment architectures for MRTA. Shown here for each architecture are the computational and communication requirements, as well as solution quality.

The intent of this taxonomy is to show how various MRTA problems can be positioned in the resulting problem space and to explain how organizational theory relates to those problems and to proposed robotics solutions. For example, in [20] we shown how ST-SR-IA problems are actually an instance of the *Optimal Assignment Problem* Gale [16], how ST-SR-TA problems are ones requiring the building of a time-extended *schedule* of tasks for each robot and minimizing total weighted cost, and how ST-MR-IA problems are studied in other domains and referred to as *coalition formation* [40]. In some cases, it will be possible to construct provably optimal solutions, while in others only approximate solutions are available. There are also some difficult MRTA problems for which there do not currently exist good approximations. When designing a MRS, it is essential to understand what kind of task allocation problem is present in order to solve it in a principled manner. Being able to identify different classes of MRTA problems and understanding how various formal methodologies, some of which are previously well-studied in other fields, may be applied to their solution is fundamental to building more principled and effective MRTA systems.

### B. Analysis of MRTA Approaches

Presumably because it is the simplest case of MRTA, the ST-SR-IA problem has received the most attention from the research community. Having developed a formal framework in which to study this MRTA problem, we can now apply it to an analysis of some of the key task allocation architectures from the literature. In this section, a summary of a more extensive analysis of six approaches to the ST-SR-IA problem are analyzed, focusing on computation and communication requirements and solution quality. Gerkey and Matarić [19] provide the full analysis.

Computational requirements, or running time, are determined in the usual way, as the number of times that some dominant operation is repeated. For the MRTA domain that operation is usually either a calculation or comparison of utility, and running time is stated as a function of  $m$  and  $n$ , the number of robots and tasks, respectively. Communication requirements are determined as the total number of inter-robot messages sent over the network. Solution quality is reported as a competitive factor, which bounds an algorithm's performance as a function of the optimal solution. for a maximization problem, an algorithm is called  $\alpha$ -competitive if, for any input, it finds a solution whose total utility is never less than  $1/\alpha$  of the optimal utility.

Tables I and II summarize the results for the iterated assignment architectures and online assignment architectures, respectively. Perhaps the most significant trend in these results is how similar the architectures look when examined in this manner. For example, the iterated architectures listed in Table I, which assign all available tasks simultaneously, exhibit almost identical algorithmic characteristics. Only the ALLIANCE architecture [37] shows any difference; in this case the decrease in communication overheard is achieved by having each robot internally model the fitness of the others, thereby

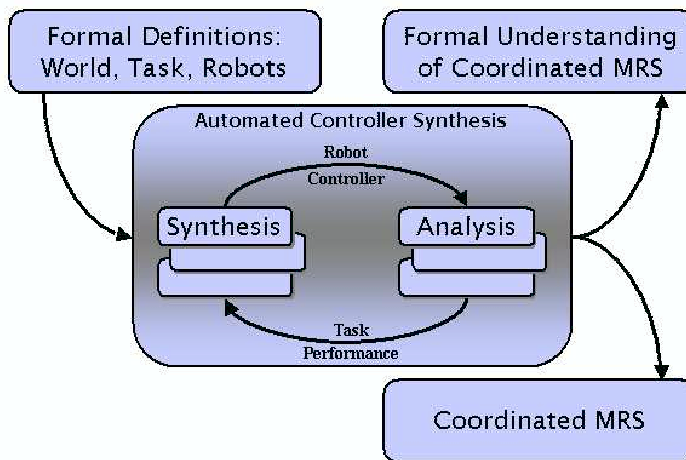


Fig. 2. A formal MRS controller design methodology multi-robot systems.

effectively distributing the utility calculations. More striking are the results in Table II, which lists architectures that assign tasks in a sequential manner. With respect to computational and communication requirements, these architectures are *identical*.

These results are particularly interesting because they suggest that there is some common methodology underlying many existing approaches to MRTA. This trend is difficult or impossible to discern from reading the technical papers describing the work, as each architecture is described in different terms, and validated in a different task domain. However, with the analysis described here, fundamental similarities of the various architectures become obvious. These similarities are encouraging because they suggest that, regardless of the details of the robots or tasks in use, the various authors are all studying a common, fundamental problem in autonomous coordination. As a corollary, there is now a formal grounding for the belief that these *ad hoc* architectures may have properties that allow them to be generalized and applied widely.

### III. FORMAL MRS CONTROLLER DESIGN METHODOLOGY

In this section we present a high-level overview of our work on the development of a principled controller design methodology for distributed homogeneous MRS performing sequential tasks. This methodology is composed of a formal foundation upon which a suite of automatic and systematic MRS synthesis and analysis methods operate to produce and optimize task-directed MRS. Figure 2 diagrams the structure of this controller design methodology. Inputs to the system are formal definitions for the world, task, and robots *sans* controller. The output of the system is a complete task-directed MRS design, including robot controllers, accompanied by a predictive formal analysis of task performance.

Typically, robot controllers are synthesized by hand, relying on expert knowledge of the designer and an intimate understanding of how the robots will interact with each other and with the world during task execution. Our approach is novel in that it represents a systematic and automated method for MRS controller synthesis. The designer's responsibility is to define the task domain using a provided formal framework. The information required to do so is inherently required to hand-design a system as well, so provides few challenges.

#### A. Related Controller Design Methodologies

Other notable work is motivated by the same desire to develop a systematic methodology to design MRS. For example, Donald [11] presented information invariants, which is aimed at defining the information requirements of a given task and ways in which those requirements can be satisfied in a specific robot design. This work put the design of MRS on a formal footing and began to identify how various robot sensors, actuators, and control strategies could be used to satisfy task requirements and how they are related through systematic reductions. The concept of information invariants was experimentally studied in a distributed manipulation task domain [10]. Also, Klavins [30] presented an interpreter called the Computation and Control Language interpreter that implements a robotics modeling and programming software tool called the Computation and Control Language (CCL). The CCL formalism was used to study the scalability of algorithms for multi-robot control, in particular the communication complexity of a variety of communication schemes [29]. Furthermore, there exist alternative approaches and methodologies to the design and optimization of MRS controllers such as evolutionary [15] and learning methods [34, 37] and a number of MRS design environments, control architectures, and programming languages which assist in the design of task-directed coordinated MRS [35, 2, 1, 20].

A distinguishing factor of our work over previous related work is that our controller design methodology is formally grounded and consists of an integrated suite of MRS synthesis and analysis methods capable of systematically and automatically constructing MRS that make use of a variety of coordination mechanisms.

### B. Formal Controller Design Methodology

The methodology is based on a MRS formalism that provides a principled framework for formally defining and reasoning about concepts relevant to MRS – the world, task definition, and capabilities of the robots themselves, including action selection, sensing, maintenance of local and persistent internal state, and broadcast inter-robot communication. Based on this formalism, the methodology utilizes an integrated set of MRS synthesis and analysis methods.

Our formal methodology includes a suite of systematic MRS synthesis methods that use the MRS formalism in a *prescriptive* manner. Synthesis is the process of constructing a specific instance of a MRS which meets design requirements such as achieving the desired level of task performance while meeting constraints imposed by limited robot capabilities. The synthesis methods take as input the formal definitions of the world, task, and robots *sans* controller and they output a robot controller. Since we focus on homogeneous MRS, our approach is to synthesize an individual robot controller, such that when every robot in the MRS executes the controller, system-level task-directed coordination is achieved. Using the MRS formalism, the synthesis methods automatically produce a complete robot controller through a unique systematic logic-induced procedure. Each of these methods is independent and produces a coordinated MRS through the use of a unique set of coordination mechanisms, including through the use of internal state [26], inter-robot communication [27], and/or deterministic and probabilistic action selection. The suite of automated synthesis methods allows for a more efficient and broader sampling of the design space. This increases the likelihood that a design will be found that meets the constraints imposed by the task, robot capabilities, and desired performance characteristics.

Complimentary to the synthesis methods, the methodology incorporates two MRS modeling approaches that use the MRS formalism in a *descriptive* manner. The design of an effective task-directed MRS is often difficult because there is not an accurate understanding of the relationship between different design options and resulting task performance. In the common trial-and-error design process, the designer will construct a MRS and then try it out in either a simulation or on actual robots. Either way, this process is resource-intensive as it requires much effort and time to evaluate many possible designs. Ideally, the designer should be equipped with an analytical tool for the analysis of a potential MRS design. Such a tool would allow for quick and efficient evaluation of different design options and likely result in more effective MRS designs as the design can more effectively be optimized with respect to some desired performance metric. To address this issue, our design methodology includes both macroscopic [28] and microscopic MRS modeling approaches that are capable of quantitatively predicting the probability a MRS will correctly execute a given sequential task.

Together, the synthesis and analysis methods provide more than just pragmatic design tools. A defining feature of this design methodology is the integrated nature of the controller synthesis and analysis methods. The fact that they are integrated allows for the capability to automatically and iteratively synthesize and analyze a large set of possible designs, thereby resulting in more optimal solutions and an improved understanding of the space of possible designs.

Furthermore, based on their common formal foundations and integrated nature, the synthesis and analysis methods provide a platform from which to formally characterize some relationships and dependencies among MRS task requirements, individual robot control and capabilities, and resulting task performance. Given the formal grounding of the methodology, it also provides a means to systematically determine limitations of different types of coordination mechanisms, to understand the inherent relationships between different mechanisms, to contribute methods to systematically reduce one mechanism type another, and to provide insight into the general requirements necessary for achieving different forms of coordination. It allows for the ability to formally answer fundamental questions such as: ‘*In what conditions is it necessary for the robots to be able to communicate?*’, ‘*In what conditions is communication alone insufficient?*’, and ‘*When are the use of internal state and communication interchangeable?*’.

This principled approach to MRS controller design has been demonstrated in a sequentially-constrained multi-robot construction task domain [26, 28, 27]. Controllers for a number of different construction tasks were synthesized and demonstrated in extensive simulation trails and in a limited number of real-robot trials. We are currently undertaking more extensive validation of the design methodology in more complex task domains and forming some formal conclusions on the uses and limitations of the different coordination mechanisms in the construction task domain.

## IV. FORMAL APPROACH TO LARGE-SCALE MRS

There are numerous challenges to be overcome before robotic systems with large numbers of robots (e.g. numbers exceeding 1000) are actively deployed. The manufacture and maintenance of this many robots is a technical problem, requiring that individual units be physically robust beyond the level commonly seen today. Questions regarding the tractable simulation of these sorts of systems mean that even simulated results of this sort are rarely seen. There are also major scientific considerations, perhaps the most obvious being the question of exactly which tasks such a system is best suited for. In the case of conventional multi-robot teams (with numbers on the order of ten or so) the class of tasks for which a given team is best suited is not properly understood; far less is known for groups of greater size. Large

groups mean that various natural intuitions and ideas about the nature of cooperation, teamwork, and synergy become malformed. One expectation is that a better understanding of the defining characteristics can be found for these notions through the study of Large Scale MRS (LS-MRS).

Exactly how to go about programming such systems in order to achieve goal directed behavior remains an open question. Just as the growth of MRS from single robots raised new issues and questioned old assumptions, so too does the evolution to LS-MRS reexamine existing MRS coordination methodologies. Expectations about reasonable classes of communications, justifiable protocols, acceptable timeouts, and various other assumptions about the general behavior of the system may prove to be invalid for LS-MRS. A large scale system may have robots attempting to communicate with others whose distance is very much larger than the communication and sensing range of the individual; mechanisms that are assumed to be axiomatic in some MRS (like global broadcast) may well be difficult to achieve efficiently, or at all, in LS-MRS.

#### A. Current Approaches

Either the current limitations in understanding of LS-MRS, or perhaps the inspiration of nature's examples of many-body societies in insects, mean that minimalist approaches are often judged as the most applicable for large numbers. Simple agents may be sufficient if complex collective capabilities of the group can arise through sheer numbers and synergy. This minimalism means that implicit coordination mechanisms are considered as most appropriate. As already described in Section I-A, implicit coordination presents challenges for formal analysis, including considerations of physical dynamics, the necessity of multiple levels of detail in models, emergent dynamics, etc. It is hoped that LS-MRS will simplify the analysis dilemma through two ways. The first is through the simplicity of the agents themselves. As numbers increase the level of competency expected of each agent generally decreases and simpler agents are more amenable to analytical approaches. The second is the anticipation that statistical methods for predicting group behavior become more applicable. Examples are approaches of Lerman and Galstyan [32] and Sato and Crutchfield [41] where mean behavior of simple agents is modeled; instantaneous fluctuations from the mean behavior can decrease in statistical likelihood as numbers are increased. Thus mean behavior predictions are more compelling for more robots.

Simply considering the mean characteristics of an arbitrary interacting system is unlikely to result in anything directly applicable to LS-MRS – simply large numbers alone is insufficient. For example, in chaotic systems small aberrant behavior may be amplified to the point of global consequences. Any analysis that must consider arbitrarily small behavioral changes in all agents in order to construct the global picture is unlikely to be useful for physical robots wherein influential local behavioral dynamics may be affected by uncertainty, noise, real world constraints, and other factors.

We hypothesize that as the number of degrees-of-freedom in the system increases with the addition of agents, those coordination mechanisms that have the least behavioral constraining assumptions are most robust. Essentially those coordination approaches that couple robot interactions as loosely as possible are favored as they are likely to scale better. Klavins [29] describes a similar notion with the words “the less coordination... the better it should scale.” Reducing inter-robot couplings makes otherwise nonlinear behavior more and more linear (and predictable), while reducing the repertoire of behavior (and hence applicable tasks). This highlights a related trade-off between the robustness and system responsiveness. The hardware, controllers and communication methodologies are intended to be robust with respect to noise and various other ‘inappropriate’ sensing and actuation variations. At the same time the system must be sensitive to control signals useful for task achievement. In an implicitly coordinated system, where there may be little or no semantic separation between task achieving and communication acts, these two requirements on the interaction dynamics are constraints in opposite directions.

#### B. Ergodic by Design

We propose a methodology that recognizes the importance of dynamics, while dealing with a class of systems which we believe are feasible for physical robots. The essential property of our approach is the focus on LS-MRS defined in order to achieve global equilibrium states with interaction properties that are achievable in the presence of noise. Our focus is on the multi-agent systems exhibiting dynamics of the class known as *ergodic*. This can be best understood in terms of the behavioral configuration space of the robots, which consists of a dimension for each variable of interest for every robot (e.g. position of the robot, internal state variable, etc). Constraints on the range and values of variables are reflected as bounds on each of respective the dimensions. At any point in time the behavior of the entire global system can be represented by a point in this space, and a time evolution of the LS-MRS is represented by a trajectory in this many dimensional space. The requirement of ergodicity translates into the requirement that all reachable parts of the configuration space are visited given sufficient time, and thus that time-average behavior (after much time) equals space-average behavior.

Focusing on ergodic dynamics is attractive because it makes prediction of macroscopic behavior feasible: only the space in which the dynamics evolve, rather than the full time-evolution, need be considered. For example the addition of a suitable (global) behavioral constraint can result in an ergodic dynamics that operates on only a subset of the configuration space. Studying that subset's topology gives important insights into the (equilibrated) system behavior.

Methods for discovering this topology rely on the simulation of “pseudo-dynamics” as performed, for example, by the famous Metropolis algorithm [36]. The stochasticity of probabilistic controllers (and real world noise) aid in the achievement of ergodicity. Also, Erdmann [14] showed that these effects may increase the class of tasks that may be executed by the system. In spite of the occurrence of natural large scale behavior that is not ergodic (e.g. Helbing [22]), it is an easily achievable design criterion for artificially synthesized systems. In a sense, we are considering systems with uninteresting dynamics, recognizing that many systems with far more interesting behavior can be amended to this basic case.

### C. Macroscopic Phenomena

This far the discussion of ergodicity has described how it is a restricted dynamics which reduces the space of achievable collective behavior. There are numerous interesting macroscopic structures that can be used in order to achieve particular tasks. Several models of natural systems studied in physics are ergodic and display a range of interesting and well understood macroscopic behavior[44]. Most examples are of extraordinarily large systems and mathematical methods consider behavior in the infinite limit. Such systems exhibit global structure in terms of equilibrium phases, phase transitions, coexistence lines and critical points. These ideas are well developed mathematically and part of a nomenclature that deals with macroscopic phenomena. Gross [21] describes the recent (and ongoing) development of ideas intended to deal with smaller numbers without taking the system size limit and particularly topological mechanisms for understanding notions of phase transitions in finite systems. Other mechanisms for understanding hysteresis and transient behavior do exist, but again are generally only considered for infinite systems.

The further question of how exactly to leverage a small set of mechanisms for producing general macroscopic phenomena remains open. We are pursuing an approach similar in a number of ways to the basis behavior [34] ideas for typical MRS. A set of algorithms, each independently known to produce stable equilibrium behavior, can be combined through sequencing, concurrent execution, and various other operations. Examples of mechanisms anticipated as useful “basis behaviors” include local sensing mechanisms for division of labor (e.g. Jones and Mataric [25]), and current ongoing work using a system roughly modeled by the Potts Model [38] for reaching consensus through induced phase-transition like behavior.

This currently ongoing work is investigating the limits of ergodic dynamics as a mechanism for synthesis of LS-MRS that are capable of performing useful task directed behavior. We aim to better understand the limits of this mechanism before considering the need for more general non-equilibrium systems and complex dynamics.

## V. CONCLUSIONS

We have presented an overview of our ongoing work on three principled multi-robot system (MRS) design methodologies. Each of these methodologies is directed toward the design of a different class of coordinated MRS, from explicitly coordinated MRS permitted through a formal approach to the study of multi-robot task allocation (MRTA), to a principled controller design methodology focusing on homogeneous MRS achieving coordination through a variety of coordination mechanisms, to a formal approach to the study of large-scale MRS. Through the development of such formal methodologies, we aim to place the design of MRS on a more formal footing and place the design process closer to the realm of a science instead of an art.

Many questions remain unanswered and much research is needed to formalize the design of MRS. For example, each of these design methodologies may be able to make some principled statements on the relationships of different design options within their local sphere of the design space. However, there is still very little understanding of the relationships between design options in very different areas of the design space. This leaves much fertile ground for the ongoing effort to formalize coordinated MRS.

## VI. ACKNOWLEDGMENTS

The research report here was conducted at the Interaction Lab, part of the Center for Robotics and Embedded Systems (CRES) at the University of Southern California. This work is supported in part by Defense Advanced Research Agency (DARPA) Grants F30602-00-2-0573 and DABT63-99-1-0015, by Office of Naval Research Grants N00014-00-1-0638 and N00014-01-1-0354, and by National Science Foundation (NSF) Grant EDAct-0121141.

## REFERENCES

- [1] R. Alur, R. Grosu, Y. Hur, V. Kumar, and I. Lee. Modular specification of hybrid systems in charon. In *Proceedings of the 3rd International Workshop on Hybrid Systems: Computation and Control*, pages 6–19, Pittsburgh, PA, Mar 2000.
- [2] R. Arkin and T. Balch. Aura: Principles and practice in review. *Journal of Experimental and Theoretical Artificial Intelligence*, 9:175–189, 1997.
- [3] R. Beckers, O. Holland, and J. Deneubourg. From local actions to global tasks: Stigmergy and collective robotics. In *Proceedings of the 4th International Workshop on the Synthesis and Simulation of Living Systems*, pages 181–189, Cambridge, MA, 1994.



- [4] K. Bohringer, R. Brown, B. Donald, and J. Jennings. Distributed robotic manipulation: Experiments in minimalism. In O. Khatib, editor, *Experimental Robotics IV, Lecture Notes in Control and Information Sciences 223*, pages 11–25. Springer Verlag (Berlin), 1997.
- [5] S. Botelho and R. Alami. M+: a scheme for multi-robot cooperation through negotiated task allocation and achievement. In *Proceedings of the IEEE International conference on robotics and automation (ICRA-99)*, pages 1234–1239, Detroit, MI, Apr 1999.
- [6] Y. U. Cao, A. S. Fukunaga, and A. B. Kahng. Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots*, 4:7–27, 1997.
- [7] L. Chaimowicz, M. Campos, and V. Kumar. Dynamic role assignment for cooperative robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA-02)*, pages 293–298, Washington, DC, Apr 2002.
- [8] M. Dias and A. Stentz. A market approach to multi-robot coordination. Technical report, CMU-RI-TR-01-26, The Robotics Institute, Carnegie Mellon University.
- [9] M. Dias and A. Stentz. A free market architecture for distributed control of a multirobot system. In *Proceedings of the 6th international conference on Intelligent Autonomous Systems*, pages 115–122, Venice, Italy, Jul 2000.
- [10] B. Donald, J. Jennings, and D. Rus. Information invariants for distributed manipulation. In K. Goldberg, D. Halperin, J.-C. Latombe, and R. Wilson, editors, *International Workshop on the Algorithmic Foundations of Robotics*, pages 431–459, Boston, MA, 1995.
- [11] B. R. Donald. Information invariants in robotics. *Artificial Intelligence*, 72(1–2):217–304, 1995.
- [12] G. Dudek, M. Jenkin, E. Milios, and D. Wilkes. A taxonomy for multi-agent robotics. *Autonomous Robots* 3, (4): 375–397, 1996.
- [13] Gregory Dudek, Michael Jenkin, and Evangelos Milios. A taxonomy of multirobot systems. In Tucker Balch and Lynne Parker, editors, *Robot Teams: From Diversity to Polymorphism*, pages 3–22. A.K. Peters, Natick, Massachusetts, Apr 2002.
- [14] M. A. Erdmann. *On Probabilistic Strategies for Robot Tasks*. PhD thesis, Massachusetts Institute of Technology, 1989.
- [15] D. Floreano. Patterns of interactions in shared environments. In *Proceedings of the Second European Conference on Artificial Life*, pages 347–366, Brussels, Belgium, May 1993.
- [16] D. Gale. *The Theory of Linear Economic Models*. McGraw-Hill Book Compan, Inc., New York, 1960.
- [17] B. Gerkey. *On Multi-Robot Task Allocation*. PhD thesis, University of Southern California, Aug 2003.
- [18] B. Gerkey and M. Matarić. Sold!: Auction methods for multi-robot coordination. *IEEE Transactions on Robotics and Automation, Special Issue on Multi-Robot Systems*, 18(5):758–768, Oct 2002.
- [19] B. Gerkey and M. Matarić. Multi-robot task allocation: Analyzing the complexity and optimality of key architectures. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 3862–3867, Taipei, Taiwan, Sep 2003.
- [20] B. Gerkey and M Matarić. A formal analysis and taxonomy of task allocation in multi-robot systems. *to appear in International journal of robotics research*, 2004.
- [21] D. H. E. Gross. *Microcanonical Thermodynamics: Phase Transitions in “Small” Systems*. World Scientific Lecture Notes in Physics - Volume 66. World Scientific, Singapore, 2001.
- [22] D. Helbing. Traffic and Related Self-Driven Many-Particle Systems. *Reviews of Modern Physics*, 4(73):1067–1141, 2001.
- [23] O. Holland and C. Melhuish. Stigmergy, self-organization, and sorting in collective robotics. *Artificial Life*, 5(2): 173–202, 2000.
- [24] A. Howard, L. E. Parker, and G. Sukhatme. The SDR experience: Experiments with a large-scale heterogenous mobile robot team. In *9th International Symposium on Experimental Robotics*, Singapore, Jun 2004.
- [25] C. Jones and M. Matarić. Adaptive division of labor in large-scale minimalist multi-robot systems. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1969–1974, Las Vegas, Nevada, Oct 2003.
- [26] C. Jones and M. Matarić. Towards a multi-robot coordinaton formalism. In *Proc. of the Workshop on the Mathematics and Algorithms of Social Insects*, pages 60–67, Atlanta, Georgia, 2003.
- [27] C. Jones and M. Matarić. Automatic synthesis of communication-based coordinated multi-robot systems. In *To appear in IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sendai, Japan, Sep 2004.
- [28] C. Jones and M. Matarić. Synthesis and analysis of non-reactive controllers for multi-robot sequential task domains. In *Proceedings of the International Symposium on Experimental Robotics*, Singapore, Jun 2004.
- [29] E. Klavins. Communication complexity of multi-robot systems. In J.-D. Boissonnat, J. Burdick, K. Goldberg, and S. Hutchinson, editors, *Algorithmic Foundations of Robotics V*, Springer Tracts in Advanced Robotics, pages 275–292. Springer, 2003.
- [30] E. Klavins. A computation and control language for multi-vehicle systems. In *Proceedings of the 42nd IEEE Conference on Decision and Control*, pages 4133–4139, Hawaii, Dec 2003.

- [31] K. Konolige, D. Fox, C. Ortiz, A. Agno, M. Eriksen, B. Limketkai, J. Ko, B. Morisset, D. Schulz, B. Stewart, and R. Vincent. Centibots: Very large scale distributed robotic teams. In *9th International Symposium on Experimental Robotics*, Singapore, Jun 2004.
- [32] K. Lerman and A. Galstyan. Mathematical model of foraging in a group of robots: Effects of interference. *Autonomous Robots*, 13(2):127–141, 2002.
- [33] Z. Mason. Programming with stigmergy: using swarms in construction. In *Artificial Life VIII: Proceedings of the Eighth International Conference on Artificial Life*, pages 371–374, Sydney, Australia, 2002. MIT Press.
- [34] M. Mataríć. Designing and understanding adaptive group behavior. *Adaptive Behavior*, 4:51–80, Dec 1995.
- [35] M. Mataríć. Issues and approaches in the design of collective autonomous agents. *Robotics and Autonomous Systems*, 16(2–4):321–331, Dec 1995.
- [36] N. Metropolis, A. W. Rosenbluth, Rosenbluth. M. N., A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 6(21):1087–1092, 1953.
- [37] L. Parker. Alliance: An architecture for fault-tolerant multi-robot cooperation. *Transactions on Robotics and Automation*, 14(2):220–240, Apr 1998.
- [38] R. B. Potts. Some generalized order-disorder transformations. In *Proceedings of the Cambridge Philosophical Society*, volume 48, pages 106–109, 1952.
- [39] M. Rahimi, R. Pon, W. Kaiser, G. Sukhatme, D. Estrin, and M. Srivastava. Adaptive sampling for environmental robotics. In *Proceedings of the International Conference on Robotics and Automation*, pages 3537–3544, New Orleans, Louisiana, Mar 2004.
- [40] T. W. Sandholm and V. R. Lesser. Coalitions among computationally bounded agents. *Artificial Intelligence*, 94(1):99–137, 1997.
- [41] Y. Sato and J. P. Crutchfield. Coupled Replicator Equations for the Dynamics of Learning in Multiagent Systems. *Physics Review E*, 1(67):40–43, 2003.
- [42] B. Werger and M. Mataríć. Broadcast of local eligibility for multi-target observation. In *Distributed Autonomous Robotic Systems 4*, pages 347–356, New York, 2001. Springer-Verlag.
- [43] B. Zhang, G. Sukhatme, and A. Requicha. Adaptive sampling for marine microorganism monitoring. In *To appear in IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sendai, Japan, Sep 2004.
- [44] Frank Zielen and Andreas Schadschneider. Broken Ergodicity in a Stochastic Model with Condensation. *Physics Review Letters* 89, 090601, 2002.