

## Multi-robot task allocation through vacancy chain scheduling

Torbjørn S. Dahl<sup>a,\*</sup>, Maja Mataric<sup>b</sup>, Gaurav S. Sukhatme<sup>b</sup>

<sup>a</sup> Robotic Intelligence Lab, University of Wales, Newport Allt-yr-yn Campus, Newport NP2 5DA, United Kingdom

<sup>b</sup> Computer Science Department and Center for Robotics and Embedded Systems, University of Southern California, 941 West 37th Place, Los Angeles, CA 90089-0781, United States

### ARTICLE INFO

#### Article history:

Received 26 September 2007

Received in revised form

26 November 2008

Accepted 2 December 2008

Available online xxx

#### Keywords:

Multi-robot task allocation

Reinforcement learning

### ABSTRACT

Modeling the effects of robot interaction in multi-robot systems, i.e., the *group dynamics*, is difficult due to the complexity of such interactions. This article formalizes the concept of group dynamics in the framework of scheduling and presents a proof that multi-robot task allocation (MRTA), in systems with significant performance effects from group dynamics, is an  $\mathcal{NP}$ -complete problem. As a way of dealing with this complexity we have developed *vacancy chain scheduling* (VCS), a new formal model of MRTA inspired by a resource distribution process commonly found in nature. VCS is also the foundation of a new MRTA algorithm which relies on optimal allocation patterns to emerge from the stigmergic effects of robot interactions. We present experimental evidence of the validity of the VCS model from high-fidelity simulations. The experimental results validate the VCS model by reliably producing the predicted allocation patterns in both homogeneous and heterogeneous groups of robots. The evidence also supports our claim that VCS is a feasible solution for a restricted class of MRTA problems.

© 2008 Elsevier B.V. All rights reserved.

### 1. Introduction

Existing task allocation algorithms generally assume *task independence*, i.e., they do not consider explicitly synergistic or interference-related effects on group performance. Task independence facilitates the identification of optimal task allocation patterns and in some problem domains it is a reasonable assumption to make. For a large number of problems, however, it is not valid to assume that tasks are independent. One particular example is problems that involve groups of mobile robots. Here the effects of interaction between robots, i.e., the *group dynamics*, commonly have a critical impact on the group's performance. Another example is the problem of allowing concurrent access to resources over a limited communication network. In such problems, the interference effects are similar to those produced by physical interference in multi-robot systems.

Using *scheduling* [1] as a formalism for task allocation we show that, in domains where the group dynamics have a significant effect on group performance, task-allocation is  $\mathcal{NP}$ -hard. This implies that identifying optimal solutions is not feasible in large systems and that we need to look for helpful heuristics and practical ways of restricting the problem considered. In the work presented here we study a restricted class of problems where tasks are *spatially*

*classifiable*. We present a formal model of this restricted multi-robot task allocation (MRTA) problem called the *vacancy chain scheduling* (VCS) model. VCS simplifies the general scheduling problem in MRTA by dividing global system performance into individual robot contributions. This division, in turn, allows VCS to be used as a basis for a distributed, communication-free MRTA algorithm. We show, through experiments in simulations, that when greedy individuals continually update local task utility estimates, solutions reliably emerge which are optimal according to VCS. We present one specific VCS algorithm and give examples of other published algorithms that fit within the VCS framework. The experimental evidence shows that the VCS algorithm can improve the performance of a group of robots on a cooperative prioritized transportation problem and that the improvement goes beyond what is possible with hand-coded solutions. We also demonstrate that VCS can be made sensitive to differences in individual performance levels, making it suitable for MRTA in heterogeneous systems.

This article is organized as follows. In Section 2 we present the motivation for this work. In Section 3 we review previous work on MRTA and the modeling of group dynamics. In Section 5 we formalize MRTA and in particular the concept of group dynamics in the framework of scheduling. Section 6 presents VCS, both the formal model and our algorithm. In Section 8 we describe a set of experiments designed to validate the VCS model and to demonstrate the reliability and applicability of the algorithm. Section 9 presents the results from our experiments. Section 10 draws conclusions from the experimental evidence. Section 10 also discusses a number of issues that limit the algorithm, in particular

\* Corresponding author. Tel.: +44 1633 432344; fax: +44 1633 432407.

E-mail addresses: [torbjorn.dahl@newport.ac.uk](mailto:torbjorn.dahl@newport.ac.uk) (T.S. Dahl), [mataric@usc.edu](mailto:mataric@usc.edu) (M. Mataric), [gaurav@usc.edu](mailto:gaurav@usc.edu) (G.S. Sukhatme).

the possibility of convergence on sub-optimal allocations and unstable solutions.

## 2. Motivation

Increasing our insight into the fundamental principles of complex systems such as multi-robot systems, is necessary in order to increase the efficiency with which we control them and hence their general applicability. Systems where multiple agents interact with the real world and with each other, e.g., multi-robot systems, are inherently complex. Better control mechanisms for such systems will make them better suited for real-world applications such as industrial automation, in particular in domains with significant dependence on collaborative work.

Behavior-based control provides a way of handling the complexity of real world interaction by producing reliable solutions to specific problems through combining simple robust behaviors such as taxis and obstacle avoidance [2]. Machine learning algorithms can be used to optimize the application criteria of the different simple behaviors [3]. The work presented here is inspired by previous work on learning efficient expression criteria for behaviors and it explores potential performance gains in groups of robots. The work considers communication-free multi-robot learning in particular. Removing explicit communication completely from multi-robot learning is an extreme restriction. As a result, this work informs us about one of the limits of multi-robot control. There are also certain real-world problem domains where communication must be minimized or avoided because it is expensive, difficult or dangerous, e.g., underwater systems or stealthy operations where robots must maintain radio silence.

A vacancy chain is a robust resource distribution mechanism [4]. In nature we find vacancy chains forming in systems analogous to multi-robot systems, i.e., systems where individuals can evaluate and exchange resources. Based on this similarity, we use vacancy chains to a model a restricted version of general multi-robot task allocation (MRTA).

This work contributes to the growing body of research on formalizing multi-robot systems. Gerkey and Mataric [5] have previously presented a taxonomy of MRTA problems over three dimensions: single-task robots vs. multi-task robots, single-robot tasks vs. multi-robot tasks, and instantaneous assignment vs. time-extended assignment. Gerkey and Mataric discussed several frameworks for formalizing MRTA problems. Here, we develop the scheduling formalism for MRTA by formalizing the concepts of group dynamics and vacancy chains.

## 3. Related work

This work impacts on two important areas of research: MRTA and complex systems modeling. It focuses in particular, on systems where complexity is a result of interaction between individuals in a group, i.e., the group dynamics. In this section we review relevant work in these two areas and relate it to our approach.

### 3.1. Multi-robot task allocation

A number of algorithms for MRTA already exists. They fall broadly into two classes. The first class is pure coordination algorithms for distributing pre-specified tasks efficiently and reliably. The second class is task utility estimation algorithms which attribute changes in global performance to task utilities and individual or sub-group activities. A number of the utility estimation methods, like VCS, use machine learning to adapt allocation patterns based on observed runtime system properties.

#### 3.1.1. Coordination algorithms

Botelho and Alami's M+ algorithm [6] uses a task allocation protocol based on the Contract Net protocol with formalized

capabilities and task costs. The need to pre-define the capabilities and costs limits the applicability of the M+ algorithm to domains where these are known. It might, however, be possible to use the M+ algorithm with local utility estimates based on machine learning methods rather than hard-coded capabilities and costs.

Gerkey and Mataric's MURDOCH system [7], although independently developed, is also equivalent to the Contract Net protocol, and is based on a set of hard-coded metrics to locally score the suitability of the participating robots for a given task. MURDOCH uses a publish/subscribe protocol for communication and an auction mechanism for task allocation. As with the M+ algorithm, the hard-coded suitability metrics have limited applicability, but the MURDOCH system can also, conceptually, use local task utility estimates based on machine learning.

Werger and Mataric's work on the Broadcast of Local Eligibility algorithm for task allocation [8] compares locally decided eligibilities to allocate tasks using Port Arbitrated Behaviors, an inter-robot coordination mechanism based on a fully connected communication network. Their work on Cooperative Multi-Robot Observation of Multiple Moving Targets used spatial proximity as a measure of eligibility, but Broadcast of Local Eligibility can also use eligibility estimates based on machine learning.

In the L-ALLIANCE work by Parker [9], each robot explicitly estimates its own performance and the performance of other robots on selected tasks and uses these values to locally reallocate tasks by taking them over or acquiescing. The L-ALLIANCE algorithm uses local utility estimates and local communication to make local allocation decision and could potentially make use of machine learning.

Other Contract-Net and auction-based systems broadly equivalent to those described above are Hoplites [10], COMETS [11] and TAEMS [12].

Our approach differs from the above and other existing work in that it uses no direct inter-robot communication. All exchange of information between robots is done through stigmergic effects on the environment. The most important of these effects is congestion, the fact that robots get in each others way when they work on the same task. In our algorithm, the individual robots use their own task performance times to estimate levels of congestion. The absence of communication makes our algorithm applicable in domains where the opportunities for communication are limited. Reduced communication means that it is harder to identify globally optimal solutions. All distributed problem solving approaches, including global exhaustive searches, heuristic methods and learning methods can benefit from communication. Here we provide evidence that it is possible to get optimal solutions to emerge reliably *without* explicit inter-robot communication.

#### 3.1.2. Utility estimation algorithms

One way of dealing with  $\mathcal{NP}$ -hard problems is to use heuristic algorithms that produce satisfactory but potentially sub-optimal solutions. Learning such heuristics in the domain of scheduling is a well-studied problem. Here we review a selection of successful uses of machine learning as a way of finding heuristics for scheduling. We also discuss the applicability of these algorithms in the MRTA domain.

Zhang and Dietterich [13] presented a reinforcement learning approach to scheduling that learned domain specific heuristics for the scheduling procedure itself. The state space consisted of possible schedules and actions were possible changes to the schedules. The system learned what changes would quickly create feasible schedules with maximized capacity utilization. The problem domain considered was space shuttle payload processing.

Zomaya et al. [14] presented another algorithm for learning scheduling heuristics. Their algorithm learned *dynamic scheduling*, i.e., scheduling when there is no *a priori* knowledge about the tasks.

It used a back-propagation neural network and a history queue that functioned like an eligibility trace to learn how to associate a set of job parameters with a set of machines.

Both above algorithms are general and applicable to all scheduling problems. However, they rely on centralized learning mechanisms, which can fail critically when robots break down. These mechanisms also rely heavily on inter-robot communication and do not scale well to large groups.

Brauer and Weiß [15] used the multi-agent learning paradigm to distribute the scheduling of set of multi-operation jobs. The operations were totally ordered and each operation could only be performed by a subsets of machines. Using reinforcement learning, each machine estimated the efficiency of the possible predecessors and together the machines learned to improve the total production rate from an initial state where all machines were assumed to be equally efficient. Brauer and Weiß also showed that the learning allowed the system to adapt to machine breakdowns. This algorithm, however, is restricted to completely ordered jobs and this limits its applicability.

Blum and Sampels [16] studied different pheromone representations in ant colony optimization of first order job shop scheduling. Each job traversed a set of machines by stochastically selecting machines in accordance with a given set of constraints on the job's operations. Throughout the traversal, the system updated a pheromone trace, effectively learning environmentally embedded solution. The different pheromone representations allowed the system to estimate utilities for different combinations of machines, including estimating the utility of a machine in isolation and the utility of a machine depending on the last machine visited. This algorithm, however, is limited to jobs that consist of multiple operations. It is difficult to see how a similar pheromone trail could be used successfully as a general scheduling algorithm.

Tangamchit et al. [17] used distributed reinforcement learning on a group of robots to allocate patrolling tasks. Each robot locally estimated the utility of a set of patrolling points and together learned to divide the set of point between them in an optimal manner. Tangamchit et al. also used local and global proximity measures as heuristics for action selection in order to speed up learning. The specificity of the spatial heuristics limits the applicability of this algorithm to general task allocation problems.

The work reviewed above demonstrates the successful application of machine learning algorithms to scheduling problems. However, the domains considered in all of the work reviewed above, except for the work by Tangamchit et al., do not feature significant effects from group dynamics. In spite of this, all the learning algorithms reviewed above would likely improve a group's performance in the presence of significant effects from group dynamics, as they in general learn where to allocate jobs according to feedback based on processing time. The advantages of VCS over existing algorithms are its general applicability and its distributed learning mechanism.

Finally, since the initial presentation of VCS [18], Low et al., [40] have presented a similar self-organizing algorithm for ant-based MRTA. The approach taken by Low et al., like ours, uses local interference estimates. Low et al. call these *encounter patterns*. The domain considered by Low et al. is sensor network coverage. They use five to fifteen robots that self-organize using stigmergy. The results presented by Low et al. support our claims about the reliability and dynamic properties of MRTA algorithms based on local task utility estimation and stigmergic communication only. Similar results have also been reported by Labella et al. [19] using a swarm-based system to do foraging with four robots. All these results show that local utility estimates are sensitive to individual robot performance levels and that this can be used to differentiate between robots as lower levels of performance will lead to lower utility estimates and lower levels of commitment to high-value

tasks. The fact that the same results have been demonstrated in three different problem domains strongly supports our claim that this is a feasible solution for MRTA in domains with limited or no communication.

#### 4. Modeling group dynamics

We refer to the interactions between members of a group of robots as the *group dynamics*. The group dynamics produce group-level features such as congestion, when the robots have to spend large amounts of time avoiding each other. Cooperation can produce group dynamics that have a positive effect on performance. When group dynamics have a major effect on performance, it becomes necessary to understand and model those dynamics in order to predict the value of different solutions. There is always a significant level of uncertainty related to interacting with the real world in terms of accuracy of perception and effects of actions. In a multi-robot systems, the number of these interactions can grow exponentially, making it very difficult to predict future system states. Models of group dynamics must reflect this inherent unpredictability and, as a result, their accuracy is limited.

Models of group dynamics have been divided into *microscopic* and *macroscopic* [20]. Microscopic models explicitly represent each robot and may or may not be *simulation-based*. Simulation-based models simulate the actions taken by each of the robots so that properties of the system can be recorded as the robots interact. Macroscopic models describe systems properties on the basis of abstract features such as the number and general distribution of robots. They do not explicitly represent each robot. Below we review the most relevant work on producing and using explicit models of group dynamics to improve group performance.

##### 4.1. Microscopic models

Game theory [21] is an abstract model of multiple interacting agents as players in formalized  $n$ -person stochastic games. Game theory represents each player in terms of his/her *strategy* for playing a given game. Each game has a related *payoff matrix* for the possible combinations of moves from each player. On the basis of these factors, game theory describes some of the resulting features of the ensuing games, e.g., their optimality and stability. A game theoretic formulation of would be a payoff matrix reflecting how different combinations of individual moves would effect individual payoffs. As the group dynamics are typically not known, the construction of such payoff matrices is not possible. To overcome the unknown payoff matrices, Littman [22] presented a learning algorithm, the Minimax-Q algorithm, which is guaranteed to find the equilibrium of any stochastic game. As a general payoff policy for  $n$  tasks and  $m$  robots is of size  $n^m$ , it is not feasible in general to estimate the payoffs for each state/action combination. Bowling et al. [23] presented experimental data demonstrating how a reinforcement learning algorithm based on policy gradient ascent and the Win or Learn Fast (WoLF) principle for adjusting the learning rate can overcome this complexity. We discuss our approach to MRTA in terms of game theory in more detail in Section 10.

Goldberg and Mataric [24] developed a microscopic model called the Augmented Markov Models, transition probability matrices with additional temporal information, to learn statistical models of interaction in a space of abstract behaviors. They also used these models to maximize reward in a multi-robot mine collection task.

Balch et al. [25] studied live insect colonies and constructed three-dimensional histograms of insect presence over a discretized area. The work was a step toward a long term goal of combining

spatio-temporal models with Behavior Hidden Markov Models for behavior recognition [26] in order to recognize colony behaviors.

Yan and Mataric [27] have presented multi-level models of group behaviors from spatial data describing both human and robot activity. Like Balch et al., they used three-dimensional histograms to recognize and describe different activity distributions as produced by underlying behaviors.

Seth [28] has pointed out the distinction in biology between ‘phenomenological’ and ‘mechanistic’ models of interference, where the former identifies mathematical relationships between intake rates and robot density in empirical data while the latter constructs individual-based models with pre-specified rules for robot behavior. This distinction corresponds closely to that between macroscopic and simulation-based models used in robotics. Mechanistic models allow a derivation of the interaction between robot density and rule application with respect to intake rates. The phenomenological models assume that robots always optimize their individual intake rates. They also assume unstructured environments. These assumptions are often invalid both in biological systems and in the multi-robot task allocation domain. As an alternative to the traditional biological models, Seth presents a simulation-based model using genetic algorithms to evolve foraging behaviors for multiple robots in spatially explicit environments. The evolved systems are able to reproduce interference functions previously described in field studies of insects, but not reproduced by phenomenological models.

#### 4.2. Macroscopic models

Mataric [29] proposed a macroscopic model of interference as an estimate based on *group density*. Group density is defined as the ratio between the robots’ *footprints* and the available *interaction space*. A robot’s footprint is its sphere of influence, including factors such as geometry, motion constraints, and sensor range. If only the number and size of the robots are used, a *mean free path* can be computed and used to estimate the *number of expected collisions* for robots executing random walks.

Lerman et al. [20] studied three different types of models of cooperation and interference in groups of robots: sensor-based simulations, microscopic numerical models, and macroscopic numerical models. They used differential equations to model the group dynamics on a macroscopic level and showed that the three different models produced corresponding results. Unlike the sensor-based simulation and the microscopic numerical model, the macroscopic model had the advantage of being very fast and independent of the number of robots modeled. To make a macroscopic model tractable, however, many simplifying assumptions were necessary.

Macroscopic or phenomenological models are not generally sophisticated enough for optimizing specific task allocation algorithms. In particular, Mataric’s model for estimating interference does not consider synergistic effects or environmental complexity. The differential equation models produced by Lerman et al. similarly assume uniform distributions of robots and tasks. The simulation-based models or mechanistic models produce problem-specific solutions, but the time needed to construct those renders them unsuitable for use in MRTA algorithms.

Currently there are no models of the effects of group dynamics with the speed, generality, and predictive accuracy necessary for specifying the effects of interaction on task processing times in MRTA problems with the aim of constructing optimal schedules. Simulation-based models are in general too slow while macroscopic mathematical models make too many simplifying assumptions to be of predictive use. VCS, as a model of group dynamics spans both the microscopic and the macroscopic levels. It does not describe individual agents. However, it does explicitly

represent global system performance in terms of possible single agent contributions. This span allows VCS to be used both to reason about global systems performance and to inform the design of individual robot control algorithms. This duality makes VCS unique among current tools for modeling group dynamics.

### 5. MRTA as scheduling

MRTA is the problem of optimizing the allocation of tasks to robots over time, with respect to some given criteria. MRTA can be formalized within the existing framework of *scheduling* [1], a well studied problem. Formalizing MRTA problems as scheduling problems allows us to establish the formal complexity of these problems and to consider what classes of algorithms are suitable for solving them. When task allocation in multi-robot systems is formalized as scheduling, the terms *job* and *task*, and the terms *machine* and *robot*, become interchangeable. We use the terms *job* and *machine* when relating the task allocation problem to other scheduling problems and the terms *task* and *robot* when we discuss specifics of MRTA.

#### 5.1. Problem definition

Scheduling problems are described in terms of a set of **machines**,  $M_j$  ( $j = 1, \dots, m$ ), that process a set of **jobs**,  $J_i$  ( $i = 1, \dots, n$ ). Jobs may consist of a set of **operations**,  $O_{i1}, \dots, O_{in_i}$ . If any job can run on any machine, the machines are called **parallel**. Parallel machines are the most general instance of a class called **multi-purpose machines**, (MPM). In MPM, a job can be processed on any machine which has the appropriate tools. If the machines in  $M_j$  are used simultaneously, the scheduling problem is called a **multiprocessor task scheduling problem**. In the simplest case where  $n_i = 1$ , each job,  $J_i$ , has a related **processing requirement**,  $p_i$ , and a cost function,  $f_i(t)$ , reflecting the cost of completing  $J_i$  at time  $t$ . The cost function,  $f_i$ , may use a **due date**,  $d_i$ , and a **weight**,  $w_i$ . Further constraints on precedence, preemptability, and batching of jobs are also common. The machine environment is subject to a number of further formalizations. Parallel machines are divided into three classes: **identical** machines,  $P$ , **uniform** machines,  $Q$ , and **unrelated** machines,  $R$ . For identical machines,  $P$ , the processing time of a job,  $J_i$ , is the same on all machines,  $M_j$ , i.e.,  $p_{ij} = p_i$ . For uniform machines,  $Q$ , each machine has a related speed,  $s_j$ , and the processing time of a job is dependent on the machine, i.e.,  $p_{ij} = p_i/s_j$ . Lastly, for unrelated machines,  $R$ , the processing time may be dependent on the machine-job combination, i.e.,  $p_{ij} = p_i/s_{ij}$ .

A common optimization criterion in scheduling is the **weighted total flow time**, denoted  $\sum w_i C_i$ , where  $C_i$  is the **finishing time** of each job,  $J_i$ . Other common optimization criteria are defined with respect to other values than the total flow time, e.g., job earliness, job tardiness, or deviation from deadline. Scheduling problems in general can be denoted by three main features: the class of machines, the attributes of the jobs, and the function to be optimized, e.g.,  $Q \mid p_i = p \mid \sum f_i$ , denotes a problem with uniform machines and uniform task processing times, where the value to optimize is the sum of the cost. Some scheduling problems are solvable in polynomial time. Other scheduling problems, such as finding the minimal weighted total flow time for identical machines,  $P \parallel \sum w_i C_i$ , have been shown to be  $\mathcal{NP}$ -hard [1].

#### 5.2. Formalizing group dynamics

In a static environment, the group dynamics depend only on what machines are allocated what jobs. An *allocation* is a function mapping machines to jobs at a given point in time. For  $m$  machines and  $n$  jobs, this function can be represented as a vector  $\mathcal{A}$  of size  $m$ ,




Fig. 1. Job-oriented Gantt chart.

where vector element  $i$  indicates what job is allocated to machine  $i$ . We use  $\emptyset$  to indicate a machine that has not been allocated a job. Informally, an allocation can be seen as a slice of a *schedule*, i.e., allocations of jobs to machines over time, as represented by a Gantt chart [1]. The dashed vertical lines in Fig. 1 illustrate this for two allocations,  $\mathcal{A}_1$  and  $\mathcal{A}_2$ . The allocations are also resented in Eqs. (1) and (2).

$$\mathcal{A}_1 = [M1, M2, M3, \emptyset] \quad (1)$$

$$\mathcal{A}_2 = [\emptyset, \emptyset, \emptyset, M1]. \quad (2)$$

We suggest that the effects of group dynamics on job processing times can be included in the formal scheduling framework by making the job processing time,  $p_i$ , from the time the job is started,  $t_s$ , to the time the job is finished,  $t_f$ , a function of the allocations during this time,  $\mathcal{A}_{t_s, t_f}$ , as shown in Eq. (3).

$$p_i = g_i(\mathcal{A}_{t_s, t_f}). \quad (3)$$

We call the function  $g$  the *interaction function*. We call the class of machines whose performance is significantly affected by group dynamics *interacting machines*. Interacting machines can be identical, uniform or unrelated, so we indicate the fact that machines are interacting using the subscript  $l$ . The interaction function, as a formalization of group dynamics, captures much of the complexity of group dynamics in otherwise static environments such as transportation between stationary locations. It is worth mentioning that this is not a suitable formalization of group dynamics in more dynamic problems, e.g., problems where the processing times are also dependent on a changing spatial distribution of jobs. An example of such a problem is cooperative tracking and pursuit [30].

### 5.3. Proof of complexity

Here we present a trivial, but formal proof that scheduling problems where job processing times are affected by machine allocation patterns, i.e., group dynamics, are  $\mathcal{NP}$ -hard.

**Theorem 1.**  $R_l \parallel \sum w_i C_i \in \mathcal{NP}$ -hard.

**Proof.** The optimization problem  $R_l \parallel \sum w_i C_i$  describes a set of unrelated, interacting machines,  $R_l$ , a set of non-preemptable jobs,  $J_i$ , with related weights,  $w_i$ , and processing times,  $p_i$ . The machine speeds,  $s_{ij}$ , are job dependent. Unrelated machines are a sub-class of interaction-dependent machines in that the function describing the task processing time,  $p_{ij} = p_i/s_{ij}$ , is a deformed interaction function that is only dependent on the task and machine in question and not on what tasks are allocated to the other machines. We can create, in polynomial time, a corresponding problem  $R_{ID} \parallel \sum w_i C_i$  with an identical set of machines,  $R_l$ , an identical set of weighted jobs, but with task processing times described by the interaction function  $g_i(\mathcal{A}_{t_s, t_f})$ .

The assumption that the machines in the original problem are unrelated, but not interaction-dependent, implies that the value of  $g_i(\mathcal{A}_{t_s, t_f})$  is only dependent on which machine  $M_j$ , the related

job,  $J_i$ , is allocated to in the relevant allocations  $\mathcal{A}_{t_s, t_f}$ . For non-preemptable jobs,  $M_j$  will be the same for all those allocations. The interaction functions  $g_i(\mathcal{A}_{t_s, t_f})$  can be produced in polynomial time, by finding which machine job,  $J_i$ , is allocated to in the initial allocation,  $\mathcal{A}_{t_s}$ . The job,  $J_i$ , and machine  $M_j$ , can then be used to find the machine speed using the given function,  $p_i/s_{ij}$ .

An algorithm solving  $R \parallel \sum w_i C_i$  problems uses sets of machines, sets of jobs, and optimization criteria that are identical to the ones used by an algorithm for solving the corresponding  $S \parallel \sum w_i C_i$  problems. The only difference between the problems is the representation of the function for job processing times. We have shown how the function  $p_i/s_{ij}$  can be accurately represented in the more general form  $g_i(\mathcal{A}_{t_s, t_f})$ , and that the transformation can be done in polynomial time. Hence, an optimal solution to the new problem would be an optimal solution to the original problem. This again shows that  $R \parallel \sum w_i C_i$ , which is  $\mathcal{NP}$ -hard, reduces to  $R_l \parallel \sum w_i C_i$  which, *ipso facto*, must also be  $\mathcal{NP}$ -hard.  $\square$

## 6. Vacancy chain scheduling

We developed the VCS formalization of MRTA to address two of the main problems in systems where the effects of the group dynamics are significant: first, the engineering problem of overcoming scheduling complexity, as discussed in Section 5.1, and, second, the scientific problem of modeling of group dynamics, as discussed in Section 4.

### 6.1. Vacancy chains

The inspiration for VCS is the *vacancy chain* process [4], through which resources are distributed to consumers. The typical example of a vacancy chain is a bureaucracy where the retirement of a senior employee creates a vacancy that is filled by a less senior employee. This promotion, in turn, creates a second vacancy to be filled, and so on. The vacancies form a chain linked by the promotions. The resources distributed in this example are the positions and the consumers are the employees.

The general process of distributing resources in this way has been recognized in many different domains. It was originally reported in human populations relating to houses and apartments as well as to jobs in bureaucracies. Chase [4] proposed that major human consumer goods such as cars, boats, and airplanes, also move through vacancy chains and that vacancy chains are common in other species such as the hermit crab, the octopus and different species of birds. In the case of the hermit crab, the empty gastropod shells they carry around as portable shelters are distributed through vacancy chains.

Chase lists three requirements for resource distribution through vacancy chains:

- (1) The resource must be reusable, discrete, and used by only one individual.
- (2) A vacancy is required before an individual takes a new resource unit, and individuals must need or desire new units periodically.
- (3) Vacant resource units must be scarce, and many individuals must occupy suboptimal units.

The vacancy chain resource distribution mechanism is both simple and powerful. It is based on *stigmery*, unintentional communication between the consumers through their effects on the environment [31]. This makes distribution through vacancy chains robust and efficient in large groups/societies.

While distribution through vacancy chains ensures that the most attractive resources are consumed, it does not, like more sophisticated resource distribution mechanisms, take into account the quality of the consumer. As such, the vacancy chain distribution

process cannot exploit the possible advantages of distributing particular resources to particular consumers. The vacancy chain distribution mechanism treats all consumers as equals. In terms of task allocation, the vacancy chain distribution algorithm does not guarantee optimal allocations as it does not consider possible differences in machine speed. Our initial study of vacancy chain distribution reflects this fact by considering only homogeneous groups of robots.

## 6.2. Spatially classifiable jobs

In order to create a common framework for scheduling theory and resource distribution through vacancy chains, we introduce the concept of *spatial classifiability*. By seeing space as a restricted resource of varying quality we can distribute jobs using vacancy chains. In certain problems, e.g., collaborative transportation, jobs can be divided into a set of classes,  $K_k$ , where there are no significant effects on a group's performance from interactions between machines working on tasks from different classes. This means that the interaction function for a given job need only consider the allocation pattern for the jobs within the jobs spatial class,  $\mathcal{A}_{t_s, t_f, k}$ . In terms of scheduling, we call such jobs *spatially classifiable* and denote them  $sc$ .

The problem of optimizing the weighted total flow time in a system with interacting machines and with spatially classifiable jobs is denoted  $P_I | sc | \sum w_i C_i$ . Our motivation for defining this class is our work on problems where robots interfere with each other when working on the same task, but where different tasks take place in spatially separate areas. For such problems, a job,  $j_i$ , has a processing time,  $p_i$ , that depends only on the interaction between the machines currently allocated to jobs in that job's class,  $\mathcal{A}_{t_s, t_f, k}$ , as expressed in Eq. (4). One problem that falls within this class is the prioritized transportation problem described in Section 7.

$$p_i = g_i(\mathcal{A}_{t_s, t_f, k}). \quad (4)$$

Although the allocation of robots to jobs in other spatial classes does not impact the processing time of a task in a given class, it remains impossible to predict accurately an individual robot's contribution to the class processing time. As a result it is not possible to predict the optimal global allocation pattern. Considering only spatially classifiable jobs is a simplification of the general MRTA problem and, by reasoning similar to that used for general interacting machines, it can be shown that MRTA on spatially classifiable jobs is still an  $\mathcal{NP}$ -hard problem. What spatial classifiability gives us is the opportunity to apply the VCS model and with it the possibility of using the VCS algorithm for efficient MRTA.

## 6.3. Homogeneous robots

Within the class of interacting machines processing spatially classifiable jobs, we initially restricted ourselves further to looking at problems where the machines are *homogeneous*. By homogeneous machines, we mean robots that are of the same type, i.e., have the same physical and behavioral capabilities, e.g., identical turning rates and traversal speeds in unobstructed environments and the same behavioral strategies for problems such as avoiding obstacles and following walls. A group of robots working on different tasks in different locations is considered homogeneous as long as their physical and behavioral capabilities are the same. In this case, the interaction function is dependent only on the *number* of robots,  $m_k$ , currently working on jobs in each class,  $K_k$ . The job processing time for this class is given in Eq. (5).

$$p_i = g_i(m_k). \quad (5)$$




Fig. 2. Example number and ranks of service-slots for two task classes.

The problem of scheduling with two or more identical machines,  $P2 || \sum w_i C_i$ , is known to be  $\mathcal{NP}$ -hard [1]. In this problem, the job processing time,  $p_i$ , is only dependent on the job,  $j_i$ . By the same reasoning as with unrelated interacting machines, it can be shown that  $P2 || \sum w_i C_i$  reduces to  $P2_I | sc | \sum w_i C_i$ , i.e., scheduling spatially classifiable jobs over two or more interacting machines. Hence,  $P2_I | sc | \sum w_i C_i$  is also  $\mathcal{NP}$ -hard. In Section 8 however, we present experimental data showing that VCS is also applicable to heterogeneous groups of robots, i.e., problems where job processing times also depend on the machine processing the job. In this case, VCS not only distributes the robots over different job classes, like a greedy solution would have done, but also allocates fast robots to high-value jobs and slow robots to low-value jobs.

## 6.4. The VCS model

Inspired by the vacancy chain process for distributing resources to consumers we developed a corresponding formal model of tasks allocation for spatially classifiable tasks. According to the VCS model, any number of robots/machines can be assigned to tasks/jobs from a given class. A particular number of homogeneous robots,  $j$ , servicing the same class of tasks,  $i$ , will have a corresponding task processing frequency,  $c_{i,j}$ . This task processing frequency depends on how efficiently  $j$  robots are able to work concurrently on task  $i$ . As an increasing number robots are assigned to a task, the task processing frequency will typically increase up to a given point where the robots start getting in each others way. When a robot is assigned to a task, the resulting difference in the task processing frequency, together with the task value,  $w_i$ , define the contribution of that robot.

When robots are homogeneous, the contribution of the last robot added is independent of which robots are in the group. Because of this, VCS abstracts away individual robots and models tasks as an ordered list of slots that can be filled by any robot. The contribution of one robot in a group of size  $j$ , in terms of change in performance from a group of size  $j - 1$ , is called the *slot-value*,  $s_{i,j}$ . When the assignment of an additional robot to a task leads to a decrease in the task processing frequency, the slot-value, correspondingly, becomes negative. The slot-values define a global ranking of the service-slots. This ranking is dependent on the group dynamics and, hence, it is not possible to establish it by analytical means. However, slot values can be estimated at runtime allowing robots to fill service slots in order of ranking. A graphical example of ranked service-slots for two classes of tasks is presented in Fig. 2.

The formal definition of slot-value in terms of difference in task processing frequency is given in Eq. (6).

$$s_{i,j} = w_i(c_{i,j} - c_{i,j-1}). \quad (6)$$

From a resource distribution perspective, service-slots are resources to be distributed among robots. Allocating robots to tasks based on service-slot values is what we refer to as *vacancy chain scheduling*. According to the VCS model, an allocation is optimal when it maximizes the sum of the filled service-slots. With spatially classifiable tasks, a group of robots can optimize the value of completed tasks over time, i.e., the *throughput*, by

allocating robots to tasks or service-slots in order of decreasing slot-value. The total system performance or throughput value,  $T$ , for  $n$  robots, is simply the sum of the individual contributions as stated in Eq. (7).

$$T_n = \sum_i \sum_j s_{i,j}. \quad (7)$$

Note that this problem has the *greedy-choice* property [41] but is non-trivial as the slot values, unlike the task-values, are not directly available to the robots. The problem of allocating tasks to robots over time based on the task values has not got the greedy-choice property due to the effects of the group dynamics.

In a scenario where the service-slots are allocated optimally, a failure in a robot will result in an empty service-slot. If the value of the vacant slot is greater than the value of one of the occupied slots, the vacant slot should be filled in order to restore optimal allocation. Expressed in vacancy chain terminology, a vacant high-value service slot is a resource to be distributed between the robots. The VCS model does not apply to problems that are not spatially classifiable. In such problems there is no clear concept of service-slot as a robot would contribute to the average task completion time of multiple job classes simultaneously.

VCS reduces the complexity of the general MRTA problem by considering only *spatially classifiable* tasks and homogeneous groups of robots. The experiments presented in Section 8, however, go beyond the VCS model in that they also evaluate the efficiency of the VCS algorithm, presented in Section 6.5, when applied to the control of heterogeneous groups of robots.

### 6.5. The VCS algorithm

Because slot values can be locally estimated and because the slot value-based task allocation problem has the greedy-choice property, a slot-value based MRTA algorithm can be distributed by letting each robot optimize the value of the service-slot it occupies based on its local estimate. The local task utility estimates depend on the effects of the group dynamics produced by the allocation of all the other robots and reflect the contribution of the individual robot to the relative ranking of the tasks. It is the effects of the group dynamics on the individual task utility estimates that make the optimal allocation patterns emerge.

We have previously studied the use of distributed reinforcement learning (RL) [32] as a way of optimizing the performance of multi-robot systems in an interaction-sensitive manner [33]. Our VCS algorithm is a specialization of that work for MRTA. For each individual robot, our algorithm handles traditional RL actions that correspond to tasks in the MRTA framework. The algorithm's state space reflects the class of the last task completed by the robot. Each robot keeps local estimates of task utilities in form of Q-values, and chooses its tasks using an  $\epsilon$ -greedy action selection function. Basic Q-learning is not sensitive to the frequency of rewards. Hence, the estimated values of actions do not necessarily correspond to their contribution to performance over time. Rather than using one of the time-sensitive reinforcement learning frameworks, such as Semi-Markov Decision Processes [34], we made the temporal aspect of performance explicit in the reward function. The reward function in terms of the last task processing time,  $t$ , and the task value,  $w_i$ , is given in Eq. (8).

$$r = w_i/t. \quad (8)$$

This reward function promotes the actions with the highest contributions to the system performance because these will on average provide a higher reward. If a robot consistently occupies a service-slot that is suboptimal due to too high levels of interference, the increased average traversal time will reduce the average reward for that slot below the average reward of the

optimal service-slots. This change in average reward will attract the robot back to an optimal slot. Eq. (9) presents the relationship between the definitions above by expressing the relationship between the individual task throughput,  $T_i$ , the slot values,  $s_{i,j}$ , the task value,  $w_i$ , the task completion frequencies,  $c_{i,j}$ , the average task processing times,  $\bar{t}$ , and the average reward,  $\bar{r}$ , for  $j$  robots servicing a given task,  $i$ .

$$T_i = \sum_j s_{i,j} = w_i \sum_j (c_{i,j} - c_{i,j-1}) = w_i \sum_j \frac{1}{\bar{t}_{i,j}} = \sum_j \bar{r}_{i,j}. \quad (9)$$

#### 6.5.1. Softmax action selection as a performance differentiator

When robots keep a set of task-related utilities, different functions can be used to select the next task to undertake based on these utilities [34]. With a Greedy- $\epsilon$  function, all the tasks, apart from the one with the highest utility, have equal probability of being selected. With a softmax function however, the probability of selecting a task is correlated with the relative estimated utility of that task. Using a softmax function on the top level of our control architecture has reliable effects on an inter-robot level, where it functions as a mechanism for allocating high-value tasks to high-performance robots. A robot that on average can service tasks in time  $\bar{p}$  will have a difference in estimated task utility that correlates with the expression  $\frac{w_h - w_l}{\bar{p}}$  where  $w_h$  and  $w_l$  are the values of high- and low-value tasks respectively. A fast robot with a lower average service time  $\bar{p}$  will have a correspondingly higher difference between the estimated utility of high- and low-value tasks. Using a softmax action selection function, this greater difference in estimated utility theoretically translates into a greater probability of servicing high-value tasks. Such persistence will lead to the fast robots servicing high-value tasks and may, depending on the group dynamics and the task values, lead to the slow robots servicing low-value tasks.

## 7. Prioritized transportation

In the basic transportation problem, a group of robots traverse a given environment in order to transport items between the sources and the sinks. We call the time taken to traverse the environment once from a sink via a source and back to a sink the *traversal time*. To perform optimally on this task the robots must maximize the number of traversals in general. The basic transportation problem [33] is one of the sub-problems of foraging [35–37]. If the locations of sources and sinks are given, the foraging problem is reduced to a problem of transportation.

Here we present the *prioritized transportation* problem, which generalizes the basic transportation problem to problems where sources and sinks may have different values. Because the robots get in each others way, an algorithm where individual robots greedily select high-value tasks may fail to find the optimal solution to this problem allocation. To optimize performance, a group of robots must strike the correct balance between different target values and levels of interference. In Section 9 we present experimental evidence showing that the VCS algorithm can solve this problem using continual individual learning.

We consider fetching and delivering one puck to constitute one *task*. In terms of scheduling, each transportation task corresponds to a job and the traversal times correspond to job processing times. The value of a task/job corresponds to its weight. The robots correspond to machines and, since their allocation has a significant effect on the performance of the system, the robots/machines are interacting. When the sources and sinks are distributed so as to make the tasks/jobs spatially classifiable, the problem of optimizing throughput in prioritized transportation is an instance of either  $P2_D|sc| \sum w_i C_i$  or  $Q_1|sc| \sum w_i C_i$  according to whether the robots are homogeneous/identical or heterogeneous/uniform.

## 8. Evaluation

To evaluate both the VCS model and the VCS algorithm we conducted a set of experiments in the domain of prioritized transportation. We first conducted experiments to validate the VCS model as a model of MRTA. These experiments also demonstrated the validity and efficiency of the VCS algorithm for homogeneous robots. Second, we conducted an experiment which demonstrated that the VCS algorithm, if using a softmax action selection function, is sensitive to different levels of performance among the robots and hence can be applied to heterogeneous teams.

### 8.1. Validating the VCS model

The goal of our validation experiments was two-fold. First, we wanted to show both that the allocations produced by the VCS algorithm satisfied the predictions of the VCS model. This would indicate that the model could be used to produce efficient solutions in spatially classifiable scheduling problems. Second, we wanted to show that the performance of our algorithm improved on the performance of any non-adaptive solutions. We designed three initial experiments:

- (1) **Base distribution:** The goal of this experiment was to demonstrate that the VCS algorithm would reliably distribute six robots over two tasks of different value in an optimal way according to the VCS model.
- (2) **Filling a vacancy:** This experiment was designed to demonstrate that the VCS algorithm could recover from a change in the optimal allocation by filling a vacancy. We created a vacancy by removing one of the robots occupying a high-value service-slot. According to the VCS model, this vacancy should be filled by one of the robots servicing a low-value service-slot. Re-establishing optimality in this way would demonstrate the VCS algorithm can improve on the performance of hand-coded solutions.
- (3) **Breakdown without vacancy:** This experiment was a control experiment showing that the removal of a robot from a low-value service-slot did not lead to a change in the allocation of the remaining robots. According to the VCS model, such a change does not create a vacancy.

Together, these three experiments demonstrate that the VCS algorithm establishes and maintains optimal allocations as defined by the VCS model. To show that the performance improvement in our experiments are due to the VCS algorithm, we also ran the three control experiments using robots that randomly chose between the available tasks. To show how the vacancy chain algorithm outperforms hand-coded MRTA algorithms, we also ran an experiment where the assumed optimal allocations for the initial setup was hand-coded in the robot controllers. We emphasize that in the general case, it is not possible to reliably identify optimal allocations.

#### 8.1.1. Learning parameters

The robots used in the initial experiments were homogeneous with respect to hardware configuration and control algorithm. Over time, the reinforcement learning differentiated the group by specializing robots on different tasks. The robots used temporal difference Q-learning [34] to associate the different states with the available actions. The Q-tables were initialized with random values between  $-0.1$  and  $0.1$ , the learning rate  $\alpha$  was set to  $0.1$ , and the discount factor  $\gamma$  was set to  $0.95$ .

The input- or state-space reflects the circuit the robot used for its last traversal. This allowed the robots to learn policies that served alternating circuits. In spite of this, the robots consistently learned policies dedicated to a single circuit. The action space

included servicing one of  $N$  available circuits. The tasks were always available, so whenever an action was complete, a robot was free to choose its next action from servicing either circuit. This setup gave each robot an  $N$ -by- $N$  sized Q-table. For the experiments validating the VCS model we used a greedy- $\epsilon$  action selection strategy [34], where  $\epsilon$  was set to  $0.1$ . Because we wanted the system to remain adaptive to changes in the environment, we did not decrease  $\epsilon$  over time, as is common.

#### 8.1.2. Reward function

We empirically estimated a range of relevant average traversal times and, using the VCS model, designed a reward function that was likely to demonstrate vacancy chain distribution. In order to produce an initial allocation where three robots serviced one circuit and three robots serviced the other, it was necessary that it was less attractive to be one of four robots servicing the high-value circuit than to be one of three servicing the low-value circuit. This constraint on the reward function is presented formally in Eq. (10).

$$\forall(x, y). \bar{r}_{x,4} < \bar{r}_{y,3}. \quad (10)$$

In order for a vacancy in the high-value circuit to be filled, it must be more attractive to be the third robot in that circuit than to be the third robot in the low-value circuit. This is expressed formally in Eq. (11), where  $p$  denotes the preferred circuit.

$$\forall(x \neq p). \bar{r}_{x,3} < \bar{r}_{p,3}. \quad (11)$$

To satisfy the constraints given these times, we set the circuit values in Eq. (8) to be  $w_1 = 2200$  and  $w_2 = 2000$ . We call the circuit with the highest reward the *high-value* circuit and, correspondingly, the circuit with the lowest reward the *low-value* circuit. Note that the reward values were set purely based the VCS model and independently estimated parameters, such as the average task completion frequencies. They were not set based on evidence from running the VCS algorithm. The fact that the values were based on the VCS model means that the experimental results would either validate or invalidate the model. The robots do not keep any information about the task values and that the optimal allocation emerges as each robot's individual utility estimates converge. The fact that the robots are oblivious to task values also allows new allocations to emerge when external factors change, e.g., the task values or the group size.

### 8.2. Controller architecture

All the robots in our demonstration used the same behavior-based controller. Importantly, however, our VCS algorithm is independent of the underlying architecture and is defined purely in terms of distributed task utility estimation. The controller used by each of the robots in our experiments had three levels. The top level was a reinforcement learning algorithm controlling the application of a set of high-level behaviors available in the middle level. On the middle level, each behavior corresponded to servicing one of the available tasks. These behaviors again, consisted of multiple lower level behaviors which were shared between the high-level behaviors. The low-level behaviors were as follows:

- **Obstacle avoidance:** avoided obstacles detected by laser range finder in the desired path
- **Visible target approach:** approached the desired target when it was visible to the laser range finder
- **Target location approach:** approached the location of the desired target when the target itself was not visible
- **Wall following:** followed the walls to the first available target when the desired target was not visible and the odometry-based localization was deemed to be inaccurate.






Fig. 3. The three-level robot control architecture.

A graphical illustration of the robot control architecture is presented in Fig. 3.

Localization was deemed to be inaccurate whenever the desired target was not visible, but should have been so according to the robot's estimated position and orientation. On encountering a target, the localization estimate was reset and again deemed to be accurate.

### 8.3. The simulated environment

The experiments were done in simulation using the *Player* robot device server and *Stage* simulator [38]. Controllers written for *Player/Stage* have been repeatedly shown to work with little or no modification on real mobile robots. The robots in the experiments were simulated Pioneer 2DX robots with SICK laser range-finders. The Pioneer 2DX platform is a common platform for mobile robot research. It is a highly holonomic differential drive robot 44 cm long, 38 cm wide and 22 cm high. The real Pioneer 2DX has two 16.5 cm diameter drive wheels and one caster at the rear. It is capable of rotating in place and can carry a payload of up to 23 kg. We used a simulated SICK laser range finder as the only sensor for the low level behaviors described in Section 8.2 and simulated binary coded laser beacons [39], i.e., bar codes that can be read by a laser range finder.

In order to verify the VCS model we designed an environment with enough interference between the robots to make a greedy approach to MRTA sub-optimal, i.e., a approach where all available robots service the high-value task. We also designed the environment in such a way that the failure of one robot would be significant enough to require a redistribution of robots in order to maintain optimal system performance. The experiments took place in a simulated 12 by 8 m environment with six robots and two sets of sources and sinks. The sources and sinks were marked with simulated laser beacons. We did not require actual objects to be carried. A proximity of less than 1 m to a source or sink was interpreted as a delivery or a pick-up. The relatively high number of robots and a relatively low number of sources and sinks, or task classes, emphasizes the fact that the dominant factor in the complexity of MRTA with spatially classifiable tasks is the number of robots, not the number of tasks. The group dynamics in this system are too complex to allow accurate predictions of optimal allocation patterns. Our design of the environment was based on rough experimental estimates of completion times using similar environments and different numbers of robots for each circuit. The experimental results presented below provide evidence indicating that the allocation pattern in Fig. 4 is optimal and that the VCS model is a valid model of MRTA for spatially classifiable tasks. Fig. 4 shows a graphical rendering of the simulated environment in which the experiments took place, with the circuits indicated by dashed arrows and the sources and sinks labeled.

### 8.4. VCS for heterogeneous robots

In addition to the experiments for validating the VCS model as a model of MRTA, we designed a second experiment to evaluate




Fig. 4. The simulation environment, with circuits indicated.

Table 1

State-time distributions for six robots.

State		0:6	1:5	2:4	3:3	4:2	5:1	6:1
VCS	$\hat{\mu}$	0.1	2.7	19.0	41.8	29.3	6.6	0.5
	$s$	0.3	1.9	7.7	6.6	9.0	3.2	0.5
R	$\hat{\mu}$	2.0	7.2	22.6	34.5	24.7	8.2	0.7
	$s$	0.6	2.9	3.4	3.1	3.3	3.1	0.4
T		1.6	9.4	23.4	31.2	23.4	9.4	1.6

the efficiency of the VCS algorithm for heterogeneous robots. We had two goals for this experiment. First, to test whether the modified VCS algorithm could produce allocations where high-value tasks are allocated to high performance, i.e., fast, robots. Second, we aimed to demonstrate that this allocation improved the performance of the system to a level significantly above the performance level of a group where tasks were allocated randomly.

For this experiment, the robots were divided into two groups. The first group was made to operate at a default speed of 0.3 m/s. The second group had a default speed of 0.2 m/s. The temperature parameter  $\tau$  was set to 0.005. With these values the experimentation rate was significant without being overwhelming. Because we wanted the system to remain adaptive to changes in the environment we did not decrease  $\tau$  over time, as is common.

## 9. Results

The following analysis of the results from the four experiments defined in Section 8 uses student-*t* tests for statistical significance with respect to a 90% confidence level. For each experiment we defined a convergence period and a stable period according to the stability of the system performance.

### 9.1. Base distribution

This experiment used six robots. It consisted of 20 10-h trials, each averaging 3000 traversals in total or 500 traversals per robot. The convergence period was 2.5 h.

To evaluate performance, we consider the last target visited by each robot. This gives seven possible system states. We refer to each state using the notation  $h : l$ , where  $h$  is the number of robots whose last target was on the high-value circuit. Correspondingly,  $l$  is the number of robots whose last target was on the low-value circuit. The rows labeled VCS in Table 1 show the mean,  $\hat{\mu}$ , and standard deviation,  $s$ , of the time the system spent in each of the states. The values are percentages of the total stable period. The rows labeled R describe the same values for a set of 20 trials using a group of robots that randomly chose between tasks.

The row labeled T lists the number of different ways to choose a sample of size  $n$  from a population of  $m$ , as percentage of all




Fig. 5. The state-time distributions of six randomly allocated robots and six VCS-controlled robots.




Fig. 6. The performance of VCS and random allocation for six robots.

possible samples, according to Eq. (12). It is worth noticing that the time distribution produced by random allocation is closely aligned with this theoretical estimate, though the differences are statistically significant. We do not suggest that a random allocation algorithm is a sensible choice for a system, we include the data here to illustrate the underlying probability distribution affecting the VCS algorithm due to the random experimentation.

$$T = \frac{m!}{n!(m-n)!2^m} \quad (12)$$

The two time distributions given in Table 1 are presented as histograms in Fig. 5 with the standard deviation indicated by the error bars for each state.

Over 20 experiments, the difference in time spent in state 3:3 is statistically significant. The time the VCS-controlled group spent in state 3:3 is also statistically higher than the time spent in any of the other states. This confirms that the group's action-selection policies have converged to promote the state defined as optimal by the VCS model, given the estimated task-processing times.

Fig. 6 presents the average performance of a group of robots controlled by the VCS algorithm over both the convergence period and the stable period. The performance is averages over the whole of each period. This group's performance is indicated by the thick, solid line. The average performance of random allocation is indicated by the dashed line. The thin solid line indicates the performance of a group of robots where three robots were hand-coded to serve the high-value circuit while the three remaining robots were hand-coded to server the low-value circuit. The performance is calculated as the sum of the delivery frequencies for each circuit weighted by the value of the task.

The performance data show that a group of robots controlled by the VCS algorithm performs significantly higher than six randomly allocated robots. Together, the time distribution data and the performance data show that the adaptive controllers improve the group's performance by adopting a dedicated service structure that conforms to the predictions of the VCS model.




Fig. 7. The state-time distributions of five randomly allocated and five VCS-controlled robots after a breakdown created a vacancy on the high-value circuit.

Table 2

State-time distributions of five robots after a breakdown created a vacancy on the high-value circuit.

State		0:5	1:4	2:3	3:2	4:1	5:0
VCS	$\hat{\mu}$	0.8	7.8	31.1	40.7	17.7	1.8
	$s$	1.1	5.1	6.3	5.3	5.9	1.6
R	$\hat{\mu}$	2.6	12.6	32.4	34.8	15.5	2.1
	$s$	0.9	2.2	4.0	2.3	3.7	0.4
C		3.1	15.6	31.3	31.3	15.6	3.1

9.2. Filling a vacancy

This experiment used five robots. We removed randomly one of the three robots that were servicing the high-value circuit at the end of the previous experiment, thus creating a vacancy on that circuit. The experiment consisted of 20 10-h trials. The convergence period was 5 h.

The state-time distribution of the converged VCS algorithm shows that the system is in state 3:2 for a significantly larger amount of time than a group of five randomly allocated robots. The state-time distribution is given numerically in Table 2 and as a bar chart in Fig. 7. This showed that the group had adapted its structure from one that promoted the 3:3 state to one that promoted the 3:2 state. The change demonstrated that a robot from the low-value circuit had filled the vacancy we had created in the high-value circuit.

The system performance is presented in Fig. 8. It shows that the removal of a robot from the high-value circuit caused the performance to drop sharply. After the re-convergence period, the performance rose again to a level that was significantly higher than the performance of random allocation and also significantly higher than the mean performance, over 20 trials, of a group of robots controlled by a hand-coded allocation algorithm optimized for six robots. The average performance of that group is indicated by the thin solid line.


9.3. Breakdown without vacancy

This experiments used five robots from the end of the first experiment. We removed randomly one robot from the low-value circuit. According to the VCS model, this did not create a vacancy and, hence, the system was expected to remain in the 3:2 state. The experiment consisted of 20 10-h trials, and the convergence time was 2.5 h.


The state-time distribution during the stable period of this experiment, presented in Table 3, was not significantly different from the distribution produced during the second experiment presented in Section 9.2.

**Table 3**  
State-time distributions after a breakdown that did not create a vacancy.

State	0:5	1:4	2:3	3:2	4:1	5:0
$\hat{\mu}$	0.3	6.7	34.6	47.1	10.5	0.7
$s$	0.3	3.7	9.2	9.4	3.8	0.4



**Fig. 8.** The performance of VCS, random allocation, and hand-coded control before and after a breakdown creating a vacancy on the high-value circuit.



**Fig. 9.** The performance of VCS and random allocation before and after a breakdown that did not create a vacancy.

As shown in Fig. 9, performance fell significantly when the robot was removed, but remained significantly higher than the performance of five random controllers. There was no significant difference in the performance during the stable period of this experiment and the stable period of the second experiment where a vacancy was created. Also, there was no significant difference in performance between the convergence and stable periods of this, third, experiment. This consistency in the performance reflects the fact that the group structure remained unchanged.

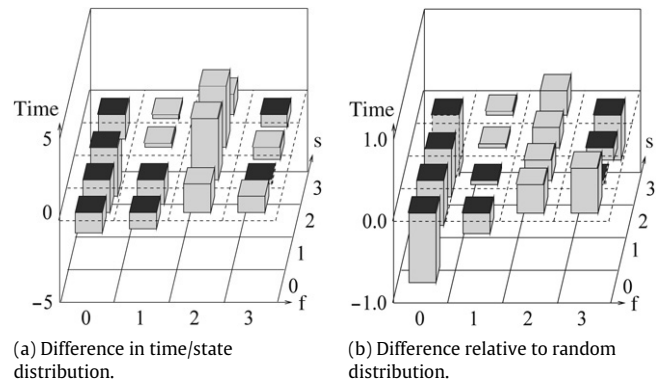
The result demonstrates that the VCS algorithm produces the allocations defined as optimal by the VCS model, independently of which robot it is that breaks down.

9.4. Heterogeneous robots

For this experiment we defined a convergence period of 15 h based on the stability of the system performance. Allocations were identified by looking at which of the robots visited the high-value circuit last. We used three fast and three slow robots, giving fifteen possible system states. We refer to each state using the notation  $f:s$ , where  $f$  is the number of fast robots whose last target was on the high-value circuit. Correspondingly,  $s$  is the number of slow robots whose last target was on the high-value circuit. The columns labeled  $\hat{\mu}_a$  and  $s_a$  in Table 4 show the mean and standard deviation of the time the system spent in each of the states while running the VCS algorithm with softmax action selection. The values are percentages of the total stable period. The columns labeled  $\hat{\mu}_r$  and

**Table 4**  
State-time dist. for heterogeneous robots.

$f:s$	$\hat{\mu}_a$	$s_a$	$\hat{\mu}_r$	$s_r$	$T$	$\hat{\mu}_a - \hat{\mu}_r$	$\frac{\hat{\mu}_a - \hat{\mu}_r}{\hat{\mu}_r}$
0:0	0.2	0.2	1.4	0.5	1.5	-1.3	-0.88
0:1	1.5	1.2	3.4	1.7	4.7	-3.2	-0.67
0:2	2.0	2.6	5.0	1.4	4.7	-2.7	-0.58
0:3	0.5	0.7	2.0	1.1	1.5	-1.1	-0.70
1:0	2.9	1.6	3.6	1.4	4.7	-1.8	-0.38
1:1	12.1	4.9	13.5	3.3	14.1	-2.0	-0.14
1:2	14.4	6.1	14.4	2.4	14.1	0.3	0.02
1:3	4.1	2.7	4.0	1.4	4.7	-0.6	-0.13
2:0	7.0	5.0	5.1	1.2	4.7	2.3	0.48
2:1	19.4	7.4	15.7	2.3	14.1	5.30	0.37
2:2	18.5	5.0	14.7	3.0	14.1	4.4	0.3
2:3	5.7	3.6	4.4	2.4	4.7	1.0	0.20
3:0	2.7	4.1	1.7	0.6	1.5	1.2	0.78
3:1	5.2	4.6	5.2	1.6	4.7	0.5	0.10
3:2	3.3	2.6	4.2	1.2	4.7	-1.4	-0.29
3:3	0.7	0.7	1.5	0.8	1.5	-0.9	-0.56



**Fig. 10.** Difference in distributions.

$s_r$  give the mean and standard deviation of the time the system spent in each of the states for a set of 15 trials using a group of robots that randomly chose between tasks.

The column labeled  $T$  lists the combinatorial probability of choosing a sample of size  $f$  from a population of  $g = 3$  fast robots as well as choosing a sample of size  $s$  from a population of  $h = 3$  slow robots. This probability is given in Eq. (13). It is worth noticing that the time distribution produced by random allocation is closely aligned with the theoretical estimate, though the differences are statistically significant. It is also worth noticing that a hand-coded solution would have spent 100% of the time in state 3:0 with a standard deviation of 0.0. However, we would like to refer to our previous argument that, in the general case, the optimal solution is not necessarily obvious, and the results from the previous experiments demonstrating that a hand-coded solution can fail spectacularly in specific cases.

$$T = \frac{100g!h!}{f!(g-f)!s!(h-s)!2^g2^h} \tag{13}$$

The difference between the state-time distribution produced by the modified VCS algorithm and the distribution produced by random allocation is presented in the column labeled  $\hat{\mu}_a - \hat{\mu}_r$ . In the last column labeled  $\frac{\hat{\mu}_a - \hat{\mu}_r}{\hat{\mu}_r}$ , this difference is presented as a percentage of the mean times from the random distribution,  $\hat{\mu}_r$ . The difference between the distributions produced by the adaptive controllers and the random controllers, i.e., the last two columns, are also presented graphically by the bar charts in Fig. 10.

Over these 15 experiments, the increase in time spent in state 0:3 is statistically significant. In Fig. 10 (b) the optimal state, 0:3, stands out as the state with the highest relative increase in time.

This confirms that the group's set of Q-tables have converged to promote the state defined as optimal according to the VCS model. The performance data also show that the performance of a group of robots controlled by the VCS algorithm, 0.081 of target value per 10 s, is significantly higher than the performance of random allocation, 0.078 of target value per 10 s. Together, the state-time distribution data and the performance data show that the modified VCS algorithm improve the group's performance by adopting a dedicated service structure that conforms to the predictions of the VCS model. The learned Q-tables confirm that, as predicted, the fast robots have higher estimated utilities for high-value tasks and higher average differences between the estimated utilities of high- and low-value tasks, than the slow robots.

## 10. Conclusions

From the experimental evidence we conclude that the VCS model is a valid model of MRTA in spatially classifiable domains. We also conclude that the VCS algorithm is an efficient and robust task allocation algorithm sensitive to differences in individual robot performance. The fact that the VCS algorithm needs only a minimal amount of information about any particular problem domain makes it applicable to a large class of problems. The evidence demonstrates that, in spite of the difficulties related to modeling group dynamics and the general complexity of scheduling, it is possible, under certain restrictions, to use the interference sensitive VCS algorithm to improve on random allocation and on hand-coded solutions.

It would be informative to compare the performance of the VCS algorithm to other known machine learning-based algorithms such as the progress estimation-based algorithm used by the L-ALLIANCE system [9] or a more general algorithm based on, e.g., Markov games [22]. However, one of the important differences between such machine learning solution and the VCS algorithm is that the VCS algorithm works without any explicit communication, while other algorithms rely on a significant communication overheads. For example, in L-ALLIANCE, each robot explicitly estimates the progress of other robots, and for Markov games, the global state has to be communicated in order to satisfy the Markov property. From the review of related work we conclude that the VCS algorithm is currently the best alternative in domains where communication is undesirable, e.g., for stealthy robotics or AUVs. We discussed such domains in detail in Section 2.

Our work on the VCS algorithm brings up several issues that do not relate directly to the experimental results. Below we discuss the most important of those in some detail.

### 10.1. Applicability

The inspiration for the VCS algorithm and the experimental work presented here are all taken from the multi-robot domain. As a general scheduling algorithm it is interesting to evaluate the applicability of the VCS algorithm outside that domain. All domains where cooperation and/or competition results in significant performance effects are possible candidate problems for the VCS algorithm. One example is network load minimization for distributed cooperating processes. In this domain different allocations imply different communication loads and a group of migrating processes using the VCS algorithm could converge on distributions that would minimize the network load. In this domain, a distributed, communication-free algorithm such as the VCS algorithm would have an advantage over communication-based task-allocation algorithms in that it would not add traffic to an already congested network. In general, in order to optimize an allocation problem using the VCS algorithm, the jobs must be spatially classifiable. Also, the machines must be able to choose between the jobs, and the value of each job must be available to provide feedback for the local utility estimates.

### 10.2. Learning efficiency

Each transportation task was completed in roughly one minute. All the experiments in homogeneous robots had convergence times of 2.5 h. This indicates that roughly 150 trials per robot is needed to adapt to a new situation. This number of trials is not surprising in a reinforcement learning context where hundreds of trials are commonly required [34]. Suitable application domains, however, would have to accept this level of sub-optimal performance. As we have argued above, in certain domains this is the best performance achievable. For the heterogeneous robot experiment, the convergence time was 15 h. This likely reflects the more complex solution sought, where particular robots have to be allocated to specific tasks. This time reflects roughly 900 trials and is again, in a reinforcement learning context, not out of the ordinary, but clearly more difficult to accommodate in real application domains.

### 10.3. Stability

A group of non-communicating, adaptive, greedy robots can be seen as participants in a multi-player game and as such, the group behavior can be analyzed using game theory [21]. To converge to a stable solution, the VCS algorithm depends on finding allocations where no single robot can benefit from unilaterally changing tasks. In game theory, such allocations are called *Nash equilibria*. It is possible, with heterogeneous robots, to construct processing time matrices that imply an ever-changing allocation. One such job processing time matrix is given in Table 5.

A game that produces involves unstable strategies are, in game theory, said to lack Nash equilibria. With the job processing time matrix presented in Table 5, there will always be an incentive for  $m_1$  to change in order to be working alone on a task. For  $m_2$ , on the other hand, it will always be advantageous to change in order to work on the same task as  $m_1$ . This will leave the system in a constant state of change, not settling on any of the optimal allocations. On problems without Nash equilibria, the VCS algorithm may not converge.

### 10.4. Optimality

Another important result from game theory is that some problems have sub-optimal Nash equilibria, and some problems do not have Nash equilibria at all, unless we consider mixed strategies. Mixed strategies are not practical as we have placed a limiting factor on switching between classes of tasks in our definition of spatially classifiable tasks in Section 5.1. For problems with sub-optimal Nash equilibria the VCS algorithm may produce sub-optimal allocations. Consider two robots,  $m_1$  and  $m_2$ , and two task classes,  $k_1$  and  $k_2$ , with identical values. Assuming identical traversal times and identical start and finishing times, the interaction function defined in Eq. (14) will lead to convergence on a suboptimal allocation. This interaction function assumes that some allocations produce synergistic effects. One example of a synergistic effect is when the robots need to continuously repair a degrading path, e.g., by removing obstacles that appear on a regular basis. In this case, the average traversal time for two robots can be higher than for one robot.

$$g_i(\mathcal{A}) = \begin{cases} 3 & \text{if } \mathcal{A} = [k_1, k_1] \\ 2 & \text{if only } m_i \text{ serves } k_1 \\ 5 & \text{if only } m_i \text{ serves } k_2 \\ 4 & \text{if } \mathcal{A} = [k_2, k_2]. \end{cases} \quad (14)$$

The average processing time matrix corresponding to the interaction function presented in Eq. (14) is presented in Table 6. This matrix shows the expected average processing time,  $g_1(\mathcal{A})/g_2(\mathcal{A})$  for  $m_1$  and  $m_2$  servicing task classes  $k_1$  and  $k_2$ .

The optimal allocation, with the minimal processing times for the interaction functions given above, is  $[k_1, k_1]$ , with the minimal

**Table 5**

An unstable processing time matrix.

$m_1$	$m_2$	
	$k_1$	$k_2$
$k_1$	2/2	4/1
$k_2$	4/1	2/2

**Table 6**

A processing time matrix leading to sub-optimal allocation.

$m_1$	$m_2$	
	$k_1$	$k_2$
$k_1$	3/3	5/2
$k_2$	2/5	4/4

total average processing time 4. The individual reward, due to the lower processing time, for a robot when unilaterally changing to task  $k_2$  or *defecting* will likely be reflected in both the robots' task utility estimates. With an exploration rate,  $\epsilon$ , below 0.5, the majority of the exploratory changes will, on average, be unilateral. This is likely to lead to one of the robots eventually changing tasks. In game theory, a problem with this structure is called a *prisoner's dilemma*, and a unilateral change away from an optimal solution for individual gain is called a *defection*. The defection of one robot will increase the remaining robot's processing time on  $k_1$  task and eventually force it to defect as well. Intuitively,  $k_2$  tasks will look more attractive due to the lower repair load, but turns out to be worse due to high levels of interference, e.g., due to less available space. An exploration rate of less than 0.5 will inhibit the robots from returning to the  $[k_1, k_1]$  allocation, as most exploratory changes are unilaterally and will yield a reward lower than the one for the  $[k_2, k_2]$  allocation.

### 10.5. Future work

The described prioritized transportation problem has a very restricted interaction function which reduces the scheduling complexity. We hope to explore problems with more complex interaction functions in order to assess the general applicability of the VCS algorithm. Vacancy chain distribution only partially describes the distribution mechanisms used in nature, such as for example the distribution of shells among hermit crabs. It is common for crabs to fight over shells [4]. Inspired by such negotiated resource exchanges, it might be possible to produce algorithms that go further in allocating high resources to high quality consumers. In the future we aim to develop distributed learning algorithms further in order exploit, on a practical level, their potential as performance-sensitive task allocation mechanisms. Demonstrating the specializing ability of distributed learning algorithms as well as their potential for performance-based task-allocation are the main contributions of the work presented in this article.

### Acknowledgments

This work was supported in part by the Department of Energy (DOE) Robotics and Intelligent Machines (RIM) Grant DE-FG03-01ER45905 and in part by the Office of Naval Research (ONR) Defense University Research Instrumentation Program (DURIP) Grant 00014-00-1-0638.

### References

- [1] P. Brucker, *Scheduling Algorithms*, 2nd ed., Springer Verlag, Berlin, Germany, 1998.
- [2] R.C. Arkin, *Behaviors-based Robotics*, MIT Press, Cambridge, MA, 1998.
- [3] M.J. Matarić, Reinforcement learning in the multi-robot domain, *Autonomous Robots* 4 (1) (1997) 73–83.
- [4] I.D. Chase, M. Weissburg, T.H. Dewitt, The vacancy chain process: A new mechanism of resource distribution in animals with application to hermit crabs, *Animal Behavior* 36 (1988) 1265–1274.
- [5] B.P. Gerkey, M.J. Matarić, A formal analysis and taxonomy of task allocation in multi-robot systems, *International Journal of Robotics Research* 23 (9) (2004) 939–954.
- [6] S. Botelho, R. Alami, M+: A scheme for multi-robot cooperation through negotiated task allocation and achievement, in: *Proceedings of the 1999 IEEE International Conference on Robotics and Automation, ICRA'99*, IEEE Press, Detroit, MI, 1999, pp. 1234–1239.
- [7] B.P. Gerkey, M.J. Matarić, Sold!: Auction methods for multi-robot coordination, *IEEE Transactions on Robotics and Automation* 18 (5) (2002) 758–768.
- [8] B.B. Werger, M.J. Matarić, Broadcast of local eligibility for multi-target observation, in: L.E. Parker, G. Bekey, J. Barhen (Eds.), *Proceedings of the 5th International Symposium on Distributed, Autonomous Robotic Systems, DARS'00*, in: *Distributed Autonomous Robotic Systems*, vol. 4, Springer, Knoxville, TN, 2000, pp. 347–356.
- [9] L.E. Parker, L-ALLIANCE: Task-oriented multi-robot learning in behavior-based systems, *Advanced Robotics*, in: *Selected Papers from IROS'96* 11 (4), 1997, pp. 305–322 (special issue).
- [10] N. Kaldra, D. Ferguson, A. Stentz, Hoplitest: A market-based framework for planned tight coordination in multirobot teams, in: *Proceedings of the 2005 IEEE International Conference on Robotics and Automation, Barcelona, Spain, 2005*, pp. 1170–1177.
- [11] T. Lemaire, R. Alami, S. Lacroix, A distributed task allocation scheme in multi-uav context, in: *Proceedings of the 2004 IEEE International Conference on Robotics and Automation, New Orleans, LA, 2004*, pp. 3622–3627.
- [12] K. Decker, TAEMS: A framework for environment centered analysis and design of coordination mechanisms, in: G. O'Hare, N. Jennings (Eds.), *Foundations of Distributed Artificial Intelligence*, Wiley Inter-Science, 1996, pp. 429–448. (Chapter 16).
- [13] W. Zhang, T.G. Dietterich, A reinforcement learning approach to job-shop scheduling, in: *Proceedings of the 14th International Joint Conference on Artificial Intelligence, IJCAI'95*, Morgan Kaufmann, Montréal, Canada, 1995, pp. 1114–1120.
- [14] A.Y. Zomaya, M. Clements, S. Olariu, A framework for reinforcement-based scheduling in parallel processor systems, *IEEE Transactions On Parallel and Distributed Systems* 9 (3) (1998) 249–259.
- [15] W. Brauer, G. Weiß, Multi-machine scheduling— A multi-agent learning approach, in: *Proceedings of the 3rd International Conference on Multi-Agent Systems, ICMAS'98*, IEEE Press, Paris, France, 1998, pp. 42–48.
- [16] C. Blum, M. Sampels, Ant colony optimization for fop shop scheduling: A case study on different pheromone representations, in: *Proceedings of the 2002 Congress on Evolutionary Computing, CEC'02*, IEEE Press, Honolulu, Hawaii, 2002, pp. 1558–1563.
- [17] P. Tangamchit, J.M. Dolan, P.K. Khosla, Learning-based task allocation in decentralized multirobot systems, in: L.E. Parker, G. Bekey, J. Barhen (Eds.), *Proceedings of the 5th International Symposium in Distributed Autonomous Robotic Systems, DARS'00*, in: *Distributed Autonomous Robotic Systems*, vol. 4, Springer, Knoxville, TN, 2000, pp. 381–390.
- [18] T.S. Dahl, M.J. Matarić, G.S. Sukhatme, Multi-robot task-allocation through vacancy chains, in: *Proceedings of the 2003 IEEE International Conference on Robotics and Automation, ICRA'03*, IEEE Press, Taipei, Taiwan, 2003, pp. 2293–2298.
- [19] T.H. Labella, M. Dorigo, J.-L. Deneubourg, Self-organized task allocation in a group of robots, in: *Proceedings of the 7th International Symposium on Distributed Autonomous Robotic Systems, DARS'04*, Toulouse, France, 2004, pp. 371–380.
- [20] K. Lerman, A. Galstyan, A. Martinoli, A.J. Ijspeert, A macroscopic analytical model of collaboration in distributed robotic systems, *Artificial Life* 7 (4) (2001) 375–393.
- [21] R.B. Myerson, *Game Theory: Analysis of Conflict*, Harvard University Press, Cambridge, MA, 1991.
- [22] M.L. Littman, Markov games as a framework for multi-agent reinforcement learning, in: W.W. Cohen, H. Hirsh (Eds.), *Proceedings of the 11th International Conference on Machine Learning, ICML'94*, Morgan Kaufmann, New Brunswick, NJ, 1994, pp. 157–193.
- [23] M. Bowling, Convergence and no-regret in multiagent learning, in: L.K. Saul, Y. Weiss, L. Bottou (Eds.), *Proceedings of the Annual Conference on Neural Information Processing Systems, NIPS'04*, December 13–18, 2004, Vancouver, Canada, in: *Advances in Neural Information Processing Systems*, vol. 17, MIT Press, Cambridge, MA, 2005, pp. 209–216.
- [24] D. Goldberg, M.J. Matarić, Learning multiple models for reward maximization, in: P. Langley (Ed.), *Proceedings of the 17th International Conference on Machine Learning, ICML'00*, Morgan Kaufmann, Stanford, CA, 2000, pp. 319–326.
- [25] T.R. Balch, Z. Khan, M.M. Veloso, Automatically tracking and analyzing the behavior of live insect colonies, in: J.P. Müller, E. Andre, S. Sen, C. Frasson (Eds.), *Proceedings of the Fifth International Conference on Autonomous Agents, Agents'01*, ACM Press, Montreal, Canada, 2001, pp. 521–528.
- [26] K. Han, M.M. Veloso, Automated robot behavior recognition applied to robotic soccer, in: *Proceedings of the 9th International Symposium on Robotics Research, ISRR'99*, Snowbird, Utah, 1999, pp. 199–204.
- [27] H. Yan, M.J. Matarić, General spatial features for analysis of multi-robot and human activities from raw position data, in: *Proceedings of the 2002 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS'02*, IEEE Press, Lausanne, Switzerland, 2002, pp. 2770–2775.
- [28] A.K. Seth, Modeling group foraging: Individual suboptimality, interference, and a kind of matching, *Adaptive Behavior* 9 (2) (2001) 67–91.
- [29] M.J. Matarić, Interaction and intelligent behavior, Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, 1994.
- [30] B. Jung, G.S. Sukhatme, Tracking targets using multiple robots: The effect of environment occlusion, *Autonomous Robots* 13 (3) (2002) 191–205.

- [31] O. Holland, C. Melhuish, Stigmergy, self-organization and sorting in collective robotics, *Artificial Life* 5 (2) (1999) 173–202.
- [32] R.S. Sutton, D. Precup, S. Singh, Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning, *Artificial Intelligence* 112 (1999) 181–211.
- [33] T.S. Dahl, M.J. Mataric, G.S. Sukhatme, Adaptive spatio-temporal organization in groups of robots, in: *Proceedings of the 2002 IEE/RSJ International Conference on Intelligent Robots and Systems, IROS'02*, IEEE Press, Lausanne, Switzerland, 2002, pp. 1044–1049.
- [34] R.S. Sutton, A.G. Barto, *Reinforcement Learning: An Introduction*, MIT Press, Cambridge, MA, 1998.
- [35] T.R. Balch, The impact of diversity on performance in multi-robot foraging, in: O. Etzioni, J.P. Müller, J.M. Bradshaw (Eds.), *Proceedings of the Third International Conference on Autonomous Agents, Agents'99*, ACM Press, Seattle, WA, 1999, pp. 92–99.
- [36] D. Goldberg, M.J. Mataric, Design and evaluation of robust behavior-based controllers for distributed multi-robot collection tasks, in: T.R. Balch, L.E. Parker (Eds.), *Robot Teams: From Diversity to Polymorphism*, A K Peters Ltd., 2001, 315–244.
- [37] E. Ostergaard, G.S. Sukhatme, M.J. Mataric, Emergent bucket brigading—A simple mechanism for improving performance in multi-robot constrained-space foraging tasks, in: *Proceedings of the 5th International Conference on Autonomous Agents, Agents'01*, ACM Press, Montreal, Canada, 2001, pp. 29–30.
- [38] B.P. Gerkey, R.T. Vaughan, K. Støy, A. Howard, G.S. Sukhatme, M.J. Mataric, Most valuable player: A robot device server for distributed control, in: *Proceedings of the 2001 IEE/RSJ International Conference on Intelligent Robots and Systems, IROS'01*, IEEE Press, Wailea, Hawaii, 2001, pp. 1226–1231.
- [39] A. Howard, Relaxation on a mesh: A formalism for generalized localization, in: *Proceedings of the 2001 IEE/RSJ International Conference on Intelligent Robots and Systems, IROS'01*, IEEE Press, Wailea, Hawaii, 2001, pp. 1055–1060.
- [40] K.H. Low, W.K. Leow, M.H. Ang, Task allocation via self-organizing swarm coalitions in distributed mobile sensor network, in: D.L. McGuinness, G. Ferguson (Eds.), *Proceedings of the 19th National Conference on Artificial Intelligence, AAAI-04*, AAAI Press/MIT Press, San Jose, CA, 2004, pp. 28–33.
- [41] T.H. Cormen, C.E. Leiserson, R.L. Rivest, C. Stein, *Introduction to Algorithms*, 2nd ed., MIT Press, Cambridge, Massachusetts, 2001.



**Torbjørn S. Dahl** is a Senior Lecturer in Computing at the University of Wales, Newport. He has a four year M.Eng. in Artificial Intelligence and Knowledge Engineering from Imperial College, London and a Ph.D. from the University of Bristol. Dr. Dahl also has work experience from Hewlett-Packard Labs, Bristol, the University of Southern California and the Norwegian Defence Research Establishment. Dr. Dahl's main research interest is in the area of Cognitive Robotics and he is the founder and coordinator of the Robotic Intelligence Lab. Dr. Dahl is the principle investigator on the research project 'Defying the rules:

How self-organising systems work' funded by the UK Engineering and Physical Sciences Research Council (EPSRC). Dr. Dahl is also an investigator on the EU

FP7 funded project 'ROBOSKIN: Skin-Based Technologies and Capabilities for Safe, Autonomous and Interactive Robots'. Dr. Dahl is a member of IEEE and ACM and has served on a number of program committees, including the IEEE 2009 International Conference on Robotics and Automation (ICRA'09) and the IEEE/RSJ 2007 International Conference on Intelligent Robots and Systems (IROS'07). He has also contributed as a reviewer to a number of scientific journals including the IEEE Transactions on Robotics and the Adaptive Behavior Journal.



**Maja Mataric** is Professor of Computer Science and Neuroscience at the University of Southern California (USC), Founding Director of the USC Center for Robotics and Embedded Systems, Co-Director of the USC Robotics Research Lab, and Senior Associate Dean for Research in the USC Viterbi School of Engineering. She received her Ph.D. in Computer Science and Artificial Intelligence from MIT in 1994, M.S. in Computer Science from MIT in 1990, and B.S. in Computer Science from the University of Kansas in 1987. She is a recipient of the Okawa Foundation Award, NSF Career Award, the IEEE Robotics and Automation Society Early Career Award, the MIT TR100 Innovation Award, and the USC School of Engineering Junior Research Award and Service Award. Her research is aimed at endowing robots with the ability to help people through individual assistive human–robot interaction (in convalescence, rehabilitation, training, and education) and human–robot team cooperation (in habitat monitoring and emergency response). The research addresses problems of intelligent control and learning in complex, high dimensional/high degree of freedom systems that integrate perception, representation, and interaction with people. Research details are found at <http://robotics.usc.edu/interaction/>.



**Gaurav S. Sukhatme** is an Associate Professor of Computer Science (joint appointment in Electrical Engineering) at the University of Southern California (USC). He received his undergraduate education at IIT Bombay in Computer Science and Engineering, and M.S. and Ph.D. degrees in Computer Science from USC. He is the co-director of the USC Robotics Research Laboratory and the director of the USC Robotic Embedded Systems Laboratory which he founded in 2000. His research interests are in multi-robot systems and sensor/actuator networks. He has published extensively in these and related areas. Sukhatme has served as PI on numerous NSF, DARPA and NASA grants. He is a Co-PI on the Center for Embedded Networked Sensing (CENS), an NSF Science and Technology Center. He is a senior member of IEEE, and a member of AAAI and the ACM. He is a recipient of the NSF CAREER award and the Okawa foundation research award. He has served on many conference program committees, and is one of the founders of the Robotics: Science and Systems conference. He is one of the program chairs of the 2008 IEEE International Conference on Robotics and Automation. He is the Editor-in-Chief of *Autonomous Robots*. He has served as Associate Editor of the IEEE Transactions on Robotics and Automation, the IEEE Transactions on Mobile Computing, and on the editorial board of IEEE Pervasive Computing.