

Investigating the Adaptiveness of Communication in Multi-Agent
Behavior Coordination
(The Adaptiveness of Communication for Coordination)

Paul Schermerhorn and Matthias Scheutz
Artificial Intelligence and Robotics Laboratory
Department of Computer Science and Engineering
University of Notre Dame
Notre Dame, IN 46556, USA
Tel. +1 (574) 631-8380
{pscherm1,mscheutz}@cse.nd.edu

July 7, 2007

The Adaptiveness of Communication for Coordination

Abstract

Some previous studies of the adaptiveness of communication for coordination have found communication beneficial, others have not. We claim that this results from the lack of a systematic examination of important variables such as communication range, sensory range, and environmental conditions. We present an extensive series of simulations exploring how these parameters effect the utility of communication for coordination in the *multi-agent territory exploration* (MATE(n)) task. MATE(n) requires agents to visit all checkpoints in the environment in as little time as possible; n agents must be at a checkpoint simultaneously for it to be counted “visited.” A comparison of the *absolute performance* of communicating and non-communicating agents on MATE(n) (i.e., performance without regard to cost) finds that communication can be beneficial. A subsequent analysis of the results establishes constraints on the cost of communication for it to provide *relative performance* benefit (i.e., absolute performance scaled by cost).

Keywords: Communicating agents, behavior coordination, cost tradeoffs

1 Introduction

Many studies have examined the utility of communication for the purpose of behavior coordination. Some find that communication is beneficial (MacLennan, 1991; Levin, 1995; Ackley and Littman, 1991; Marocco et al., 2003; Cangelosi et al., 2004), others conclude that communication does not aid coordination (Noble, 1999; Grim et al., 2002). One possible explanation for these inconclusive findings is that there is not a simple “yes or no” answer to whether communication is beneficial; the answer may be “yes” for certain agent configurations and “no” for others. Yet, interactions between sensory range and communication range each can substantially impact the benefit of communication, as can the structure of the environment (e.g., how objects of potential interest are distributed, randomly or otherwise), but their effects are not examined systematically in previous studies. Here, “communication range” refers to the maximum distance across which two agents can communicate, and “sensory range” refers to the maximum distance objects of potential interest can be from the agent and still be detected. Moreover, the “cost of communication”—the expense of having and using additional architectural, sensory and effector mechanisms for communication—is often ignored in previous studies, while those studies that do include cost fail to examine how it is affected by sensory and communication ranges.

We revisit the question of the utility of communication for behavior coordination with a systematic study that varies sensory range, communication range, and environmental conditions, comparing the performance of communicating and non-communicating agents that are identical except with regard to their ability to share information. Agents’ performance is evaluated in the *multi-agent territory exploration* task (MATE(n), for short). In a MATE(n) task, a group of agents A of kind K is required to visit a certain number of locations (“checkpoints”) in an environment in the least amount of time, with a “visit” requiring n agents to be simultaneously present at a checkpoint. This last requirement creates a substantial need for coordination between agents, which otherwise might never choose to target the same checkpoint for visitation (e.g., if the nearest checkpoint to each agent is not the same). MATE(n) is a good target task for an investigation of the utility of communication because it encapsulates essential components of a wide variety of biological (e.g., foraging) and engineering (e.g., planetary exploration) tasks.

To evaluate utility of communication, we employ *absolute* and *relative* performance measures. The *absolute performance* of a group of agents on MATE(n) is the time it takes the group to complete the task (i.e., to visit all the checkpoints). When communicating agents take less time to complete the task than non-

communicating agents, communication provides an *absolute performance advantage*. A group's *relative performance* is its time to completion scaled by agent cost (i.e., the costs of construction, maintenance, and operation of the agent). Communication provides a *relative performance advantage* when communicating agents have an absolute performance advantage that is large enough to offset the cost of communication.

The remainder of the paper is an exploration of the utility of communications for MATE(n) based on absolute and relative performance. Section 2 provides the task specification for MATE(n). Section 3 outlines the experimental approach to investigating the utility of communication, and presents the basic agent model used throughout the paper. Section 4 presents a detailed examination of agent performance without communication, where opportunities for performance improvement are identified. The basic agent model described in Section 3 is extended in Section 5 to include communication, as well as another, less expensive alternative that improves coordination without the use of communication: the *wait timer* mechanism that increases the probability that agents will be at the same checkpoint simultaneously. In addition, a minimally structured environment model is introduced that should increase the value of communicated information, and should, therefore, contribute to the value of communication.

A large-scale set of experiments exploring each of these additional dimensions is analyzed in Section 6. We find that communication does provide an absolute performance advantage for MATE(n) under certain conditions. In those cases where there is an absolute performance advantage, the question of whether there is a relative performance advantage to communication depends on the particular cost assumptions. While it is not possible to fill in the details of those assumptions (e.g., because communication can be implemented in many ways, some of which are more costly than others), we do present an analysis of the cost constraints that will need to be met in order for communication to provide a relative performance advantage, even in the most favorable case (i.e., where communication's absolute performance advantage is greatest). We find that communication costs must be very small relative to other agent costs to yield sufficiently high relative performance. Communicating artificial agents must be designed to satisfy this constraint. It is unlikely that biological agents can achieve such a ratio, so we conclude in Section 7 that communication is very unlikely to have a relative performance benefit (and hence is unlikely to evolve) for biological agents in MATE(n) tasks. Section 8 presents a brief summary of the conclusions drawn from this research.

2 The Multi-agent Territory Exploration Task

An appropriate task for the study of communication's utility will require agents to work directly together; if agents can complete the task individually without ever interacting with others, it will be unsurprising to find that communication is not useful, whereas if they require the aid of others (even if it is only to improve performance), communication may be adaptive. Ideally, the task should be parameterized to allow the degree of reliance on other agents (and hence the amount of coordination required) to be varied, making it possible to probe the limits of communication's utility (e.g., what is the lowest coordination requirement at which communication proves beneficial?). However, care must be taken to ensure that the task is not tailored too specifically to the search for communication. Here we list four *potential pitfalls* for investigations like ours that use agent-based models, or "synthetic ethology" (MacLennan, 1991):

1. *Communication is a part of the task requirement.* Tasks must not intrinsically require communication for successful task completion. For otherwise the "deck is stacked" in favor of communication and it will be neither surprising nor very telling if communication can subsequently be shown to be of benefit.
2. *The fitness function rewards agents for communication.* Even if communication is not necessary for task completion, explicitly rewarding agents for communication (e.g., as part of the fitness function in a genetic algorithm) will still bias an investigation in favor of communication. Rather than rewarding agents separately for communicating, the reward should come directly from the task performance—if communication helps the agents achieve their goals, then it will lead to higher task performance and does not have to be rewarded explicitly. Especially in biologically plausible settings, communication must provide a benefit for the accomplishment of some task that is germane to the survival and prosperity of the agent (e.g., foraging). Agreeing on the "meaning" of a signal should not benefit the agent in and of itself; rather that agreement must facilitate performance of some essential survival task.
3. *The cost of communication is ignored.* It is often important to determine whether communication can lead to better absolute performance in a task (e.g., Ackley and Littman (1991); Marocco et al. (2003); Cangelosi et al. (2004)). However, for an argument about the relative performance of communication, or the evolvability of communication, one must also consider the associated costs of construction, maintenance, and operation of the communication mechanisms; communication will yield a relative

performance advantage only if its benefits are greater than its costs.

4. *High-level coordination is required for success.* Especially for investigations of the utility of communication, it is important to avoid tasks requiring coordination that can *only be reliably achieved* via communication (it would be unsurprising to find that communication is beneficial, even essential, to the task of landing a rover on Mars). A class of tasks with variable complexity, as described above, enables the experimenter to select a degree of complexity appropriate to the problem under consideration.

Foraging is a straightforward example of a task that avoids these pitfalls and is well-represented in biological agents (e.g., among social insects; see Capaldi and Dyer (1999); Menzel et al. (1998); Nicholson et al. (1999); Shen et al. (1998)). It can be accomplished without communication, whatever rewards communication may provide are the result of improved task performance, the cost is built in as an exchange of energy (for communication acts) for energy (in the form of food), and high-level coordination is not necessary for successful foraging. Foraging can also serve as a model for important artificial tasks, such as planetary exploration by rovers (Rabideau et al., 1999; Estlin et al., 2000; Clement and Barrett, 2003), making results from foraging tasks applicable also to those domains. Performance in foraging tasks could be defined in terms of the efficiency with which agents locate food. Hence, we can define a *time-bounded multi-agent territory exploration task*:

Definition 1: A MATE(t) task $T(C, A, R_S, D)$ requires a group of identical agents A , each with sensory range R_S , to visit all members of set of checkpoints C , where checkpoints are placed in a 2D environment according to a probability distribution D , and the agents are allotted a period of time t within which to visit as many checkpoints as possible.

If checkpoints are defined such that they can be visited only once (such that checkpoints can be taken to represent food items), MATE(t) provides an accurate evaluation of foraging efficiency for such environments. MATE(t) is easily extended to a form that requires a greater degree of coordination by requiring n agents present at a checkpoint in order for it to be “visited” (i.e., a single agent cannot gather the food source alone, it takes n agents to do so), yielding the more general MATE(n, t) task.

MATE(n, t) can be viewed as a time slice of the actual (ongoing) biological task in which the agents are foraging for food. As such, MATE(n, t) can serve as a measure of *foraging efficiency* for those biological

agents. Given that foraging efficiency is critical to the survival and propagation of biological agents, visitation rate (the average number of checkpoints visited by the group in the given timespan) is predictive of success in biological populations (Scheutz and Schermerhorn, 2005). Hence, $MATE(n, t)$ is well-suited for exploring questions regarding the utility of communication for biological agents. However, it can be very difficult to determine appropriate values for the parameters of $MATE(n, t)$ (i.e., the number of checkpoints in the environment and the time allowed); if there are too many checkpoints, the task will be too easy and agents will be able to locate checkpoints even with poor foraging skills. Conversely, if there are too few checkpoints, very good foragers may visit them early on, allowing poorer foragers to “catch up” in the remaining time. In either case, there will be insufficient variability in performance to distinguish between agent kinds. Of course, the duration t interacts with these tradeoffs as well, contributing to the difficulty of finding appropriate parameters. Moreover, parameters that are appropriate for groups of one size may prove useless for another group size (e.g., placing checkpoints with a sufficiently low initial density to detect differences between agent kinds in groups of size ten will make the task virtually impossible for groups of size two, leading to the lack of performance variation mentioned above); while it may be possible to hit upon useful parameters for multiple group sizes, doing so will make comparisons across group sizes difficult.

What is needed is a task that is not so sensitive to its initial parameters, to allow us to compare a wide variety of agent kinds, group sizes, and environments under the same task conditions. Fortunately, the task described below (i.e., $MATE(n)$) is not subject to these difficulties. The formal definition of $MATE(n)$ is presented next, followed by an argument demonstrating that results from the (less biologically plausible) $MATE(n)$ task are just as relevant to the utility (and evolvability) of communication as the (more biologically plausible) $MATE(n, t)$ task.

2.1 $MATE(n)$ Tasks

We define $MATE(n)$ as follows:

Definition 2: A $MATE(n)$ task $T(C, A, R_S, D)$ requires a group of identical agents A , each with sensory range R_S , to visit all members of set of checkpoints C , where each checkpoint requires $n \geq 2$ agents present to remove its mark and agents and checkpoints are placed in a 2D environment according to a probability distribution D .

In a $MATE(n)$ task, there must be n agents present at a checkpoint *at the same time* in order for that checkpoint's mark to be removed. If fewer than n agents are located at the checkpoint, it is not considered visited and the mark remains in the environment. The stipulation that $n \geq 2$ ensures that coordination (but not necessarily communication) is beneficial for each individual subtask (i.e., bringing n agents together at a checkpoint), rather than only for the overall task (e.g., coordinating which agent will visit which checkpoint). Removing the time constraint t makes it much easier to determine an appropriate number of checkpoints $|C|$ to detect performance differences in a wide range of physical, architectural, and environmental configurations.

The main difference between the $MATE(n)$ and $MATE(n, t)$ tasks is that in the former agents have to find all checkpoints in the environment, whereas in the latter agents have to visit as many checkpoints as they can within the allotted time t . Note, however, that these tasks are intrinsically related as they both measure *foraging efficiency*, i.e., the efficiency with which agents find and visit checkpoints. Suppose an agent kind K_1 has a higher foraging efficiency than another agent kind K_2 ($FE_1 > FE_2$), where the FEs are given in terms of “average number of visited items per time unit.” Then we would expect K_1 , on average, to perform better than K_2 in both $MATE(n)$ and $MATE(n, t)$ tasks. Moreover, we would not expect any agent kind K_1 to ever perform better on average than some other agent kind K_2 in $MATE(n)$, but worse than K_2 in $MATE(n, t)$, and vice versa, given that better foragers will find checkpoints more quickly, regardless of whether their performance is measured in terms of the time it takes to find all checkpoints or the number of checkpoints they find within a given time period. However, it is possible that performance differences between two agent kinds in $MATE(n)$ are *not preserved* in $MATE(n, t)$ (and vice versa), e.g., because the allotted time in $MATE(n, t)$ is so long that even bad foragers with poor performance can visit all checkpoints in C within time t , so that the difference in foraging efficiency does not come to bear. Therefore, we cannot conclude necessarily that an agent kind that performs better than another in $MATE(n)$ will also perform *better* in $MATE(n, t)$, only that it will not perform *worse*. Given that $MATE(n)$ and $MATE(n, t)$ are mutually predictive in the sense that an agent kind K_1 that performs better than another kind K_2 in one of the tasks performs at least as well as K_2 in the other, we use $MATE(n)$ tasks (instead of $MATE(n, t)$ tasks) for the investigations in this paper, due to the difficulties described above in determining appropriate parameters for $MATE(n, t)$ experiments.

2.2 Solving MATE(n)

There are two general classes of solutions to MATE(n) tasks: *offline* and *online*. If the locations of all agents and all checkpoints are known in advance, an offline solution can be computed via an exhaustive search of the solution space (i.e., which agents target which checkpoints and when). This solution is, however, exponential in the number of agents and checkpoints, and is, therefore computationally intractable. Moreover, optimal solutions require full knowledge of the environment, which may not be practical. Practically speaking, then, online approximations are the only viable solutions to MATE(n) tasks. While MATE(n) is similar to the Traveling Salesman Problem (TSP, Cormen et al. (1990)) and TSP variants such as the the Multi-Traveling Salesman Problem (MTSP) and the delivery-scheduling problem, it differs critically in that agents do not begin at a common position and are not required to return to their initial positions. Hence, the online approximations (Golden et al. (1980); Johnson and McGeoch (1997) for the TSP, Gavish and Srikanth (1986); Bugera (2004); Wang and Xue (2002); Applegate et al. (2002); Montemanni et al. (2003); Thompson and Psaraftis (1993) for the MTSP) are not appropriate for MATE(n).

Performance on MATE(n) can be defined in many ways. One goal could be to minimize the total distance traveled by all agents A . Or, one could balance the load among all agents by minimizing the difference in length between all agents' paths. For the experiments described below, we define the performance measure to be the time required to visit all checkpoints (sometimes referred to the "makespan"). Thus, agent groups that visit all checkpoints in the least amount of time have the best absolute performance.

MATE(n) is a task that requires multiple agents to work together in order to achieve good performance, but it avoids all four pitfalls (connected to tasks for the study of the utility of communication) described above. The criteria for task completion are not dependent on communication or explicit coordination (pitfall 1); although communication and explicit coordination may be helpful, it is possible to complete a MATE(n) task without using either. There is no explicit reward for communication (pitfall 2). The measure of performance is an objective count of how long it takes each group to complete the task; communication may help, or it may not. Moreover, while the comparisons presented below are of *absolute performance*, or performance without regard to cost, a subsequent analysis of the results in terms of *relative performance*, or scaled performance where benefits are offset by costs addresses pitfall 3. The task complexity and thus the degree of coordination required for efficient task completion is variable (pitfall 4), so it should be possible to determine the smallest n for which communication pays off in MATE(n) (if such an n exists).

The degree of coordination among agents required to do well in $MATE(n)$ increases as n increases because more agents must find a way to be at the same place at the same time, something that is increasingly unlikely to happen by chance as n rises. While it is possible to achieve this goal without the aid of communication, we conjecture that sharing information should significantly increase the likelihood of agents visiting checkpoints together. Specifically, it should be beneficial for an agent in a $MATE(n)$ task to communicate the location of a checkpoint to another agent and to attract that agent to the checkpoint, as two agents are required to visit it. Hence, we expect communication to be useful for $MATE(n)$ because communication can serve to coordinate the behavior of two or more agents to the extent that it brings them together in the vicinity of a checkpoint. The experiments below use $MATE(2)$ as the target task; as mentioned above, we are interested in exploring the limits on the degree of complexity required for communication to be beneficial for $MATE(n)$. $MATE(2)$ has the lowest degree of complexity, while still requiring agents to actively coordinate behavior in order to achieve good performance. We argue in Section 6.3 that the results of the $MATE(2)$ experiments provide an adequate basis for generalizing our conclusions to $MATE(n)$ (for $n \geq 2$).

2.3 Related Work

Many projects have examined the potential benefit or evolvability of communication. Game-theoretic studies point to the benefit of communication (Miller et al., 2002; Miller and Moser, 2003; Cooper et al., 1992), however, embedded games such as these ignore many of the practical issues confronting embodied agents. Simulation studies commonly find that communication is beneficial. However, communication is often directly rewarded by the fitness function rather than used to improve performance on an independent fitness function (MacLennan, 1991; Levin, 1995; Noble and Cliff, 1996), or it is free of cost (Ackley and Littman, 1991; Marocco et al., 2003; Cangelosi et al., 2004). When communication is rewarded directly (i.e., when agents are rewarded specifically for successful communication), there is no possibility of other, non-communicating solutions, to succeed. Similarly, cost is an important component of utility, so ignoring the cost of communication biases a study in favor of the utility of communication. Studies that include considerations of cost find that communication is too expensive to evolve for the tasks studied (Noble, 1999; Grim et al., 2002).

Communication has been used to aid coordination in robotic tasks. Parker et al. (2004) use communication for explicit coordination in a robotic sensory net deployment task. Leader robots lead columns of subordinates through an environment. When the column arrives at a deployment point, the leader orders

a subordinate to navigate into position and begin scanning the room. While this project uses communication only to transmit orders, a use of communication more relevant to MATE(n) tasks is presented in Tang and Parker (2005), where robots use communication to share perceptual information in a deployment task. In this schema-based approach, information is shared between robots via communication schemas, which can have as their inputs perceptual schemas or other communication schemas, and as their outputs perceptual schemas (potentially on another robot), motor schemas, or other communication schemas. This allows groups of robots to “share” sensors.

Others have explored means of coordination without communication. Werger (1999) found that communication is unnecessary for coordinating offensive and defensive group formation in RoboSoccer. Gervasi and Precipe (2004) also find that communication is unnecessary for robotic formation tasks. And Reynolds’ Boids project (Reynolds, 1987) was an early exploration of multi-agent flocking behavior using simple rules and no communication. Coordination on common classes of tasks (e.g., robotic formations) does not appear to require communication. Scheutz and Schermerhorn (2003) find that a simple lookahead mechanism which tries to avoid conflicts with other agents in a collection task outperforms a more complex deliberative system that includes memory and route planning, even though the cost of the lookahead mechanism is much less than the cost of deliberation.

2.4 Relative Utility

Relative performance represents an important aspect of overall agent performance. While it is common to find simulation (Ackley and Littman, 1991; Marocco et al., 2003; Cangelosi et al., 2004) and even robotic (Parker et al., 2004; Tang and Parker, 2005) studies that do not take into account the costs of the agent architectures being proposed, the cost of an adaptation can be just as important as its benefits in the real world, outside the simulation environment or laboratory. There are surely cases in which we are willing to accept high costs in order to eke out tiny performance improvements, however, there are also many cases in which efficiency is important. Extraterrestrial surface exploration robots are very expensive to build and transport to their targets, making it imperative that researchers maximize the benefits obtained within the constraints of potentially limited battery or solar power. An expensive strategy that returns results hours earlier than the alternative but leads to battery failure days earlier is probably not worth pursuing.

For biological agents, the importance of cost is at least as great. Potential adaptations that do not provide a relative performance advantage are actually maladaptive. They will, on average, make survival more dif-

ficult, thereby reducing the chances that the “adaptation” will be passed on to offspring. In this context, the only useful understanding of the term “beneficial” must imply a relative performance advantage; communication is not beneficial to the agent merely in virtue of attracting a mate more quickly if the effort exhausts the agent to such a degree that it is unable to participate in the mating ritual.

The results presented below are analyzed with respect to the absolute performance of the agents. In this way, we are able to identify the agents with the best possibility of maintaining a relative performance advantage over their opponents. Section 6.2 then examines the prospects for communication to yield a relative performance advantage in MATE(2), providing a detailed analysis of the cost constraints that must be satisfied by an implementation of the agent architecture if communication is to be beneficial with respect to relative performance.

3 Methodology

Our approach to evaluating the utility of communication for MATE(2) begins with determining what effect (if any) communication has on absolute performance (i.e., average cycles to completion). Base agent models are presented that do not use communication, and extensive simulation sets are used to determine their absolute performance. Communication is then included in the model, and the absolute performance of the new agents is evaluated in another series of simulations. The addition of communication leads to an increase in cost. To determine the relative performance difference, the absolute performance difference is compared with this additional cost. If there is no absolute performance difference, or if communicating agents have *worse* absolute performance, relative performance must be worse with communication for these agents on MATE(2). When absolute performance with communication is *better*, the relative performance difference depends on the cost of communication (i.e., the cost of building, maintaining, and using communication facilities). When the cost is very low and the absolute performance benefit of communication is very high, communication is likely to have a relative performance advantage. Conversely, when the cost is high and the benefit low, communication is unlikely to have a relative performance advantage for these tasks. When cost and benefit are of similar magnitude, a detailed cost-benefit analysis is required to determine whether communication provides a relative performance advantage.

Communication may not be the only (or the best) coordination mechanism for MATE(n). To gain a fuller understanding of the adaptiveness of communication, we also compare its effects with those of

another, lower-cost coordination mechanism: the *wait timer* (described below). For, once again, although communication may have an absolute performance advantage over the wait timer for MATE(n), the lower cost of the latter may give it a relative performance advantage, giving us a better understanding of the adaptiveness of communication for MATE(n).

The remainder of this section describes the base agent model used and the employed experimental setup.

3.1 Basic Agent Model

Sensory requirements for agents in the MATE(n) task are as follows: agents must be able to sense the presence of checkpoints within the given sensory range R_S . The sensor model is kept purposefully simple, abstracting over sensor modalities (e.g., vision vs. sonar) and assuming that agents can sense all checkpoints within sensory range in a full 360° radius. The agent's sensory input S combined with its inner states I produces behaviors according to the *agent function* F :

Agent function: An agent function F is a mapping $F : 2^S \times I \rightarrow D$, where 2^S is the power set of sensory data S given as pairs $\langle d, \alpha \rangle$ of distances $d \in [0, R_S]$ and angles $\alpha \in [0, 2\pi]$, I is a set of inner states (possibly empty), and $D \in [0, 2\pi]$ is the set of possible directions the agent will move in.

These rules constitute a model for *target-closest* agents in MATE(n) tasks:

- *Rule 1:* if no checkpoint is perceived ($S = \emptyset$), perform a random walk $RW(rwd, \beta)$ (i.e., move in the direction of the current heading θ for rwd cycles, then change heading randomly to some value in $[\theta - \beta, \theta + \beta]$)
- *Rule 2:* if some checkpoint C is within visiting distance (i.e., the agent is at the checkpoint) and there are n agents within visiting distance of C , remove the mark (if it is still there—one of the other $n - 1$ agents may remove it first)
- *Rule 3:* if some checkpoint C is within visiting distance and there are fewer than n agents within visiting distance of C , wait at C
- *Rule 4:* if some checkpoints are perceived ($S \neq \emptyset$) and no checkpoint is within visiting distance, go directly towards the closest checkpoint (the direction is given by α such that $\min_d \{ \langle d, \alpha \rangle \mid \langle d, \alpha \rangle \in S \}$)

These simple agents will move directly to the nearest checkpoint and wait there indefinitely for a total of n agents to arrive. They are called “target-closest” because they target the closest perceived checkpoint to visit and select a heading directly toward that checkpoint. They employ a random walk mechanism that changes their heading after rw_d cycles without perceiving a checkpoint. The value of rw_d was determined for each combination of “number of agents” and “sensory ranges” by systematically exploring the performance space based on a sampling of values for these two parameters and selecting those with the best performance.¹ Additionally, the environment is bounded by an 800 x 800 area (in comparison, agents have a size of 8). To prevent agents from leaving the area, they “bounce” off the edge of the environment with a small amount of random error analogous to a billiard ball bouncing off the cushion.

Physical factors (such as sensory range) and *social* factors (such as the number of agents participating in the task) should effect performance on the task regardless of the agent model. Increasing either sensory range or the number of agents participating should improve performance, with the following reservations: (R1) that it is likely that increasing sensory range of agents beyond a certain limit (related to the environment size) will not increase group performance (although it should not lower it either), and similarly (R2) that there is likely to be group size beyond which increasing the number of agents no longer improves performance.

Each experiment reported in this paper consists of 40 experimental runs using different randomly generated initial conditions in the same 800 by 800 square region of the continuous 2D world. The same set of 40 checkpoint distributions ($|C| = 10$) is used for all experiments, in which checkpoints are distributed uniformly throughout the environment at random locations. Moreover, the same 40 initial agent distributions (also randomly placed throughout the environment) is used for all experiments with identical size of A (e.g., experiments with $|A| = 2$, $|A| = 3$, etc.). This allows us to compare directly between agent types and sensory ranges. The results reported here are the average cycles to completion of each experiment set for each architectural configuration.

4 Target-closest Agents

The agent model in Section 3 represents the minimal model that can reasonably be expected to succeed in the task. However, because n agents are unlikely to arrive at a checkpoint simultaneously, agents often spend a large number of cycles “hovering” at a checkpoint. Therefore, we do not expect the model to perform well without additional coordination mechanisms, as waiting indefinitely can cause “deadlocks” in which each

agent waits for the others at separate checkpoints.

The experiments for this study examine the performance of agents in the MATE(2) task. Sensory range R_S was varied from 25 to 250 in steps of 25 and from 300 to 800 in steps of 50. The group size $|A|$ was varied from 2 to 10. This yields a total of 189 experiments, each consisting of 40 experimental runs. In each of the 40 runs, all agents share the same fixed sensory range, and the group size is also fixed throughout each experimental run.

4.1 Results

[Figure 1 about here.]

Figure 1 presents the results of these experiments for agents with sensory ranges from 25 to 800. The results are in line with the predictions: increasing sensory range improves performance, up to a point. At around 300, the improvement levels off and subsequent increases in sensory range do not affect performance substantially. Adding more agents also improves performance, producing a clear performance ordering based on group size. The incremental benefit of increasing group size appears to be diminishing by the time $|A| = 10$, indicating reduced efficiency with increased group size for this environment size and number of checkpoints.

A two-way 9x21 ANOVA with *group size* (2-10) and *sensory range* (25-800) as randomized variables and *average cycles to completion* as the dependent variable shows highly significant main effects for group size ($F(8,7371) = 1356.609$, $p < 0.001$) and range ($F(20,7371) = 107.600$, $p < 0.001$), confirming that increasing group size and sensory range improves absolute performance. There is also a highly significant interaction between group size and sensory range ($F(160,7371) = 3.706$, $p < 0.001$); this confirms that larger groups increase performance faster as sensory range increases, reaching earlier the point at which added sensory range is of no benefit. These effects are exactly what one would expect, and will continue to be seen throughout the series of experiments presented below.

4.2 Analysis

The two main effects for performance (group size and sensory range) can be clearly seen in Figure 1. The results support the prediction that physically extending sensory range will increase performance, but, in line with (R1) above, only up to a certain point. Increasing sensory range beyond about 300 affords little benefit

and will certainly increase the cost of operation for the agents, making agents with $R_S > 300$ inefficient relative to agents with $R_S \leq 300$. For $R_S \leq 200$, however, performance improves rapidly as R_S increases, making it likely that these improvements will be beneficial even with cost taken into account.

Performance differences for sensory ranges higher than 300 to 400 are due almost entirely to the increases in group size. Increasing the group size increases performance on the MATE(2) task, as predicted. The benefit of additional agents is quite good initially, with substantial drops in average cycles to completion (*acc*) as $|A|$ increases. If we examine the region where $R_S > 300$ (i.e., the area in which differences are attributable only to $|A|$ and not R_S), we see that performance improvements tend to get smaller as group size increases, lending support to (R2). In terms of efficiency, however, the ordering is altered substantially. If we compare the performance of different group sizes scaled by the number of agents in the group, the results demonstrate a downward trend in total cycles (starting with 17834.90 for $|A| = 2$, down to 8042.0 for $|A| = 10$). This indicates an increase in efficiency as $|A|$ increases, seemingly refuting (R2) for this limited number of agents, although it is still likely that (R2) will be true for larger groups.

The work of an agent can be divided into two parts: finding a checkpoint and visiting it. Once a target-closest agent detects a checkpoint, it moves there and waits for another agent. Until it finds a checkpoint, however, it is wandering randomly through the environment looking for one. Agents with low sensory range will have more difficulty detecting checkpoints; they will need to explore more of the environment on average than agents with high sensory range, which will be able to perceive most of the checkpoints in the environment regardless of their own positions. We predict that these cycles spent foraging comprise the difference between agents with low sensory range and those with high sensory range.

4.3 Discussion

Given the random placement of agents and checkpoints in the environment, it seems difficult to improve on a random walk scheme to reduce foraging cycles. We do, however, see performance increases as sensory range increases; when agents see more, they spend less time wandering around unproductively. Hence, to the extent that communication can be considered an extension of sensory range, it should have effects similar to extending sensory range. We refer to this as the “Communication Used as Sensory Prosthesis” (“CUSP”) thesis, formulated as follows:

Communication Used as Sensory Prosthesis (CUSP): Communication can serve the functional role of extending an individual agent’s sensory range and thus improve absolute performance for agents with shorter than optimal sensory ranges.

CUSP will be true in situations where two agents are within communication range of one another, but only one is within sensory range of a checkpoint. If the other agent would otherwise have moved in a different direction (i.e., instead of toward the checkpoint), then communication will serve the role of extending its sensory range.

The fact that there is no improvement in performance when $R_S > 400$ suggests that it is unlikely that communication will be useful at these ranges—even if CUSP were true, communication would not help for that reason, as increasing sensory range does not help. It is likely that better decision-making mechanisms are needed to realize performance improvements at these ranges. For lower sensory ranges where increasing sensory range does improve performance, however, the additional information provided by communication may be beneficial. We add communication to the agent model in Section 5.

5 Extended Agent Models and Environments

The results and analysis presented in the previous section for the timeout agent experiments suggest multiple possible avenues for improving agent performance, including the addition of communication. Several of these possibilities were explored, and the results are presented below. However, the number of experiments conducted and possible comparisons between them preclude the sort of detailed presentation given of target-closest agents above. Therefore, the results and analysis in Section 4 are given only as a concrete example of the detailed examination possible, and the remaining results are presented in summary form via a single large ANOVA (see Section 6). This section presents the extended agent models used in those additional experiments, as well as a modified environment model in which all variants of the agent model were tested.

Communicating Target-Closest Agents For communicating agents, the agent model is extended by adding a new source of information: other agents. Communication is accomplished via the following actions:

- *Sending*: any sensed checkpoint is communicated to any agent within communication R_C as the tuple $\langle d, \alpha \rangle$ of distances $d \in [0, R_S]$ and angles $\alpha \in [0, 2\pi]$ relative to the receiving agent's position (i.e., every agent receives the checkpoint information as if it had sensed it).
- *Receiving*: the set of communicated checkpoint locations CS (i.e., the union of all such sets received from all agents within communication range R_C) is combined with the agent's own set of sensory inputs S to form the overall set of perceptions ($S' = S \cup CS$).

At each cycle, agents communicate the locations of every checkpoint they detect to all other agents within communication range. This information is included in the receiving agent's sensory input S' and is treated as if it were directly sensed by the agent. This additional information does not effect the behavior of target-closest agents unless there is no checkpoint within sensory range. In that case, non-communicating agents will begin a random walk according to *Rule 1*. Communicating agents, on the other hand, will move directly toward the nearest *communicated* checkpoint when there is nothing directly sensed, according to *Rule 3*. Note that agent locations are not communicated. Target-closest agents do not make use of agent locations, so this limitation does not effect them.

The Wait Timer Observation of experimental runs with target-closest agents highlight the large proportion of time they spend at checkpoints waiting for other agents to arrive. The compromise we propose to reduce that waiting time is the *wait timer*: a timeout mechanism that allows an agent to leave a checkpoint after it has waited for a certain number (W) of cycles. By itself, the wait timer would not be effective, because an agent would return immediately to the checkpoint it had just left, essentially waiting indefinitely at that checkpoint. To support the wait timer, we introduce another mechanism: the *filter set*. Checkpoints included in an agent's filter set are excluded from from the agent's sensory input S , forcing the agent to ignore them as long as they remain in the filter set; a counter with an initial value of F is associated with each checkpoint in the filter set. This counter is decremented every cycle until it reaches zero, at which point the checkpoint is removed from the filter set. Given these additions, the agent model is modified as follows:

- *Rule 1*: if no checkpoint is perceived ($S = \emptyset$), perform a random walk $RW(rwd, \beta)$ (i.e., move in the direction of the current heading θ for rwd cycles, then change heading randomly to some value in $[\theta - \beta, \theta + \beta]$)

- *Rule 2*: if some checkpoint C is within visiting distance and there are n agents within visiting distance of C , remove the checkpoint's mark and reset w to 0
- *Rule 3*: if some checkpoint C is within visiting distance, there are fewer than n agents at C , and $w = W$, reset w to 0 and add C to the agent's filter set
- *Rule 4*: if some checkpoint C is within visiting distance, there are fewer than n agents at C , and $w < W$, increment w
- *Rule 5*: if no checkpoint is within visiting range but some checkpoints are perceived ($S \neq \emptyset$), go directly towards the closest checkpoint (the direction is given by α such that $\min_d [\langle d, \alpha \rangle | \langle d, \alpha \rangle \in S]$)

This modified ruleset constitutes the model for *timeout* agents. In the presence of a checkpoint a timeout agent bases its action partly on its internal state I , specifically, on the value of the wait timer w . The timer allows agents to wait at a checkpoint for other agents to arrive when there are insufficient numbers of agents at the checkpoint to mark it visited. The wait timer mechanism works in conjunction with the filter mechanism. When the timer expires, the checkpoint is added to a set of checkpoints to be filtered from the agent's sensory input and the agent leaves. Checkpoints remain in the filter set for F cycles, during which time the agent does not perceive the checkpoint. When F cycles have elapsed, the checkpoint is removed and the agent can once again perceive the checkpoint when it is within sensory range. The filter set allows Rule 5 to remain in effect and still allow the agent to target a checkpoint that is farther away than the one just vacated, or invoke the random walk if no other checkpoint is within sensory range. The value of W was determined for each number of agents and sensory range, analogous to $rawd$, by testing a range of value combinations and choosing the one with the best performance.² We fix the value of F at 200 cycles for all experiments described here in order to reduce the size of the parameter space.

With the addition of the *Sending* and *Receiving* actions above, communicating variants of timeout agents are also possible. Communicating timeout agent performance is included in the analysis below.

Clustered Checkpoint Environments All previous agent models were tested in environments where checkpoints were randomly distributed uniformly throughout. However, we reasoned that clustering checkpoints together should allow agents to “amortize” their foraging cycles over a number of checkpoints, because once one checkpoint in a cluster has been located, the other checkpoints in the cluster will be easier

to find nearby. Communicating agents should receive even more benefit than non-communicating agents, because (potentially) only one agent needs to discover each cluster, attracting other agents to the checkpoints they may not have found as quickly otherwise. Similarly, communication should be more beneficial in structured environments than in random environments, as the value of communicated information should increase—recipients are not being informed of just a single checkpoint’s location, but rather (potentially) the locations of many checkpoints. In the modified environments the same total number of checkpoints needs to be visited ($C = 10$), but they are contained in two clusters C_i ($i \in \{1, 2\}$), each with an equal number of checkpoints ($|C_i| = 5$). Within each cluster, checkpoints are placed randomly according to a Gaussian distribution. The checkpoint clusters are centered at (200,200) and (-200,-200) with a radius of 150 units and a standard deviation of 75 units (recall that size of the environment is 800 by 800). Agents, however, do not have any *a priori* knowledge of where or how clusters are distributed. Moreover, note that agents are still placed randomly in the environment (at the same locations as above), and must search for checkpoints just as in random environments.

The results presented below are from a total of 2,268 experiments, each consisting of 40 experimental runs. The target-closest agent is tested in non-communicating, limited-range communicating (i.e., communication range identical to sensory range), and unlimited-range communicating (i.e., agents can communicate regardless of where they are in the environment) variants. Similarly, results from non-communicating, limited-range communicating, and unlimited-range communicating timeout agents are also presented below. Finally, each agent kind is tested with 21 different sensory ranges (25–800), nine different group sizes (2–10), and in each of the two environment models.

6 Summary and Implications

To summarize our findings, we conducted a five-way $9 \times 21 \times 3 \times 2 \times 2$ ANOVA with *group size* (2-10), *sensory range* (25-800), *communication* (none, limited range, and unlimited range), *wait timer* (without and with), and *environment* (random and Gaussian) as randomized variables and *average cycles to completion* as the dependent variable. Highly significant main effects were found for all five main factors, group size ($F(8,88452) = 19369.32$, $p < 0.001$), sensory range ($F(20,88452) = 1240.75$, $p < 0.001$), communication ($F(2,88452) = 937.89$, $p < 0.001$), wait timer ($F(1,88452) = 860.31$, $p < 0.001$), and environment type

($F(1,88452) = 2586.33, p < 0.001$), confirming that each of these has a significant impact on agent performance. In addition, most two-way and some three-way interactions were significant (and discussed below), but none of the four- or five-way interactions was significant.

[Figure 2 about here.]

Sensory Range vs. Group Size ($F(160,88452) = 29.96, p < 0.001$) Larger groups take better advantage of increasing sensory range when sensory range is low (Figure 2, top left). There is no clear difference between groups as to when increasing sensory range no longer helps performance, but larger groups start out with steeper and deeper reductions in average cycles to completion compared to smaller groups.

Sensory Range vs. Communication ($F(40,88452) = 90.41, p < 0.001$) Communication is most effective at low sensory ranges (i.e., exactly when information is most scarce), with performance converging as sensory range reaches 600 (Figure 2, top right). Unlimited communication range shows the best improvement, with the characteristic dip at low sensory ranges indicating that coordination effects are highest when information is harder to come by directly.

Sensory Range vs. Wait Timer ($F(20,88452) = 4.57, p < 0.001$) Target-closest and timeout agents perform similarly at very low sensory range (i.e., 25-50), but as sensory range increases, agents with the wait timer begin to outperform those without (Figure 2, bottom left).

Sensory Range vs. Environment ($F(20,88452) = 10.25, p < 0.001$) When agents' sensory range is low, performance is similar in random and Gaussian environments; it is nearly as hard to find checkpoint clusters as it is to find individual checkpoints when agents have to be almost at the cluster to sense the checkpoints. With higher sensory range, agents in Gaussian environments perform better (Figure 2, bottom right).

[Figure 3 about here.]

Group Size vs. Communication ($F(16,88452) = 51.04, p < 0.001$) Communication has more impact in medium-sized groups than in small or large ones (Figure 3, top left). When group size is too small, limited range communicating agents tend to be within range of one another only rarely. Unlimited range communicating agents receive transmissions from other agents, but are often diverted to directly sensed

checkpoints. Additional agents increase the chance that another agent will be nearby, or will be able to make it to the remote checkpoint without being sidetracked. In large groups, the task is just too easy; coordination is unnecessary, because when there are many agents in the environment, the probability of two agents selecting the same target are higher than when there are fewer agents.

Group Size vs. Wait Timer ($F(8,88452) = 86.57, p < 0.001$) The wait timer is also less effective in small and large groups than in medium groups (Figure 3, top right), for similar reasons. The timeout mechanism is effective only when agents leave a checkpoint and subsequently visit another with an agent present. When there are few agents in the environment, it is common for agents to explore different regions of the environment, seldom coming into contact. The wait timer can do little to affect this. In large groups, again, there is little need for coordination; chance is often sufficient to bring agents together at checkpoints.

Group Size vs. Environment ($F(8,88452) = 152.59, p < 0.001$) Gaussian environments have a similar effect on the coordination mechanisms: they do not directly coordinate agents, instead they decrease the degree of coordination needed by agents to perform well. It is unsurprising, then to find that the interaction between group size and environment is similar to other group size interactions (Figure 3, bottom). Small and large groups perform similarly in random and Gaussian environments; medium-sized groups show the biggest performance differences between random and Gaussian environments.

Environment vs. Communication ($F(2,88452) = 5.10, p = 0.006$) Non-communicating and limited communicating agents show nearly identical average performance improvements between random and Gaussian environments. Average cycles to completion drops by 692.6622 for non-communicating agents and by 696.4406 for limited communicating agents. However, unlimited communicating agents improve their performance by 791.323 when moving from random to Gaussian environments. Non-communicating agents must find the cluster on their own, and limited communicating agents must be near it in order for communication to be helpful, whereas unlimited communicating agents can draw others from far away, leading to better performance increases.

Environment vs. Wait Timer ($F(1,88452) = 33.05, p < 0.001$) The utility of the wait timer is less in Gaussian environments because the coordination requirements of these environments are less. Moreover, given the close proximity of checkpoints in clusters, once two agents have visited a checkpoint, there is a

strong tendency for them to move in lockstep to the remaining checkpoints in the cluster, eliminating the possibility of the wait timer expiring in timeout agents.

[Figure 4 about here.]

Sensory Range vs. Group Size vs. Communication ($F(320,88452) = 7.91, p < 0.001$) The top left graph in Figure 4 presents only 2, 3, 5, and 10 agent groups for the sake of clarity. These four group sizes display the major trends found in all sizes. First, there is little difference between non-communicating and limited communicating performance in any group size. The interesting differences are found between these two and unlimited communicating agents. Second, differences are found mostly at low sensory ranges. Third, differences begin small in small groups, get larger to a point as group size increases, after which they decrease again. The overall tendency, then, is for the largest differences to be found between unlimited communicating agents and the others in medium-sized groups with low sensory range.

Sensory Range vs. Group Size vs. Environment ($F(160,88452) = 2.05, p < 0.001$) Again, we present a reduced graph of these interactions in Figure 4 (top right). Overall, we find little performance difference in small agent groups, but as group size increases, agents become better able to take advantage of increased sensory range in Gaussian environments. For large groups, however, while there is a performance difference between random and Gaussian environments, it is fairly constant, indicating that increasing sensory range in these groups is no more effective in Gaussian than in random environments.

Sensory Range vs. Communication vs. Environment ($F(40,88452) = 1.54, p = 0.016$) We find that the benefit of increasing sensory range extends to higher ranges in Gaussian than in random environments, and that unlimited communicating agents make better use of sensory range increases in both environments (Figure 4, bottom). Non-communicating and limited communicating agents perform similarly in both environments when sensory range is low, because the difficulty of finding a cluster is similar to that of finding a checkpoint when little information is available.

[Figure 5 about here.]

Group Size vs. Communication vs. Wait Timer ($F(16,88452) = 1.73, p = 0.034$) We find once again that small and large groups exhibit less variation in performance. Non-communicating and limited com-

communicating agents perform similarly with and without the wait timer, with unlimited communicating agents showing substantial advantages in both configurations. In fact, unlimited communicating agents without the wait timer average fewer cycles to completion than either non-communicating or limited communicating with the wait timer. These differences are most pronounced in medium-sized groups.

Group Size vs. Wait Timer vs. Environment ($F(8,88452) = 12.15, p < 0.001$) In both environments, agents with the wait timer outperform agents without it. Small groups of timeout agents in random environments average fewer cycles to completion than target-closest agents in Gaussian environments, but that trend reverses as group size increases. Target-closest agents approach the performance of timeout agents in smaller groups in Gaussian faster than in random environments, confirming the effect of environmental structure on coordination requirements for MATE(2).

6.1 Discussion

It is interesting to note that for smaller group sizes, unlimited communicating agents perform better at small sensory ranges than at large ones. This can be attributed to a kind of “tunnel vision” effect: when a receiving agent A_R senses no checkpoint directly, but is attracted to a remote checkpoint via a transmitting agent A_T 's communication, it will begin to move toward the remote checkpoint. In the meantime, A_T will move to the checkpoint and wait for another agent to arrive. In the absence of communication from another agent about nearer checkpoints, A_R will continue to A_T 's checkpoint unless it directly senses some other checkpoint (which, by definition is closer than A_T 's checkpoint). If A_R has a small sensory range, it is less likely to sense some other checkpoint on its way to A_T than if it has a large sensory range. Moreover, the density of checkpoints in the environment will contribute to whether A_R senses another checkpoint en route to A_T . The probability of A_R sensing another checkpoint between its original location and A_T 's checkpoint can be given in terms of the expected number of checkpoints found in the area of the territory covered by A_R 's sensors on its way from its original location to A_T . This value (the “probability of distraction”, or P_D) is the area covered multiplied by the environment's checkpoint density:

$$P_D = (d - R_S) \cdot 2R_S \cdot \frac{|C|}{Area_{Env}} \quad (1)$$

where d is the distance A_R must travel to arrive at the checkpoint and $Area_{Env}$ is the overall area of the

environment. When P_D is less than one, it can be taken as the probability of sensing another checkpoint on the way, whereas when it is greater than one, the agent is likely to sense more than one checkpoint. Thus, the checkpoint density, A_R 's sensory range, and distance traveled determine how likely A_R is to make it all the way to A_T . If A_R is “distracted,” it will move to the newly discovered checkpoint and wait indefinitely, just as A_T is doing. In many cases, both would be better off if A_R never found another checkpoint on the way to A_T 's, hence smaller sensory range is *helpful*. As group size increases, there is a greater chance of some agent making it all the way to a remote checkpoint without sensing another checkpoint (due both to the increased probability of a nearby agent and the increased number of agents “answering the call”), so large groups do not have better performance at low sensory ranges than at high ones (Figure 2, upper right).

We have shown that communication does provide absolute performance advantages to target-closest agents in MATE(2) tasks. Communicating agents perform as well as or better than non-communicating agents. In some cases, this advantage can be attributed to a reduction in forage cycles—communication makes it easier to find checkpoints when sensory range is low, indicating that CUSP is true for MATE(2) tasks. However, communication has an additional coordinating effect beyond CUSP. This is particularly evident in unlimited communicating agents at low sensory range. At high sensory ranges, agents can see most of the environment, so all of the information is available. Yet, unlimited communicating agents at low sensory range have better performance than if they perceived almost everything. For example, when there are four agents with sensory range R_S , agents can potentially sense things that are all the way across the environment, similar to high sensory range, but the area they perceive is *at most* $4\pi R^2$. This is equivalent to extending sensory range to $\sqrt{4R_S^2}$ (i.e., doubling it). There are many cases in which unlimited communication pays off more than doubling sensory range, indicating that additional effects are at work other than just increasing the information available to the agent. Communicated data is not just random information, it is in many cases biased toward information that will be useful. Unlimited communication is biased in another way: transmitting agents do not need to be near the receiving agent, so information is more likely to be about areas where the receiving agent has not recently been. Thus, unlimited communication may have the effect of drawing agents away from regions they have already explored. The coordinating effects of communication lead us to accept a stronger version of CUSP, call it CUSP⁺: *CUSP is true and communication has additional coordination effects that improve (absolute) performance beyond that of extended sensory range*. Consequently, it is possible for communication to provide a relative performance advantage for MATE(2) tasks *if* the cost of communication is less than the benefit of the absolute performance difference it affords,

i.e., if relative performance with communication is better than without communication.

6.2 Prospects for Relative Performance Benefits in MATE(2)

The above results demonstrate that communication does provide an absolute performance advantage in some configurations; for the most part, low sensory range and medium-size groups are required for communication to be advantageous, and communication is more likely to be advantageous in random than in Gaussian environments. Again, when there is no absolute performance advantage to communicating, there can be no relative performance advantage, because agents must pay for communication. This rules out a relative performance advantage for communicating agents with sensory range above 400, for beyond this range, communicating and non-communicating performance have converged for both limited and unlimited communication.

Unlimited communication does provide substantial absolute performance advantages at low sensory ranges in both environments, but are they sufficient to constitute a relative performance advantage? This question cannot be answered definitively without a detailed cost analysis for each agent type. However, the absolute performance results do allow us to derive the relationship between the costs of the two agent types required for communication to have a relative performance advantage.

Relative performance is absolute performance scaled by cost, making the relationship between the relative performance of non-communicating and communicating agents $P(R_S, k) \cdot F(R_S) = d \cdot Q(R_S, R_C, k) \cdot G(R_S, R_C)$, where $P(R_S, k)$ is the absolute performance of k non-communicating agents with sensory range R_S , $Q(R_S, R_C, k)$ the absolute performance of communicating k target-closest agents with sensory range R_S and communication range R_C , $F(R_S)$ and $G(R_S, R_C)$ the cost functions for non-communicating and communicating agents with sensory range R_S and (in the latter case) communication range R_C , respectively, and d the constant that specifies the relation between the two. When communicating agents have a relative performance advantage over non-communicating agents $d > 1$, so $Q(R_S, R_C, k) \cdot G(R_S, R_C) < P(R_S, k) \cdot F(R_S)$. Hence, for communication to provide a relative performance benefit, the following inequality must hold: $G(R_S, R_C) < \frac{Q(R_S, R_C, k)}{P(R_S, k)} \cdot F(R_S)$. This allows us to determine, for any two agent configurations from the results above, the constraints on the cost of communication that must be satisfied to show a relative performance advantage for communication.

In the experiments above, $P(R_S, k)$ ranged from 1 to 2 times $Q(R_S, R_C, k)$. In the worst case, where the ratio is 1:1 (i.e., performance was identical), communicating agents would have to have lower cost

than non-communicating agents ($G(R_S, R_C) < F(R_S)$). As we have said, this is not possible, given that communication adds costs, so communication cannot have a relative performance advantage at the low end. In the best case, the cost of communicating agents could be nearly twice the cost of non-communicating agents ($G(R_S, R_C) < 2 \cdot F(R_S)$). At first glance this may seem enough to cover the cost of communication, but recall that the best-case results were from agents with unlimited communication range, which must certainly be very expensive. Hence, it is still unclear whether these results represent a relative performance benefit. To gain a better understanding, more details of the cost functions are needed.

For the purposes of this analysis, we make some simplifying assumptions about agent cost. We consider only operating, not construction costs, and we aggregate most of the operating costs (e.g., the cost of moving, processing costs, etc.) into a generic *base cost*, treating separately only sensory and communication cost. Many components of base cost are not accrued every cycle (e.g., an agent does not move every cycle), in which case their component of the base is the average per cycle. Moreover, we assume that the cost of communication (and sensing) is limited to the physical cost of *receiving* the information each cycle (i.e., of operating the sensors required to receive shared information, such as listening for a call, or monitoring a wireless signal); other costs of communication (e.g., the physical costs of transmitting the information, the processing costs incurred to translate coordinates from one agent-centric system to another agent-centric system, etc.) are rolled into the base cost. As a final simplification, to facilitate comparison, we assume that the base cost for communicating and non-communicating agents is identical, although in reality differences in their inputs will lead to cost differences (e.g., because one kind travels more than the other, communicating agents will have more inputs to process, etc.).

The simplified cost functions for non-communicating target-closest agents ($F(R_S)$) and communicating target-closest agents ($G(R_S, R_C)$) are as follows:³

$$F(R_S) = s \cdot R_S^2 + b \quad (2)$$

$$G(R_S, R_C) = s \cdot R_S^2 + c \cdot R_C^2 + b \quad (3)$$

The constants s and c are the unit costs for sensory range and communication range, respectively, and b is the base cost. Signal quality drops quadratically with distance, so receiving distant signals requires additional effort, making the cost of receiving quadratic in the communication (or sensory) range. For communication to have a relative performance advantage given these cost functions, the inequality $s \cdot R_S^2 + c \cdot R_C^2 + b <$

$\frac{Q(R_S, R_C, k)}{P(R_S, k)} \cdot (s \cdot R_S^2 + b)$ must hold.

The most promising prospect for a relative performance advantage (i.e., the configuration with the largest absolute performance advantage for unlimited communication over non-communicating target-closest agents) is in random environments with group size $|A| = 4$ and sensory range $R_S = 50$ (Figure 4, upper left), where non-communicating agents require just over double the number of cycles on average, making $\frac{Q(R_S, R_C, k)}{P(R_S, k)} \approx 2$. As for R_C , although we term it “unlimited,” it is best analyzed as limited communication that extends to the whole of the limited environment, in this case slightly more than $R_C = 1100$ (i.e., $\sqrt{2} \cdot 800$ to reach from corner to corner), so the cost is fixed based on R_C .⁴ Hence, for communication to yield a relative performance advantage in this configuration, the inequality $c \cdot R_C^2 < b + s \cdot R_S^2$ must hold. Either the unit cost of communication c must be very small ($c < \frac{b+s \cdot 2500}{1210000}$), or the base cost b must be very large ($b > c \cdot 1212500$, if we assume that the unit costs $c \approx s$). We leave open this possibility (e.g., perhaps some inexpensive wireless communication), but note again that, not only do these calculations represent the best-case scenario, they are also biased in favor of communication by ignoring some of its costs, as described above.

Although relative performance for limited communication can be analyzed in the same way, it is probably not necessary to do so. The cost of limited communication is much less than for unlimited communication, but its benefit is also much smaller. Moreover, in those configurations where limited communication holds its greatest absolute performance advantage over target-closest agents, timeout agents perform better absolutely. Given the low cost of the wait timer,⁵ this is likely to translate to better relative performance. This, again, leaves only a narrow set of circumstances in which limited communication would have a relative performance advantage over timeout agents: communication must be very inexpensive.

6.3 Implications for MATE(n)

What do the MATE(2) results tell us about MATE(n)? We claim that the conclusions transfer directly to MATE(n): the cost of communication must be very low in order for communication to have a relative performance advantage over non-communicating strategies in MATE(n). To make this case based only on the MATE(2) results, we must demonstrate that there is nothing specific to MATE(n) tasks of higher degree that would lead to a different outcome than in MATE(2). For, one might accept the claim for MATE(2), but maintain that its degree of complexity is not sufficiently high to allow communication to pay off. In that case, communication could have positive utility for some MATE(n) task with larger n .

We acknowledge that increasing the complexity of the task is likely to increase the benefit of communication. However, there is no reason to believe that the same mechanisms that work for MATE(2) to coordinate agent behavior such that two agents can be at the same checkpoint at the same time will not generalize to MATE(n). Thus, mechanisms for coordination as part of the agent function *per se* will not change (modulo the generalization), even though the actual performance (based on the distribution of agents and checkpoints) might differ between MATE(n) and MATE(k) tasks ($n \neq k$). Communication may contribute more to coordination in MATE(3) than in MATE(2), but other coordination mechanisms that work well for MATE(2) (e.g., the wait timer) will also contribute more to coordination in MATE(3). There is no evidence that the performance ordering between coordination strategies would change simply based on an increase in the degree of a MATE(n) task.

7 Biological Implications

MATE(n) tasks are closely related to tasks performed by biological agents (e.g., in the context of foraging or mating). If communication were to evolve in these agents, it could be a powerful tool to aid in the performance of these tasks. This suggests the *Communication Evolved for Coordination Hypothesis* (CECH), according to which communication evolved in agent kind K for the purpose of coordinating behavior among agents in task T . Based on the results above, what can we conclude about CECH for biological variants of MATE(n) tasks?

We chose to examine MATE(2) here because it is the MATE(n) task of lowest complexity that still requires agents to actively coordinate behavior in order to achieve good performance. MATE(2) is also the most appropriate choice for studying coordination in biological agents because, whereas higher-degree MATE(n) tasks are difficult to motivate biologically, tasks in which two agents must detect and meet at some location are biologically plausible. Animals that have to come together at an appropriate mating location in the environment (e.g., for safety from predators or proximity to food sources) to lay and fertilize eggs (e.g., like bullfrogs Emerson and Boyd (1999)) are examples of animals performing MATE(2) tasks. There are likely MATE(3) tasks in nature (e.g., a food source requiring three agents to harvest), but they are certainly not common. For communication to evolve for coordination on a rarely occurring task, the benefit would need to be very high to outweigh the cost of maintaining the communication facilities the rest of the time. It is more likely that communication evolved for some other purpose, and is subsequently applied to those

few MATE(3) tasks with which the agents are confronted.

Having established that performance in MATE(n) is relevant to biological agents (i.e., predictive of the performance of biological populations; see Section 2), and having demonstrated the stringent constraints on the cost of communication that must be met if it is to be beneficial (see Section 6.2), we conclude that CECH is likely to be false for MATE(n) tasks and closely-related variants: communication is unlikely to have evolved for the purpose of coordination in MATE(n)-like biological foraging tasks. It is at least possible to envision some engineering solution wherein the cost of communication is sufficiently small relative to the remaining costs, but biological agents are constrained by their “flesh and blood” construction, making it very difficult to imagine how such cost ratios could be achieved. We know, for example, that in some cases communication is actually more expensive than the other costs we categorized as base costs above (McLister, 2001).

It is important to note that CECH is a statement about the evolution of communication for coordination in general, not about any particular, actual evolutionary history on earth. We claim that CECH is likely false for MATE(n) tasks *in general*, but this does not preclude the possibility that particular circumstances might have favored the evolution of communication in a particular species in a MATE(n) task. For example, the actual environmental conditions on earth could have been such that, while CECH is false in most reasonable conditions, communication *is* beneficial in a limited range of conditions, one of which which happened to be the case in the actual evolutionary history on earth. Without the details of the initial distributions of resources and agents, however, we can make comments only about distributions in general. Moreover, note that while the statistical analyses of the results of many MATE(n) task experiments demonstrate that it is unlikely that CECH is true for MATE(n), this does not rule out the possibility that communication could have evolved for other purposes and that it was only subsequently used for coordination in those tasks.

Finally, note that we do *not* claim that CECH is false for *every* foraging task. After all, communication has evolved many times in many species—bees (Capaldi and Dyer, 1999; Menzel et al., 1998), for example, and, perhaps to some extent, ants (Nicholson et al., 1999; Shen et al., 1998) communicate the locations of resources to their peers. So there may be aspects of these natural foraging tasks are *not* captured by MATE(n) tasks. For example, bees and ants return to their nests either to deposit the resources they have gathered or to inform their peers about where the resources are. Hence, it could be that departing from and returning to a common location makes a difference in foraging patterns such that communication would be more efficient (e.g., because agents do not need to be able to communicate over long distances in order to

share information about distant food sources).

Interestingly, there is biological evidence for the claim that communication is critical for the survival of colony creatures like bees *as well as* its negation—that bees are capable of surviving *without communication*—depending on the food distribution in the environment (Dornhaus and Chittka, 2004). Hence, it may well be that the same effects of communication’s benefit being contingent upon the distribution of food in the environment can be observed in $MATE(n)$ tasks modified to include the additional constraint of having to bring resources back to the hive (i.e., always having to return to a common “base checkpoint” after visiting a “field checkpoint”). In fact, we have evidence from preliminary simulations that this is the case (Schermerhorn and Scheutz, 2005). However, note that such tasks require additional premises about agent distributions and visiting patterns that were not part of our investigation here and thus do not support CECH for $MATE(n)$. If communication proved essential for such a $MATE(n)$ -like task, augmented with “home base” constraints given a particular distribution of checkpoints, then it would be easy to point to the property that ultimately favored communication in it.

8 Conclusion

In this paper we investigated the adaptiveness of communication for coordination in biologically plausible $MATE(2)$ tasks. A systematic examination of the effects of varying sensory and communication range and environmental conditions revealed that there is no simple “yes or no” answer. Although communication is not adaptive throughout most of the space defined by these variables, there are some regions in which communication does provide an absolute performance advantage. In particular, medium-sized groups of agents with low sensory range and unlimited communication range gain a substantial absolute performance benefit from communication, and even those with limited communication range enjoyed small absolute performance advantages over non-communicating agents.

The analysis of relative performance demonstrates a method for finding the relationships that must hold between different components of agent cost to demonstrate a relative performance advantage. Applying this method to the configuration in which communication’s absolute performance advantage was greatest (unlimited communicating vs. non-communicating agents in groups of size 4 with sensory range 50), we demonstrated that even in the best case, the cost of communication must be much smaller than the remaining costs if it is to be relatively beneficial. Limited communication incurs a much smaller cost than unlimited

communication, but also a much smaller absolute performance benefit. Moreover, the wait timer mechanism outperforms unlimited communication absolutely, despite its significantly lower cost.

The relative performance analysis establishes constraints on the cost of communication that engineers of communicating agents must satisfy if a relative performance benefit is desired. Artificial agents may be able to achieve the required ratio, but it is very unlikely that a solution that preserves a relative performance advantage for communication can be found for biological agents. CECH (i.e., the claim that communication evolved for the purpose of coordination), then, is likely to be false for this task; the cost of communication is simply too great to establish a relative performance advantage in biological agents, so communication is unlikely to evolve for coordination on MATE(2) tasks. We also claimed that the results for MATE(2) generalize to all MATE(n), because there is no reason to believe that increasing n (and hence the complexity of the task) would change the performance ordering established in MATE(2). Therefore, we conclude, where communication has evolved in agents performing MATE(n)-like foraging tasks, it has evolved for reasons other than coordination (although it may subsequently be used for that purpose). Similarly, where communication *has* evolved for coordination in foraging, there must be characteristics of those foraging tasks that are importantly different from MATE(n) (e.g., as may be the case when agents must return found resources to a central hive location). Detailed analyses and results to explore these possibilities will be left for future work.

Notes

¹The experiments and results of these *random walk spaces* are available at <http://www.nd.edu/~airolab/publications/schermerhornscheutz07mate2-calibration.pdf>.

²Like the random walk spaces, experiments and results of these *wait timer duration spaces* are available at <http://www.nd.edu/~airolab/publications/schermerhornscheutz07mate2-calibration.pdf>.

³The cost functions presented here should not be taken to represent accurate cost models for realistic agents. The actual cost models for two competing agent types would need to be determined before a precise comparison could be made. However, that level of detail is unnecessary to establish a sense of the relationship between communication cost and other costs necessary for a relative performance advantage.

⁴Similarly, extraplanetary landers may have virtually unlimited communication range (i.e., anywhere on the surface) via the use of satellite communication, in which case the cost is fixed according to the distance at which the satellites orbit.

⁵The wait timer can be implemented by a simple connectionist unit with recurrent inhibitory connections. Functionally, the unit receives activation whenever an agent arrives at a checkpoint, which it subsequently suppresses via its self-inhibitory connection. The output of the unit, in turn, suppresses the agent's navigation module. Over time, the activation will become smaller and smaller. Eventually, it will be insufficient to suppress the navigation module any longer, at which point the agent will resume its foraging behavior. Moreover, the "filter set" in which the checkpoint is included can be implemented by a simple "avoid past" mechanism (Balch and Arkin, 1993; Scheutz and Andronache, 2004).

List of Figures

1	Average performance of target-closest agents for 2 to 10 agents and sensory ranges from 25 to 800.	39
2	Two-way interactions between sensory range and other factors	40
3	Two-way interactions between group size and other factors	41
4	Three-way interactions between sensory range and other factors	42
5	Three-way interactions between group size and other factors	43

References

- Ackley, D. and Littman, M. (1991). Altruism in the evolution of communication. In Brooks, R. and Maes, P., editors, *Artificial Life IV: Proceedings of the Fourth Workshop on Artificial Life*, pages 40–48, Cambridge, MA. MIT Press.
- Applegate, D., Cook, W., Dash, S., and Rohe, A. (2002). Solution of a min-max vehicle routing problem. *INFORMS Journal on Computing*, 14(2):132–143.
- Balch, T. and Arkin, R. (1993). Avoiding the past: a simple but effective strategy for reactive navigation. In *Proceedings of the 1993 IEEE International Conference on Robotics and Automation*, pages 678–685, Atlanta, GA. IEEE Computer Society Press.
- Bugera, V. (2004). Properties of no-depot min-max 2-traveling-salesmen problem. In Butenko, S., Murphey, R., and Pardalos, P., editors, *Recent Developments in Cooperative Control and Optimization*, pages 45–60. Kluwer Academic Publishers, Norwell, MA.
- Cangelosi, A., Riga, T., Giolito, B., and Marocco, D. (2004). Language emergence and grounding in sensorimotor agents and robots. In *First International Workshop on Emergence and Evolution of Linguistic Communication*, Kanazawa, Japan.
- Capaldi, E. A. and Dyer, F. C. (1999). The role of orientation flights on homing performance in honeybees. *The Journal of Experimental Biology*, 202:1655–1666.
- Clement, B. J. and Barrett, A. C. (2003). Continual coordination through shared activities. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems (AAMAS '03)*, pages 57–64, Melbourne, Australia. ACM Press.
- Cooper, R., DeJong, D. V., Forsythe, R., and Ross, T. W. (1992). Communication in coordination games. *The Quarterly Journal of Economics*, 107(2):739–771.
- Cormen, T. T., Leiserson, C. E., and Rivest, R. L. (1990). *Introduction to Algorithms*. MIT Press, Cambridge, MA.
- Dornhaus, A. and Chittka, L. (2004). Why do honey bees dance? *Behavioral Ecology and Sociobiology*, 55(4):396–401.

- Emerson, S. and Boyd, S. (1999). Mating vocalizations of female frogs: control and evolution mechanisms. *Brain, Behavior and Evolution*, 53:187–197.
- Estlin, T., Rabideau, G., Mutz, D., and Chien, S. (2000). Using continuous planning techniques to coordinate multiple rovers. *Electronic Transactions on Artificial Intelligence*, 4:45–57.
- Gavish, B. and Srikanth, K. (1986). An optimal solution method for large-scale multiple traveling salesmen problems. *Operations Research*, 34(5):698–717.
- Gervasi, V. and Prencipe, G. (2004). Coordination without communication: The case of the flocking problem. *Discrete Applied Mathematics*, 143(1-3):203–223.
- Golden, B., Bodin, L., Doyle, T., and Jr., W. S. (1980). Approximate traveling salesman algorithms. *Operations Research*, 28(3):694–711.
- Grim, P., Kokalis, T., Tafti, A., and Kilb, N. (2002). Evolution of communication with a spatialized genetic algorithm. *Evolution of Communication*, 3(2):105–134.
- Johnson, D. S. and McGeoch, L. A. (1997). The traveling salesman problem: A case study in local optimization. In Aarts, E. H. L. and Lenstra, J. K., editors, *Local Search in Combinatorial Optimization*, pages 215–310. John-Wiley and Sons, Ltd., Hoboken, NJ.
- Levin, M. (1995). The evolution of understanding: A genetic algorithm model of the evolution of communication. *BioSystems*, 35:167–178.
- MacLennan, B. (1991). Synthetic ethology: An approach to the study of communication. In Langton, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II: Proceedings of the Second Workshop on Artificial Life*, pages 631–658, Redwood City, Calif. Addison-Wesley.
- Marocco, D., Cangelosi, A., and Nolfi, S. (2003). The emergence of communication in evolutionary robots. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, 361(1811):2397–2421.
- McLister, J. (2001). Physical factors affecting the cost and efficiency of sound production in the treefrog *Hyla versicolor*. *Journal of Experimental Biology*, 204(1):69–80.
- Menzel, R., Geiger, K., Joerges, J., Müller, U., and Chittka, L. (1998). Bees travel novel homeward routes by integrating separately acquired vector memories. *Animal Behavior*, 55(1):139–152.

- Miller, J. H., Butts, C. T., and Rode, D. (2002). Communication and cooperation. *Journal of Economic Behavior and Organization*, 47(2):179–195.
- Miller, J. H. and Moser, S. (2003). Communication and coordination. Working Paper 03-03-019, Santa Fe Institute, Santa Fe, New Mexico.
- Montemanni, R., Gambardella, L., Rizzoli, A., and Donati, A. (2003). A new algorithm for a dynamic vehicle routing problem based on ant colony system. In *Second International Workshop on Freight Transportation and Logistics*.
- Nicholson, D. J., Judd, S. P. D., Cartwright, B. A., and Collett, T. S. (1999). Learning walks and landmark guidance in wood ants (*Formica rufa*). *The Journal of Experimental Biology*, 202(13):1831–1838.
- Noble, J. (1999). Cooperation, conflict and the evolution of communication. *Adaptive Behavior*, 7(3-4):349–370.
- Noble, J. and Cliff, D. (1996). On simulating the evolution of communication. In Maes, P., editor, *Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 608–617, Cambridge, MA. MIT Press.
- Parker, L. E., Kannan, B., Tang, F., and Bailey, M. (2004). Tightly-coupled navigation assistance in heterogeneous multi-robot teams. In *Proceedings of IEEE International Conference on Intelligent Robots and Systems (IROS)*. IEEE Computer Society Press.
- Rabideau, G., Estlin, T., Chien, S., and Barrett, A. (1999). Working together: Centralized command sequence generation for cooperating rovers. In *Proceedings of the IEEE Aerospace Conference (IAC)*.
- Reynolds, C. W. (1987). Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21(4):25–34.
- Schermerhorn, P. and Scheutz, M. (2005). The effect of environmental structure on the utility of communication in hive-based swarms. In *IEEE Swarm Intelligence Symposium 2005*, pages 440–443. IEEE Computer Society Press.

- Scheutz, M. and Andronache, V. (2004). Architectural mechanisms for dynamic changes of behavior selection strategies in behavior-based systems. *IEEE Transactions of System, Man, and Cybernetics Part B*, 34(6):2377–2395.
- Scheutz, M. and Schermerhorn, P. (2003). Many is more but not too many: Dimensions of cooperation of agents with and without predictive capabilities. In *Proceedings of IEEE/WIC IAT-2003*, pages 378–384. IEEE Computer Society Press.
- Scheutz, M. and Schermerhorn, P. (2005). Predicting population dynamics and evolutionary trajectories based on performance evaluations in alife simulations. In *Proceedings of GECCO 2005*, pages 35–42. ACM Press.
- Shen, J.-X., Xu, Z.-M., and Hanks, E. (1998). Direct homing behavior in the ant *Tetramorium caespitum* (formicidae, myrmicinae). *Animal Behavior*, 55(6):1443–1450.
- Tang, F. and Parker, L. E. (2005). Asymtre: Automated synthesis of multi-robot task solutions through software reconfiguration. In *Proceedings of 2005 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1501–1508. IEEE Computer Society Press.
- Thompson, P. M. and Psaraftis, H. N. (1993). Cyclic transfer algorithms for multivehicle routing and scheduling problems. *Operations Research*, 41(5):935–946.
- Wang, H. and Xue, D. (2002). An intelligent zone-based delivery scheduling approach. *Computers in Industry*, 48(2):109–125.
- Werger, B. B. (1999). Cooperation without deliberation: a minimal behavior-based approach to multi-robot teams. *Artificial Intelligence*, 110(2):293–320.

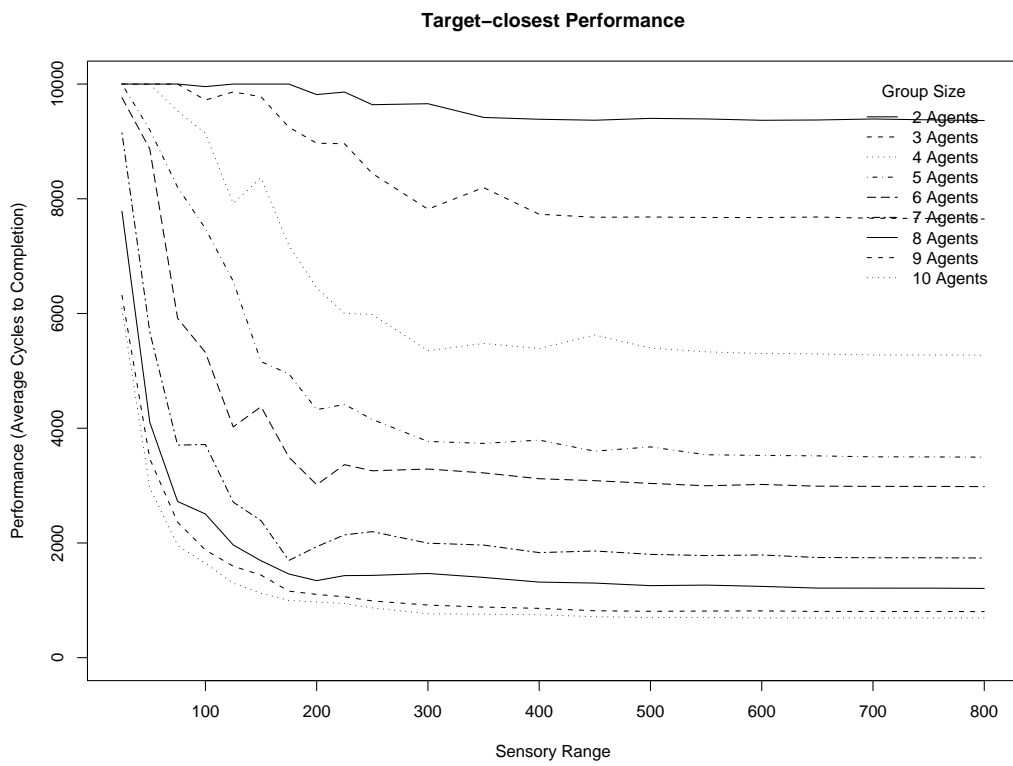


Figure 1

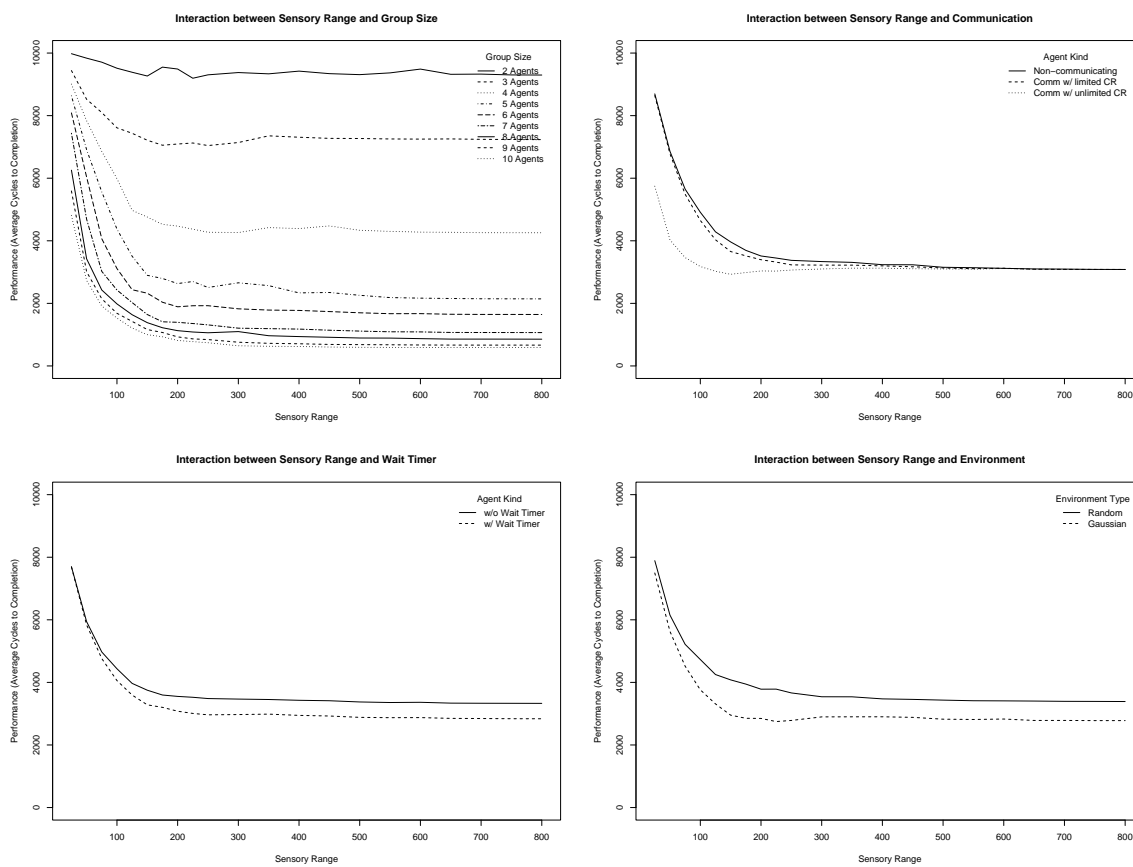


Figure 2

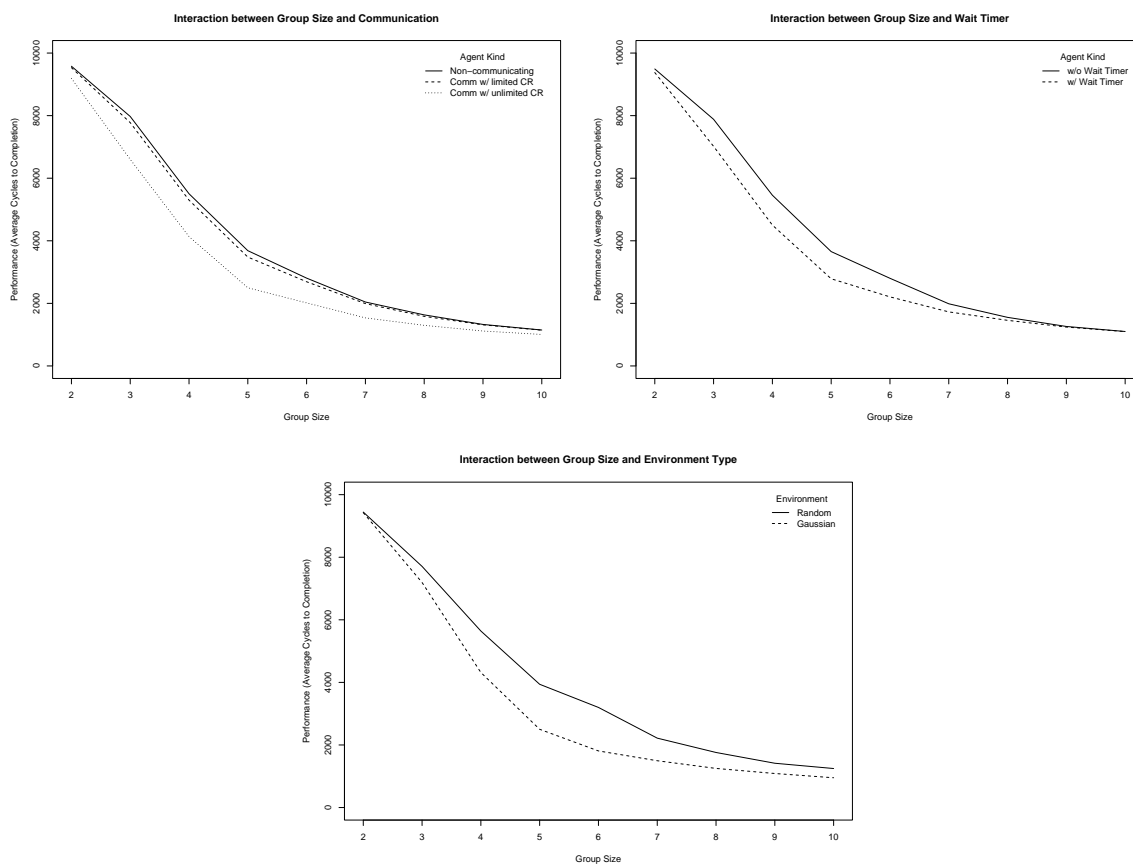


Figure 3

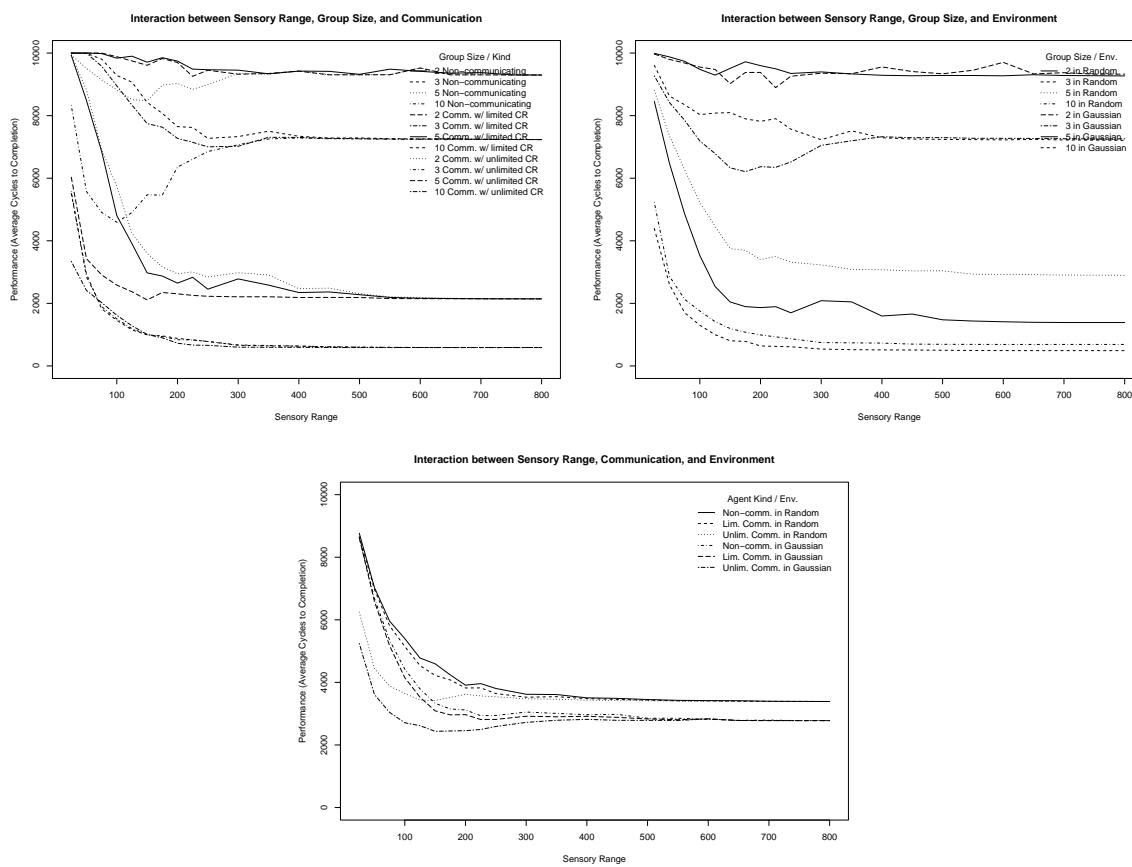


Figure 4

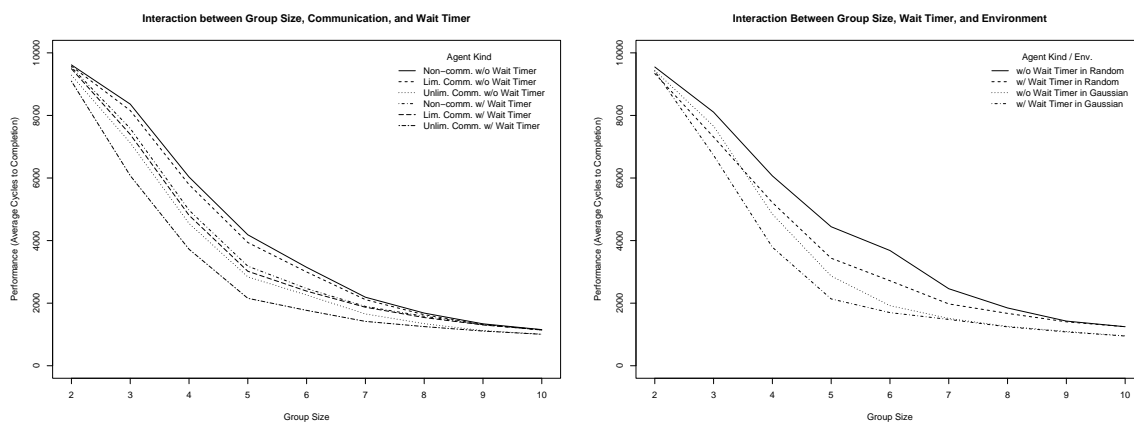


Figure 5

Paul Schermerhorn received the M.A. in philosophy at Northern Illinois University in 1999 and the M.S. and Ph.D. in computer science and engineering at the University of Notre Dame in 2002 and 2006, respectively. He is currently a post-doctoral researcher in the Artificial Intelligence and Robotics Lab at Notre Dame. His primary research interests include affective robot control systems and agent-based modeling of social interactions in biological agents. Address: 384 Fitzpatrick Hall, University of Notre Dame, Notre Dame, IN, 46556; Email: pscherm1@cse.nd.edu

Matthias Scheutz received the Ph.D. in philosophy at the University of Vienna in 1995 and the joint Ph.D. in cognitive science and computer science from Indiana University in 1999. He is an assistant professor in the Department of Computer Science and Engineering at the University of Notre Dame. He has over 80 peer-reviewed publications in artificial intelligence, artificial life, agent-based computing, cognitive modeling, foundations of cognitive science, and robotics. His current research interests include complex cognitive and affective architectures for human-robot interaction, computational models of language processing in mono- and bilinguals, distributed agent architectures, and interactions between affect and cognition. Address: 384 Fitzpatrick Hall, University of Notre Dame, Notre Dame, IN, 46556; Email: mscheutz@cse.nd.edu