Cooperative Decision-Making in Decentralized Multiple-Robot Systems: the Best-of-N Problem

Chris A. C. Parker, *Member, IEEE,* Hong Zhang, *Member, IEEE*

***Abstract*—Multiple-robot systems (MRS) that are decentrally organized have many benefits over centralized systems. Decentralized systems are less affected by computational and communicative bottlenecks, and they are more robust to the loss of individual member robots. System-level cognitive operations, though, are much more difficult to implement in decentralized systems. One example is the best-of-N decision-making problem, in which a team attempts to unanimously select a single alternative from a list that maximizes a given metric. This is a valuable operation, since many system-level operations can be expressed in this form. Optimal best-of-N decision-making, however, is intractable in large decentralized systems. The contribution of this work is a biologically inspired algorithm that enables a decentralized MRS composed of very simple robots to make good, unanimous decisions. In a series of physical experiments using real robots, the best decision was made at least 80% of the time. 100% of the decisions achieved perfect consensus, which prevented the MRS from becoming fragmented. The decisions are made using only anonymous, local communication, with no direct comparisons of the available alternatives by the individual robots.**

***Index Terms*—Decentralized, multi-robot system, decision-making.**

I. Introduction

The inability of decentralized multiple-robot systems (dec-MRS) to make intelligent collective decisions is a significant obstacle to their deployment in the real world. A dec-MRS is a robotic team which lacks any centralized (or hierarchical) organization [8]. Robustness to individual robot failures and scalability are two advantages of dec-MRS over their centralized counterparts. A centralized MRS would fail if its central controller failed, but this is not the case for a dec-MRS since they contain no such critical individuals. Their absence also allows dec-MRS to be scaled up to very large population sizes (such systems sometimes are referred to as swarms), since communicative and computational bottlenecks largely are absent. In this work, we propose a solution to the decentralized best-of-N decision problem, in which a dec-MRS must unanimously select one of N alternatives to some problem facing it. Empowering a system to make collective decisions as though it were a single intelligent entity enables designers to abstract away from the low-level complexities of decentralized control and instead focus on the high-level details of the task at hand.

Dec-MRS enjoy the advantages that they do because their control is bottom-up (emergent), which also is the source of the difficulty in their development, since there is no "top" from which control signals might emerge, nor where information could be collected for processing. Because of this, much of the research to date has concentrated coordination mechanisms for specific tasks, or the simultaneous execution of multiple tasks that do not interfere with each other. Often, these algorithms are tightly coupled to specific environments [3]. Abstract cognitive operations like best-of-N decision-making are domain-independent, and thus can be applied to a wide variety of dec-MRS and reused in many situations.

Naturally occurring decentralized multiple-agent systems like ant and bee colonies have served as valuable examples in the past for dec-MRS applications, and we follow this model by presenting a biologically inspired solution to the decentralized best-of-N decision-making problem in this work. Social insects are constrained similarly to the robots of a dec-MRS. In both systems, communication tends to be local and thus largely one-to-one. The individuals also tend to be simple, with noisy sensors and incomplete global knowledge. Because of these limitations, individual robots, just like their insect counterparts, should not be relied upon to determine the appropriate response to their environment beyond the most basic situations. Otherwise, the global behavior of decentralized systems would be dominated by interference [11] due to the conflicting actions of erroneous individuals; a problem that tends to worsen as population size increases.

In response to these concerns, this work presents a decentralized algorithm that enables a dec-MRS to make group decisions cooperatively, that takes advantage of a dec-MRS's redundancy and highly parallel structure instead of being penalized by it. We demonstrate our work with real robots in a collective relocation domain, but our algorithm easily could be applied to any problem that could be expressed as a best-of-N decision. In the next section, we discuss research related to our problem of interest. In Section III, we formally define the decentralized best-of-N decision-making problem and describe the strategy used by honeybees and ants to solve it. Inspired by their solution, we propose our best-of-N decision-making algorithm for dec-MRS in Section IV. We also present quorum testing, a new mechanism in dec-MRS that plays a central role in our algorithm. In Sections V and VI, we describe a series of experiments that we carried out to demonstrate best-of-N decision-making by a real dec-MRS. We generalize the results of these experiments with further discussion in Section VII, and then close with some conclusions and suggest the next steps for research in this area. We first introduced the use of social insect behavior to guide consensus-based [9] decisions in dec-MRS in [21], [22]. Some parts of those works are reproduced herein.

Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada e-mail: parkerca@ieee.org.

II. Related Work

Interest in cooperative teams of mobile robots [14], [32] has increased in recent years. Often, dec-MRS are composed of small, simple robots such as those described in [6], [14]. As Campbell et al. have pointed out, the cost of computation is independent of robot size [7], but the ability to compute is not. Therefore, smaller autonomous robots, which due to their small size cannot carry sophisticated on-board controllers, nor the batteries necessary to power them are more computationally constrained than larger robots. These are the sorts of robots likely to compose many dec-MRS in the future, and so cooperative decision-making algorithms intended for these systems must take their limitations into consideration. To date, however, most of the research concerning dec-MRS has concentrated on their coordination when carrying out specific tasks, rather than general purpose system-level cognitive operations.

Examples of dec-MRS applications in the literature include object sorting [13], cooperative transportation [19], formation movement, and the stick-pulling experiment [15]. These studies illustrate that dec-MRS are capable of a wide variety of tasks, but each of these dec-MRS were single-purpose systems. A more complicated task like cooperative construction [30] can be expressed as the composition of foraging, transport and sorting [24], but in order to organize it as such, a MRS would have to be able to make collective decisions regarding where to build a desired structure, when to begin, where to gather the materials, and when each subtask had been completed so that the next could begin [25]. Otherwise, the group task likely would fail as each robot might determine a different course of action on its own. Each of these questions can be posed as a best-of-N decision.

There has been work in artificial intelligence studying team decision-making, but much of it focuses on smaller teams of relatively sophisticated agents, rather than the low-cost resource-constrained and expendable robots that often are envisioned composing large dec-MRS. Its extension to large dec-MRS therefore is questionable. One example of such work is Pynadath and Tambe's communicative multiagent team decision problem (COM-MTDP) in [27]. Similar work has cast the cooperative decision-making problem as a decentralized partially observable Markov decision process, or dec-POMDP [4], suggesting that a general purpose solution to the decentralized best-of-N problem will have to be heuristic, since an optimal solution to this problem would be intractable for decentralized systems containing many robots.

International competition in robotic soccer also has driven cooperative decision-making research. If a robotic soccer team is to compete effectively, its players must coordinate their actions. This coordination often takes the form of a group decision regarding which play to follow. Plays commonly are fixed *a priori* in a playbook and the robots must agree which play best suits the circumstances in which they find themselves. Therefore, the manner in which they deliberate constitutes a solution to the best-of-N decision problem. Bowling et al. developed an effective play-selector for their team in the small-size league that was able to adapt to different opponents' strategies in [5]. Like many small-size league techniques, however, this was a centralized algorithm that dictated roles to the individual robots. Strategies in the mid-sized league tend to be more decentralized, since these robots cannot take advantage of centralized off-field hardware. Kok et al. presented several truly decentralized algorithms for play selection in [17], [18]. Because soccer teams have small population sizes, many of the cooperative decision-making algorithms used in this domain are optimized for speed in small dec-MRS and take advantage of communicative behavior that would be impractical in for larger populations. It is unlikely that these algorithms would scale up to large population populations. Also, in both of these works, the set of alternatives over which the systems deliberated were known beforehand. This is not a general property of a best-of-N decision.

An area of research that is somewhat similar to that of swarm-style dec-MRS with respect to population size and agent complexity is sensor networks. A sensor network consists of a large number of simple, expendable sensor nodes that are spread over a region to monitor some phenomenon within it [1]. Decentralized decision-making often is performed in sensor networks to reduce the amount of data that must be passed to the central sink node of the network, such as the decision-fusion proposed in [16], but because the individual sensor nodes tend to be immobile, these decisions usually are local, combining the independent decisions of nearby nodes. More distant nodes are unaffected by these decisions beyond a reduction in the volume of network traffic that they might have to carry. The resulting behavior might better be described as self-organizing rather than decision-making as we intend the term in this article. System-wide decisions in sensor networks are more likely to be made by a centralized sink node and dictated to the rest of the network. Although there has been some interest in mobile sensor nodes, the majority of the research in this field concerns stationary nodes, and thus the topology of a sensor network is quite stable when compared to that of a dec-MRS.

To our knowledge, the only other work that has investigated cooperative decision-making in large dec-MRS is that of Wessnitzer and Melhuish in [31]. In that work, a dec-MRS was tasked with pursuing and immobilizing two "prey" in a series of simulated experiments. The robots, each of which possessed minimal sensing capability and short communication ranges used majority voting and a hormone-inspired approach to cooperatively decide which prey to follow. They then pursued the other one once they agreed that first had been immobilized. Our work here is somewhat similar, but unlike that of Wessnitzer and Melhuish, we present an approach that is only very loosely coupled to the specific decision being made, and thus is more general in its approach.

III. The Decentralized Best-of-N Decision-Making Problem

In this section we formally introduce the decentralized best-of-N decision-making problem and describe the emergent behavior used by honeybees and ants to solve it. In the next section, we present our algorithm for best-of-N decision-making in a dec-MRS which is patterned after the insects' behavior.

The best-of-N decision-making problem arises when a single alternative that maximizes some value function must be selected from a set of candidate alternatives. The set of alternatives could be anything: locations at which to perform some operation, heuristically determined solutions to an intractable problem, etc. A solitary robot need only select the best one in its list of known alternatives in order to make a best-of-N decision. In a MRS, the robots must unanimously select one alternative, since the goal is for the *system* to make the best-of-N decision. The solution for a centralized system is much the same as that of a solitary robot. Making a best-of-N decision in a decentralized system, however, is more complex. Because of its structure, knowledge is spread across the individual members, meaning that none of them are likely to know of the entire set of candidate alternatives. It is the cost of sharing each robot's knowledge with every other member of a dec-MRS that makes optimal best-of-N decision-making intractable. Even if this information could be shared, only one of the alternatives ultimately will be selected by the end of a decision, so much time and energy would be wasted spreading knowledge of alternatives that will not be selected system-wide. In the presence noisy sensors, even if each robot could be made aware of each candidate alternative, it cannot be assumed that each robot would make the same best-of-N decision as every other robot, since the values ascribed by the various robots to each of the alternatives likely would differ.

A conceptually simple approach to the decentralized best-of-N decision-making problem would take advantage of global broadcast communication. Following an initial search of some predetermined length, each robot would broadcast its own local solution to the rest of its team. Once all of the broadcasts had been made, each robot would then select the alternative that had received the most support, completing the collective decision. There are several problems with this approach, though, that make it undesirable. First of all, it is an implicitly centralized algorithm. It strives to treat every member of a dec-MRS as though it was a central controller, giving it complete knowledge in order to make an omnipotent decision. Second, it relies heavily on global broadcast communication. Not only is this a relatively expensive operation, it can be unreliable in a decentralized system, and the success of the group decision depends critically on every robot receiving every one of its teammates' broadcasts. Furthermore, each robot is required to know precisely how many robots compose the system, since they can only make their decisions about which alternative to select once every teammate has broadcast its opinion. Even if this number was known initially, recall that robustness to individual failure is one of the main arguments in favor of dec-MRS. The simple, disposable robots [14] that are likely to compose these systems might fail during a dangerous mission. It is unlikely that the surviving robots would be able to maintain an accurate list of their functional teammates over the course of a mission. In the end, this approach is not even guaranteed to make good best-of-N decisions, since each robot's vote for the "best" alternative would be based on its own, likely incomplete, list of alternatives. Two rounds of each robot broadcasting would be required: one to share all of the known alternatives, and one for the robots to cast their votes. Clearly, a better strategy is required.

Instead of trying to adapt the tools of centralized control (e.g. global broadcast communication) to dec-MRS, a decentralized algorithm for best-of-N decision-making should take advantage of decentralized structure. We now present just such an algorithm that social insects have evolved to solve this very problem.

A. A Social Insect Solution to the Best-of-N Decision Problem

Social insects live in decentrally organized colonies, similar in many ways to the idealized notion of a large dec-MRS. Some species periodically relocate their nests, and this operation represents a best-of-N decision, since a relocating colony is presented with the problem of selecting the single best site from those that its members are able to find. Recent research [20], [26], [28] has described an elegant emergent behavior that honeybees and *Temnothorax* ants use in order to solve this problem.

Once a colony determines that it must relocate, scouts search the surrounding environment for a new site. Each scout that finds a candidate site contributes an alternative to a decentralized list from which the colony must select only one. The individual insects measure the quality of candidate sites on an absolute scale, albeit noisily. When a scout returns to the colony, it recruits other insects to its site. The rate at which a scout recruits is determined by its opinion of its site's quality. The better it believes its favored site to be, the more frequently it will recruit. The recruits form their own opinions of site quality before returning to the colony to recruit even more insects to it. Over time, the differences in site quality will become apparent in the number insects recruiting to each one.

While visiting a candidate site, insects estimate its popularity by measuring the rate at which they encounter their nestmates. The recruiting insects periodically visit their favored sites, so the size of a site's recruiting population - and thus its popularity - can be inferred from the size of its visiting population. Once an insect determines that its site is popular enough, it commits to it, which alters its recruiting behavior. A committed individual recruits its teammates as rapidly as it can, inducing the rest of its colony to commit to its favored site, completing the collective best-of-N decision. The best site, able to attract the most vigorous recruitment, is the one most likely to induce commitment first, and thus will tend to be the one selected by the colony as a whole. Note that only local communication is required in this approach, and no direct comparisons of the candidate sites by individual robots need ever occur.

IV. Multiple-Robot Decision-Making from Social Insect Behavior

The nest-site selection behavior of the insects described in the last section forms the basis of our dec-MRS best-of-N decision-making algorithm. All inter-robot communication is local, which allows many one-on-one exchanges

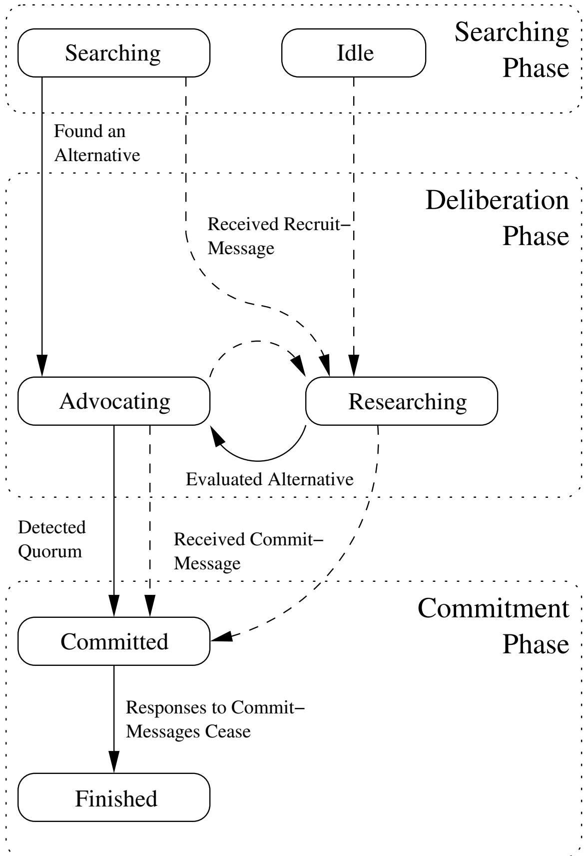

Fig. 1. This flowchart illustrates our proposed collective decision-making algorithm. Decisions are made in three phases, identified by the dotted boxes. The state transitions between the individual robot states are indicated by the arrows. Solid arrows indicate state transitions by a robot that are self-motivated, while dashed arrows correspond to those induced by messages received from teammates. Refer to Section IV for a detailed explanation of this behavior.

of information to occur in parallel. From the standpoint of network contention, using very short-range communication significantly reduces the chance of two robots interfering with each other [12]. In a dec-MRS, in which there is no central arbiter to coordinate the robots' transmissions, this is an important consideration. Its inclusion makes our algorithm ideally suited to the locally-interactive nature of dec-MRS.

Refer to Figure 1 for the following discussion. The individual robots' behaviors during the decision-making process are arranged into six states: *searching*, *idle*, *advocating*, *researching*, *committed*, and *finished*, organized into three phases: *searching*, *deliberation*, and *commitment*. Robots begin a decision in either the searching state or the idle state. Those in the searching state search for candidate alternatives, while the idle robots wait to be recruited into the process. When a robot finds an alternative, it determines its quality, and then enters the advocating state and rejoins its teammates.

It is the deliberation phase, composed of the advocating and researching states, that allows a dec-MRS to determine the best alternative from those that have been found. Advocating robots periodically send recruit-messages to their teammates, the frequency of which increases with their opinions of a favored alternative's quality. Robots in the searching, idle and advocating states are recruitable, and when such a robot receives a recruit-message, it enters the researching state. Here, it evaluates the quality of the specified alternative, and then enters the advocating state favoring it. Over time, the proportion of robots advocating in favor of the better alternatives will tend to increase, since it is more likely for robots to be recruited from poorer to better alternatives than *vice versa*.

In order to estimate the popularity of a favored alternative, advocating robots send vote-queries to the teammates that they encounter in between their sending of recruit-messages. The response to a vote-query is a vote-message, which is either "agree" or "disagree", indicating whether or not the queried robot favors the same alternative as the sender of the vote-query[1]. The robots use the proportion of the most recently received n votes that agree to estimate the popularity of their favored alternative. The specific manner in which this is done is described in Section IV-B.

Once a robot determines that its alternative is sufficiently popular, that it has satisfied a quorum, the robot enters the committed state. It is in this state that a robot begins the process of completing the group decision and from which unanimity emerges. Robots in the committed state send commit-messages to those that they encounter. When a robot receives such a message, it commits to the specified alternative. If it was not already committed to that alternative, it also responds with an acknowledgment. As more robots commit to the same alternative, the acknowledgment messages will become less frequent, disappearing altogether once complete consensus has been achieved. When committed robots no longer receive acknowledgments to their commit-messages they enter the finished state, completing the group decision and begin whatever created the need for a decision in the first place (e.g. deploy a solar array at a selected location).

A. Individual Knowledge, Direct Comparisons and Recruitment

Conspicuously absent from our proposed algorithm are direct comparisons of alternatives by individual robots. A recruit-message induces its recipient to advocate for the specified alternative regardless of the quality of any alternative previously favored. At first, it might seem as though group decision-making would be improved were the robots to ignore recruitments to alternatives of lesser quality, but this actually is not the case. It is important to understand that it is not the individual robots that make the best-of-N decision - it is their collected interactions that do so. Individual robots might make errors when evaluating the quality of an alternative, but the average of several robots' evaluations is less likely to be erroneous. The rate of recruitment to a specific alternative is based on the collected opinions of its advocating population, which reduces the impact of individual robot errors. If a robot were to overestimate an alternative's quality and were to stubbornly refuse to accept recruitment to another, this would

[1]Robots also respond to recruit-messages with vote-messages. This ensures that robots that recruit frequently are not deprived of votes due to their proportionally less frequent vote-queries.

amplify the effect of its error leading to potential stagnation. In a previous study, we showed that even in the presence of perfect sensing, direct comparisons of alternatives by the individual robots do not improve the performance of the group decision-making, and can even degrade the overall group decision-making in certain cases [23]. Of course, under our proposed model of quality-dependent stochastic recruitment, a robot might occasionally be recruited from a better alternative to a poorer one, but the reverse is far more likely.

The value of this sort of selfless behavior in a large decentralized system anecdotally is supported by the foraging strategies employed by different species of ant. Beckers *et al.* examined foraging in 98 different species of ant and found that the larger a species' average colony population size is, the less self-reliant the individual foragers tend to be [2]. Small-colony species commonly forage individually and rely on their own abilities to navigate between the colony and a food source, whereas members of species that live in colonies with larger populations rely more on emergent group navigation, such as pheromone trail-following. In a large decentralized system, it is more efficient to leverage the redundancy of the swarm, whereas smaller systems are better off taking advantage of the intelligence and unique individual skills of their members in the absence of numerical strength. The selfless iterative recruitment employed by our decision-making algorithm follows the example set by larger colony population insects, and so we believe that our algorithm will scale well to very large dec-MRS.

B. Quorum, its Evaluation, and the Probability of Commitment

In this section, we describe the manner in which the individual robots test quorum for an alternative that they favor and explain how our algorithm minimizes the likelihood of selecting a suboptimal one. It is important to remember that best-of-N decisions are made in our algorithm through two separate processes (refer to Figure 1). First, alternatives are compared in the deliberation phase through iterative, decentralized recruitment. Second, the system coalesces around the single alternative in the commitment phase that attracted the most substantial recruitment. The best solution is identified via a quorum test, and whichever alternative the quorum test identifies first will be the one chosen by the system as a whole.

Once quorum is detected and the commitment phase has begun, alternative quality no longer plays any role in the process. The commitment phase unanimously selects the alternative favored by the first robot to conclude that quorum has been met. Therefore, the algorithm's ability to select the best available alternative depends critically on the ability of the individual robots to accurately determine whether or not quorum has been satisfied. Specifically, we must minimize the chance of a robot concluding prematurely that its favored alternative has satisfied quorum. We refer to this phenomenon as a *false-positive quorum test*.

The individual robots test quorum by comparing the proportion of the n most recently received vote-messages that agree to the *quorum threshold*, $Q \in [0, 1]$. If the proportion of agreeing vote-messages is greater than or equal to Q, the test concludes that quorum has been met. Because the messages are anonymous, some robots' opinions might be over-represented in a particular set of n samples, introducing errors into the result of the test. Our goal in this section is to identify values for the quorum threshold, Q, and the sample size, n, such that the probability of a false-positive is reduced.

The first task is to determine the probability of a quorum test returning a positive result (including both false- and true-positives) in terms of Q and n and then minimize the conditional probability of a false-positive. The minimum number of agreeing votes that must be received to produce a positive result is $\lceil Q \cdot n \rceil$. When at least this many of the most recent n vote-messages are agreeing votes, the test will conclude that quorum has been met. If the MRS is *well-stirred*[2], then each vote received by a robot is equally likely to have come from any of its teammates. Assume that a MRS contains N robots, N_a of which are advocating for the same alternative. When one of these robots receives a vote-message, the probability of that vote agreeing is given by the Equation 1. C_a stands for *apparent consensus*, and we use this quantity since it is the consensus for a particular alternative that is apparent to a robot favoring it. The robots do not count their own opinions in their estimates of consensus, since they would need to know their system's population size in order to do so. As we explained earlier, it is unrealistic for individual robots in a dec-MRS to keep track of this quantity, nor is it necessary, since apparent consensus is very nearly equal to true consensus for all but the smallest systems.

$$C_a = \frac{N_a - 1}{N - 1} \tag{1}$$

Each vote can be viewed as a Bernoulli trial with a probability of being true equal to C_a. The probability of receiving i agreeing votes in a set of n therefore follows the binomial distribution, Equation 2:

$$P(i \text{ of } n \text{ votes are agreeing}) = \binom{n}{i} C_a^i (1 - C_a)^{n-i} \tag{2}$$

The overall probability of a positive quorum test is obtained by summing Equation 2 over all of the values of i such that a positive result will be returned: $i \in [\lceil Q \cdot n \rceil, n]$. This is given by Equation 3 which also is the probability that a robot will commit to its alternative, since commitment immediately follows a positive quorum test.

$$P_{commit}(C_a, n, Q) = \sum_{i=\lceil Q \cdot n \rceil}^{n} \left[\binom{n}{i} C_a^i (1 - C_a)^{n-i} \right] \tag{3}$$

To calculate the probability of a false-positive occurring, Equation 3 is summed over the values of apparent consensus less than quorum, $C_a \in [0, \lceil n \cdot Q \rceil - 1]$. Finally, Q and n are chosen to reduce this sum. Graphically, when P_{commit} is plotted as a function of apparent consensus, the probability of

[2] In a well-stirred system, a robot is equally likely to encounter any of its teammates and each encounter is an independent event.

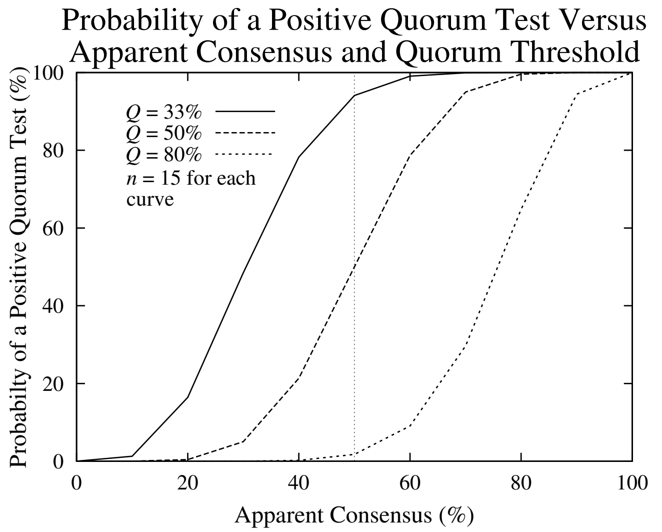

Fig. 2. This figure plots the probability of a quorum test concluding that quorum has been met versus the apparent consensus present in a system (refer to Section IV-B) for three values of quorum threshold, Q, with $n = 15$. As consensus increases, so does the probability of a positive test. To prevent premature commitment by the individual robots, the parameters n and Q should be chosen to reduce the area under the curve to the left of $C_a = 50\%$.

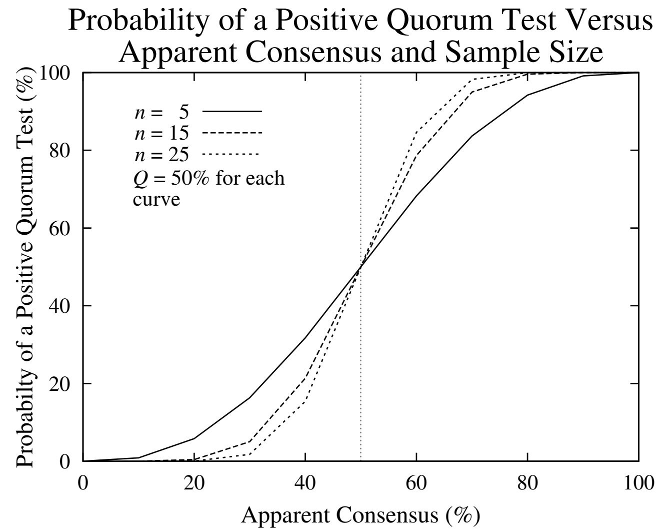

Fig. 3. This figure plots the probability of a quorum test concluding that quorum has been met versus the apparent consensus present in a system (refer to Section IV-B) for three sample sizes, n, with $Q = 50\%$. As consensus increases, so does the probability of a positive test. Increasing n makes the curve approach the shape of a step, decreasing the area under it to the left of Q, thus decreasing chance of a premature commitment.

a false-positive is represented by the area under the curve to the left of $C_a = Q$

Figure 2 plots P_{commit} versus apparent consensus for three values of Q. Changing Q does not reduce the likelihood of a false-positive, since it simply shifts the curve to the left or the right. The area under the curves to the left of $C_a = Q$ remains relatively constant.

Figure 3, illustrates the effect of changing n. As n is increased, the curve becomes more step-like, reducing the area to the left of $C_a = Q$ and thus the probability of a false-positive. However, a quorum test with a larger n will take longer to compute, since more votes will have to be received. This will slow down a group decision.

Ultimately, our goal is to prevent robots from committing until at least half of their teammates agree with them, as this guarantees that that only one site will trigger commitment. The quorum test can be tailored to this criterion by reducing the probability of a positive quorum test when $C_a < 50\%$. Returning to Figures 2 and 3, this corresponds to the area under the curves to the left of $C_a = 50\%$, while keeping n sufficiently small that C_a does not take too long to compute. The $Q = 80\%$, $n = 15$ curve in Figure 2 achieves this goal. By using these values for n and Q, the likelihood of an individual robot prematurely committing is greatly reduced, improving the overall accuracy of the group decision-making process. This result is independent of MRS population size, and the computational cost of quorum testing to the individual robots depends neither on the population size of their system, nor the number of alternatives over which a dec-MRS deliberates.

V. Experiments

The behavior of our decision-making algorithm was examined in a site-selection domain using real robots. The purpose of our experimentation was to demonstrate that unanimous best-of-N decision-making using our algorithm by a real dec-MRS is feasible despite noisy, real-world sensors. Our experimental environment (see Figure 4) was an enclosed arena that contained two sites located on opposite sides. Each site was represented by an illuminated spot on the floor, the brightness of which could be controlled. A colored beacon was located on the floor next to each site, which made them uniquely identifiable. In all of our experiments, the robots equated the quality of a site with the brightness of its overhead light. The brighter site was considered better than the dimmer one, and so the goal of the robots was to collectively choose the brighter site.

All of the experiments were carried out with an 11-robot MRS, each robot following the decision-making algorithm illustrated by Figure 1. At the beginning of an experimental trial, the robots began in either the idle or searching states. Those in the searching state wandered the environment looking for sites. Idle robots remained motionless, waiting to be recruited into the decision-making process. The better in quality a robot believed a favored site to be, the more frequently it would send recruit-messages to teammates. The frequency of recruit-messages increased linearly with perceived site quality, ranging from once per minute to once every 10 minutes. Quorum was tested by the robots using the method outlined in the last section, using the $n = 15$ most recently received vote-messages to estimate C_a. Quorum thresholds (Q) of 33%, 53%, and 80% (corresponding to 5, 8, and 12 agreeing votes out of 15) were implemented.

A photograph of one of our robots is given in Figure 5. Each robot carried an upward-pointing set of light sensors, as well as colored-light sensors on its front. The robots used the overhead light sensors to search for sights while wandering their environment. Once a robot found a site, it would move in the direction of the overhead light's gradient in order to

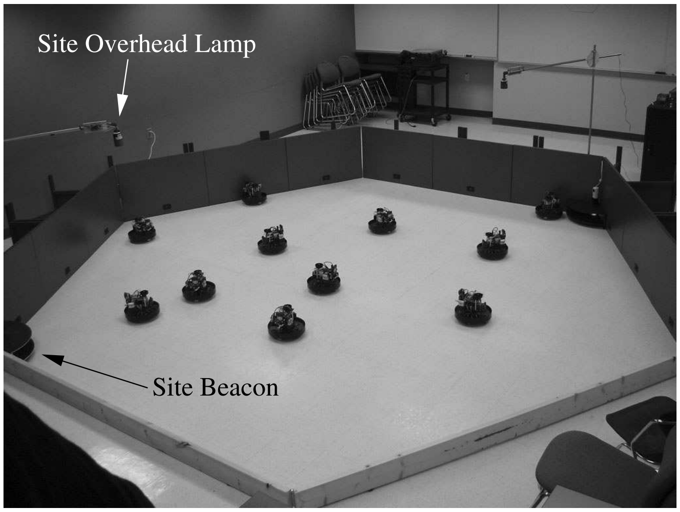

Fig. 4. Our experiments were carried out in an enclosed area containing two sites from which the robots had to select the best one. The sites were represented by illuminated spots on the ground with a colored beacon next to each one to permit unique identification. The quality of each site was determined by its brightness, which could be precisely controlled.

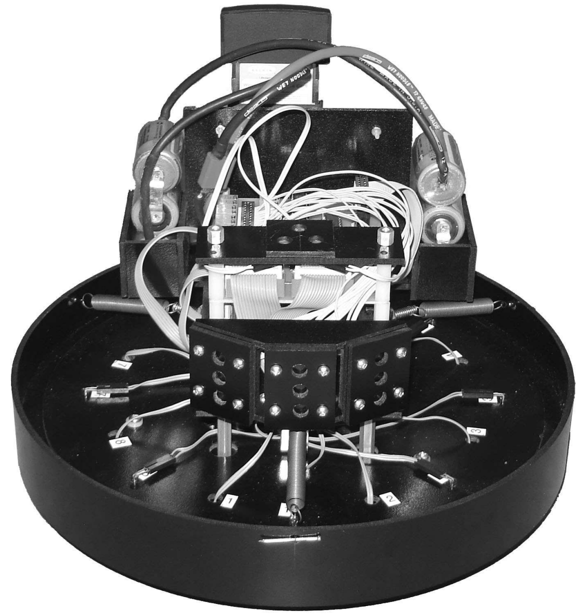

Fig. 5. This figure illustrates one of the robots used in this study. Each robot possessed overhead light sensors that allowed it to find candidate sites and measure their quality. Each robot also carried forward-facing colored light sensors that were used to search for the beacons that uniquely identified each site in the environment. Rising above the robot at the rear is its 802.11B wireless Ethernet interface, with which it communicated with its teammates.

position itself at the center of the site before measuring its quality. Next, the robot would rotate, scanning for colored beacons. The brightest beacon found during this rotation would be the one that it would associate with its found site.

When a robot recruited a teammate or requested a vote-message, it included the color of the beacon of its favored site in its message. Robots that received a vote-query responded positively if they favored the specified site and negatively otherwise. Robots would search for sites nearby the specified beacon while in the researching state. However, until the specified site had been found, a researching robot would continue to favor any site that it had favored prior to being recruited, and it would respond to vote-queries accordingly. A recruited robot would favor the site specified in the recruit-message only once it had successfully found it. If a researching robot was unable to find the site to which it had been recruited within three minutes, it would revert to whatever state it had been in prior to the reception of the recruit-message (searching for a site as a scout, advocating for a different site, or sitting idle). The robots' beacon sensors had a relatively short range, and they searched for a site in the researching state by following a random walk while scanning for the site's beacon. Less than 8% of recruited robots were unable to find a site once recruited, so the three minute time limit rarely was reached.

Once a robot determined that quorum had been satisfied, it would enter the committed state, in which it would instruct the teammates that it encountered to commit to its favored site. If the recipient of a commit-message was not already committed to the site specified in the commit-message, it would respond with an acknowledgment and then commit to the site. Committed robots entered the finished state (exiting the decision-making process) once they had been in the committed state for one minute without receiving an acknowledgment to a commit-message.

Our decision-making algorithm assumes that all of the inter-robot communication is very short range. Robots are meant to communicate only when they encounter a teammate as they wander about their environment. However, our robots possess only 802.11B wireless Ethernet interfaces, which have a global range in our experimental environment. In order to simulate local one-to-one communication, the robots were provided with a list of their teammates' IP addresses. When a robot encountered a teammate, it would select an address from its list at random and would send that robot a message via unicast. The recipient of such a message would respond via unicast back to its original sender. A simulation that we conducted with the Teambots [29] package showed that the communication behavior of our randomized approach is similar enough to that of true short-range peer-to-peer communication as to make any differences between them insignificant. This simulation also demonstrated that the well-stirred assumption is a valid one.

In order to examine its impact on the decision-making process, we also varied the number of robots that began each trial in the searching state. Any robots that did not begin a trial in the searching state began in the idle state and needed to receive a recruit- or commit-message in order to be brought into the decision-making process. Each experimental configuration was repeated with 4 and 11 robots participating in this initial search for sites.

The quality (brightness) of the two sites was set such that one tended to be perceived as better than the other, yet similar enough that the noise in the individual robots' site evaluations resulted in some overlap between the two. Each of the six experimental configurations (three quorum thresholds and two

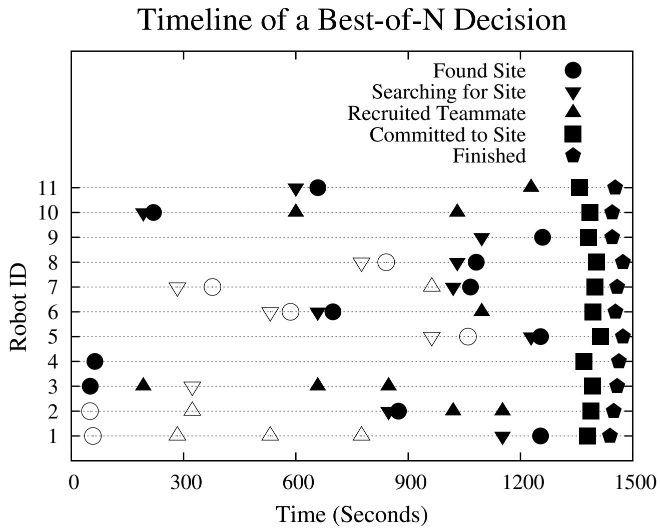

Fig. 6. This figure depicts a timeline of a single decision-making trial. Each horizontal sequence of symbols records the actions of a different robot in the decision. Solid shapes correspond to the better site, whereas hollow ones correspond to the poorer site. In this trial, the first four robots searched for sites, while the remaining seven began the decision in the idle state. Refer to the Section VI for a more complete explanation.

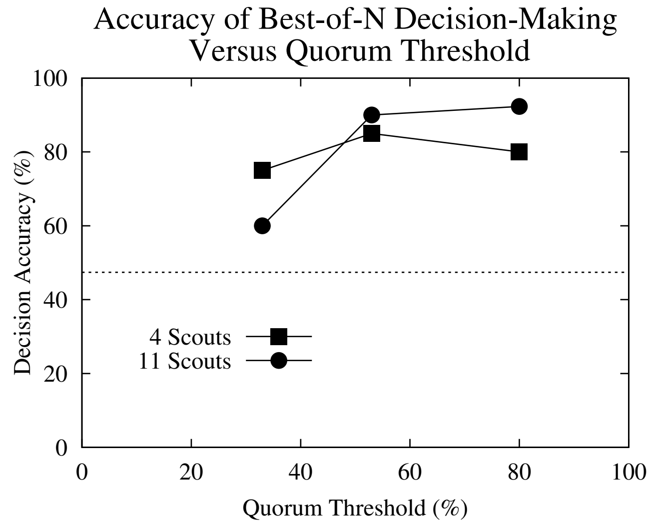

Fig. 7. This figure illustrates the accuracy of the proposed decision-making algorithm as a function of the quorum threshold used by the deliberating robots. A decision was considered to be accurate if the best site found by the scouts during the initial search was selected unanimously by the system by a trial's end. The horizontal dotted line is the accuracy of a system that favored all known sites equally. Using our decision-making algorithm, even when quorum was less than 50%, the robots were able to make accurate decisions when selecting amongst alternatives. Increasing quorum tended to increase the decision-making accuracy of the robots.

searching populations) was repeated at least 20 times, with the robots logging all of their actions and communications. Figure 6 graphically depicts a timeline of a typical decision-making trial, and a video of a trial can be viewed online[3].

A control experiment also was run in which the advocating robots recruited at a constant rate (every 90 seconds) regardless of their favored sites' measured qualities. Other than their constant, site-quality independent rate of recruitment, these robots behaved identically to those outlined above.

VI. Results

In this section, we present the results of our physical experiments, and discuss the effects of the parameters that were varied on the decision-making behavior of the robots. In the next section, we generalize these results and discuss our algorithm as a general-purpose best-of-N decision-making framework.

Figure 6 presents a timeline for one of our experimental trials. In this trial, two of the scouts find the good site (filled symbols) and two find the poor site (hollow symbols). Recruitment is indicated by triangles. When a robot recruits a teammate, the sending of a recruit-message is indicated by an upward pointing triangle, and the reception of this message is indicated by a downward pointing triangle in the recruited robot's timeline at the same time. For example, at approximately 200 seconds, robot-3 recruits robot-10 to the better site, and robot-10 robot finds the site shortly thereafter[4].

[3]http://www.cs.ualberta.ca/~parker/movies/decision.mov

[4]In this timeline, it appears as though robot-4 never attempts to recruit any of its teammates despite having found one of the sites. This is not actually the case. It sent a recruit-message to robot-11 at $t = 628$ seconds, and another at $t = 1237$ seconds to robot-1. However, both of these robots already favored the same site as robot-4, and so the messages were treated queries for votes, and thus they are not plotted on the timeline.

The better site induces the most frequent recruitment, and the population favoring it increases rapidly. Eventually, robot-11 determines that the apparent consensus for its favored site (the better one) has met quorum, and it commits. The rest of the system rapidly commits to the better site as commit-messages flood the MRS. Approximately one minute after the last commitment, the robots enter the finished state having unanimously selected the better of the two sites.

Every one of the decision-making trials in this work achieved unanimity. That is, every decision ended with all of the robots in the finished state favoring the same site. This result demonstrates that the commitment phase of the algorithm was effective. The sole purpose of this phase is to ensure unanimity regardless of which site is selected. Site quality plays no role in a decision once quorum is satisfied. It is in the preceding deliberation phase that good decisions are promoted over bad ones. The remainder of this section therefore is devoted to the performance of the deliberation phase in the experimental trials.

A. Decision Accuracy

A decision was considered accurate if the robots selected the best site that the scouts were able to find during the searching phase[5]. The relationship between decision accuracy and quorum for of each experimental configuration is presented in Figure 7. In all six of the experimental configurations investigated, the robots were able to collectively choose the better site much more frequently than they would be able to simply by chance (indicated by the performance of the control

[5]In other words, a trial that selected the poor site was not labeled inaccurate if the better site was never found.

trial). In general, increasing the quorum threshold of a decision tended to increase its accuracy, as earlier work suggests [21]. However, decision accuracy appears to decrease for the 4-scout system when quorum was increased from 53% to 80%. Why?

The drop in performance of the 4-scout system is within experimental error, as it corresponds to results of just one best-of-N decision out of the 20 that were run. However, the 11-scout system *should* tend to outperform the 4-scout system when quorum is higher and *vice versa*. Our experimental environment contained only two candidate sites. If every one of the scouts in the 11-scout system found a site, then quorum would have been satisfied without any recruitment when quorum was 33%, since C_a would have been at least 50% for one of the two sites. It is random which of the sites this would have been, which explains the decreased decision-making accuracy of this configuration. Because the robots did not find sites at the same time, there still was time for the deliberation phase to promote the better site somewhat, which explains why the 11-scout, 33% quorum system performed better than the control trial's performance suggests it should have. It is impossible for the searching phase alone to satisfy quorum when only four of the robots scouted, making deliberation necessary for commitment to occur. We therefore should expect this system to out-perform the 11-scout system when quorum is low, which it did.

When quorum is higher, necessitating deliberation for both scouting populations, the 11-scout system is the more accurate of the two. This is due to the relative impact of stochastic effects in the early stages of the deliberation phase. A hypothetical example best illustrates why. Consider a 4-scout system in which three of the scouts find the poorer site and only one finds the better site. The chance of this occurring is 25%, so it is not unlikely. If one of the poorer-site favoring robots recruited the lone better-site favoring robot before it was able to recruit, the better site would be forgotten by the dec-MRS, since from a system-level point of view, only those sites that are favored by at least one of its robots are remembered. This outcome would be classified as an inaccurate decision, since the best site found would not ultimately be selected. Increasing the scouting population increases the likelihood that more than one robot will find each of the sites during the initial search, making this stochastic effect less likely to have an impact. In general, a greater scouting population will tend to improve the accuracy of decision making as long as the initial search is unlikely to satisfy quorum on its own. This conclusion agrees with the data plotted in Figure 7.

B. Observed Quorum

Consensus cannot be observed directly by the individual robots. Instead, they can only estimate its value by sampling their anonymous teammates' opinions through the reception of vote-messages. We use the term *observed quorum* to denote the actual consensus present in a dec-MRS at the time of commitment, and this is plotted against the quorum threshold (the threshold to which the robots compare their estimates of apparent consensus when testing quorum) in Figure 8. As one would expect, the observed quorum increases with the quorum threshold, but also note that the former consistently is greater than the latter.

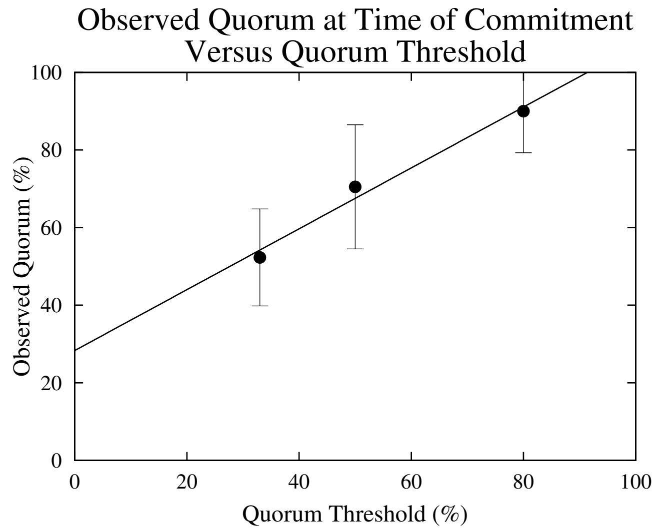

Fig. 8. Observed quorum is the actual consensus present in a dec-MRS when the commitment phase begins. This figure illustrates that the observed quorum clearly depends on the quorum threshold (Q) used by the robots in their quorum tests. The observed quorum consistently is greater than Q, which is to be expected given the conservative nature of the test as described in Section VI-B.

Observed quorum is a measure of true consensus ($\frac{N_a}{N}$), so it is strictly greater than apparent consensus ($\frac{N_a-1}{N-1}$), but this difference is only a few percent in an 11-robot system. It does not account for the degree to which the observed quorum exceeds the quorum threshold in our experimental results.

In Section IV-B, we explained the trade-off between the speed and accuracy of the individual robots' quorum tests. Measuring apparent consensus takes time. During the time required by the quorum test to measure it, the robots' deliberation simultaneously is changing the proportion of the dec-MRS that favors each of the known alternatives. As a result, the estimates of apparent consensus made by each robot are averages of C_a during the period in which the vote-messages were received.

The consensus in favor of the best alternative will tend to increase over time, so the estimates of C_a for it will be underestimates. Because the best alternative is the most likely to trigger commitment, we should expect the observed quorum in a dec-MRS at the time of commitment to exceed the quorum threshold, which is what we observe in Figure 8.

The quorum test outlined in Section IV-B will be conservative in nature, tending to underestimate consensus in a best-of-N decision. This is a good property, since the role of the quorum test is to indicate reliably when quorum has been met, and not to do so prematurely. The data of Figure 8 demonstrates that our algorithm's quorum test meets this goal.

C. Time Required for Deliberation

The accuracy of a decision increases with quorum, and there is an inherent tradeoff between speed and accuracy when making group decisions [9]. This is evident in the duration of the deliberation phase of the trials. Figure 9 plots the length

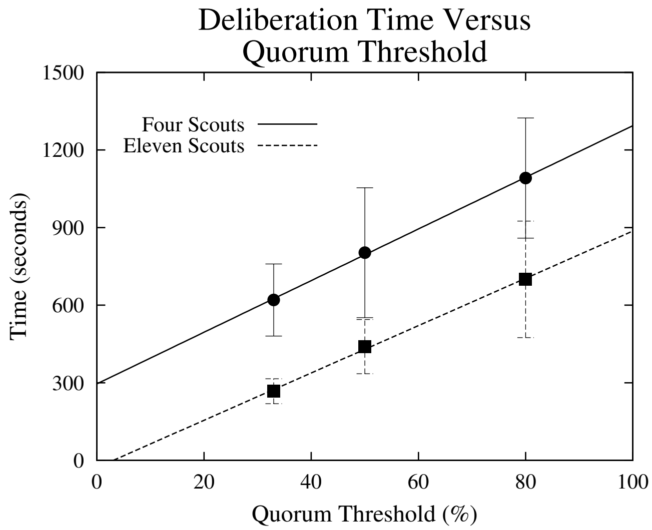

Fig. 9. The length of the deliberation phase, the phase in which good decisions are promoted over bad ones, increases with quorum. This occurs because more recruitment is necessary in order to build up an advocating population large enough to satisfy quorum. Increasing the population that executes the initial search for sites decreases the deliberation period because fewer robots need to be recruited into the process.

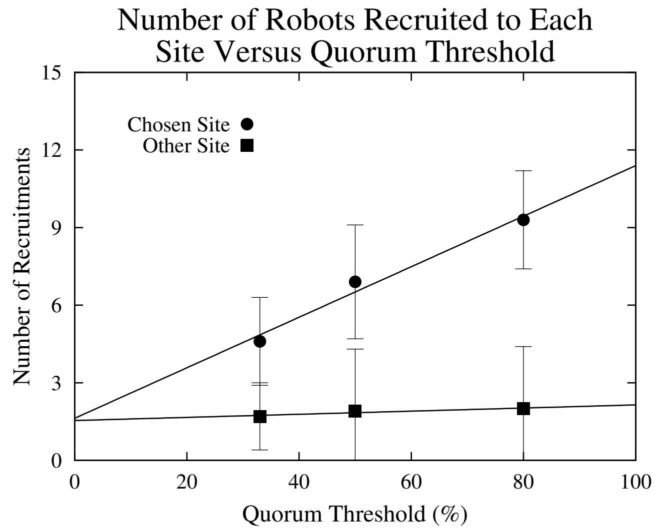

Fig. 10. The deliberation phase of our emergent decision-making process is able to focus its effort on the site that ultimately is chosen by the dec-MRS. As quorum increases, more attention is paid to the better site through increased recruitment, but recruitment to the site not chosen is minimal and relatively constant.

of the deliberation phase versus quorum threshold. Both the 4- and 11-scout systems required more time to deliberate as the quorum threshold was increased. This occurred because more recruitment was necessary to build sufficient support in order for quorum to be satisfied. From the system-level point of view, we can say that the group was more careful in its comparisons of the known sites as quorum was increased. The cost of this increased accuracy is the increased deliberation time.

The incremental cost of increasing quorum, indicated by the slopes of the two regression lines, is very nearly equal for both scouting populations. Increasing the scouting population speeds up the process because fewer robots need to be recruited into the decision, which accounts for the 11-scout data being less than that of the 4-scout system. Refer to our earlier discussion of decision accuracy for a more detailed explanation of the effect of scouting population size on decentralized best-of-N decision-making.

D. The Focus of Deliberation

The ideal decision-making algorithm would waste as little effort as possible considering alternatives that ultimately will not be selected. Sites are considered by our algorithm through deliberation. Better ones induce recruitment more often than those less likely to be selected, and so the collective deliberation tends to focus on them. In practice some time must be spent considering each alternative, and so some recruitment should be expected towards each one found during the initial searching phase, but our algorithm keeps this to a minimum, illustrated by Figure 10. The site ultimately chosen by the system is responsible for the majority of the recruitment carried out and this increases with quorum. Extra effort is not wasted on the other site. Our dec-MRS best-of-N decision-making algorithm is able to discard lesser alternatives quickly and focus its attention on the best one known.

VII. Discussion

The experiments described in this work all were conducted in an environment containing exactly two alternatives. Earlier simulated work has demonstrated that our algorithm performs equally well when more than two alternatives are available [23]. Ultimately, a dec-MRS can make a best-of-N decision and compare as many alternatives as there are scouts that participate in the initial searching phase. Because all of the operations are decentralized, our algorithm could be utilized by very large dec-MRS without encountering computational or communicative bottlenecks that a more centralized approach would.

In the introduction, we stated that algorithms for use by dec-MRS should be tolerant of individual robots' faults, as these should be expected if low-cost mass-produced robots are used to compose dec-MRS. Individual robots can make errors by either over- or under-valuing candidate alternatives. In both of these cases, our algorithm should be self-correcting, just like the original insects' behavior upon which our approach is based [10]. If an individual under-values an alternative, it will recruit slowly, but those robots that it does recruit are likely to make more accurate assessments of its quality and recruit more rapidly. What if an alternative is over-valued? In this case, the robot that made the error would recruit others more frequently than it should. Again, those that it recruited would be tend to recruit at a rate more in line with the alternative's true quality, and so they would be unlikely to further contribute to the error. Nonetheless, the initial erroneous individual would continue to recruit too frequently. However, robots favoring a better site also would be recruiting quickly, and their recruits would do tend to do so as well. After only a generation or two, the cumulative effect of recruitment to better sites would tend to render the erroneous individual's counter-productive recruitment insignificant, and it likely would be recruited to favor a better site as well. This behavior is visible in Figure 6. Robot-1 finds the poorer site, overestimates its quality, and quickly recruits robots 3 and 6. These robots, however, recruit slowly, and robot-1 eventually is recruited to the better site by robot-2, correcting the error. Noise in the individual robots' sensing will make the overall group decision somewhat noisy, and so individuals' errors might negatively impact the outcome of a given decision. However, as the population size of a dec-MRS increases, the relative contributions of each robot to the deliberation phase decrease. This tends to reduce the impact of individual robots' errors on a decision, increasing the system-level reliability of the process.

VIII. Conclusions and Future Work

The contribution of this work is a best-of-N decision-making algorithm for use by dec-MRS. Central to this algorithm is the use of quorum to link two emergent behaviors. The first of these, a period of iterative deliberation compares a set of alternatives and encourages the best one to attract the most support. The detection of quorum by one of the robots triggers the second collective behavior, in which the alternative that satisfied quorum is adopted by every member of the system. By preceding the deliberative behavior with a search for alternatives, an intelligent decentralized best-of-N decision-making algorithm is produced. The manner in which quorum is detected was calibrated to reduce the chance of adopting an alternative before the support for it exceeds 50%, increasing the decision-making accuracy. Because this algorithm utilizes only local interaction and communication, and because the costs of iterative recruitment and quorum testing to the individual robots are independent of MRS population size, our proposed algorithm is well-suited to very large dec-MRS. An empirical study demonstrated the performance of our algorithm in a series of physical experiments in which the robots were able to select the best alternative in their environment at least 80% of the time when the recommended value was used for the quorum threshold. This is good performance given the intractable nature of the decentralized best-of-N problem and the noisy individual sensing of our robots. Unanimity was achieved by every trial.

Because robots are recruited asynchronously into decisions by their teammates, decisions could be initiated by individual robots as they identified the need for one. These robots would serve as the initial scouts, and the commitment phase of the algorithm would synchronize a dec-MRS around the outcome of the decision. We would like to investigate the use of our algorithm in this way to enable a dec-MRS to react cohesively to a changing environment. Furthermore, because this work presents an empirical study, only a limited portion of our algorithm's parameter space could be investigated. We intend to expand our analysis and application of iterative recruitment, consensus estimation, and quorum testing to dec-MRS in future works, as well as compare the performance of our algorithm to other decision-making strategies.

References

[1] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayi rci. Wireless sensor networks: A survey. *Computer Networks*, 38:393–422, 2002.

[2] R. Beckers, S. Goss, J. L. Deneubourg, and J. M. Pasteels. Colony size, communication and ant foraging strategy. *Psyche*, 96:239–256, 1989.

[3] R. Beckers, O. E. Holland, and J.L. Deneubourg. From local actinos to global tasks: Stigmergy and collective robotics. In R. Brooks and P. Maes, editors, *Artificial Life IV*, pages 181–189. MIT Press, 1994.

[4] Daniel S. Bernstein, Shlomo Zilberstein, and Neil Immerman. The complexity of decentralized control of markov decision processes. In *UAI '00: Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence*, pages 32–37, San Francisco, CA, USA, 2000. Morgan Kaufmann Publishers Inc.

[5] Michael Bowling, Brett Browning, and Manuela Veloso. Plays as effective multiagent plans enabling opponent-adaptive play selection. In *Proceedings of the Fourteenth International Conference on Automa ted Planning and Scheduling*, pages 376–383, 2004.

[6] Benjamin H. Brown Jr., J. Michael Vande Weghe, Curt A. Bereton, and Pradeep K. Khosla. Millibot trains for enhanced mobility. *IEEE/ASME Transactions on Mechatronics*, 7(4):452–461, December 2002.

[7] Jason D. Campbell, Padmanabhan Pillai, and Seth Copen Goldstein. The robot is the tether: Active, adaptive power routing for modular robots with unary inter-robot connectors. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2005)*, pages 4108–15, Edmonton, Alberta Canada, August 2005.

[8] Y. Uny Cao, Alex S. Fukunaga, and Andrew B. Kahng. Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots*, 4:1–23, 1997.

[9] Larissa Conradt and Timothy J. Roper. Consensus decision making in animals. *TRENDS in Ecology and Evolution*, 20(8):449–456, 2005.

[10] Nigel R. Franks, Stephen C. Pratt, B. Mallon, Eamonn, Nickolas F. Britton, and David J. T. Sumpter. Information flow, opinion polling and collective intelligence in house-hunting social insects. *Philosophical Transactions of the Royal Society of London Series B*, 357:1567–1583, 2002.

[11] Dani Goldberg and Maja J. Matarić. Interference as a tool for designing and evaluating multi-robot controllers. In *Proceedings of AAAI 1997*, pages 637–642, 1997.

[12] Piyush Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, 2000.

[13] Owen Holland and Chris Melhuish. Stigmergy, self-organisation, and sorting in collective robotics. *Journal of Adaptive Behaviour*, 5(2):173–202, 1999.

[14] Jianyang Huang, Shane M. Farritor, Ala' Qadi, and Steve Goddard. Localization and follow-the-leader control of a heterogeneous group of mobile robots. *IEEE/ASME Transactions on Mechatronics*, 11(2):205–215, 2006.

[15] Auke Jan Ijspeert, Alcherio Martinoli, and Aude Billard. Collaboration through the exploitation of local interactions in autonomous collective robotics: The stick pulling experiment. *Autonomous Robots*, 11(2):149–171, 2001.

[16] Natallia Katenka, Elizaveta Levina, and George Michailidis. Local vote decision fusion for target detection in wireless sensor networks. *IEEE Transactions on Signal Processing*, 56(1):329–338, 2008.

[17] Jelle R. Kok, Matthijs T. J. Spaan, and Nikos Vlassis. Multi-robot decision making using coordination graphs. In *Proceedings of the 11th International Conference on Advanced Robotics*, pages 1124–1129, June 2003.

[18] Jelle R. Kok and Nikos Vlassis. Distributed decision making of robotic agents. In *Proceedings of the 8th Annual Conference of the Advanced School for Computing and Imaging*, pages 318–325, 2003.

[19] C. Ronald Kube and Eric Bonabeau. Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, 30(1/2):85–101, 2000.

[20] E. B. Mallon, S. C. Pratt, and N. R. Franks. Individual and collective decision-making during nest site selection by the ant *Leptothorax albipennis*. *Behavioral Ecology and Sociobiology*, 50:352–359, 2001.

[21] Chris A. C. Parker and Hong Zhang. Biologically inspired decision making for collective robotic systems. In *Proceedings of the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2004)*, pages 375–380, September 2004.

[22] Chris A. C. Parker and Hong Zhang. Collective decision making: A biologically inspired approach to making up all of your minds. In *Proceedings of the 2004 IEEE International Conference on Robotics and Biomimetics (ROBIO 2004)*, August 2004.

[23] Chris A. C. Parker and Hong Zhang. Passive versus active expression of preference in the control of multiple-robot decision-making. In *Proceedings of the 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2005)*, pages 915–920, 2005.

[24] Chris A. C. Parker and Hong Zhang. Collective robotic site preparation. *Adaptive Behavior*, 14(1):5–20, 2006.

[25] Chris A. C. Parker and Hong Zhang. Consensus-based task sequencing in decentralized multiple-robot systems using local communication. In *Proceedings of the 2008 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2008)*, pages 1421–1426, 2008.

[26] Stephen C. Pratt. Quorum sensing by encounter rates in the ant *Temnothorax albipennis*. *Behavioral Ecology*, 16(2):488–496, 2005.

[27] D. V. Pynadath and M. Tambe. The communicative multiagent team decision problem: Analyzing teamwork theories and models. *Journal of Artificial Intelligence Research*, 16:389–423, 2002.

[28] Thomas D. Seeley and P. Kirk Visscher. Group decision making in nest-site selection by honey bees. *Apidologie*, 35:101–116, 2004.

[29] Teambots home page. http://www.teambots.org, April 2005.

[30] Jens Wawerla, Gaurav S. Sukhatme, and Maja J. Matarić. Collective construction with multiple robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2002)*, pages 2696–2701, 2002.

[31] Jan Wessnitzer and Chris Melhuish. Collective decision-making and behaviour transitions in distri buted ad hoc wireless networks of mobile robots: Target hunting. In W. *et al.* Banzhaf, editor, *European Conference on Artificial Life (ECAL) 2003*, pages 893–902. Springer-Verlag, 2003.

[32] Erfu Yang and Dongbing Gu. Nonlinear formation-keeping and mooring control of multiple-autonomous underwater vehicles. *IEEE Transactions on Mechatronics*, 12(2):164–178, 2007.

Chris A. C. Parker received his B.Eng. (Electrical Engineering) from McMaster University in Hamilton (ON), Canada and his M.Sc. in Computing Science from the University of Alberta in Edmonton (AB), Canada. He currently is a Ph.D. Candidate at the University of Alberta, where he studies collective intelligence in multiple-robot systems. In particular, his research focuses on the application of social insect behavior to the cooperative decision-making problem. His other research interests include the development of novel robot actuators and the connection between local interaction and group behavior.

Hong Zhang received his B.Sc. from Northeastern University, Boston, USA and his Ph.D. from Purdue University in West Lafayette (IN), USA, both in Electrical and Computer Engineering. He currently is a Professor in the Department of Computing Science, University of Alberta, and Director of the Center for Intelligent Mining Systems. His research interests include robotics, computer vision, image processing, and intelligent systems.