A Distributed Consensus Model for House-Hunting in *Temnothorax* Ant Colonies

Emily Y. Zhang Jiajia Zhao Nancy Lynch

Massachusetts Institute of Technology

Abstract

We study the problem of house-hunting in ant colonies, where ants reach consensus on a new nest and relocate their colony to that nest, from a distributed computing perspective. We propose a house-hunting algorithm that is biologically inspired by *Temnothorax* ants. Each ant is modelled as a probabilistic agent with limited power, and there is no central control governing the ants. We show a $\Omega(\log n)$ lower bound on the running time of our proposed house-hunting algorithm, where n is the number of ants. Further, we show a matching upper bound of expected $O(\log n)$ rounds for environments with only one candidate nest for the ants to move to. Our work provides insights into the house-hunting process, giving a perspective on how environmental factors such as nest qualities or a quorum rule can affect the emigration process.

1 Introduction

Recently, there has been an interest in the distributed computing community on studying biologically-inspired algorithms [1]. Tissues found within the human body and insect colonies of ants and bees are good examples of naturally occurring systems where there are many agents with limited power, a global goal, and no central control. Interestingly, an ant colony as a whole exhibits a high level of collective intelligence and is able to achieve global goals, such as relocating to new nests. It is puzzling how the distributed system is able to quickly reach consensus through local communications, especially given the high noise levels observed in nature.

The house-hunting process in *Temnothorax* ant colonies is a naturally occurring algorithmic task that is closely related to consensus, a fundamental problem in distributed computing theory. The goal of the ants is to relocate the colony of ants to a new nest with superior quality. During the house-hunting process, a colony is able to reach consensus on a new nest and execute the move of the entire colony, even though each individual actively-scouting ant has information about only a small subset of the new candidate nests.

In 2015, Ghaffari et al. [3] modeled the ant colony house-hunting process as a distributed algorithm on independent random agents. They also showed theoretical guarantees on the number of rounds required for various house-hunting algorithms under their model to converge. Recently, Zhao et al. [13] developed a simulator that closely mimics the ants' behaviors on both individual and colony levels. The simulator is based on the agent-based model of ants' house-hunting process that Pratt et al. [10] created by studying videotaped behavior of ants. Zhao et al. showed that their simulator is biologically plausible in that it accurately reflects many behaviors observed in real ant colonies, and their simulator is also useful for predicting some of the behaviors of ants that are harder for biologists to directly study in experiments. The algorithm presented by Zhao et al.

is more biologically plausible than those studied by Ghaffari et al., but one remaining challenge is that there are no theoretical bounds on the convergence speed of the algorithm by Zhao et al.

We present an algorithm for the house-hunting emigration process that is structured like the model considered by Ghaffari et al. [3] and additionally takes into account the more diverse behavior of ants in different phases and states observed by Pratt et al. [10] in ant colonies and modelled by Zhao et al. [13] in computer simulations. As a result, our algorithm is more biologically plausible while still tractable to rigorous analysis. Our model is a mathematical agent-based model. We show that the theoretical guarantees on the running time of our algorithm are similar to those of the algorithms considered by Ghaffari et al.

Our work has many implications for both the biology community and the computer science community. Natural algorithms have evolved over time to have many advantageous properties. For example, algorithmic tasks carried out by collections of living beings are usually highly adaptive to different types of environments, robust to noise, and also optimized in terms of their speed and accuracy. Thus, insights from these biological algorithms can inspire more robust, efficient algorithms for distributed computer systems, such as robot swarms [5]. Additionally, using mathematical tools to analyze the house-hunting algorithm can allow for a better understanding of the qualities of ant colonies that are harder for biologists to directly observe, such as the dependence on various environmental parameters.

1.1 The House-Hunting Process

Temnothorax ants often search for and move to new nests, as living in favorable nests is important to the survival of their colony. Their moving process is highly distributed, as each individual ant has limited information and communication, and there is no central control governing the emigration process.

Ant colonies are typically composed of active and passive ants [10]. Active ants execute the emigration while passive ants, such as brood items or inactive adult ants, are transported to new nests by active ants.

Biologists have observed that the house-hunting process involves several stages. Active ants search for nests, assess nests, recruit other ants, and also transport other ants. Once an active ant has found a new nest of satisfactory quality, it moves on to the recruitment phase, where it recruits other active ants to the new nest via *tandem runs* [6, 11]. Should the population of active ants in a new nest surpass a *quorum threshold*, then active ants can commit to the new nest and begin *transporting* (i.e. picking up and carrying) other ants from the old nest to the new nest [9]. These transports speed up emigration to the new nest.

1.2 Main Results and Organization

Our main results are a biologically plausible house-hunting algorithm that is tractable to analysis, a lower bound on the number of rounds that the algorithm takes to converge with high probability, and an expected upper bound on the number of rounds required in single-nest emigrations (when the environment contains the original home nest and one new nest). We incorporate biological insights from [10] and [13] into our house-hunting algorithm, which is presented in Section 2.

In Section Section 3, we show a lower bound on the run-time for our model. In [3], Ghaffari et al. showed a lower bound of $\Omega(\log n)$ on the number of rounds required for any house-hunting algorithm under their model of house-hunting to converge with high probability. In Section 3, we explain a subtle difference between our algorithm and the model considered by Ghaffari et al. While this difference affects the lower bound proof, we show that the lower bound of $\Omega(\log n)$ holds for our model as well.

In Section 4, we consider the special case where there are two nests in the environment. We show that the expected running time for our house-hunting algorithm is $O(\log n)$ rounds assuming that the quorum threshold falls within a certain range. By our lower bound result, this upper bound is tight up to a constant factor. It would be interesting to extend this expected upper bound result to environments with more nests in future work.

Finally, in Section 5, we discuss possible modifications to the algorithm and directions for future research.

2 Model

We present a model of *Temnothorax* ants' house-hunting process that is both tractable to analysis and biologically plausible. This algorithm is primarily inspired by the agent-based model for house-hunting in ant colonies introduced in [13]. Like the model in [13], our algorithm has many parameters that can be tuned to reflect the changing environmental conditions and varied behaviors of ants observed in nature. Our model differs in that we reduce the number of internal variables stored for each ant and the number of possible states that the ants can be in, thus simplifying the rules for how ants change locations. These simplifications make the model tractable for proofs of theoretical guarantees.

In our algorithm, active and passive ants in the colony play different roles in the emigration process. Active ants transition through many states, including searching for a new nest, evaluating a nest, and recruiting other ants to the nest by lead forward or transport runs. Passive ants, on the other hand, change location only when transported by an active ant. The biological insights for these design decisions come from [10].

2.1 Framework

The environment consists of at least two nests, one of which is the original home of the colony. Each nest has an associated *quality*, which is a nonnegative real number. The ants are modeled as identical finite state machines that execute computations synchronously in discrete rounds. In each round, an active ant performs at most one call to each of the functions **select_action**(), **select_ant**(), and **transition**(), which are defined in Section 2.2.

We let n denote the total number of ants, and we let n_a and n_p denote the number of active ants and passive ants, respectively, with $n_a = \theta(n)$, $n_p = \theta(n)$, and $n_a + n_p = n$. The location of an ant a is denoted $a.location$, which is one of the nests in the environment. Every active ant a has an associated state, denoted $a.state$, which is one of 9 possible states: At Nest_i, Search_i, Quorum Sensing, Lead Forward, and Transport, for $i \in \{E, C, T\}$. The subscripts E, C, and T stand for Exploration, Canvassing, and Transport, three different phases of active ants described in [10]. For every ant a, the value of $a.state$ begins as At Nest_{E} and the value of $a.location$ begins as the original home nest of the colony. These values are updated by calls to the helper function **transition**().

Active ants can probabilistically select one of two possible actions to take: *advance* or *hold* (see Figure 1). The action that an ant takes determines how its state and location are changed. There are two instances where an action involves two ants. We say that an ant is *committed* to the nest that it is in if it is in the At Nest_{T}, Search_{T}, or Transport state. An ant that is committed to the nest that it is in can *transport* other ants to that nest by advancing from the Transport state. Similarly, an ant can recruit other active ants to the nest that it is in via a *tandem run* by advancing from the Lead Forward state. Transports and tandem runs can fail to move a second ant if that second ant has already transitioned on that round.

Like the model from [13], our model is parameterized by many adjustable constants. The parameters μ_q and μ_p are the *quality coefficient* and *population coefficient*, respectively. They represent the relative weight that ants give to the quality and population of a nest when evaluating that nest. We denote by θ the *quorum threshold*, or the fraction of all active ants that must be in a nest before an active ant can commit to that nest and begin transporting ants to that nest. If the fraction of active ants in the original home nest ever drops below the quorum threshold, we say that that nest *drops out of competition*, and no active ant can commit to that nest after that point.

The constants c_s, c_f, c_ℓ, and c_t denote the *search constant*, the *follow constant*, the *lead forward constant*, and the *transport constant*. These constants parameterize the probability that the corresponding type of action succeeds. Finally, λ controls for how noisy individual ants' decision-making is, with higher λ values corresponding to lower individual noise level. All constants other than λ range between 0 and 1; λ ranges from 1 to 16 [13]. See Table 1 for an example of how these values are set for experiments performed in [13].

Parameter	Value	Source
quality coefficient μ_q	0.25	trial-and-error from [13]
population coefficient μ_p	0.35	trial-and-error from [13]
quorum threshold θ	0.15	[9, 2]
search constant c_s	0.025	trial-and-error from [13]
follow constant c_f	0.4	[4, 8]
lead forward constant c_ℓ	0.6	trial-and-error from [13]
transport constant c_t	0.7	[10]
λ	8	trial-and-error from [13]

Table 1: An example of a set of parameter values that are biologically plausible.

2.2 Helper Functions

This subsection defines the helper functions that are called in the house-hunting algorithm, which is given in Section 2.3.

select_action(a): The input is an ant a. Let n' be a nest chosen uniformly at random from all of the nests other than *a.location*. Let q be the quality of *a.location*, and let p and p_a be the number of ants and active ants in that nest, respectively. Finally, let q' and p' denote the quality and population of nest n'. The ant probabilistically chooses an action u, sampled from a Bernoulli random variable $\mathbf{u} \in \{\text{advance}, \text{hold}\}$ with parameters that depend on *a.state* as shown in (1). The

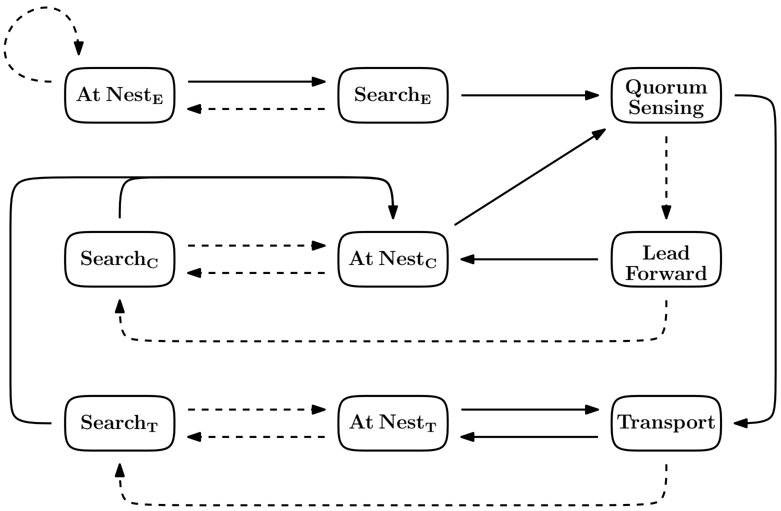

Figure 1: The State Transition Diagram. The solid arrows denote the ant choosing to *advance*; the dashed arrows denote the ant choosing to *hold*. The transition probabilities are given in (1).

function **select_action(a)** returns (u, n').

$$
\begin{aligned}
&\Pr[\mathbf{u} = \text{advance} \mid a.state = \text{At Nest}_E] = 1 - \left(1 + e^{-\lambda\left(\mu_q \cdot q + \mu_p \cdot \frac{p}{n}\right)}\right)^{-1}\\
&\Pr[\mathbf{u} = \text{advance} \mid a.state = \text{At Nest}_i] = \left(1 + e^{-\lambda\left(\mu_q \cdot q + \mu_p \cdot \frac{p}{n}\right)}\right)^{-1} \quad \text{for } i \in \{\text{C, T}\}\\
&\Pr[\mathbf{u} = \text{advance} \mid a.state = \text{Search}_i] = c_s \cdot \left(1 + e^{-\lambda\left(\mu_q \cdot (q'-q) + \mu_p \cdot \frac{p'-p}{n}\right)}\right)^{-1} \quad \text{for } i \in \{\text{E, C, T}\}\\
&\Pr[\mathbf{u} = \text{advance} \mid a.state = \text{Quorum Sensing}] = \begin{cases} 1 & \text{if quorum has been met, that is } p_a > \theta \cdot n_a \\ & \text{and } a.location \text{ has not dropped out of competition} \\ 0 & \text{otherwise} \end{cases}\\
&\Pr[\mathbf{u} = \text{advance} \mid a.state = \text{Transport}] = c_t\\
&\Pr[\mathbf{u} = \text{advance} \mid a.state = \text{Lead Forward}] = \begin{cases} c_\ell & \text{if } q' > q \\ 0 & \text{otherwise} \end{cases}
\end{aligned}
\tag{1}
$$

$\mathbf{select\_ant(a, n', action)}$: This function takes as input an ant a, a nest n', and an action *action*. If *action* is hold, then this function immediately returns *null*. First, we set a' to null. If $a.state$ is Lead Forward and there is at least one active ant in nest n', then let a' be an active ant chosen uniformly at random from n'. If $a.state$ is Transport and there is at least one passive ant in nest n', then let a' be a passive ant chosen uniformly at random from n'; if there are only active ants in n', let a' be an active ant chosen uniformly at random from n'. We return a' with probability c_f and *null* otherwise. Note that the returned ant cannot be a because a is not in nest n'.

$\mathbf{transition(a, a', n', action)}$: This function takes as input ants a and a', a nest n', and an action *action*. If $a.state$ is Search$_i$ for any $i \in \{E, C, T\}$ and *action* is advance, then we set $a.location$ to n'. If $a.state$ is Lead Forward or Transport and *action* is advance, then we set $a'.location$ to $a.location$. We set $a.state$ to the state obtained by starting from $a.state$ and following the arrow corresponding to *action* in the state transition diagram (Figure 1). If a' is not null, we set $a'.state$ to At Nest$_E$.

2.3 Algorithm

With the helper functions defined, we are ready to present the house-hunting algorithm. Let P be a permutation of all of the active ants chosen uniformly at random for every round.

Algorithm 1: One Round of the HouseHunting Algorithm

```
$M$: a set of ants, initially $\emptyset$
for $i = 1$ to $|P|$ do
    if $a_{P(i)} \notin M$ then
        $action, n' :=$ select_action($a_{P(i)}$)
        $a' :=$ select_ant($a_{P(i)}, n', action$)
        if $a' \in M$ then
            $a' \leftarrow null$
        transition($a_{P(i)}, a', n', action$)
        $M := M \cup \{a_{P(i)}\} \cup \{a'\}$
```

Algorithm 1 shows one round of the house-hunting algorithm. We repeat the procedure given in Algorithm 1 until the algorithm *converges*. The house-hunting algorithm converges when all of the passive ants have moved from the original home nest to the *winning nest*, a nest different than the original home nest and has better quality. We note that not all of the active ants are required to be in the winning nest when the algorithm converges since the active ant population is more mobile.

In Algorithm 1, the set M serves to make sure that each ant transitions at most once: either actively by initiating an action or passively by getting recruited. If there are conflicts among the actions involving two ants (i.e. transports or tandem runs), then P serves as a tie breaker between the conflicting actions. In short, for every active ant a_i that has not been recruited via a tandem run or transported by the time that all ants preceding a_i in the permutation have transitioned in a given round, the ant a_i will probabilistically select an action to take using the **select_action** function and select an accompanying ant a' using the **select_ant** function. If a_i is not successfully transporting or leading forward, then a' will be set to *null*. Finally, the **transition** helper function updates the state and location of ants a_i and a' to reflect the changes caused by the action that ant a_i executed.

3 A Lower Bound

In this section, we asymptotically lower bound the number of rounds required for a colony of ants to complete the house-hunting process with high probability. We use ideas from the lower bounds on spreading a rumor in a graph from [7] and proof methods from [3]. We also use terminology from [3]: an *informed* ant knows the id of the winning nest, and an *ignorant* ant does not know.

Ghaffari et al. [3] proved an asymptotic lower bound on the convergence rate of their algorithm that grew logarithmically in the number of ants. There is a subtle difference between our algorithm and that of Ghaffari et al. that affects how the proof of the lower bound proceeds. In the model used by Ghaffari et al., if an active ant located in a nest Nest 1 is trying to perform a transport to another nest Nest 2, then she randomly selects any other ant in Nest 1 to transport. In their model, the probability that any passive ant gets transported on any given round can be upper bounded. That is, if a minority of the ants in Nest 1 are active ants trying to transport, then the probability that any given passive ant in Nest 1 gets transported is small. On the other hand, if the majority of the ants in Nest 1 are trying to transport, then they have a significant chance of transporting each other, rather than a passive ant in Nest 1.

In contrast, in our algorithm, each active ant has a bit more computation power in a single round. An active ant located in Nest 2 can move to Nest 1 and transport an ant from Nest 1 to Nest 2 in one round. Thus, the probability that a given passive ant in Nest 1 gets transported on any given round cannot be upper bounded: when there are many active ants in Nest 2 trying to transport ants from Nest 1 to Nest 2 and few passive ants in Nest 1, a passive ant in Nest 1 gets transported with high probability.

Even though we are not able to lower bound the probability that any ignorant ant remains ignorant in any given round of the simulation, we are able to lower bound this quantity for a fraction of the passive ants, as shown in the following lemma. Our proof makes use of the quantity $m = \frac{n_a}{n_p}$; we note that m is independent of n because $n_a = \Theta(n)$ and $n_p = \Theta(n)$.

Lemma 3.1. *Let $m = \frac{n_a}{n_p}$. If $n_p \geq 4$, then for any given round where there are at least $\frac{n_p}{2}$ ignorant passive ants in the original home nest, an ignorant passive ant in the original home nest remains ignorant at the end of that round with probability at least*

$$\min\left(\frac{1}{4}, \frac{1}{16^m}\right).$$

Proof. Let $B \geq \frac{n_p}{2} \geq 2$ be the number of ignorant passive ants in the original home nest at the beginning of a given round. The only way for an ignorant passive ant to become informed is for that ant to be transported to the best nest by an active ant. There are at most n_a active ants transporting passive ants from the home nest to the best nest on any given round. While there are passive ants in the original home nest, an active ant transporting from that nest chooses one of those passive ants to transport uniformly at random. Thus, the probability that any given passive ant does not get transported is at least

$$\left(\frac{B-1}{B}\right)^{n_a}.$$

We define $f(x) := \left(\frac{x-1}{x}\right)^x$. Note that the function $f(x)$ is monotonically increasing when $x > 1$. We consider two cases and lower bound the probability that a passive ant in the original home nest remains ignorant after a round in both cases.

Case 1: $B \geq n_a$.

In this case, we have

$$\begin{aligned}\left(\frac{B-1}{B}\right)^{n_a} &\geq f(B) \\ &\geq f(2) \\ &= \frac{1}{4}.\end{aligned}$$

Case 2: $B < n_a$.

In this case, we have

$$\begin{aligned}\left(\frac{B-1}{B}\right)^{n_a} &\geq \left(\frac{\frac{n_p}{2}-1}{\frac{n_p}{2}}\right)^{m\cdot n_p}\\ &= f\left(\frac{n_p}{2}\right)^{2m}\\ &\geq f(2)^{2m}\\ &= \frac{1}{16^m}.\end{aligned}$$

Putting the two cases together finishes the proof of the lemma. ☐

In [3], Ghaffari et al. proved a lower bound of $\Omega(\log n)$ on the running time of their house-hunting algorithm. Their proof depended on a lemma that stated that all ignorant ants stay ignorant with constant probability during each round; we showed in Lemma 3.1 that this result is true for $\frac{n_p}{2}$ of the passive ants. This discrepancy by a constant factor does not affect the asymptotic behavior of the algorithm. More formally, Lemma 3.1, together with the proof of Theorem 3.2 from [3], imply the following theorem:

Theorem 3.2. *If $n_p \geq 4$, then for any constant $c > 0$, our proposed house-hunting algorithm requires $\Omega(\log n)$ rounds for all of the passive ants to move from the original home nest to the winning nest with probability at least $\frac{1}{n^c}$.*

Proof. Let $\rho = \min\left(\frac{1}{4}, \frac{1}{16^c}\right)$, which is a constant independent of n. Let S be the set of the first $\frac{n_p}{2}$ passive ants to become informed. We will investigate the number of rounds required for the ants in S to become informed with high probability. By Lemma 3.1, with probability at least ρ, an ant in S that is ignorant at the beginning of round r remains ignorant at the end of round r.

For every ant $a \in S$ and every round r, let Y_r^a be a Bernoulli random variable with $\Pr[Y_a^r = 1] = \rho^r$. Let $Y_r = \sum_{a\in S} Y_r^a$. We define random variable S_r to be the number of ants in S that are ignorant after r rounds. Note that for every $a \in S$ and every round r, $\Pr[Y_a^r = 1] = \rho^r$ lower bounds the probability that ant a is ignorant after r rounds. Thus, we have

$$\Pr\left[S_r < x\right] \leq \Pr\left[Y_r < x\right] \text{ for any } x \leq n. \tag{2}$$

For $r = \frac{1}{2}\log_{\rho^{-1}}\left(\frac{n_p^2}{4n}\right) - \log_{\rho^{-1}}(8c)$, we have $\mathbb{E}[Y_r] = |S|\rho^r = 8c\sqrt{n}$, and

$$\begin{aligned}\Pr\left[S_r < \frac{\mathbb{E}[Y_r]}{2}\right] &\leq \Pr\left[Y_r < \frac{\mathbb{E}[Y_r]}{2}\right] && \text{(Using (2))}\\ &\leq e^{-\frac{\mathbb{E}[Y_r]}{8}} && \text{(Chernoff bound)}\\ &= e^{-c\sqrt{n}}\\ &< \frac{1}{n^c}. && \left(e^{-\sqrt{x}} < \frac{1}{x} \text{ for } x > 0\right)\end{aligned}$$

Therefore, with probability at least $1-\frac{1}{n^c}$, at least $4c\sqrt{n}$ passive ants are ignorant after $\frac{1}{2}\log_{\rho^{-1}}\left(\frac{n_p^2}{4n}\right) - \log_{\rho^{-1}}(8c) = \Theta(\log n)$ rounds. Thus, the number of rounds required for all of the passive ants to move to the winning nest with probability at least $\frac{1}{n^c}$ is $\Omega(\log n)$. ☐

4 An Expected Upper Bound for Single Nest Emigration

In this section, we consider an environment with only two nests: Nest 0 (with quality q_0), and Nest 1 (with quality $q_1 > q_0$). All of the ants are in Nest 0 in the At_Nest$_{\mathrm{E}}$ state at the beginning of Round 1, and we investigate how long it takes for the house-hunting algorithm to converge as the number of ants n varies.

Since an active ant will change locations if it advances from a Search state, the frequency at which ants are in Search states is of interest to us.

Definition 4.1. For every integer $k > 0$, random variable $R_0^{(k)}$ is 1 if there are no active ants in Nest 0 at the beginning of round k; otherwise, $R_0^{(k)}$ is the fraction of active ants in Nest 0 that are in one of the Search states (Search$_{\mathrm{E}}$, Search$_{\mathrm{C}}$, or Search$_{\mathrm{T}}$) at the beginning of round k. Similarly, for every integer $k > 0$, random variable $R_1^{(k)}$ is 0 if there are no active ants in Nest 1 at the beginning of round k; otherwise, $R_1^{(k)}$ is the fraction of active ants in Nest 1 that are in one of the Search states at the beginning of round k. Let $f_0 := \min_{k>1} \mathbb{E}\left[R_0^{(k)}\right]$ and $f_1 := \max_k \mathbb{E}\left[R_1^{(k)}\right]$.

We do not explicitly compute f_0 and f_1 since this would require intensive computations. Instead, we make the following useful observation:

Observation 4.2. We have $f_0, f_1 > \epsilon$ for some constant $\epsilon > 0$ that is independent of n.

Proof. Since all ants start out in the At_Nest$_{\mathrm{E}}$ state in Nest 0, we have $R_0^{(1)} = 0$. For $k > 1$, $\Pr\left[R_0^{(k)} > 0\right] > 0$ since there is the possibility that an active ant in Nest 0 stays in the At_Nest$_{\mathrm{E}}$ state until it moves to the Search$_{\mathrm{E}}$ state during the $(k-1)^{\mathrm{th}}$ round. Thus, $\mathbb{E}\left[R_0^{(k)}\right] > 0$ for $k > 1$, so $f_0 > 0$.

If $f_1 = 0$, then that means that either no active ants ever enter Nest 1 or the expected fraction of active ants in Nest 1 that are in a Search state is 0 for every round. Both of these statements are clearly not true, so we have $f_1 > 0$ by contradiction.

Because the state transition diagram has an aperiodic recurrent class, the states are associated with steady-state probabilities. The states Search$_{\mathrm{C}}$ and Search$_{\mathrm{T}}$ are recurrent states, so their steady-state probabilities are positive. The values $\mathbb{E}\left[R_i^{(k)}\right]$ for $i \in \{0, 1\}$ only depend on the transition probabilities of the model and k, and they converge to the corresponding steady-state probabilities, which also only depend on the transition probabilities. All of the transition probabilities of the model can be upper and lower bounded by positive constants independent of n. Thus, f_0 and f_1 can also be lower bounded by a constant independent of n. ☐

Recall from Section 2 that we say that the algorithm converges when all of the passive ants are in the winning nest. In this section, we prove the following main theorem, which says that the algorithm converges in expected $O(\log n)$ rounds in an environment with two nests, given some restrictions on the parameters.

Theorem 4.3. *We consider an environment with two nests: Nest 0 (with quality q_0), and Nest 1 (with quality $q_1 > q_0$). All n ants start in Nest 0 in the At_Nest$_{E}$ state.*

We define function $a : \mathbb{R} \to \mathbb{R}$ as follows:

$$a(x) = n_a \left(\frac{\frac{f_0}{f_1} - x}{\frac{f_0}{f_1} + e^{\alpha}} \right),$$

where $\alpha = -\lambda(\mu_q(q_1 - q_0) - \mu_p)$.

Let $\epsilon > 0$ *be any small constant less than* $\frac{f_0}{f_1}$, *and let random variable* R_ϵ *denote the number of rounds required for at least* $a(\epsilon)$ *of the active ants and all of the passive ants to move from Nest 0 to Nest 1.*

If the quorum threshold satisfies $1 - \frac{a(\epsilon)}{n_a} < \theta < \frac{a(\epsilon)}{n_a}$, *then* $\mathbb{E}[R_\epsilon] = O(\log n)$.

We note that a condition of Theorem 4.3 is that the quorum threshold must fall within a certain range. In particular, we will see in Proposition 4.9 that the lower bound on the quorum threshold guarantees that backward transports from Nest 1 to Nest 0 cease to happen after $O(\log n)$ rounds. The upper bound on the quorum threshold simply guarantees that quorum will be met at Nest 1, allowing for transports to that nest. When we plug in the biologically plausible values $\lambda = 8$, $\mu_q = .25$, $\mu_p = .35$, $q_1 = 2$, $q_0 = 0$, let $\epsilon = .00001$, and use the estimate $\frac{f_0}{f_1} = 1$, we find that the condition becomes $\theta \in (.2315, .7685)$. If we instead plug in $q_1 = 3$, then the condition becomes $\theta \in (.0392, .9608)$. We thus see that the bounds on the quorum threshold required by the condition of the theorem are reasonable and not overly restrictive.

We prove Theorem 4.3 by separately examining the emigration of active and passive ants. We examine active ant emigration in Section 4.1, and we examine passive ant emigration in Section 4.2.

4.1 Active Ant Emigration

In this subsection, we focus on the emigration of active ants. We use the quantities f_0 and f_1 to study the expected number of ants that move between the nests during each round.

Lemma 4.4. *Suppose that there are* x *active ants in Nest 0 at the beginning of a given round* $r > 1$. *The expected number of ants that move from Nest 0 to Nest 1 during that round is at least*

$$\frac{c_s \cdot f_0 \cdot x}{1 + e^{-\lambda(\mu_q(q_1 - q_0) - \mu_p)}}.$$

Proof. We can see that active ants change location only when they advance from a Search state or get recruited to a new nest by another active ant via a tandem run. In our model, a tandem run can only cause an active ant to be recruited from an inferior nest to a nest with better quality. Thus, tandem runs will only speed up the rate at which ants move from Nest 0 to Nest 1, so we can disregard the population changes caused by tandem runs without loss of generality and focus on ants that change locations by advancing from the Search state.

We denote by p_0 and p_1 the population in Nest 0 and Nest 1 at the beginning of this round, respectively. The probability p that an active ant in a Search state at Nest 0 advances, thus moving to Nest 1, is

$$\begin{aligned} p &:= \frac{c_s}{1 + e^{-\lambda(\mu_q(q_1 - q_0) + \mu_p(p_1 - p_0)/n)}} \\ &\geq \min_{p_0, p_1 \in \{0,1,\ldots,n\}^2} \frac{c_s}{1 + e^{-\lambda(\mu_q(q_1 - q_0) + \mu_p(p_1 - p_0)/n)}} \\ &= p' := \frac{c_s}{1 + e^{-\lambda(\mu_q(q_1 - q_0) - \mu_p)}}, \end{aligned}$$

where the last line follows because p is minimized when $p_1 - p_0$ is minimized, which is achieved when $p_0 = n$ and $p_1 = 0$.

The expected number of ants in a Search state at Nest 0 at the beginning of this round is at least $f_0 \cdot x$. Thus, the expected number of ants that go from Nest 0 to Nest 1 during this round is at least $f_0 \cdot x \cdot p'$. $\square$

We can obtain the following lemma using the same reasoning that we used in the proof of Lemma 4.4.

Lemma 4.5. *Suppose that there are x active ants in Nest 0 at the beginning of a given round. The expected number of ants that move from Nest 1 to Nest 0 during that round is at most*

$$\frac{c_s \cdot f_1 \cdot (n_a - x)}{1 + e^{-\lambda(\mu_q(q_0 - q_1) + \mu_p)}}.$$

Proof. As in the proof of Lemma 4.4, we only focus on ants that move nests by advancing from a Search state.

We denote by p_0 and p_1 the population in Nest 0 and Nest 1 at the beginning of this round, respectively. We upper bound the probability q that an active ant in a Search state at Nest 1 advances, thus moving to Nest 0:

$$\begin{aligned} q &:= \frac{c_s}{1 + e^{-\lambda(\mu_q(q_0 - q_1) + \mu_p(p_0 - p_1)/n)}} \\ &\le \max_{p_0, p_1 \in \{0,1,\ldots,n\}^2} \frac{c_s}{1 + e^{-\lambda(\mu_q(q_0 - q_1) + \mu_p(p_0 - p_1)/n)}} \\ &= q' := \frac{c_s}{1 + e^{-\lambda(\mu_q(q_0 - q_1) + \mu_p)}}, \end{aligned}$$

where the last line follows because q is maximized when $p_0 - p_1$ is maximized, which is achieved when $p_0 = n$ and $p_1 = 0$.

The expected number of ants in a Search state at Nest 1 at the beginning of this round is at most $f_1 \cdot (n_a - x)$. Thus, the expected number of ants that go from Nest 1 to Nest 0 during this round is at most $f_1 \cdot (n_a - x) \cdot q'$. $\square$

Now, we put Lemma 4.4 and Lemma 4.5 together to show that the number of rounds required for a constant fraction of the active ant population to move to Nest 1 is independent of the size of the ant colony.

Lemma 4.6. *Let function $a(\cdot)$ and constant α be defined as in Theorem 4.3. Let random variable A_r denote the number of active ants in Nest 1 after r rounds. There exists a constant c, independent of n, such that we have $\mathbb{E}[A_c] \ge a(\delta)$ for any constant $\delta \in \left(0, \frac{f_0}{f_1}\right)$.*

Proof. Let random variable Y_x denote the number of ants that Nest 1 gains during a single round $r > 1$ that begins with x active ants in Nest 0. Putting Lemma 4.4 and Lemma 4.5 together, we have

$$\begin{aligned} \mathbb{E}[Y_x] &\ge \frac{c_s \cdot f_0 \cdot x}{1 + e^{\alpha}} - \frac{c_s \cdot f_1 \cdot (n_a - x)}{1 + e^{-\alpha}} \\ &= \Theta\left(\frac{f_0}{f_1} x \left(1 + e^{-\alpha}\right) - (n_a - x)\left(1 + e^{\alpha}\right)\right) \\ &= \Theta\left(x\left(\frac{f_0}{f_1} \cdot e^{-\alpha} + 1\right)\left(1 + e^{\alpha}\right) - n_a\left(1 + e^{\alpha}\right)\right) \\ &= \Theta\left(x\left(\frac{f_0}{f_1} \cdot e^{-\alpha} + 1\right) - n_a\right). \end{aligned}$$

If $x \geq n_a - a(\delta) = n_a \left(\frac{e^\alpha + \delta}{\frac{f_0}{f_1} + e^\alpha} \right)$, we have

$$\begin{aligned}
\mathbb{E}[Y_x] &= \Omega \left(n_a \left(\frac{e^\alpha + \delta}{\frac{f_0}{f_1} + e^\alpha} \right) \left(\frac{f_0}{f_1} \cdot e^{-\alpha} + 1 \right) - n_a \right) \\
&= \Omega \left(n_a \left(\frac{\delta}{e^\alpha} \right) \right) \\
&= \Omega(n).
\end{aligned}$$

Thus, while there are at least $n_a - a(\delta)$ active ants in Nest 0, Nest 0 loses $\Omega(n)$ active ants and Nest 1 gains $\Omega(n)$ active ants in expectation during every round. As a result, the expected number of ants in Nest 0 drops down to at most $n_a - a(\delta)$ in a constant number of rounds. ☐

Now that we have shown that the number of rounds required for an expected $a(\epsilon)$ active ants to move to Nest 1 is independent of the size of the ant colony, we will use that result to determine the expected number of rounds required for $a(\epsilon)$ active ants to move to Nest 1. In particular, now our random variable is the number of rounds, rather than the number of ants in Nest 1.

Proposition 4.7. *Let function $a(\cdot)$ be defined as in Theorem 4.3. Let random variable R_ϵ denote the number of rounds required for $a(\epsilon)$ of the active ants to move from Nest 0 to Nest 1. We have that $\mathbb{E}[R_\epsilon] = O(1)$.*

Proof. Let $0 < \delta < \epsilon < \frac{f_0}{f_1}$. Let c be the number of rounds required for an expected $a(\delta)$ of the active ants to move to Nest 1. By Lemma 4.6, c is a constant independent of n.

We denote by the random variable A the number of active ants that Nest 1 gains in c rounds. By definition, $\mathbb{E}[A] = a(\delta) > a(\epsilon)$. Using the fact that $A \leq n$, we upper bound $\Pr[A \leq a(\epsilon)]$:

$$\begin{aligned}
&\mathbb{E}[A] \leq \Pr[A \leq a(\epsilon)]\, a(\epsilon) + (1 - \Pr[A \leq a(\epsilon)])\, n \\
\implies &\Pr[A \leq a(\epsilon)] \leq \frac{n - a(\delta)}{n - a(\epsilon)}.
\end{aligned}$$

For any integer $r \geq 1$, we have that $\Pr[R_\epsilon = r]$ is upper bounded by

$$\Pr[R_\epsilon \geq r] \leq \Pr[A \leq a(\epsilon)]^{\lfloor \frac{r}{c} \rfloor} \leq \left(\frac{n - a(\delta)}{n - a(\epsilon)} \right)^{\lfloor \frac{r}{c} \rfloor}$$

because in order for $R_\epsilon \geq r$, Nest 1 must have gained less than $a(\epsilon)$ ants every consecutive c rounds since the start of the algorithm.

We thus have

$$\begin{aligned}
\mathbb{E}[R_\epsilon] &= \sum_{r=1}^{\infty} \Pr[R_\epsilon = r] \cdot r \\
&\leq \sum_{r=1}^{\infty} \left(\frac{n - a(\delta)}{n - a(\epsilon)} \right)^{\lfloor \frac{r}{c} \rfloor} \cdot r \\
&= \Theta(1).
\end{aligned}$$

The last equality follows because $\frac{n - a(\delta)}{n - a(\epsilon)}$ is less than 1 and has no dependency on n since $a(\delta), a(\epsilon) = \Theta(n_a) = \Theta(n)$, so the factors of n in the numerator and denominator cancel out. ☐

4.2 Passive Ant Emigration

Passive ants only change location when they are transported by active ants. Thus, we first examine the changes in the number of transports as the algorithm proceeds. We start by proving a lemma which will be useful in converting bounds on expected number of ants to bounds on expected number of rounds.

Lemma 4.8. *Consider n ants that probabilistically decide to leave on discrete rounds and never come back. Let random variable R denote the number of rounds that it takes for all of the ants to leave. Let $0 < c \leq 1$ be some constant independent of n. If the expected number of ants that leave during any round is at least $c \times$ [the number of ants remaining at the beginning of that round], then $\mathbb{E}[R] = O(\log n)$.*

Proof. Let $r^* = \log_{1-c} \frac{1}{2n}$. After r^* rounds, the expected number of ants remaining is at most $n(1-c)^{r^*} = \frac{1}{2}$.

Let random variable A denote the number of ants remaining after r^* rounds. By Markov's Inequality,

$$\Pr[A \geq 1] \leq \frac{1}{2}.$$

Let random variable R_1 denote the number of rounds that it takes for all but at most 1 ant to leave. For any integer $t \geq 0$, we have that

$$\begin{aligned}
\Pr\left[t \cdot r^* \leq R_1 \leq (t+1) \cdot r^*\right] &\leq \Pr\left[R_1 \geq t \cdot r^*\right] \\
&\leq \Pr\left[A \geq 1\right]^t \\
&\leq \left(\frac{1}{2}\right)^t.
\end{aligned}$$

We use this inequality to bound $\mathbb{E}[R_1]$.

$$\begin{aligned}
\mathbb{E}[R_1] &= \sum_{r=1}^{\infty} \Pr[R_1 = r] \cdot r \\
&\leq \sum_{t=0}^{\infty} \Pr\left[t \cdot r^* \leq R_1 \leq (t+1)r^*\right] \cdot (t+1)r^* \\
&\leq r^* \sum_{t=0}^{\infty} \left(\frac{1}{2}\right)^t \cdot (t+1) \\
&= \Theta(\log n).
\end{aligned}$$

Now, we investigate $R - R_1$, that is, the number of additional rounds it takes for all of the ants to leave after all but at most 1 ant already left.

$$\begin{aligned}
\mathbb{E}[R - R_1] &\leq \sum_{r=1}^{\infty} \Pr[R - R_1 \geq r] \cdot r \\
&\leq \sum_{r=1}^{\infty} (1-c)^{r-1} \cdot r \\
&= \Theta(1).
\end{aligned}$$

Thus, $\mathbb{E}[R] = \mathbb{E}[R_1] + \mathbb{E}[R - R_1] = O(\log n) + O(1) = O(\log n)$. □

Proposition 4.9. *Let $a(\cdot)$ be defined as in Theorem 4.3. Let random variable R denote the number of rounds that the algorithm runs for before no more ants perform transports from Nest 1 to Nest 0 for the rest of the algorithm. If the quorum threshold satisfies $\theta > 1 - \frac{a(\epsilon)}{n_a}$, then $\mathbb{E}[R] = O(\log n)$.*

Proof. By Proposition 4.7, after an expected constant number of rounds, there will be at most $n_a - a(\epsilon)$ active ants in Nest 0, which is below the number of active ants needed for the quorum threshold to be reached in Nest 0. After the fraction of active ants that are in Nest 0 drops below the quorum threshold θ, Nest 0 drops out of competition, and no more active ants will commit to Nest 0 (i.e. advance from the Quorum Sensing state while in Nest 0). After this point, the number of ants committed to Nest 0 cannot increase.

By the same reasoning from Lemma 4.4, a constant fraction of the ants committed to Nest 0 are expected to move to Nest 1 (thus becoming no longer committed to Nest 0) on any given round. Hence, in expectation, the number of ants committed to Nest 0 decreases by a constant factor on each round after the number of active ants in Nest 0 drops below quorum. Using Lemma 4.8, we can conclude that there will be no more ants committed to Nest 0 after expected $O(\log n)$ rounds, which means that there will be no more ants that can perform transports from Nest 1 to Nest 0. ☐

By Proposition 4.9, we have that the effect of transports from Nest 1 to Nest 0 is negligible after expected $O(\log n)$ rounds. We can now bound the expected number of rounds required for all of the passive ants to be moved to Nest 1.

Proposition 4.10. *Let random variable R denote the number of rounds required for all of the passive ants to be transported to Nest 1. If the quorum threshold satisfies $1 - \frac{a(\epsilon)}{n_a} < \theta < \frac{a(\epsilon)}{n_a}$, we have $\mathbb{E}[R] = O(\log n)$.*

Proof. We consider any given round of the algorithm; we let a_1 denote the number of active ants transporting from Nest 0 to Nest 1 at the beginning of that round, and we let p_0 denote the number of passive ants in Nest 0 at the beginning of that round. We assume that $a_1 \geq \theta \cdot n_a$ and break into two cases depending on how a_1 and p_0 compare to each other.

Case 1: $a_1 \leq \frac{1}{2}p_0$.
Each active ant in Nest 1 that is transporting (i.e. advancing from the Transport state) has at least a 50% chance of successfully moving a passive ant to Nest 1 on this round. This is because less than half of the passive ants in Nest 0 could have already been transported by a different active ant in that round. Thus, in this case, the expected number of passive ants that get transported to Nest 1 in this round is at least $\frac{1}{2}\mathbb{E}[a_1]$ in this case. When the fraction of active ants that are in Nest 1 exceeds the quorum threshold, we have $\frac{1}{2}\mathbb{E}[a_1] = \Theta(n_a) = \Theta(n) = \Omega(p_0)$.

Case 2: $a_1 > \frac{1}{2}p_0$.
The probability that a given passive ant does not get transported is at most

$$\begin{aligned}\left(\frac{p_0 - 1}{p_0}\right)^{a_1} &< \left(\frac{p_0 - 1}{p_0}\right)^{\frac{1}{2}p_0} \\ &\leq \left(e^{-\frac{1}{p_0}}\right)^{\frac{1}{2}p_0} \\ &= e^{-\frac{1}{2}}.\end{aligned}$$

Since every passive ant in Nest 0 has at least a constant probability of getting transported to Nest 1 in this case, the expected number of ants that get transported to Nest 1 in this round is $\Omega(p_0)$.

By Proposition 4.7, at least $a(\epsilon) > \theta \cdot n_a$ active ants will be in Nest 1 after an expected constant number of rounds, and by Proposition 4.9, there will be no more ants performing backward transports from Nest 1 to Nest 0 after $O(\log n)$ rounds. After that, the expected number of passive ants that are transported from Nest 0 to Nest 1 during each round will be $\Omega(p_0)$ in both Case 1 and Case 2. Using Lemma 4.8, we can conclude that all passive ants will be in Nest 1 after expected $O(\log n)$ rounds. □

Finally, putting Proposition 4.7 and Proposition 4.10 together proves Theorem 4.3.

5 Discussion and Future Work

We note that Theorem 4.3 involves a lower bound on the quorum threshold. Without such a lower bound on the quorum threshold, we would not be able to show Proposition 4.9, and the emigration progress may be delayed due to backward transports from the superior nest to the inferior nest. There is a lot of work in the biology community studying the role of the quorum threshold in the house-hunting process. As we see with our result, mathematical analyses such as ours can provide possible explanations for why ant colonies have evolved to use a quorum threshold.

For the sake of simplicity and for our analysis, we made many simplifications to the model from [13], on which our algorithm is based. Though our model consists of fewer states in the state transition diagram than the model in [13], our analysis is robust to changes to the state transition diagram, as long as all of the transition probabilities can be bounded by constants independent of the number of ants and the ways that ants change location are unaffected. The other key difference is that in the model from [13], each ant consists of more internal variables (in addition to the state and location of the ant, which are the two internal variables of ants that our model uses). For example, in addition to location, the home nest, candidate nest, and previous candidate nest are also stored for each ant. Because of this, the rules for location changes in the model in [13] are also more fine-grained, making analysis more complicated. It would be interesting to analyze a house-hunting algorithm that is more similar to the one given in [13] as the algorithm in [13] has been shown to be biologically plausible by comparison with data from biologists [10].

In our model, recruitments via tandem runs can only happen from inferior nests to superior nests, and the original home nest drops out of competition once the fraction of active ants in that nest drops below the quorum threshold. One potential direction for future research is to analyze our house-hunting algorithm with the modifications that the direction of tandem runs are unrestricted and the original home nest never drops out of competition. These two modifications would make the algorithm more general.

In Section 4, we obtain a running time result on single-nest emigrations. It would be interesting to apply the ideas from Section 4 to derive a corresponding result for the general multi-nest setting. This analysis may be more involved, but it would be interesting to see how including competing nests in the environment affects running time.

See [3] for other possible extensions to the algorithm that may be included, potentially at the expense of the runtime or simplicity of the algorithm. Some potential extensions include fault tolerance and asynchrony, where the synchronous round assumption is relaxed.

References

[1] *Biological Distributed Algorithms (BDA) 2019*, snl.salk.edu/∼navlakha/BDA2019/.

[2] Nigel R. Franks, Jonathan P. Stuttard, Carolina Doran, Julian C. Esposito, Maximillian C. Master, Ana B. Sendova-Franks, Naoki Masuda, and Nicholas F. Britton. How ants use quorum

sensing to estimate the average quality of a fluctuating resource. *Scientific Reports*, **5(1)** (2015), 1–12.

[3] Mohsen Ghaffari, Cameron Musco, Tsvetomira Radeva, and Nancy Lynch. Distributed house-hunting in ant colonies. *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing*, (2015), 57–66.

[4] Simone M. Glaser and Christoph Grüter. Ants (temnothorax nylanderi) adjust tandem running when food source distance exposes them to greater risks. *Behavioral Ecology and Sociobiology*, **72(3)** (2018), 1–8.

[5] Michael J. B. Krieger, Jean-Bernard Billeter, and Laurent Keller. Ant-like task allocation and recruitment in cooperative robots. *Nature*, **406(6799)** (2000), 992–995.

[6] Michael Möglich. Social organization of nest emigration in Leptothorax (Hym., Form.). *Insectes Sociaux*, **25(3)** (1978), 205–225

[7] Richard Karp, Christian Schindelhauer, Scott Shenker, and Berthold Vocking. Randomized rumor spreading. *In Proceedings 41st Annual Symposium on Foundations of Computer Science*, IEEE, (2000), 565–574.

[8] Stephen C. Pratt. Efficiency and regulation of recruitment during colony emigration by the ant temnothorax curvispinosus. *Behavioral Ecology and Sociobiology*, **62(8)** (2008), 1369–1376.

[9] Stephen C. Pratt, Eamonn B. Mallon, David J. Sumpter, and Nigel R. Franks. Quorum sensing, recruitment, and collective decision-making during colony emigration by the ant Leptothorax albipennis. *Behavioral Ecology and Sociobiology*, **52(2)** (2002), 117–127.

[10] Stephen C. Pratt, David J. Sumpter, Eamonn B. Mallon, and Nigel R. Franks. An agent-based model of collective nest choice by the ant Temnothorax albipennis. *Animal Behaviour*, **70(5)** (2005), 1023–1036.

[11] Thomas O. Richardson, Philippa A. Sleeman, John M. McNamara, Alasdair I. Houston, and Nigel R. Franks. Teaching with evaluation in ants. *Current biology*, **17(17)** (2007), 1520–1526.

[12] Lili Su, Martin Zubeldia, and Nancy Lynch. Collaboratively learning the best option on graphs, using bounded local memory. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, **3(1)** (2019), 1–32.

[13] Jiajia Zhao, Nancy Lynch, and Steven Pratt. The Power of Social Information in Ant-Colony House-Hunting: A Computational Modeling Approach. *bioRxiv*, (2021), 2020–10.