A Bio-Inspired Transportation Network for Scalable Swarm Foraging

Qi Lu[1], G. Matthew Fricke[1,*], Takaya Tsuno[2], and Melanie E. Moses[1,3,4]

Abstract— **Scalability is a significant challenge for robot swarms. Generally, larger groups of cooperating robots produce more inter-robot collisions, and in swarm robot foraging, larger search arenas result in larger travel costs. This paper demonstrates a scale-invariant swarm foraging algorithm that ensures that each robot finds and delivers targets to a central collection zone at the same rate regardless of the size of the swarm or the search area. Dispersed mobile depots aggregate locally collected targets and transport them to a central place via a hierarchical branching transportation network. This approach is inspired by ubiquitous fractal branching networks such as tree branches and animal cardiovascular networks that deliver resources to cells and determine the scale and pace of life. We demonstrate that biological scaling laws predict how quickly robots forage in simulations of up to thousands of robots searching over thousands of square meters. We then use biological scaling to predict the capacity of depot robots that overcome scaling constraints to produce scale-invariant robot swarms. We verify the claims for large swarms in simulation and implement a simple depot design in hardware.**

I. Introduction

Foraging is a canonical swarm robotics task with applications in search and rescue, construction, transportation, and exploration [1]–[4]. Agricultural harvesting [5], planetary exploration [6] and mining [6], [7] in particular require large numbers of robots to efficiently gather and transport resources over potentially large areas. We propose a hierarchical transportation network for thousands of robots to efficiently collect targets dispersed over large foraging arenas.

Central-place foraging (CPF) is the problem of searching for resources and transporting them to a central collection zone [8]–[10]. Here we design a scale-invariant foraging system in which the foraging rate is linear in both the number of robots and the search area. Biological swarms such as colonies of social insects, flocks of birds, and schools of fish have long served as inspiration for swarm robotics, in part due to the scalability of their solutions [11]–[15]. Biological systems also illustrate how collective systems can be adaptable and robust to individual failures, particularly in swarm foraging where multiple robots are advantageous for collecting spatially dispersed targets [9], [11]–[13], [16].

Two problems affect scalability, the first is a function of swarm size and the second is a function of the size of the search area. First, large swarms with many robots produce more inter-robot collisions both during the search process and during the return of targets to a relatively small, centralized collection zone. These collisions result in diminishing returns [17] as the swarm grows. Second, large foraging arenas require robots to travel further distances to find targets and transport them to the central collection zone.

We demonstrate the utility of applying biological scaling theory [18], [19] to swarm robotics by analysing the efficiency of central place foraging algorithms (CPFAs). We then demonstrate how to achieve near-linear scaling in two ways 1) by modifying the multiple place foraging algorithm (MPFA) [20], [21] to have variable capacity depots and 2) by creating a hierarchical transport multiple place foraging algorithm (MPFA_H). Using biological scaling theory we are able to show that *in theory* the MPFA_H and MPFA reach almost perfect scale-invariant foraging, but simulations that include collisions among depots show that only the MPFA_H approaches linear scaling. The MPFA_H uses a hierarchical network of mobile-depots of increasing capacity to achieve this scaling. We present a proof-of-concept hardware implementation of a robotic mobile-depot which is the foundation of the MPFA and MPFA_H.

I-A. Related Work

Scalability is a fundamental challenge in computer science [22]. Though swarm robot foraging has been studied for decades [14], [16], [23], analysis of the scalability of large swarms is limited. Where scalability has been analyzed, most studies find that large swarms are less efficient. For example, [24]–[26] all show dramatic reductions in per robot foraging rates for even modest increases in swarm size.

This work is related to task partitioning which has previously been shown to improve scalability. Pini et al. [27] demonstrated that a static partitioning strategy can provide a scalable and robust foraging robot swarm. Buchanan et al. [28] improved the scalability of robot swarms using a dynamic partitioning strategy which mitigates dead reckoning error. The work in [29] describes a leafcutter ant inspired foraging algorithm. The robot swarm achieves maximum foraging per-

This work was supported by a James S. McDonnell Foundation Complex Systems Scholar Award, funding from the UNM Office of the Vice President for Research, and NASA MUREP #NNX15AM14A funding for the Swarmathon.

[1]Department of Computer Science, University of New Mexico, [2]Department of Mechanical Engineering, Mie University, Tsu, Mie, Japan, [3]Department of Biology, University of New Mexico, Albuquerque, NM, USA. [4]Santa Fe Institute, Santa Fe, NM, USA.

[*]mfricke@cs.unm.edu

formance by dividing foraging and delivering tasks automatically using Grammatical Evolution [30]. In [31] the authors demonstrate the importance of tuning robot carrying capacity vs. swarm size in designing scalable non-central-place foraging algorithms. We build on this earlier work by introducing mobile depots with capacity tuned to optimize the transportation task.

II. Scaling Analysis

Biological scaling theory predicts that the cardiovascular transportation rate is proportional to animal volume (or mass) to the $3/4$ power and more generally to the power $D/(D+1)$ where D is the system's dimensionality. This constraint arises from the need to transport resources over greater distances in larger animals and limitations on the space available to devote to transportation routes. In essence, as body size increases the proportion of resources in transit, rather than in active use by cells, increases non-linearly. As a result, the resources available to each cell are reduced by a $1/4$ power of body mass [32]. For example, biological rates and times (e.g. fertility rates, heart rate and lifespan) are predicted and observed to be 30 times slower in an elephant which is a million times larger than a mouse. Below we derive how scaling produces analogous diminishing returns for CPFAs in 2D.

The transportation of resources through the cardiovascular system from the heart to dispersed cells (filling the 3D space of a mammal body) is the inverse problem of (2D) transportation of dispersed targets to a central collection zone in robot foraging. Scaling laws explain how the $3/4$ power scaling of delivery rates, and $-1/4$ power scaling of per-cell biological rates with body size result from a hierarchical transportation network which minimizes energy dissipation [18] and resource delivery time [19]. We have previously used this theory to predict how power and performance scale with microprocessor size [33].

To translate the biological scaling theory into a model of scalable robot foraging, we first identify similarities between the delivery of blood, which carries resources to cells through cardiovascular networks, and the robot foraging task, which carries targets to collection zones. In biology, scaling theory considers delivery of resources in 3D bodies divided into 'service volumes' which are the volume of tissue supplied by one capillary that delivers blood. Our ground-based robot foraging take places in a 2D area divided into "search regions" surrounding each depot that serves as a local collection zone. Both require transportation between a central site (heart or collection zone) and a service region. The venous system which conversely transports cellular waste such as CO_2 to the central pulmonary system is directly analogous to CPF. The analogies between these systems are listed in Table I.

We derive scaling predictions for foraging robot swarms using the following definitions and simplifying premises translated from [19] which derives maximal scaling by "matching supply and demand" so that there is no waiting or delay in pickup or delivery of targets. For simplicity, we omit the constants of proportionality.

TABLE I: Similarities between cardiovascular systems and robot swarms

Organisms	Robot swarms
3D bodies	2D arenas
Blood cells	Robots
Heart	Central collection zone
Service volumes	Search regions
Resources	Targets
Metabolic rate	Foraging rate
Resource delivery	Target collection

II-A. Premises and Definitions

Definition 1: The foraging problem is defined as collecting targets from an arena divided into N_r search regions. Each region has a specified area, A_r so that total arena area $A = N_r A_r$, which allows the size of search regions to vary with total arena size. Each region has a collection zone at its center. Searching robots operate within a designated search region and deliver targets to the collection zone at the center of its foraging region. Depots collect targets from their collection zones at level i and transport them to collection zones at level $i-1$ via the route specified by the network. The ant-inspired central place foraging algorithm (AntCPFA) [9] has zero such levels, MPFA [20] has 1 such level, and here we introduce the MPFA_H which has a number of levels determined by the size of the foraging area.

Definition 2: Target density, D_t, is the number of targets, M, in arena A; $D_t = M/A$. Simulation experiments have M targets distributed uniformly in A. To maintain constant D_t and simplify the analysis, as targets are collected, they are replenished in a randomly chosen new location.

Definition 3: Foraging rate, F, is the number of targets collected at the top-level central collection zone per unit time. The foraging rate in a region i is F_r^i, the number of targets collected and transported to the regional collection zone per unit time. Thus $F = \sum_{i=1}^{N_r} F_r^i$. Since the number of searching robots in each region is constant, we have $F \propto N_r$ and $A_r \propto A/N_r \propto A/F$.

Premise 1: Searching robot velocity, v_f, is constant across all experiments. Depot velocity, v_d, can vary for each foraging model and experimental setup.

Premise 2: The capacity of searching robots is always one target. The capacity of depots, C, can vary for each foraging model and experimental setup.

Premise 3: The number of targets in transit is at a steady state. The foraging rate F of targets being collected in all regions matches the rate targets that are delivered to a central collection zone.

Premise 4: The number of targets in transit is proportional to arena area: $N_t \propto A$ and thus, the density of targets in transit is the same across arena sizes. This premise is analogous to the biological scaling theory assertion that the fraction of blood (that transports resources) is constant across animal sizes.

Premise 5: The delivery rate matches the foraging rate. The system minimizes the time that collected targets are stored in regional collection zones waiting for a depot to pick them up; and no depot arrives at a collection zone and has to wait for a searching robot to drop off targets (i.e., if a depot has capacity 4, then there should be exactly 4 targets ready for pickup when it arrives at its collection zone). This design minimizes the delay in delivery. Ideally, this efficient pipelining means that targets do not unnecessarily wait to be picked up, and depots do not unnecessarily wait for targets to arrive at the collection zones. Thus, the rate of the dropoff in a region equals the rate of the pickup in that region, and therefore the delivery rate to the central collection zone equals the total foraging rate summed over all regions, F.

II-B. Sublinear Scaling in Transportation Networks

The explosion network (Figure 1a) models the MPFA which transports targets from each region directly to the central collection zone. The hierarchical branching network models the transportation routes of the MPFA_H. In the fully scalable implementation of the MPFA_H (described below in Theorem 3, and in Simulation Set III in Experimental Methods) targets are aggregated in smaller numbers of larger capacity depots along these routes.

Lemma 1. *The number of targets being transported is proportional to the mean transport route length times the foraging rate divided by the depot velocity, i.e.* $N_t \propto \bar{l}_{rt} F / v_d$.

Proof: We begin by relating foraging rate, F, to arena area, A, following [19]. The number of targets in transit is $N_t \propto n N_{rt}$, where N_{rt} is the number of routes from every region to the central collection zone and n is the average number of targets in transit per route. Therefore, $N_{rt} \propto N_r \propto F$ and $n = \bar{l}_{rt}/v_d$, where $\bar{l}_{rt}$ is the average length of a route: $\bar{l}_{rt} \propto A^{1/2}$. It follows that $N_t \propto \bar{l}_{rt} F / v_d$.

Theorem 1. *When robots follow an "explosion network" strategy with fixed robot velocity and depot capacity*[1]*, the foraging rate, F, will be proportional to the square root of the foraging arena area, i.e., $F \propto A^{1/2}$.*

Proof: From Lemma 1 we have $N_t \propto \bar{l}_{rt} F / v_d$. Since v_d is kept constant, $N_t \propto F A^{1/2}$. From Definition 3, $N_t \propto A$ and therefore $F \propto A^{1/2}$ □. Theorem 1 is the trivial scaling case in which targets spend time in transport proportional to the radius, $A^{1/2}$, which slows per capita delivery and therefore total foraging rates to also be proportional to $A^{1/2}$. To improve the scaling of the explosion transport network, following [19], we allow v_d to increase with arena size by setting v_d to the maximum value that allows consistently matching collection rate to delivery rate at all collection zones within an arena.

[1]For CPFAs without mobile depots the capacity is trivially zero.

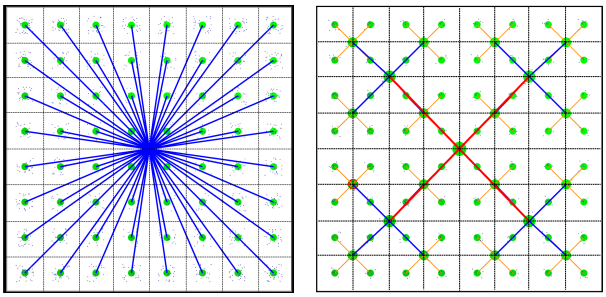

(a) Explosion (MPFA) (b) Hierarchical (MPFA_H)

Fig. 1: Transport paths in (a) explosion (MPFA) and (b) hierarchical (MPFA_H) networks. Each small square is a region with 4 searching robots (tiny dots). Mobile depots (tiny dots) transport targets to intermediate collection zones in next level (green circles) or to the central collection zone (green circle in the center) via routes (colored lines).

Theorem 2. *If the mobile-depot velocity is allowed to scale to the maximum beneficial value, then the foraging rate scales with foraging-arena area to the 2/3 power and velocity scales with area to the 1/6 power, i.e., when v_d is maximized subject to $D = F$ then $F \propto A^{2/3}$ and $v_d \propto A^{1/6}$.*

Proof: As shown in [19], the maximum velocity is proportional to the length of the shortest route (or region length) l. Exceeding this velocity forces depots to wait to be loaded which does not improve delivery rate. From Premise 3 we have $F \propto N_r$. Because the foraging arena is divided into N_r regions of length l we have $A = l^2 N_r$, which is $l = (A/N_r)^{1/2}$. Therefore $v_d \propto l \propto (A/F)^{1/2}$. It follows that $N_t \propto F A^{1/2}/(A/F)^{1/2} \propto F^{3/2}$. Thus, by maximizing velocity, we have derived the classic biological $D/(D+1)$ scaling theory for 2-dimensions: $F \propto A^{2/3}$ when $v_d \propto l \propto A^{1/6}$ □.

The hierarchical network differs from the explosion network by aggregating target transportation onto paths of increasing length and capacity. Figure 1b shows a network composed of 4 branches at each level from the central collection zone (level 0, red) to 4 regional depots (level 1, blue), to 16 sub-regional depots (level 2, yellow) and finally to 64 search regions with depots (green dots) at their center. We define a to be the ratio of route length at levels i and $i-1$ and b to be the in-degree for vertices in the directed tree that forms the hierarchical transportation network. In Figure 1b the network is a quad-tree with $a = 2$ and $b = 4$.

Lemma 2. *The minimum number of depots, N_d, required*

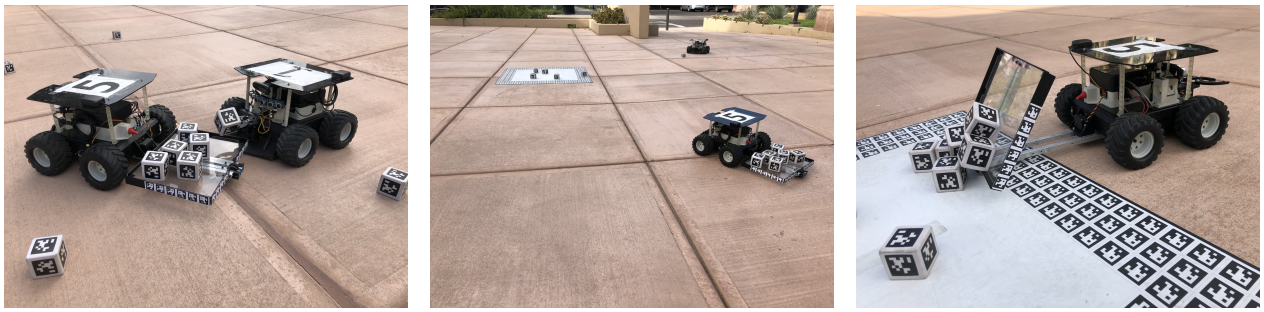

(a) Collection (b) Transport (c) Drop Off

Fig. 2: Proof of concept implementation of the MPFA_H using mobile depots. The three stages of depot behaviour are illustrated with Swarmie robots [6] and depots, collection zones labeled with April Tags: (a) searching robots place targets in the depot and send a message to the depot indicating a target was deposited, (b) once the depot reaches its pre-specified capacity, it travels to a collection zone, (c) the depot robot then rotates 180° and deposits the targets. Finally the mobile-depot returns to its original location and repeats.

TABLE II: The experimental parameters used to evaluate predictions in Theorems 1, 2 and 3. Set I has fixed mobile-depot capacity, fixed depot velocity and variable region sizes following the premises of Theorem 1. In Set II velocity of mobile-depots scales with arena area, which consequently changes region length according to Theorem 2. In Set III, the capacity of mobile-depots scales and the velocity is held constant according to Theorem 3. The number of targets (M) in the arena is set to ensure constant target density, and the number of searching robots (N_s), mobile depots (N_d), depot capacities (C), and collection zones are determined according to the premises and theorems in Section II.

	Shared configuration								MPFA	MPFA_H	
Set	Delivery Velocity (v_d)	Capacity (C)	Region Length (l)	Region (N_r)	Arena Dim. (m)	Targets (M)	Searching (N_s)	Depots (N_d) Eqn. 1	Collection Zones	Collection Zones (N_c) Eqn. 3	Level (L)
I	0.16	4	1	1	1×1	1	4	0	0	0	0
			2	4	4×4	16	16	4	4	4	1
			4	16	16×16	256	64	48	16	20	2
			8	64	64×64	4096	256	896	64	84	3
II	$\propto A^{1/6}$	4	1	1	1×1	1	4	0	0	0	0
			2	16	8×8	64	64	48	16	20	2
			4	256	64×64	4096	1024	3840	256	340	4
III	0.16	vary	5	4	10×10	100	16	16	4	4	1
				16	20×20	400	64	80	16	20	2
				64	40×40	1600	256	336	64	84	3
				256	80×80	6400	1024	1360	256	340	4

for the intake rate at the central collection zone, D, to match the total collection rate, F, is as follows:

$$N_d = \begin{cases} \frac{\sqrt{2}F_r}{v_d C}(A - A^{3/4}) & \text{if } v_d \text{ is constant} \\ \frac{\sqrt{2}F_r}{C}(A - A^{2/3}) & \text{if } v_d \propto A^{\frac{1}{6}} \end{cases} \quad (1)$$

Proof: The required number of levels, L, to connect all regions is $\log_b N_r$. N_d is minimized subject to the condition that $D = F$ and in accordance with Premise 5. The time for a depot at level i to make a round trip from the center of its region to its destination collection zone or depot is $T_d^i = 2d_i/v_d$, where d_i is the distance from the center of a region at level i to the center of the destination region at level $i-1$. The number of depots N_d^i is equal to the number of collected targets in T_d^i at level i divided by the capacity C: $N_d^i = 2b^{i+1}F_r b^{L-i-1}d_i/(v_d C)$, where $F_r b^{L-i-1}$ is the rate targets are collected in zones at level i and b^{i+1} is the number of branches at level i.

In the following summation over levels we first substitute b^L for N_r, then d_i for $d_0 a^i$, and d_0 for $\frac{\sqrt{2}}{2}l$, and then a^L for $N_r^{\frac{1}{2}}$.

$$\begin{aligned} N_d &= \sum_{i=0}^{L-1} \frac{2F_r b^L d_i}{v_d C} = \frac{\sqrt{2}F_r N_r l}{v_d C}\left(\frac{a^L - 1}{a - 1}\right) \\ &= \frac{\sqrt{2}F_r l}{v_d C}(N_r^{3/2} - N_r) \end{aligned} \quad (2)$$

Finally we rewrite Equation (2) in terms of A using Theorem 1 and Theorem 2. □

II-C. Scale-Invariant Transportation Network

In biology, the explosion and hierarchical networks are both limited to sublinear scaling: $F \propto A^{(D-1)/D}$ when v_d is constant and $F \propto A^{D/(D+1)}$ when v_d scales at its maximum value. However, we can use scaling theory to design a scale-invariant foraging swarm

which escapes these constraints so that the total foraging rate is linear with arena size and swarm size and per capita foraging rates are constant. If we increase the depot capacity C_i at level i by ab from level $i+1$, ($C_i = c(ab)^{L-i-1}$, where c is the capacity of depots in regional collection zones at level $L-1$), Equation (1) will have a constant N_d^i in each collection zone at each level i. With this depot capacity, transport keeps pace with the foraging in each identically sized region and Theorem 3 follows:

Theorem 3. *A CPFA employing a hierarchical transportation network using mobile-depots with capacity that scales with the level of the hierarchy they service will have a foraging rate that scales linearly with the foraging-arena area, i.e., $F \propto N_r \propto A$.*

For the quad-tree implementation, there are 4 searching robots in each region and 4 depots for each collection zone, and each depot has a systematically larger capacity on more central routes. So, the total number of searching robots is $4N_r$ and the total number of depots is $4N_c$, where N_c is the number of collection zones (excluding the central collection zone). Thus:

$$N_c = \sum_{i=0}^{L-1} \frac{N_r}{b^i} = \frac{4(N_r - 1)}{3} \tag{3}$$

III. Experimental Methods

We conducted three sets of experiments using the Autonomous Robots Go Swarming (ARGoS) [34] robot simulator to test our theoretical predictions (see Youtube video[2]). The experimental setup for each set is summarized in Table II. For each Set, we implement the MPFA explosion network and the MPFA_H hierarchical branching network. In Set I, we test the 1/2 scaling with constant depot velocity and capacity. In Set I, we also implement the Ant-inspired CPFA without depots for comparison.

In Set II, we test the 2/3 scaling predicted by Theorem 2, here $C = 4$ and $v_d \propto l \propto A^{1/6}$. In Set III, we test the linear scaling predicted by Theorem 3. As described in Section II-C for the MPFA_H the depot capacity C is scaled by ab on each level, and the number of depots, N_d, is $4N_c$ as predicted in Equation (3). We also implemented a new variant of the MPFA that uses variable numbers and capacities of robots (analogous to the MPFA_H). The experimental parameters listed in Table II were selected carefully to meet multiple constraints. Arena dimensions are chosen to ensure that the dependent variables are integers and divisible by 4 when necessitated to achieve a balanced quad tree. This avoids fractional regions and robots.

To provide fair comparisons within each Experimental Set, we use the same number of robots for

[2] https://youtu.be/tgSTCz264cU

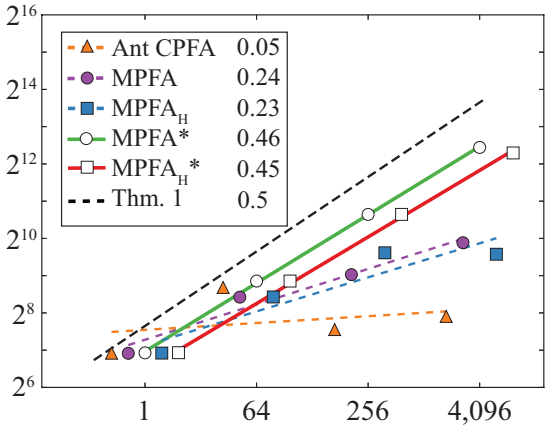

(a) Comparison between Theorem 1 and Simulation Set I. All algorithms in Set I use depots with constant capacity and velocity.

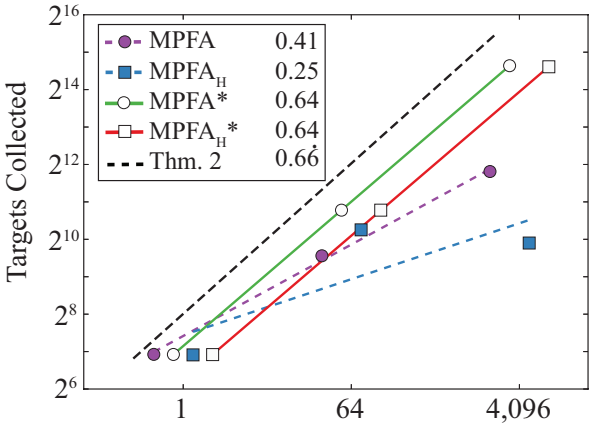

(b) Comparison between Theorem 2 and Simulation Set II. All algorithms in Set II have depots with velocity proportional to the distance they must cover (labelled "vary" in Table II).

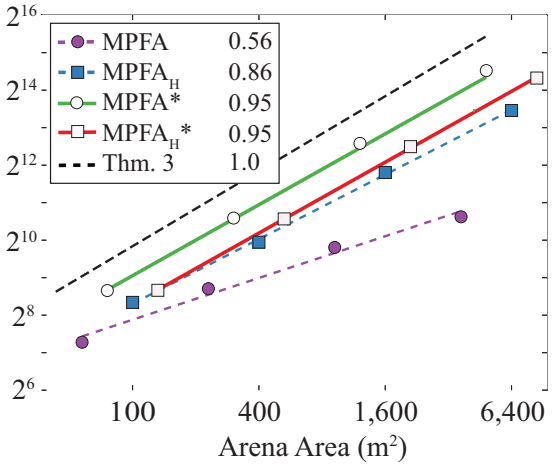

(c) Comparison between Theorem 3 and Simulation Set III. All algorithms in Set III use constant velocity and variable capacity depots.

Fig. 3: Scaling in simulation vs theory. Scaling exponents are the slopes (listed in the legends) on log-log plots of the regression line for each algorithm. Each data-point is the mean of 60 30-minute trials. Note the change in x-axis scale in panel (c).

the explosion and hierarchical implementations and the same number of depots for the MPFA and the MPFA_H. In the MPFA, we distributed depots on each delivery route proportional to the length of that route. In the AntCPFA, the depots are replaced with searching robots that directly deliver targets to the central collection zone. The AntCPFA is only included in Set I where we demonstrate that it is not scalable, consistent with results from prior work [20].

IV. Experimental Results

Generally, scaling relationships are analyzed as $Y = \alpha X^\beta$ where α is a conversion factor, and β is an exponent which indicates the scaling relationship between X and Y. We display and analyze log-transformed data so that $\log_2 Y = \beta \log_2 X + \log_2 \alpha$. The slope of the regression between the $\log_2 Y$ and the $\log_2 X$ gives an estimate of the scaling exponent β.

In Figure 3 simulations marked with an asterisk were run without depot collisions (but forager collisions were kept in the simulation). These runs agree closely with theory (p-value< 0.001 in all cases). The degree to which experiments that include depot collisions deviate from theory illustrates how sensitive each algorithm is to depot collisions.

Figure 3a highlights the degradation in performance due to depot collisions when depots have fixed capacity and travel with a fixed velocity. The slopes of the MPFA* and the MPFA_H* approach the $1/2$ slope predicted by Theorem 1, with slopes of 0.46 and 0.45 respectively. However, the implementations with depots collisions are less scalable. Searcher collisions reduce the scalability of the AntCPFA dramatically, and since that algorithm does not include depots, there is no way to improve its scaling with a transportation network.

In Figure 3b scaling of the MPFA_H is worse than the MPFA, and both are well below the $2/3$ slope predicted by Theorem 2. However, that scaling is nearly achieved by depot collision-free algorithms. This is because with large numbers of low-capacity depots, the MPFA_H has high collision rates because there are many depots crowded onto a few routes. Collisions are actually worse in Set II with faster depot velocity, compared to Set I with slower, fixed velocity. Too many depots cause more problems for the MPFA_H than for the MPFA, because targets aggregated at mid-level collection zones in the MPFA_H cannot reach the central collection zone via the four crowded routes available.

Figure 3c shows that scaling depot capacity, C, is the key to scalable foraging, even with collisions in the MPFA_H. This approach of aggregating more targets in a few larger depots (rather than in many small depots) reduces collisions. The slope of the MPFA_H, 0.86, is closer to the theoretical optimum than the MPFA, 0.56, and is the highest slope of all algorithms that consider depot collisions.

V. Discussion

Real-world applications like harvesting, mining and planetary exploration require that robots find and collect as many objects as possible from large areas in the least amount of time. We propose the MPFA_H as a highly scalable foraging algorithm, designed from biological scaling principles, tested in simulation, and with mobile depots and foraging robots implemented in hardware as a proof of concept.

We predict how to achieve scale-invariant foraging by aggregating target transport into depots, with the capacities, routes, and numbers of depots determined by biological scaling theory. The MPFA_H achieves nearly scale-invariant foraging even when collisions are included in simulations. Our proof of concept shows that depots can collect from foraging robots in one region and deliver to the next hierarchical level of our transportation architecture. Since the MPFA_H algorithm allows each foraging region to operate independently, and it effectively coordinates delivery between hierarchical levels, the feasibility we demonstrated in one region is evidence that the approach could be implemented at larger scales in physical hardware.

Our three sets of experiments validate the mathematical claims of several variants of scaling theory. Using the MPFA* and the MPFA_H* that eliminate collisions among depots, we show that both explosion algorithms and hierarchical branching algorithms approach the claims of each theorem ($1/2$, $2/3$, and linear scaling). Performance in most implementations is substantially impaired when collisions are included in simulations. For example, with collisions both the MPFA and MPFA_H are several orders of magnitude worse in the largest arena in Sets I and II (see Figure 3, panels (a) and (b); compare filled squares and circles to hollow ones). However, the scale-invariant implementation of the MPFA_H is very close to the predicted slope (0.86 vs 1.0), even with collisions, and it outperforms the MPFA by several orders of magnitude. This highlights an important feature of the hierarchical network. By aggregating targets into the sizes and numbers of depots predicted by scaling theory (analogous to the way blood is aggregated in larger vessels like the aorta in cardiovascular networks) the MPFA_H reduces both the number of paths to each collection zone, and the crowding on each path. The branching factor b results in a constant density of depots throughout the arena, even as the swarm and arena scale to large sizes. The robot foraging transportation network can increase capacity (up to the largest possible depot) to accommodate the increase in transport required in larger arenas to ensure that delivery capacity keeps up with a constant per-forager collection rate. This produces a scale-invariant swarm in theory and a nearly scale-invariant swarm when collisions are considered.

References

[1] A. F. T. Winfield, "Towards an engineering science of robot foraging," *Distributed autonomous robotic systems*, vol. 8, pp. 185–192, 2009.

[2] R. Groß and M. Dorigo, "Towards group transport by swarms of robots," *International Journal of Bio-Inspired Computation*, 2009.

[3] S. K. Yun and D. Rus, "Adaptive coordinating construction of truss structures using distributed equal-mass partitioning," *IEEE Transactions on Robotics*, 2014.

[4] W. Fink, J. M. Dohm, M. A. Tarbell, T. M. Hare, and V. R. Baker, "Next-generation robotic planetary reconnaissance missions: A paradigm shift," *Planetary and Space Science*, 2005.

[5] C. W. Bac, E. J. Henten, J. Hemming, and Y. Edan, "Harvesting Robots for Highvalue Crops: Stateoftheart Review and Challenges Ahead," *Journal of Field Robotics*, vol. 31, no. 6, pp. 888–911, 2014.

[6] S. M. Ackerman, G. M. Fricke, J. P. Hecker, K. M. Hamed, S. R. Fowler, A. D. Griego*, J. C. Jones*, J. J. Nichol*, K. W. Leucht, and M. E. Moses, "The Swarmathon: An Autonomous Swarm Robotics Competition," in *International Conference on Robotics and Automation (ICRA), 2018 IEEE/RSJ International Conference on*, IEEE.

[7] C. Brown, "Autonomous vehicle technology in mining," *Engineering and Mining Journal*, vol. 213, no. 1, p. 30, 2012.

[8] Wenguo Liu, A. F. Winfield, Jin Sa, Jie Chen, and Lihua Dou, "Towards energy optimization: Emergent task allocation in a swarm of foraging robots," *Adaptive Behavior*, 2007.

[9] J. P. Hecker and M. E. Moses, "Beyond pheromones: evolving error-tolerant, flexible, and scalable ant-inspired robot swarms," *Swarm Intelligence*, vol. 9, no. 1, pp. 43–70, 2015.

[10] E. Castello, T. Yamamoto, F. D. Libera, W. Liu, A. F. Winfield, Y. Nakamura, and H. Ishiguro, "Adaptive foraging for simulated and real robotic swarms: the dynamical response threshold approach," *Swarm Intelligence*, 2016.

[11] B. Webb, "Swarm Intelligence: From Natural to Artificial Systems," *Connection Science*, 2002.

[12] J. Kennedy, R. C. Eberhart, Y. Shi, C. Jacob, J. R. Koza, F. H. B. Iii, D. Andre, and M. a. Keane, *Swarm Intelligence The Morgan Kaufmann Series in Evolutionary Computation*. 2001.

[13] E. Sahin, "Swarm robotics: From sources of inspiration to domains of application," in *Swarm robotics*, pp. 10–20, Springer, 2005.

[14] J. C. Barca and Y. A. Sekercioglu, "Swarm robotics reviewed," 2013.

[15] Y. Khaluf, E. Ferrante, P. Simoens, and C. Huepe, "Scale invariance in natural and artificial collective systems: A review," 2017.

[16] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo, "Swarm robotics: a review from the swarm engineering perspective," *Swarm Intelligence*, vol. 7, no. 1, pp. 1–41, 2013.

[17] S. L. Brue, "Retrospectives: The Law of Diminishing Returns," *Journal of Economic Perspectives*, 1993.

[18] G. B. West, J. H. Brown, and B. J. Enquist, "A general model for the origin of allometric scaling laws in biology," *Science*, vol. 276, no. 5309, pp. 122–126, 1997.

[19] J. R. Banavar, M. E. Moses, J. H. Brown, J. Damuth, A. Rinaldo, R. M. Sibly, and A. Maritan, "A general basis for quarter-power scaling in animals," *Proceedings of the National Academy of Sciences*, vol. 107, no. 36, pp. 15816–15820, 2010.

[20] Q. Lu, J. P. Hecker, and M. E. Moses, "The MPFA: A Multiple-Place Foraging Algorithm for Biologically-Inspired Robot Swarms," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2016.

[21] Q. Lu, M. E. Moses, and J. P. Hecker, "A Scalable and Adaptable Multiple-Place Foraging Algorithm for Ant-Inspired Robot Swarms," Workshop on On-line decision-making in multi-robot coordination, 2016 Robotics Science and Systems Conference, arXiv:1612.00480, 2016.

[22] D. H. Ackley, "Indefinite scalability for living computation," in *30th AAAI Conference on Artificial Intelligence, AAAI 2016*, 2016.

[23] A. F. T. Winfield, "Foraging Robots," in *Encyclopedia of complexity and systems science* (R. A. Meyers, ed.), pp. 3682–3700, New York: Springer, 2009.

[24] J. P. Hecker and M. E. Moses, "Beyond Pheromones: Evolving Adaptable, Flexible, and Scalable Ant-Inspired Robot Swarms." 2014.

[25] G. M. Fricke, P. H. Joshua, D. G. Antonio, T. T. Linh, and E. M. Melanie, "A Distributed Deterministic Spiral Search Algorithm for Swarms," in *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*, pp. 4430–4436, Copyright 2016, IEEE reprinted with permission from the authors., 2016.

[26] A. Font Llenas, M. S. Talamali, X. Xu, J. A. Marshall, and A. Reina, "Quality-sensitive foraging by a robot swarm through virtual pheromone trails," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2018.

[27] G. Pini, A. Brutschy, A. Scheidler, M. Dorigo, and M. Birattari, "Task partitioning in a robot swarm: Object retrieval as a sequence of subtasks with direct object transfer," *Artificial Life*, 2014.

[28] E. Buchanan, A. Pomfret, and J. Timmis, "Dynamic task partitioning for foraging robot swarms," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2016.

[29] E. Ferrante, A. E. Turgut, E. Duéñez-Guzmán, M. Dorigo, and T. Wenseleers, "Evolution of Self-Organized Task Specialization in Robot Swarms," *PLoS Computational Biology*, 2015.

[30] E. Ferrante, E. Duéñez-Guzmán, A. E. Turgut, and T. Wenseleers, "GESwarm: Grammatical evolution for the automatic synthesis of collective behaviors in swarm robotics," in *GECCO 2013 - Proceedings of the 2013 Genetic and Evolutionary Computation Conference*, 2013.

[31] A. Schroeder, B. Trease, and A. Arsie, "Balancing robot swarm cost and interference effects by varying robot quantity and size," *Swarm Intelligence*, 2019.

[32] V. M. Savage, J. F. Gillooly, W. H. Woodruff, G. B. West, A. P. Allen, B. J. Enquist, and J. H. Brown, "The predominance of quarter-power scaling in biology," *Functional Ecology*, vol. 18, no. 2, pp. 257–282, 2004.

[33] M. Moses, G. Bezerra, B. Edwards, J. Brown, and S. Forrest, "Energy and time determine scaling in biological and computer designs," *Philosophical Transactions of the Royal Society B: Biological Sciences*, 2016.

[34] C. Pinciroli, V. Trianni, R. OGrady, G. Pini, A. Brutschy, M. Brambilla, N. Mathews, E. Ferrante, G. Di Caro, and F. Ducatelle, "ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems," *Swarm intelligence*, vol. 6, no. 4, pp. 271–295, 2012.