Evolving Multirobot Excavation Controllers and Choice of Platforms Using an Artificial Neural Tissue Paradigm

Jekanthan Thangavelautham, Nader Abu El Samid, Paul Grouchy, Ernest Earon
Terence Fu, Nagina Nagrani, Gabriele M. T. D'Eleuterio

Abstract— **Autonomous robotic excavation has often been limited to a single robotic platform using a specified excavation vehicle. This paper presents a novel method for developing scalable controllers for use in multirobot scenarios and does not require human defined operations scripts nor extensive modeling of the kinematics and dynamics of the excavation vehicles. Furthermore, the control system doesn't require specifying an excavation vehicle type such as a bulldozer, front-loader or bucket-wheel and it can evolve to select for an appropriate choice of excavation vehicles to successfully complete a task. The "Artificial Neural Tissue" (ANT) architecture is used as a control system for autonomous multirobot excavation and clearing tasks. This control architecture combines a variable-topology neural-network structure with a coarse-coding strategy that permits specialized areas to develop in the tissue. Training is done in a low-fidelity grid world simulation environment and where a single global fitness function and a set of allowable basis behaviors need be specified. This approach is found to provide improved training performance over fixed-topology neural networks and can be easily ported onto different robot platforms. Aspects of the controller functionality have been tested using high fidelity dynamics simulation and in hardware. An evolutionary training process discovers novel decentralized methods of cooperation employing aggregation behaviors (via synchronized movements). These aggregation behaviors are found to improve controller scalability (with increasing robot density) and better handle robot interference (antagonism) that reduces the overall efficiency of the group.**

I. INTRODUCTION

It is remarkable that an army of ants can march lock-step carrying food, construction material or excavating a network of tunnels. The collective behaves as one large entity, with the actions of the individual synchronized towards completing a global goal. Such behavior in the natural world can both awe and inspire the roboticist [1]. But how does the roboticist approach the control of such artificial systems?

Ants and other insects appear to act locally without knowledge of the greater goal. Yet how global consensus emerges from local behaviors is not well understood. Designing controllers by hand is a daunting task. Previous work in the field such as [2], [3], [4] rely on task-specific human knowledge to develop simple "if-then" rules to solve multirobot tasks. One might well wonder whether devising a control system by the same means that insects came to acquire theirs is not a more effective strategy.

Jekanthan Thangavelautham is with the Mechanical Engineering Dept., Massachusetts Institute of Technology, 77 Massachusetts Ave., Cambridge, MA, 02139. Nader Abu El Samid, Paul Grouchy, Ernest Earon, Terence Fu, Nagina Nagrani and Gabriele M. T. D'Eleuterio are with the Institute for Aerospace Studies, University of Toronto, 4925 Dufferin St., Toronto, ON M3H 5T6, Canada

We investigate here a controls architecture inspired by neural networks and evolved using an artificial Darwinian process called the "Artificial Neural Tissue" (ANT) framework [5], [6]. This approach exploits machine learning techniques to obtain the necessary coordination rules to solve a prespecified task.

The focus of this paper is the extension of the ANT framework to facilitate concurrent selection of excavation platforms and controllers to solve a multirobot excavation task. The controllers can exploit self-organized *aggregation*, a mode of cooperative behavior that can limit the effects of interference among robots in a multiagent setting. Previous work in this field has shown the need for evolving controllers that can limit the effects of interference (antagonism) and perform well with increased density of robots. Aggregation is an insect phenomenon and results in efficient, well coordinated actions among individuals. The ANT controllers have access to an extended palette of sensors and generic basis behaviors. These are merely building blocks that may be used to acquire aggregation capabilities.

Evolution is driven only by a global fitness function that measures (global) excavation success. It is up to the ANT controllers to determine what excavation platform to use and how best to use a palette of sensors and basis behaviors, from which the aggregation capability can emerge. Aggregation is not shaped or guided using a fitness function as in other work [7], [8]. Instead, this capability is expected to emerge through self-organization. ANT controllers with aggregation capability are compared with previously published results that lack this functionality.

Two of the major concerns in developing multirobot controllers is scalability and *antagonism* [9], where multiple robots trying to perform the same task interfere with one another and reduce the overall efficiency of the group or, worse, unable to solve the task. As previously shown, in a decentralized configuration, increased density of robots beyond a critical point, reduces overall efficiency due to antagonism. In nature, there are other strategies that limit the effects of antagonism. Aggregation via choreographed movement, as shown here, facilitates minimizing robot-robot interactions as obstacles to be avoided, thus increasing robot productivity. Furthermore, this feat is possible without requiring centralized control and coordination. By minimizing antagonism, it is possible to increase further the density of robots in a given workspace. This in turn provides important benefits, notably increased redundancy and parallelism.

We shall begin with a short background on robotic exca-

vation followed by details on the implementation of ANT for the task at hand. This is followed by discussion of the training process and analysis of the controller in simulation and in hardware.

II. BACKGROUND

In terrestrial settings, autonomous excavation has often been limited to a single robotic platform, with support vehicles for loading and unloading [10]. Digging is performed using hand coded operational scripts that decompose a specified task into repetitive excavation and truck loading cycles. Other approaches are devoted to modeling kinematics and/or dynamics of the excavation vehicles [11] and simulating their performance. The ANT controllers presented here are used in a multirobot scenario and are trained in a low-fidelity grid world simulation environment.

The ANT framework allows for scalable control of multiple excavation platforms towards achieving a specified global goal and is a synthesis of several collective robotics concepts. Collective robotic controllers typically mimic the mechanisms used by social insects. These include the use of templates, stigmergy, aggregation and self-organization. Templates are environmental features perceptible to the individuals within the collective [12]. In robotic applications, template-based approaches include the use of light fields to direct the creation of linear walls [13]. Stigmergy is a form of indirect communication mediated through the environment [14]. Stigmergy has been used extensively in collective-robotic construction, including blind bulldozing [3], box pushing [2], heap formation [15] and tiling pattern formation [16]. Aggregation is a process of clustering of units into a mass or whole. This functionality uses two basic mechanisms, positive and negative feedback, that result in attraction towards a source or repulsion [7], while self-organization describes how local or microscopic actions give rise to a macroscopic structure in systems that are not in equilibrium [17].

The collective robotic and autonomous excavation works cited earlier excluding [16] rely on either user-defined, deterministic 'if-then' rules, or on stochastic behaviors. In both cases, designing these controllers is an ad-hoc process that relies on the experimenter's or operator's knowledge of the task at hand. However, for collective robotic tasks, the global effect of local interactions is often difficult to gauge, and the specific interactions required to achieve a global consensus may be counterintuitive. Thus at this stage of the field's development, hand-coded controllers are derived through trial and error (manually).

One approach to reducing the amount of trial and error done by hand is to use Evolutionary Algorithms to train the robot controller. Fixed topology neural networks are often used as the scheme allows for plasticity and generalization unlike lookup-table approaches [16]. Variable-length neural-network topologies like ANT provide added benefits such as allowing for improved generalization of sensory input, improved scalability over fixed-network topologies, and allowing more extensive task decomposition capability [5], [6]. More important, for excavation, ANT allows for both sensor and behavior *extensibility* and is not constrained to a specific digging platform.

We also argue that increased controller flexibility may be obtained by allowing for aggregation behaviors. In nature, aggregation behaviors are used by a wide range of organisms including bacteria, insects, birds and mammals to increase the chance of survival (by herding) or through improved efficiency in locomotion. Earlier works in the field use fitness functions that explicitly encourage aggregation to perform clustering, foraging [7] and collective transport [18]. A few of these experiments were realized in hardware using s-bots [19] that have grappling actuators, enabling multiple robots to be physically linked [20], [8].

III. ARTIFICIAL NEURAL TISSUE MODEL

The ANT architecture (Figure 1) presented in this paper consists of a developmental program, encoded in the "genome," that constructs a three-dimensional neural tissue and associated regulatory functionality. The tissue consists of two types of neural units, decision neurons and motor-control neurons. Decision neurons are fed sensory input and dynamically activate and inhibit (mask) groups of motor neurons using a simulated chemical signalling scheme. Groups of motor neurons activated by the decision neurons are fully connected to the sensory input and in turn produce controller outputs, that can trigger a set of basis behaviours. Further details on the ANT framework can be found in [5], [6].

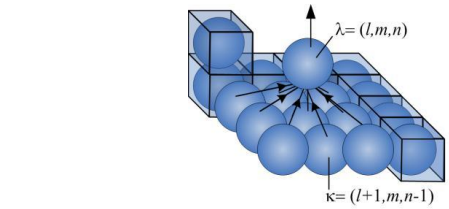

(a) Synaptic connections

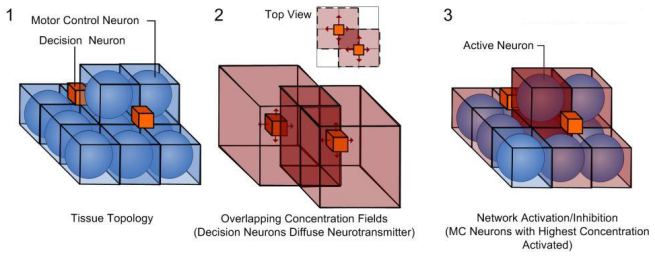

(b) Coarse-coding

Fig. 1. Synaptic connections between motor neurons and operation of neurotransmitter field.

IV. EXPERIMENTAL SETUP

The excavation task is intended to demonstrate the feasibility of autonomous teams of robots digging holes and clearing landing pads for lunar base construction. It can be argued that emergent task decomposition is necessary to accomplish the task given a global fitness function. A typical layout of the simulation experiment area used for training

is shown in Figure 2. The experiment region or workspace is modeled as a two-dimensional grid environment with each robot occupying four grid squares. For this task, it is argued that the controller needs to possess a number of capabilities to complete the task, including interpreting excavation 'blueprints', performing layered digging, avoid burying or trapping other robots, clearing and maintaining excavation routes. Each robot controller has access to a goal map that defines the location of the dump area and target depth of the excavation area (Figure 2). The fitness function f for the task is given as follows:

$$f = \frac{\sum_{j=1}^{J}\sum_{i=1}^{I}\vartheta_{i,j}\cdot e^{-2|g_{i,j}-h_{i,j}|}}{\sum_{j=1}^{J}\sum_{i=1}^{I}\vartheta_{i,j}} \quad (1)$$

where I and J are the dimensions of the entire area and $\sum_{j=1}^{J}\sum_{i=1}^{I}\vartheta_{i,j} > 0$ and $\vartheta_{i,j} = 1$ if grid square (i,j) is to be excavated and 0 otherwise; $g_{i,j}$ is the target depth and $h_{i,j}$ is the current regolith depth.

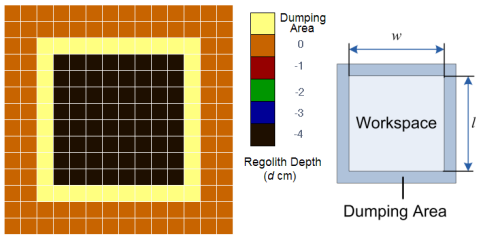

Fig. 2. An example goal map consisting of an excavation area surrounded by a dump area.

A. Excavation Robot Model

For these experiments, the controllers are homogenous (i.e., a single controller is replicated for use on each robot), decentralized and make use of local sensor data to complete the task. The inputs to the ANT controller are shown in Table I and Figure 3. The robots have access to current position (x, y) from localization scans performed in simulation. The variable z is computed through integration of changes in depth values. The discretized x and y coordinates are used to look up the goal depth $g_{x,y}$ of each grid square region in front of the robot.

All raw input data are discretized and fed to the controller. $Z_1 \ldots Z_4$ and $FZ_1 \ldots FZ_2$ are obtained using simulated ground scans. A set of simulated sensors are used to detect obstacles at the front and back (S_1, S_2). Robot orientation, heading and whether a robot is stuck or not are represented using TL_1, D_1 and ST_1 respectively. In addition, the controllers can perform aggregated group actions that we describe below. Alignment sensors ($A_1 \ldots A_4$) are used to determine if neighboring robots are in alignment position. The robots also have access to one memory bit (M_1), that can be manipulated using some of the basis behaviors.

Table II lists the basis behaviors the robot can perform (in order) within a single timestep. These prespecified behaviors are generic and can be used for many different tasks, with

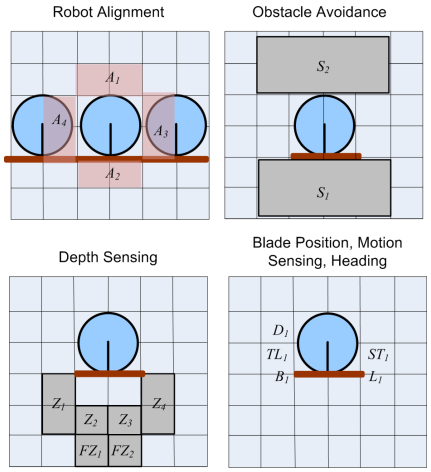

Fig. 3. Robot input sensor mapping (with bulldozer implement) for the simulation model.

various excavation implements. It is up to the ANT controller to determine when to execute these basis behaviors (using sensory input) to solve the overall task. Furthermore, the excavation behaviors allow for actuation of an implement. For the bulldozer implement, the blade can be raised or lowered to one of four positions. Fine adjustments are made using PID control to ensure a constant force is met. While for the front-loader and bucket-wheel, the lower most implement position actuates a 'dig and scoop' action that fills a bucket with regolith. For the bucket-wheel, the rotating buckets when lowered will dig, scoop and accumulate the regolith into an attached tank. Behavior 4 is also implemented differently according to the attached implement. For the bulldozer implement, moving backwards results in pushed regolith being 'dumped.' However for the front-loader, the bucket is raised and tilted downwards to 'dump' accumulated regolith. While for the bucket-wheel, the attached tank is opened and tipped over much like a dump truck towards the back. The capacity of the bucket-wheel tank, front-loader bucket and volume of soil pushed by the bulldozer is set to the same value (26 units), to make experimental comparisons between the implements meaningful.

For the aggregation capable robots, once two robots are in alignment position and behavior 2 is triggered and behavior 3 is not triggered (Table II), then the AS_1 value for both robots is set to 1. A group of robots detach from aligned state ($AS_1 = 1$) once all members of the group trigger behavior 3 and behavior 2 is not triggered. This capability allows for positive feedback, in which the robots choose to aggregate and negative feedback whereby they vote to detach and part ways. When a group of robots are in aligned state ($AS_1 = 1$), motor behaviors such as 'move forward', 'turn left 90^o' and 'turn right 90^o' are performed through consensus (Figure 4). Each member of the group triggers the 'move forward' behavior for the aggregated group to move forward otherwise the 'move forward' triggered by some members are vetoed.

TABLE I

ANT CONTROLLER SENSOR INPUTS

Variable	Function	States
$Z_1 ... Z_4$	Depth Sensing Relative to Goal Depth	Level, Above, Below, Don't Care
FZ_1, FZ_2	Depth Sensing Relative to Ground	Above, Below, Level
B_1	Excavation Implement Position	Above, Level, Below Ground, Home
S_1	Front Obstacle Detection	Obstacle, No Obstacle
S_2	Back Obstacle Detection	Obstacle, No Obstacle
D_1	Heading	North, East, West, South
TL_1	Robot Tilted Downwards	True, False
ST_1	Robot Stuck	True, False
$A_1 ... A_4$	Robot Alignment Position	Alignment Position, Not in Alignment Position
AS_1	Robot Alignment Status	Not Aligned, Aligned
M_1	Memory bit	0, 1

TABLE II

ANT CONTROLLER BASIS BEHAVIORS

Behavior	Function	States
1	Set Throttle	Set rover throttle to high otherwise remain nominal
2	Want to Align	Intention of rover to go into "align state"
3	Not Want to Align	Intention of rover to exit out of "align state"
4	Move Forward	Move one grid square forward
5	Dump and Move Backward	Dump regolith and move one grid square backward
6	Random Turn	Randomly Turn 90° right or Turn 90° left
7	Turn Right	Turn 90° right
8	Turn Left	Turn 90° left
9	Implement position: Above	Set excavation implement above ground d cm
10	Implement position: Below	Set bulldozer implement below ground d cm or use bucket-wheel/front -loader implement to scoop regolith.
11	Implement position: Level	Set excavation implement level to ground
12	Implement position: Home	Set implement to home position (makes no contact with soil)
16	Bit Set	Set memory bit 1 to 1
17	Bit Clear	Set memory bit 1 to 0

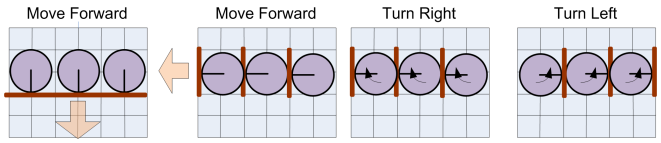

Fig. 4. Robots (shown with bulldozer implement) in aligned state performing motor behaviors. Note that the 'move forward behavior can be performed in two configurations (side by side and one behind another).

B. Simulation Parameters

Darwinian selection is performed based on the fitness value of each controller averaged over 100 different initial conditions. The population size for the experiments is $P = 100$, with crossover probability $p_c = 0.7$, mutation probability $p_m = 0.025$ and a tournament size of $0.06P$. The fitness f for each controller is calculated for an excavation area to be dug to a goal depth of d below ground and area spanning $l \times w$ squares. A 100 different scenarios are used during training, with excavation areas varying from 4×4 up to 12×12 squares and with $d = 1$ or 2. The dump site either surround the work site as in Figure 2 or is located in a line on the eastern side of the work site. Initially, the robots are placed in one of two configurations during training. In one configuration, the robots are randomly dispersed over the work site, while in the other, the robots are placed in a row next to a corner.

C. LEGO® Hardware Experiments and Environment

The ANT controllers are trained (evolved) in a grid world simulation environment written in C/C++ (Figure 5 left). The best individuals from several evolutionary runs are then tested in Digital Spaces™, a high-fidelity simulation environment and on LEGO® Mindstorm robots. The LEGO® hardware experiments are intended to verify robot-robot interactions and are not sufficient to validate excavation performance on regolith. Regolith excavation performance is tested within the Digital Spaces™ simulation environment. For these hardware experiments, three LEGO® Mindstorms™ NXT Robots have been used (Figure 5 right). Each robot consists of a pair of servo drive motors in a differential drive configuration, a gyro that records the rate of change of heading, a compass, a sonar unit that measures distance to obstacles and a one degree of freedom blade. The ANT controller for each robot is run on a dedicated laptop and it communicates using a wireless Bluetooth interface.

After initialization, the robots enter a four-step loop until the experiment is completed. For the first step, input data is collected from the sensors onboard each robot and each robot's local map is updated by an overhead camera system. This is followed by the second step, where ANT chooses a series of basis behaviors to perform. In the third step the chosen behaviors are executed, and in the final step heading is measured using the compass and corrected using proportional control. The move forward and backwards behaviors adapted for the LEGO® robots uses odometry to measure

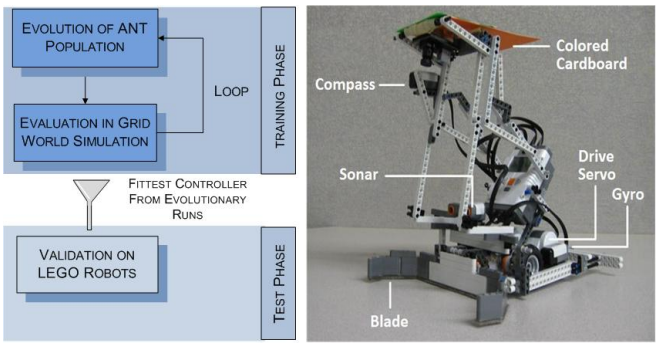

Fig. 5. (Left) Hardware validation process of the evolved ANT controllers. (Right) LEGO® Robot Configuration

distance traveled and proportional control using the onboard gyro to maintain a straight path. The turn 90° to the left or right behaviors also uses proportional control with the onboard gyro.

An overhead camera system was set up using a Logitech® webcam with 640×480 resolution. Colored Styrofoam peanuts is used to represent regolith and colored cardboard was mounted on top of the robots to distinguish each one for use by the overhead system. With this configuration, a naïve Bayes classifier is simultaneously able to detect and distinguish resources and robots. Using overhead cameras, each robot gets an updated reading of $Z_1 ... Z_4$, FZ_1, FZ_2, ST_1, S_2 and $A_1 ... A_4$. Algorithms were written that could take advantage of this color detection to provide robot localization, robot heading and styrofoam location. The overhead system can also be used to maintain alignment between the aggregated robots.

D. Digital Spaces™ Simulation

The ANT-based excavation controllers are demonstrated using high-fidelity simulations within Digital Spaces™, an off-the-shelf commercial 3-D Lunar simulation tool. Once the controllers have acquired the necessary traits from the evolutionary training procedure, they are ported onto the Digital Spaces™ simulation environment for further testing. Robotic hardware is represented with increased fidelity and it includes sensor interfacing, and deformable terrain modeling. Deformable terrain modeling allows for realistic regolith-tool and wheel-track-terrain interactions. This physics-based virtual environment facilitate prototyping and testing of alternative digging concepts and can potentially reduce hardware experiment costs. The goal map dimensions used for these simulations is a volume $8 \times 8 \times 1.5$ m deep. The robots modeled in Digital Spaces™ are four-wheeled skid steered (holonomic) systems with two separate drive motors on each side. Each rover can deliver up to 300 W of power. Three different implements, namely front-loader, bulldozer blade and bucket-wheel have been mounted on a common vehicle chassis to make comparisons meaningful (Figure 6).

The simulated lunar terrain is divided into rectangular boxes, representing regolith containers, that rise and fall to reflect interaction with the excavation tool. When the bucket shape in the simulator comes in contact with one of these boxes, a volumetric friction procedure that computes the Balovnev forces acting on the blade is executed [21]. The effects of wheel slippage, sinkage, cutting resistance and lunar gravity are all accounted for in the simulation [22].

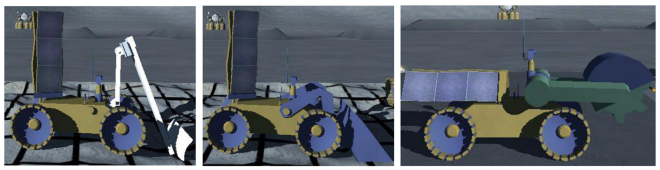

Fig. 6. Digital Spaces™ simulation models of a rover equipped with a bulldozer blade (left), front-loader (center) and bucket-wheel (right).

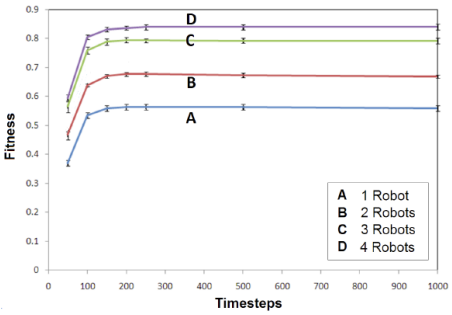

Fig. 7. Fitness performance vs. time of the evolved ANT controllers (bulldozer implement) within the grid world environment.

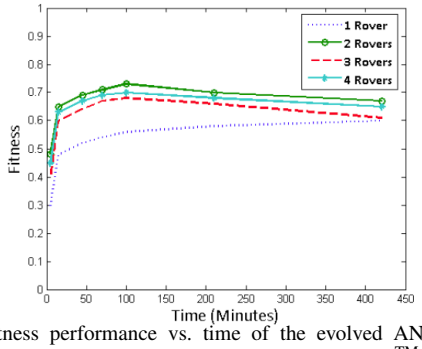

Fig. 8. Fitness performance vs. time of the evolved ANT controllers (bulldozer implement) from Figure 7 using Digital Spaces™.

In Digital Spaces™, it is assumed that an external localization system provides the rovers with x, y, z position and orientation. However, the relative depth of regolith in front of the rover is obtained using simulated infrared range sensors that provides a sampling of depth value within predesignated regions as shown from Figure 3. To ensure consistency between the training simulation and the Digital Spaces™ environment, fitness is monitored for 1, 2, 3, and 4 bulldozer simulations. The results are shown in Figures 7 and 8. The two simulations largely agree but differences should also be noted. It should be observed in the gridworld simulations, the fitness plateaus once the robots reach the goal depth. While within Digital Spaces™, the robots continue to compress and/or scrape off the regolith even after reaching the goal depth. This is because the blade is kept level once the controller reaches the goal depth and, in the Digital Spaces™

simulation environment, slight errors in the contact modeling results in some excavation of the regolith over prolonged periods. This difference can be minimized with more precise tuning of blade position in Digital Spaces™ to match the ideal conditions of the grid world simulation.

V. RESULTS AND DISCUSSION

We first analyze the evolutionary training performance of the ANT controllers using the bulldozer blade implement. Figure 9 shows the fitness (population best) of the system evaluated at each generation of the training process. For the evaluations, the time provided is $l \times w \times h$ timesteps, where l, w, h is the length, width and depth respectively. One notices an improvement in fitness with increased number of robots. With more robots, each robot has a smaller area to cover in performing excavation. The evolutionary runs also shows a point of diminishing returns is reached. Beyond this point, additional robots have a minimal effect on system performance for a constant work area.

The performance of randomly initialized fixed-topology, fully connected networks with up to 40 hidden and output neurons is shown in Figure 10. While this is not intended to be a thorough comparison, in a fixed-topology network there tends to be more 'active' synaptic connections present (since all neurons are active), and thus it takes longer for each neuron to tune these connections to the sensory inputs. ANT is an improvement as the topology is evolved and decision neurons learn to mask out noisy neurons. The net result is that ANT requires fewer genetic evaluations to evolve desired solutions compared to standard neural networks.

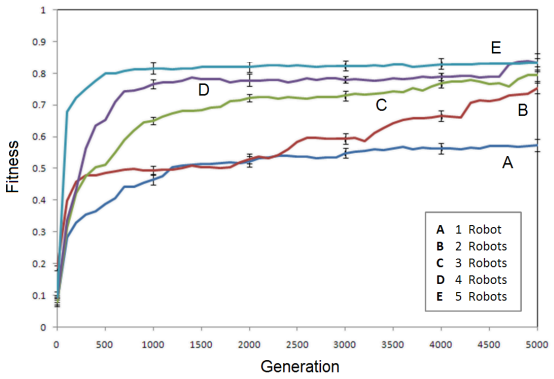

Fig. 9. Evolutionary performance comparison of ANT based solutions (aggregation capable using bulldozer implement) for between 1 and 5 robots, population best averaged over 30 evolutionary runs.

A. Implement Comparison

Figure 11 shows the fitness (population best) of the system evaluated at each generation of training. The time provided for evaluation is $l \times w$ timesteps, where l and w is the length and width respectively. This time is shorter than the previously described experiments that require multiple excavation passes. For comparison, we consider the controller performance with a preset vehicle implement (i.e., bucket-wheel or front-loader) and a situation where the vehicle implement selection is coevolved with the controller. In this latter scenario, an additional parameter within the ANT genome specifies vehicle implement type. Comparison of the preset vehicle implements show that the bucket-wheel option has a substantial performance advantage over the other two options. This advantage is attributed to the fact that the bucket-wheel has a tank to store and carry excavated regolith while the bulldozer does not. At times, it may unavoidable that a bulldozer pass be abandoned due to obstacles obstructing a pathway and thus results in leaving behind excavated regolith. Bulldozer robots require having to excavate an area in a layered fashion to avoid having to accumulate too much regolith and stalling the vehicle. In contrast, with a bucket-wheel, the vehicle can remain stationary and the rotating blade can churn and deposit the regolith into its tank. This enables the vehicle to reach the goal depth without having to do multiple layered passes. These factors simplify control actuation and facilitates finding efficient solutions within fewer genetic evaluations. The front-loader shares some functionality with the bucket-wheel, namely digging and holding regolith. However, the front-loader cannot match the bucket-wheel in excavating to a goal depth while remaining stationary. The comparison indicates that the ability to store regolith after being excavated improves fitness performance. This is expected because otherwise additional time is spent interrupting an excavation pass by dumping the material and then resuming the excavation process. Interestingly, ANT controllers with a gene-specified excavation implement match the performance of the controllers with the bucket-wheel implement after 4,000 generations. This is despite the fact the search space has increased with one additional parameter, namely choice of vehicle implement. The additional parameter allows effective coevolution of the best implement and controller. This result is particularly useful when it is not obvious what excavation implement is optimal for a given task.

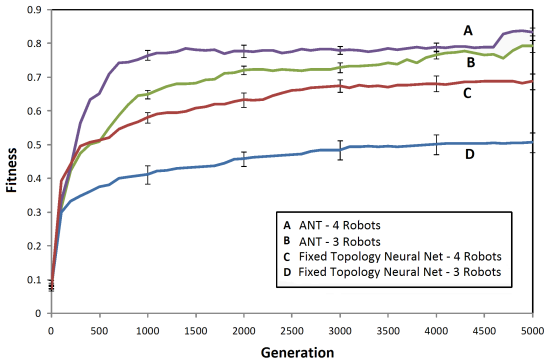

Fig. 10. Evolutionary performance comparison of ANT based solutions with randomly initialized fixed-topology neural networks for 3 and 4 robot configurations, population best averaged over 30 evolutionary runs.

B. Controller Scalability

It is found that with aggregation capability, the system performance improves with increased density of robots per digging area (Figure 12). This is in contrast to our results with aggregation functionality turned off as shown in

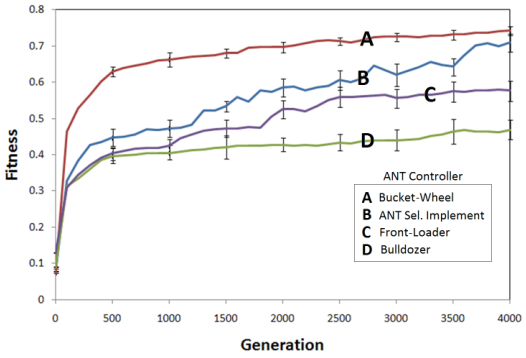

Fig. 11. Evolutionary performance comparison of ANT-based solutions for various excavation implements. The bucket-wheel implement shows significant performance improvement over the other implements. The ANT selected implement also tends towards selecting bucket-wheel in excess of 80% of the time at 4,000 generations.

Figure 13. Beyond a certain robot density, the problem of antagonism reduces the overall system performance. It is evident that the evolutionary search process can find the necessary excavation solutions without explicit *shaping* of the fitness function and furthermore it can fine tune these coordination behaviors depending on evolved robot densities.

The ANT controllers with aggregation functionality (Figure 12) show better performance with increased robot density compared to the ANT controllers without this functionality (Figure 13). However, these aggregation controllers do show decreased performance for lower densities, due impart to acquired traits from training. In this scenario, the robots have acquired skills to work efficiently in larger groups and when faced with fewer robots, these evolved skills are less efficient. Controllers trained at lower densities show a significant decrease in fitness for higher rescaled densities because these controllers lack the experience of using excess robots. For a system designer, it is not obvious as to how many excavation robots are needed to complete the task in the shortest time. Using the aggregation capable rovers, one has the advantage of overestimating the number of robots without facing reduced system performance. Having an excess number of robots allows for the system to better handle erroneous maneuvers by a single individual. As the training proceeds, the population learns to correct erroneous maneuvers, hence making better use of the each robot.

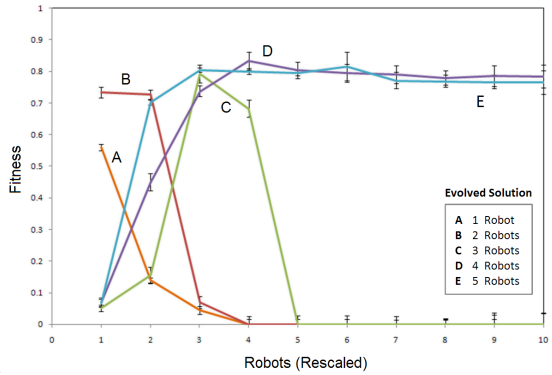

Fig. 12. ANT based solutions evolved from 1 to 5 robots in aggregation mode and evaluated for between 1 and 10 robots (8×8 excavation area, average $1.5d$ cm depth, 100 runs).

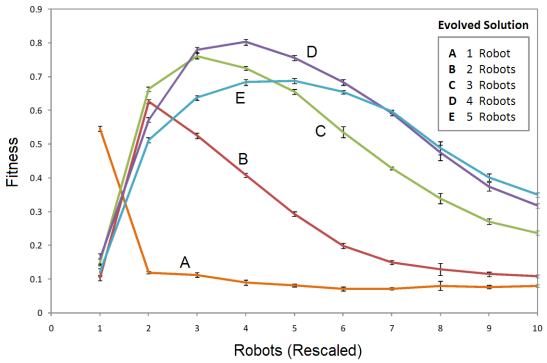

Fig. 13. ANT based solutions evolved from 1 to 5 robots (without aggregation capability) and evaluated for between 1 and 10 robots (8×8 excavation area, average $1.5d$ cm depth, 100 runs)

C. Behavioral Adaptations

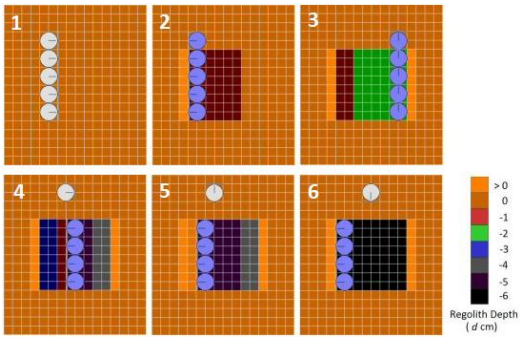

Fig. 14. Simulation snapshots of an excavation task simulation using aggregation capable ANT controllers (5 robots) after 0, 16, 25, 45, 47, 65 timesteps. Shaded robots are in aggregation mode.

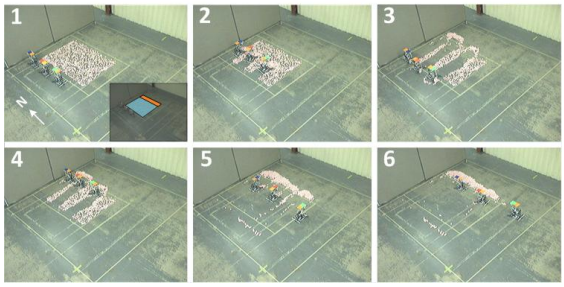

Fig. 15. Snapshots of three robots performing the excavation task using an ANT controller in aggregation mode (14×8 excavation area, dump site on eastern side). Frame 1 (inset) shows the excavation area in light blue and dump area in orange. Frame 3 shows robots turning at the boundary of the excavation area before facing the dump site. Frame 4 shows the robots making a left turn after dumping the peanuts (regolith).

It is found that the aggregation capable ANT controllers reduce the effects of antagonism by permitting an excess number of robots to act in a 'neutral' fashion (Figure 14). These excess robots avoid being in the way of other robots, or detach and remain outside the excavation area. In both cases, the system self-organizes to reach a desired robot

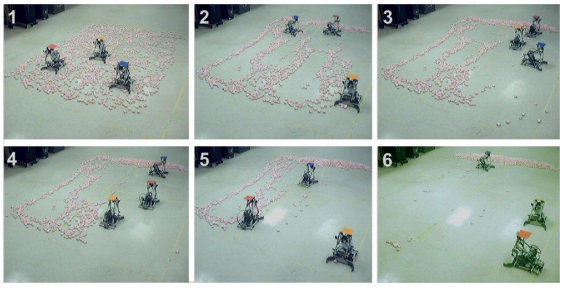

Fig. 16. Snapshots of three robots performing the excavation task using an ANT controller in a random initial position (10×15 excavation area, dump site on eastern side). Frame 2 shows the robots having pushed the styro peanuts (regolith) to the eastern side. Frame 3 shows robots clearing an excavation path parallel to the dump site. Frames 4, 5 shows the robots systematically clearing areas missed from previous excavation passes.

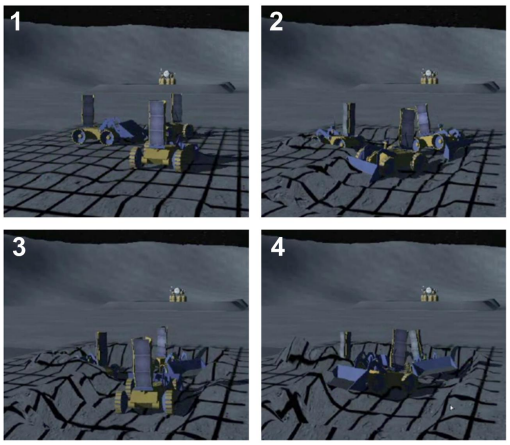

Fig. 17. Digital Spaces^{TM} simulation snapshots showing three randomly positioned front-loader robots performing excavation of a rectangular region (surrounded by dump sites) along with an attached ramp.

density. Communication between the robots occur in a number of stigmergic ways. Similar to the resource gathering experiments, individuals communicate by manipulating the environment. The aggregated robots indirectly communicate via their group movements (by choosing to veto or support aggregated group decisions). The controllers also learn to exploit templates by correctly interpreting the goal maps (excavation blue prints). This is used to determine when to lower, level or raise the blade. Once a robot senses a dump area in front, the controller executes a combination of 'move backward' followed by a 'turn left' or 'turn right' to offload the excavated material. Ongoing experiments using the LEGO® robots (Figure 15, 16) and Digital Space^{TM} simulations (Figures 17) show the robots correctly interpreting a given goal map and dump regolith in the designated areas.

VI. CONCLUSIONS

The paradigm of an "Artificial Neural Tissue" (ANT) framework has been successfully applied towards the development of controllers for a simulated excavation task. These controllers have been developed and initially evaluated in computer simulation using a model of the excavation problem and then tested on hardware in a 2-D laboratory environment. The training of the ANT controllers requires a global fitness function that measures the performance of the controller and a generic set of basis behaviors. Because less knowledge is used a priori in the training and controller development cycle, an ANT architecture may select a suitable excavation implement and discover novel solutions that might otherwise be overlooked by a human supervisor. The ANT controllers have also been compared with fixed-topology neural network architectures and produce fitter solutions in fewer genetic evaluations. It is evident that the evolutionary search process can indeed find efficient excavation solutions using self-organized aggregation behaviors.

REFERENCES

[1] C. Melhuish, J. Welsby, and C. Edwards, "Using templates for defensive wall building with autonomous mobile ant- like robots," in *TIMR99 Toward Intelligent Mobile Robots*, 1999.

[2] M. J. Matarić, M. Nilsson, and K. T. Simsarian, "Cooperative multi-robot box-pushing," in *IEEE/RSJ IROS*, 1995, pp. 556–561.

[3] C. A. Parker, H. Zhang, , and C. R. Kube, "Blind bulldozing: Multiple robot nest construction," in *IEEE/RSJ Int. Conference on Intelligent Robots and Systems*, 2003, pp. 2010–2015.

[4] J. Werfel, Y. Y., and R. Nagpal, "Building patterned structures with robot swarms," in *IJCAI*, 2005, pp. 1495–1502.

[5] J. Thangavelautham and D. G. M. T., "A coarse-coding framework for a gene-regulatory-based artificial neural tissue," in *Proc. of the 8th European Conference on ALife*, 2005, pp. 67–77.

[6] J. Thangavelautham, A. D. S. Smith, D. Boucher, J. Richard, and G. M. T. D'Eleuterio, "Evolving a scalable multirobot controller using an artificial neural tissue paradigm." in *ICRA*. IEEE, 2007, pp. 77–84.

[7] V. W Trianni, R. Gro, T. Labella, E. Sahin, and M. Dorigo, "Evolving aggregation behaviors in a swarm of robots," in *Proc. of the 7th European Conference on Artificial Life*, vol. 2801, 2003, pp. 865–874.

[8] E. Bahceci and E. Sahin, "Evolving aggregation behaviors for swarm robotic systems: A system case study," in *Swarm Intelligence Symposium, 2005*. IEEE, 2005, pp. 333–340.

[9] F. Chantemargue, T. Dagaeff, M. Schumacher, and B. Hirsbrunner, "Implicit cooperation and antagonism in multi-agent systems," University of Fribourg, IIUF-PAI, Internal Report, Tech. Rep., 1996.

[10] A. Stentz, J. Bares, S. Singh, and P. Rowe, "A robotic excavator for autonomous truck loading," in *Proceedings of IEEE/RSJ International Conference on Intelligent Robotic Systems*, vol. 1, 1998, pp. 175–186.

[11] M. Dunbabin and P. Corke, "Autonomous excavation using a rope shovel," in *Journal of Field Robotics*, vol. 23, 2006, pp. 379–394.

[12] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm Intelligence: From Natural to Artificial Systems*. New York: Oxford Univ. Press, 1999.

[13] J. Wawerla, G. Sukhatme, and M. Matari, "Collective construction with multiple robots," in *IEEE/RSJ Int.l Conference on Intelligent Robots and Systems*, 2002, pp. 2696–2701.

[14] P. Grassé, "La reconstruction du nid les coordinations interindividuelles; la theorie de stigmergie," in *Insectes Sociaux*, 1959, pp. 41–84.

[15] R. Beckers, O. E. Holland, and J. L. Deneubourg, "From local actions to global tasks: Stigmergy and collective robots," in *Fourth Int. Workshop on the Syntheses and Simulation of Living Systems*. MIT Press, 1994, pp. 181–189.

[16] J. Thangavelautham, T. Barfoot, and G. M. T. D'Eleuterio, "Coevolving communication and cooperation for lattice formation tasks," in *Proc. of the 7th European Conference on ALife*, 2003, pp. 857–864.

[17] E. Bonabeau, G. Theraulaz, J.-L. Deneubourg, S. Aron, and S. Camazine, "Self-organization in social insects," in *Trends in Ecology and Evolution*, vol. 12, May 1997, pp. 188–193.

[18] C. Kube and H. Zhang, "Collective robotic intelligence." in *Second International Conference on Simulation of Adaptive Behaviour*. MIT Press, 1993, pp. 460–468.

[19] R. Groß, E. Tuci, M. Dorigo, M. Bonani, and F. Mondada, "Object transport by modular robots that self-assemble," in *Int. Conf. on Robotics and Automation*. IEEE, 2006, pp. 2558–2564.

[20] R. Groß and M. Dorigo, "Evolving a cooperative transport behavior for two simple robots," in *Artificial Evolution*, vol. 2936, 2004.

[21] V. Balovnev, "New methods for calculating resistance to cutting of soil," in *Amerind Publishing (Translation)*. TIS, 1983.

[22] N. Abuelsamid, "Lunar excavation methods and evaluation," in *Masters Thesis*. U. of Toronto Institute for Aerospace Studies, 2008.