Multirobot Lunar Excavation and ISRU Using Artificial-Neural-Tissue Controllers

Jekanthan Thangavelautham[1], Alexander Smith[1], Nader Abu El Samid[1], Alexander Ho[1], Dale Boucher[2], Jim Richard[3], Gabriele M. T. D'Eleuterio[1]

[1]Institute for Aerospace Studies, University of Toronto, 4925 Dufferin St., Toronto, ON M3H 5T6
[2]Northern Centre for Advanced Technology, 1400 Barrydowne Rd., Sudbury, ON P3A 3V8
[3]Electric Vehicle Controllers Ltd, 2200 ValleyView Rd., Val Caron, ON P3N 1L1
416-667-7731, thanga@ecf.utoronto.ca

Abstract. Automation of site preparation and resource utilization on the Moon with teams of autonomous robots holds considerable promise for establishing a lunar base. Such multirobot autonomous systems would require limited human support infrastructure, complement necessary manned operations and reduce overall mission risk. We present an Artificial Neural Tissue (ANT) architecture as a control system for autonomous multirobot excavation tasks. An ANT approach requires much less human supervision and pre-programmed human expertise than previous techniques. Only a single global fitness function and a set of allowable basis behaviors need be specified. An evolutionary (Darwinian) selection process is used to 'breed' controllers for the task at hand in simulation and the fittest controllers are transferred onto hardware for further validation and testing. ANT facilitates 'machine creativity', with the emergence of novel functionality through a process of self-organized task decomposition of mission goals. ANT based controllers are shown to exhibit *self-organization*, employ *stigmergy* (communication mediated through the environment) and make use of *templates* (unlabeled environmental cues). With lunar in-situ resource utilization (ISRU) efforts in mind, ANT controllers have been tested on a multirobot excavation task in which teams of robots with no explicit supervision can successfully avoid obstacles, interpret excavation blueprints, perform layered digging, avoid burying or trapping other robots and clear/maintain digging routes.

Keywords: ISRU, collective robotics, neural networks, evolutionary algorithms, developmental systems.
PACS: 45.40.Ln, 84.35.+i, 07.05.Mh.

INTRODUCTION

The use of multipurpose teams of robots for construction of key elements of a human habitat on the moon offers many benefits. It is very likely an imperative due to concerns of health and safety limiting the productivity of human astronauts. These teams of autonomous robots could work continuously in harsh environments making them very productive and appealing for extended duration tasks without the need for teleoperation infrastructure. One has only to observe the effectiveness of a colony of ants excavating tunnels or termites building towering cathedral mounds with internal heating and cooling shafts (Bristow and Holt, 1997) to see why approaches to such tasks in robotics are often biologically inspired and multiagent oriented (Melhuish, Welsby and Edwards, 1999).

An earth-based teleoperation infrastructure would require robust operational procedures with the ability to handle intermittent communication interruptions, bandwidth disruptions/traffic limitations and latency issues. Such a system has been demonstrated successfully with the Lunakhod 1 and 2 rover missions; however, operator fatigue is of concern (Miller and Machulis, 2005) especially when coordinating actions with teams of robots over extended duration missions.

Alternatively, a lunar-based teleoperation system will still require a dedicated human habitat infrastructure for on-site operation. Operation of multiple robots will either require multiple operators or single operator using an automated scheduling system (Mau and Dolan, 2007). A scheduling system would be less efficient than continuous operation since its limited by *saturation* (when a human operator cannot attend to any more robot tasks) in contrast to an autonomous robotic system that is not limited by these concerns. These attributes make an autonomous robotic system more appealing, with the possibility of having a base deployed and operational in time for astronauts to arrive from Earth. There currently exist two major approaches to developing autonomous control systems: human knowledge/model-based controllers and systems based on machine learning techniques. Human knowledge/model-based behavior control strategies rely on human input in the form of ad-hoc control rules, task-specific assumptions, and human experience and knowledge. In contrast, machine learning systems presented here perform task decomposition through 'emergent' self-organization behavior. In lunar and planetary environments, task-specific assumptions may not always be valid in situ and may require operational reassessment during a mission. There is also a growing necessity for the development of generic teams of multipurpose 'utility' robots that can facilitate in-situ resource utilization and perform specific tasks that may never have been envisioned during mission planning and modeling stages.

The approach outlined in this paper involves the application of a machine learning approach called the Artificial Neural Tissue (ANT) framework (Thangavelautham and D'Eleuterio, 2005) to solve a multirobot excavation task, where teams of robots need to excavate a hole in which to bury a nuclear power source. With minimal task-specific assumptions and limited supervision, ANT can produce controllers which can interpret and follow excavation blueprints, can successfully avoid obstacles, perform layered digging, avoid burying or trapping other robots and clear/maintain excavation routes. These innovative behaviors are achieved through the exploitation of unlabeled blueprint cues and templates, *stigmergy* (indirect communication mediated through the environment), and self-organization. Since little preprogrammed knowledge is given, ANT may discover novel solutions. In this paper, we compare performance characteristics of an ANT-based neural network controllers for a multirobot excavation task.

BACKGROUND

Previous work into autonomous excavation (Stentz et al., 1998) has been limited to single robotic excavation platforms and separate loading/unloading vehicles for terrestrial applications. Digging operations is performed through use of human defined automated scripts (behaviors) that simplify repetitive excavation/truck loading cycles. These are a more elaborate form of 'if-then' rules, defined specifically for the task at hand. The scripts are used to position and unload an excavator bucket relative to a dump truck, based on a suite of sensors onboard the vehicles. The scripts are developed based on input from expert human excavator operators and model vehicle specific limitations such as load handling capacity. These scripts also incorporate a coarse and refined planner to sequence digging operations within a localized area. Such systems have been comparable in efficiency to human operated systems. Kinematic modelling-based techniques have been used to automate the digging operations of an electric rope shovel (Dunbabin and Corke, 2006). Sensing of the digging terrain is done using laser ranger sensors. Both systems are explicitly designed for specific vehicle platforms and lack any longer term task planning capability. A control system such as LUCIE (Bradley and Seward, 1998), apart from identifying and automating cyclic excavation related subtasks, incorporates a whole sequence of intermediate goals that need to be achieved to complete a trench digging task. The task is decomposed and prioritized by a human operator and the controller attends to automatic sequencing of subtasks to achieve each intermediate goal.

Autonomous collective robotic tasks typically employ some of the same mechanisms used by social insects. These include the use of templates, stigmergy, and self-organization. *Templates* are environmental features perceptible to the individuals within the collective (Bonabeau, Dorigo and Thereaulaz, 1999). In robotic applications, template-based approaches include the use of light fields to direct the creation of circular (Stewart and Russell, 2003) and linear walls (Wawerla, Sukhame and Mataric, 2002) and planar annulus structures (Wilson et al., 2004). *Stigmergy* is a form of indirect communication mediated through the environment (Grasse, 1959). Stigmergy has been used extensively in collective-robotic construction tasks, including blind bull dozing (Parker, Zhang and Kube, 2003), box pushing (Mataric et al., 1995), heap formation (Beckers, Holland and Deneubourg, 1994) and tiling pattern formation (Thangavelautham, Barfoot and D'Eleuterio, 2003). *Self-organization* describes how local or microscopic behaviors give rise to a macroscopic structure in systems which are not in equilibrium (Bonabeau et al., 1997). With decentralized self-organized systems, no individual possesses knowledge of the overall environment or the end goal. Individuals merely react to local sensor data. Our approach in this paper uses a Darwinian selection process to evolve robotic controllers for performing a desired task. The evolved controllers make use of templates and stigmergy in order to achieve the level of self-organization necessary to achieve the global goals.

The collective robotic and autonomous excavation works cited earlier excluding (Thangavelautham, Barfoot and D'Eleuterio, 2003) rely on either user-defined, deterministic 'if-then' rules, or on stochastic behaviors. In both cases, designing these controllers is an ad-hoc procedure 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 even be counterintuitive. Thus, at this stage of the field's development at least, designing successful controllers by hand is a process of trial and error. With autonomous excavation, only one vehicle is considered and thus the impact of digging behaviors when introducing multiple interacting vehicles is also bound to the same problems encountered for collective robotic tasks.

One approach to reducing the amount of trial and error done by hand is to encode controllers as behavioral look-up tables, and allow a genetic algorithm to evolve the values in the table. This approach was demonstrated for collective heap formation (Barfoot and D'Eleuterio, 1999) and 2×2 tiling pattern formation tasks (Thangavelautham, Barfoot and D'Eleuterio, 2003). The limitations to this approach are poor sensor scalability and lack of generalization. An increased number of sensors leads to a combinatorial explosion in the size of the look-up table, resulting in premature search stagnation (the 'bootstrap' problem). For excavation, poor sensor scalability imposes severe constraints on the choice of digging vehicle and number of sensors allowed. As an action must be encoded for each combination of sensor inputs, the controller does not generalize from one state to another one with similar inputs. Neural-network controllers can often overcome this second limitation by effectively implementing a compressed representation of the problem space. A neural-network controller was able to solve the harder 3×3 tiling formation task (Thangavelautham and D'Eleuterio, 2004). Other fixed-topology neural controller approaches have been used to build walls, corridors and briar patches (Crabbe and Dyer, 1999).

Fixed-topology neural networks present an additional problem: The size and structure of the network must be fixed ahead of time. Inappropriate choices may lead to a network that is unable to solve the problem. The ANT framework is able to overcome this problem. This variable-length neurocontroller model allows for the generalization of sensory input, for improved scalability over fixed-network topologies, and for both stochastic and deterministic arbitration schemes. More importantly for excavation, ANT allows for both sensor and behavior extensibility and is not constrained to a specific digging platform.

ARTIFICIAL NEURAL TISSUE MODEL

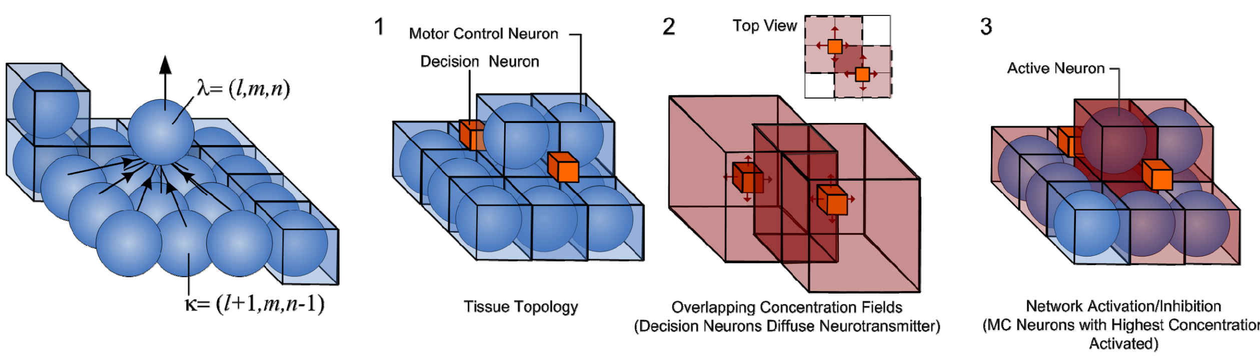

(a) Synaptic Connections. **(b)** Coarse Coding.

FIGURE 1. (a) Synaptic Connections between Motor Neurons from Layer l+1 to l. (b) Activated Decision Neurons Diffuse Neurotransmitter Concentration Field Resulting in Activation of Motor Control Neurons with Highest Activation Concentration.

The ANT architecture (Figure 1) used 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, or simply motor neurons. Regulation is performed by the decision neurons that dynamically exhibit or inhibit motor-control neurons within the tissue based on a coarse-coding framework. Further details on ANT can be found in Thangavelautham et al. (2007).

SIMULATION EXPERIMENT SETUP

The excavation task is intended to demonstrate the feasibility of team of robots in digging a hole to bury a nuclear reactor into lunar regolith. This is intended as the first step in setting up nuclear power source for a lunar base. It could be argued that *emergent* task decomposition may be necessary to accomplish the task given a global fitness function. A 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 the size of four squares in the grid being just able to accommodate one robot. For this task, the controller need to accomplish a number of subtasks including interpret excavation blue prints, perform layered digging, avoid burying or trapping other robots and clear/maintain excavation routes. Each robot controller has access to a goal map that defines the location of the dumping area and target depth of the excavation area. The global fitness function f for the task is given as follows:

$$f = \frac{\sum_{j=1}^{J} \sum_{i=1}^{I} p_{i,j} \cdot e^{-2\left|g_{i,j} - h_{i,j}\right|}}{\sum_{j=1}^{J} \sum_{i=1}^{I} p_{i,j}} \tag{1}$$

I and J are the dimensions of the workspace, where $\sum_{j=1}^{J} \sum_{i=1}^{I} p_{i,j} > 0$ and $p_{i,j} = 1$ 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. For the evolutionary runs, fitness f is calculated after T=100 timesteps, for an excavation area of 8×8 squares surrounded by the dumping area.

TABLE 1. Sensor Inputs.

Variables	Function	States
$Z_1 \ldots Z_3$	Depth Sensing relative to Goal Depth	Level, Above, Below, Don't Care
FZ_1, FZ_2	Depth Sensing relative to Ground	Above, Below or Level
B_1	Blade Position	Above, Level, Below ground, Home
BL_1	Blade Load	0-4
S_1	Front Obstacle Detection	Obstacle, No Obstacle
D_1	Separation Distance From Nearest Robot	0-3
H_1	Heading from nearest robot	North, East, West, South
TL_1	Robot Tilted Downwards	True, False
ST_1	Robot Stuck	True, False

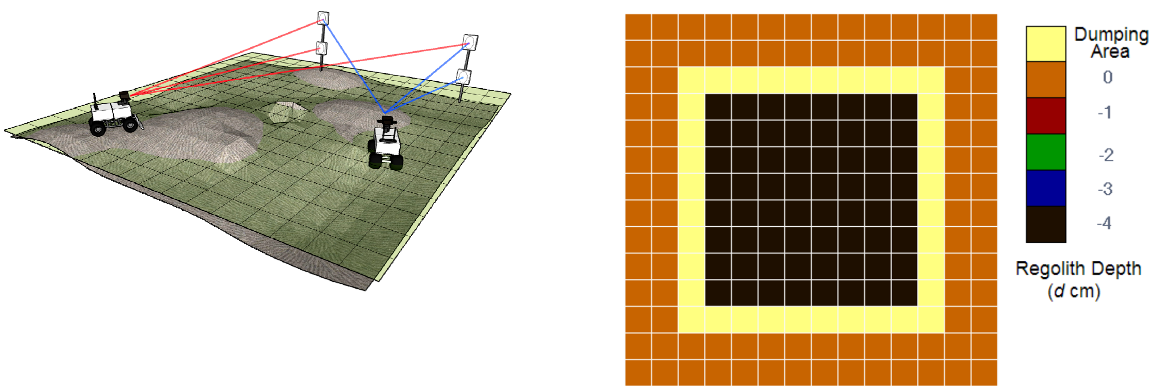

(a) Localization **(b)** Goal Map

FIGURE 2. Localization setup for hardware demonstrations and an example goal map.

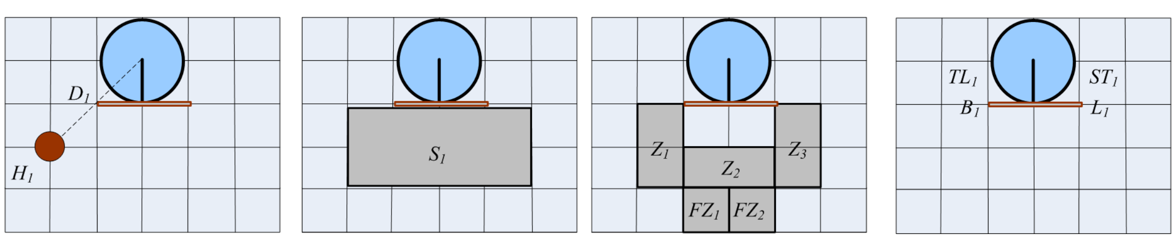

(a) Virtual Radar **(b)** Sonar **(c)** Depth Sensing **(d)** Blade, Motion Sensing

FIGURE 3. Robot Input Sensor Mapping for the Simulation Model.

TABLE 2. Basis Behaviors.

Order	Behavior	Description
1	Set Throttle	Set rover throttle to high otherwise remain nominal/
2	Move Forwards	Move one grid square forwards.
3	Move Backwards	Move one grid square backward.
4	Random Turn	Randomly Turn 90^{o} right or Turn 90^{o} left.
5	Turn Right	Turn 90^{o} right.
6	Turn Left	Turn 90^{o} left.
7	Blade position: Above	Set blade above ground *d* cm
8	Blade position: Below	Set blade below ground *d* cm
9	Blade position: Level	Set blade level to ground
10	Blade position: Home	Retract blade to home position (makes no contact with regolith)
11	Bit Set	Set memory bit 1 to 1
12	Bit Clear	Set memory bit 1 to 0

For this task, inputs to the ANT controller are shown in Table 1 (right). The robots have access to a pair of webcams and a laser range finder mounted on a pan tilt unit. The webcams are used to detect and center the pan tilt units on up to four localization targets consisting of ellipses with contrasting patterns (Figure 2a). A laser range finder is used to measure the distance from each target to robot and triangulation is performed to determine (x, y, z) coordinates of the robot within the target workspace. The discretized x and y coordinates are use to lookup the goal depth $g_{x,y}$ of each grid square region in front of the robot (Figure 2b). The ground topology is discretized into d cm increments, where d is dependent on blade height and actuation parameters. For use on the Argo class rovers d is set to 1 cm. In addition, a laser scan of the ground in front of the robot is performed to measure depth of regolith in relation to the wheel depth to determine $Z_1..Z_3$ and $FZ_1..FZ_2$ (Figure 3c). The current x and y positions of each robot is used to determine the relative position of the nearest robot and its relative heading through map sharing. All raw input data are discretized. The sonar sensors are used to determine the values of S_1. A pair of load cells on the blade is used to determine BL_1 and onboard tilt sensor is used to determine TL_1. Frame differencing of two consecutive webcam images of the ground is used to determine ST_1 (whether the rover is stuck or not). The robots also have access to one memory bit, which can be manipulated using some of the basis behaviors.

Table 2 lists the basis behaviors the robot can perform (in order) within a single timestep. Darwinian selection is performed based on the fitness value of each controller averaged over 50 different initial conditions, within an 8×8 excavation area (Figure 2b). The Evolutionary Algorithm population size for the experiments is $P = 100$, crossover probability $p_c = 0.7$, mutation probability $p_m = 0.025$ and tournament size of 0.06 P.

RESULTS AND DISCUSSION

Figure 4 shows the fitness (population best) of the overall system evaluated at each generation of the artificial evolutionary process. It is apparent that the system performance is affeced by the density of robots per digging area (8×8 squares). A single robot is not as efficient as 4 robots, as the excavation can be performed in parallel, with each robot having a smaller area to cover. However, with more than 4 robots, the problem of *antagonism* arises, when multiple robots trying to perform the same task interfere with one another and reduce the overall efficiency of the group (see Figures 6, 7). The emergent solutions indicate that the individual robot exploits *templates* by learning to sense depth relative to the specified goal map and determine whether to lower, level or raise the blade (Figure 12). Once the robot senses a dumping area in front, the controller executes are combination of 'move backward' followed by a 'turn left' or 'turn right' to offload the excavated material. As with other multiagent systems, communication between robots occurs through manipulation of the environment in the form of *stigmergy*. The manipulation of the environment involves either excavating a region or dumping off excavated material.

Controllers also exploit the ability to sense the depth of soil relative to wheel depth (see Figure 3c). This enables each robot controller to sense whether it is excavating deeper or backfilling at the current depth. The ability to backfill, although useful in some situations, can also undo the effort of other robots excavating at different depths within the system. The robots also have the ability to sense the relative position of a nearby robot much like a virtual form of radar (Figure 3a). This sensing capability appears to be exploited particularly to avoid collisions when a series of output behaviors such 'move forward' and 'turn left' is applied in sequence. Although obstacles can be detected using sonars directly in front of the robot, there exists blind spots to the extreme right and left making it difficult to detect and react to obstacles when a sequences of behaviors are executed.

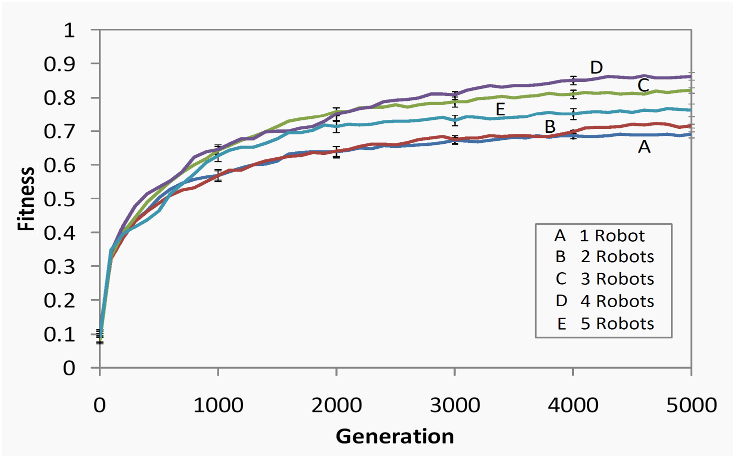

FIGURE 4. Evolutionary Performance Comparison of ANT Based Solutions for between 1 and 5 Robots.

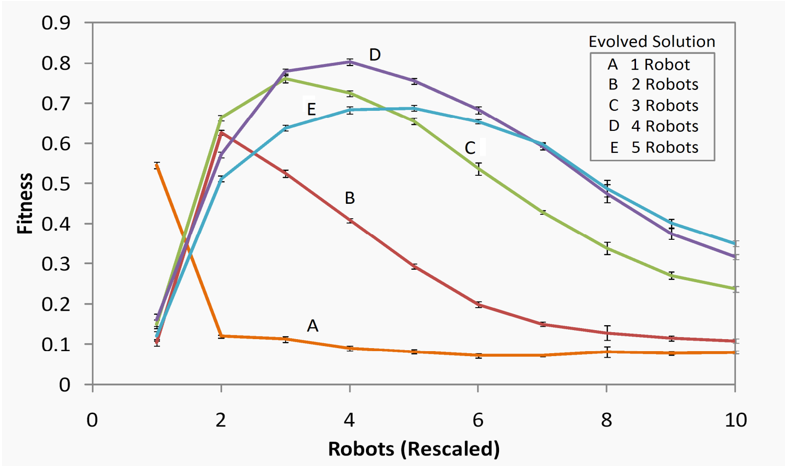

FIGURE 5. Scaling of ANT based Solutions from 1 to 5 robots (8×8 excavation area, average 1.5*d* cm depth).

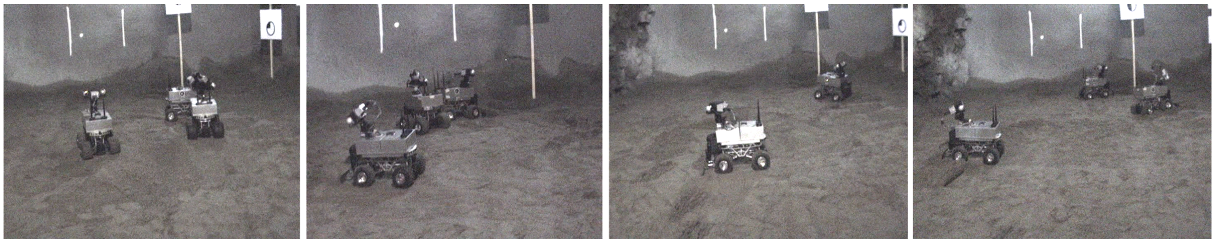

(a) Frame 1. **(b)** Frame 2. **(c)** Frame 3. **(d)** Frame 4.

FIGURE 6. Movie frames of Argo Rover with ANT Controller Performing an Excavation Task Conducted at the Planetary and Mining Sciences Symposium, Sudbury, Canada (June, 2007).

We examine the fittest solutions from the simulation runs shown in Figure 4 for scalability in the number of robots while holding the size of the digging area constant (Figure 5). Taking the controller evolved for a single robot and running it on a multirobot system shows a steep degradation in performance. This is expected since the single-robot controller lacks the cooperative behavior necessary to function well within a multirobot setting, showing similarity to the resource gathering task (Thangavelautham *et al.*, 2007). For example, such controllers fail to develop ‘robot collision avoidance’ behaviors. Similarly, a multirobot system scaled down to a single robot setting also shows a degradation in system performance. With the multirobot systems, controllers have evolved to exploit and depend on cooperative actions to complete the task; thus when the environment is abruptly changed to exclude such a possibility the controller performs poorly.

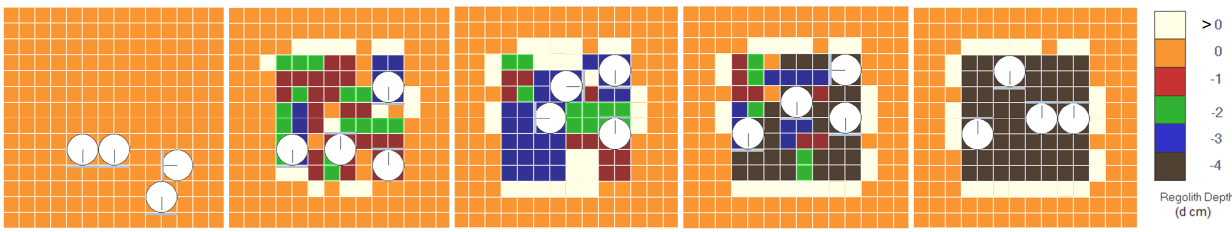

(a) Timestep 0. **(b)** Timestep 50. **(c)** Timestep 75. **(d)** Timestep 100. **(e)** Timestep 170.

FIGURE 7. (a)-(e) Snapshots of an a Excavation Task Simulation (4 robots).

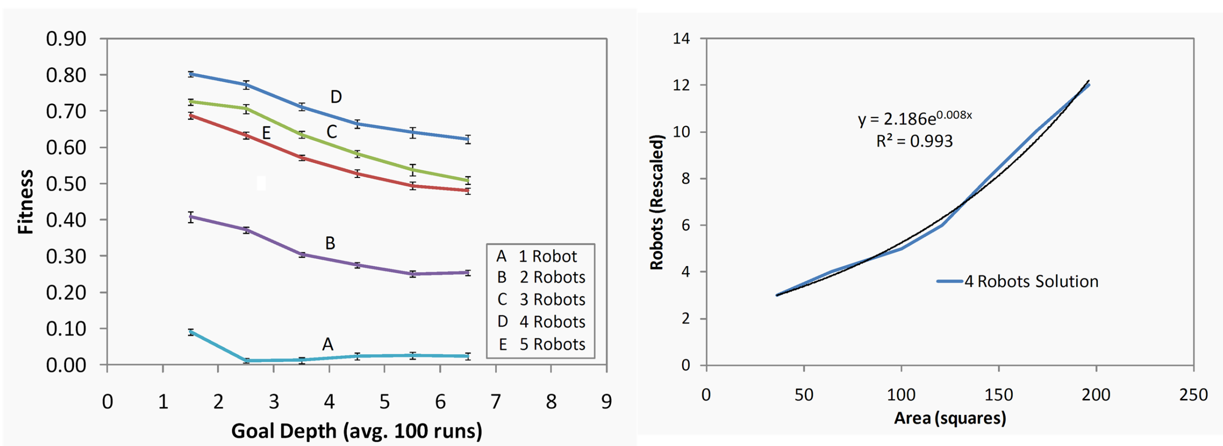

(a) Varying Excavation Depth **(b)** Varying Excavation Area

FIGURE 8. Scaling of ANT based solutions for varying depth (a) and 4 robot solution for varying excavation area (b).

It is interesting that the controllers trained with 4 robots for an 8×8 digging area perform considerably better overall than solutions trained for other densities. It is also apparent that this solution scales better for increased goal depth (see Figure 8a) and was found to scale better than other solutions for increased excavation area. The optimal ratio (best fit) of robots to digging area using solution trained with 4 robots is shown in Figure 8b. However, when the number of robots is set higher than the optimal number, solutions trained under higher robot densities perform better, although the system performance falls short compared to the optimal setting (Figure 5). It is apparent from these simulation experiments that there exists an optimal set of training conditions under which solutions show improved scalability. Although the controllers maybe better adapted to antagonism under higher training densities with improved obstacle avoidance techniques, these behaviors may not be as well tuned to completing the overall objectives (excavation) effectively.

CONCLUSIONS

A developmental Artificial Neural Tissue (ANT) architecture has been demonstrated for a multirobot excavation task in support of lunar in-situ resource utilization efforts. ANT controllers require only a global fitness function that merely measures the performance of the controller for a given task and a generic set of basis behaviors. ANT controllers are shown to exploit a number of mechanisms known to be used in multiagent systems in unsupervised manner, including environmental templates, stigmergy and self-organization. ANT is shown to generalize and interpret user defined excavation blueprints. Since little preprogrammed knowledge is given, an ANT architecture may permit novel solutions that might otherwise be overlooked by a human supervisor. The scalability of the ANT controllers under varying training conditions indicate existence of optimal set of training parameters. The controllers are shown to be scaleable for both increased excavation depth and area.

ACKNOWLEDGMENTS

This work was partially funded by the Natural Sciences and Engineering Research Council of Canada. We also wish to thank the reviewers for their helpful comments.

REFERENCES

Barfoot, T. D. and D'Eleuterio G.M.T., "An Evolutionary Approach to Multiagent Heap Formation," in proceedings of *IEEE Congress on Evolutionary Computation (CEC)*, edited by Angeline et al., IEEE, Washington DC, 1999, pp. 427-435.

Beckers, R., O. E. Holland, and J. L. Deneubourg, "From Local actions to Global Tasks: Stigmergy and Collective Robots," in *Artificial Life IV*, edited R. Brooks, P. Maes, MIT Press, Cambridge, MA, 1994, pp. 181-189.

Bonabeau, E., et al., "Self-organization in social insects," *Trends in Ecology and Evolution*, **12**, 188-193, (1997).

Bonabeau, E., Dorigo, M., and Theraulaz G., *Swarm Intelligence: From Natural to Artificial Systems*, Oxford University Press, New York, NY, 1999, pp. 183-205.

Bradley, D. A., and Seward, D., "The Development, Control and Operation of an Autonomous Robotic Excavator," *Journal of Intelligent and Robotic Systems*, **21**, 73-97, (1998).

Bristow, K. L. and Holt, J. A. "Can Termites Create Local Energy Sinks to Regulate Mound Temperature?" *Journal of Thermal Biology*, **12**, 19-21, (1997).

Crabbe, F., Dyer, M., "Second-order networks for wall-building agents," in proceedings of *IEEE International Joint Conference on Neural Networks*, IEEE, Washington DC, 1999, pp. 2178– 2183.

Dunbabin, M., and Corke, P., "Autonomous Excavation using a Rope Shovel," *Journal of Field Robotics*, **23**(6/7), 379-394 (2006).

Grasse, P., "La reconstruction du nid les coordinations interindividuelles; la theorie de stigmergie,'' *Insectes Sociaux*, **35**, 41–84 (1959).

Mau, S., and Dolan, J., "Scheduling for Humans in Multirobot Supervisory Control," in proceedings of *IEEE/RSJ International Conference on Intelligent Robotic Systems*, edited by K. Balavanis et al., IEEE, Washington DC, 2007.

Mataric, M. et al. "Cooperative Multi-Robot Box Pushing." in proceedings of *IEEE International Conference on Intelligent Robots and Systems*, edited by S. Tzafestas et al., IEEE, Washington DC, 1995, pp 556-561.

Melhuish, C., Welsby, J. and Edwards, C., *Using Templates for Defensive Wall Building with Autonomous Mobile Ant-Like Robots*, Manchester Tech. Report/UMCS-99-3-1, Univ. of Manchester, Manchester, UK, 1999.

Miller, D. and Machulis, K., "Visual Aids for Lunar Rover Tele-operation," in proceedings of *8th International Symposium on Artificial Intelligence, Robotics and Automation in Space*, edited by R. Battrick, ESA Publishing, Netherlands, 2005.

Parker, C., Zhang, H., and Kube, R., "Blind Bulldozing: Multiple Robot Nest Construction," in proceedings of *IEEE International Conference on Intelligent Robots and Systems*, edited by S. Tzafestas et al., IEEE, Washington DC, 2003, pp. 2010- 2015.

Stentz, A., Bares, J., Singh, S., and Rowe, P., "A Robotic Excavator for Autonomous Truck Loading," in proceedings of *IEEE/RSJ International Conference on Intelligent Robotic Systems*, edited by S. Tzafestas et al., IEEE, Washington DC, 1998, pp. 175- 186.

Stewart, R. and Russell, A., "Emergent structures built by a minimalist autonomous robot using a swarm-inspired template mechanism," in proceedings of *The First Australian Conference on ALife (ACAL2003)*, Springer, Berlin, 2003, pp. 216-230.

Thangavelautham, J. and D'Eleuterio G. M. T., "A Neuroevolutionary Approach to Emergent Task Decomposition" in proceedings of *8th Parallel Problem Solving from Nature Conference*, edited by X. Yao et al., Springer, Berlin, 2004, pp. 991-1000.

Thangavelautham, J., Barfoot, T.D. and D'Eleuterio G.M.T., "Coevolving Communication and Cooperation for Lattice Formation Tasks (Updated)," in *Advances In Artificial Life: Proceedings of the 7th European Conference on ALife*, edited by W. Banzhaf et al., Springer, Berlin, 2003, pp. 857-864.

Thangavelautham, J. and D'Eleuterio G. M. T., "A Coarse-Coding Framework for a Gene-Regulatory-Based Artificial Neural Tissue" in *Advances In Artificial Life: Proceedings of the 8th European Conference on ALife*, edited by M. Capcarrere et al., Springer, Berlin, 2005, pp. 67-77.

Thangavelautham, J., Smith, A., Boucher, D., Richard, J. and D'Eleuterio, G. M. T., "Evolving a Scalable Multirobot Controller Using an Artificial Neural Tissue Paradigm," in *Proceedings of IEEE International Conference on Robotics and Automation*, edited by S. Hutchinson et al., IEEE, Washington DC, 2007, pp. 77-84.

Wawerla, J., Sukhatme, G. and Mataric, M., "Collective construction with multiple robots," in proceedings of *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'02)*, edited by S. Tzafestas et al., IEEE, Washington DC, 2002, pp. 2696-2701.

Wilson, M., Melhuish, C., Franks, A. and Scholes, S., "Algorithms for building annular structures with minimalist robots inspired by brood sorting in ant colonies," *Autonomous Robots*, **17**, 115–136 (2004).