Proceedings of DSCC2008
2008 ASME Dynamic Systems and Control Conference
October 20-22, 2008, Ann Arbor, Michigan, USA

DSCC2008-2284

RESILIENT SELF-REPRODUCING SYSTEMS

Amor A. Menezes*
Department of Aerospace Engineering
University of Michigan
Ann Arbor, Michigan 48109
Email: amenezes@umich.edu

Pierre T. Kabamba
Department of Aerospace Engineering
University of Michigan
Ann Arbor, Michigan 48109
Email: kabamba@umich.edu

ABSTRACT

This paper is motivated by the need to minimize the payload mass required to establish an extraterrestrial robotic colony. One approach for this minimization is to deploy a colony consisting of individual robots capable of self-reproducing. An important consideration once such a colony is established is its resiliency to large-scale environment or state variations. Previous approaches to learning and adaptation in self-reconfigurable robots have utilized reinforcement learning, cellular automata, and distributed control schemes to achieve robust handling of failure modes at the modular level. This work considers self-reconfigurability at the system level, where each constituent robot is endowed with a self-reproductive capacity. Rather than focus on individual dynamics, the hypothesis is that resiliency in a collective may be achieved if: 1) individual robots are free to explore all options in their decision space, including self-reproduction, and 2) they dwell preferentially on the most favorable options. Through simulations, we demonstrate that a colony operating in accordance with this hypothesis is able to adapt to changes in the external environment, respond rapidly to applied disturbances and disruptions to the internal system states, and operate in the presence of uncertainty.

INTRODUCTION

Recent scientific research in self-reproduction has raised the prospect of advances in such diverse areas as space colonization, bioengineering, evolutionary software and autonomous manufacturing. Inspired by the work of John von Neumann [1], extensive study of self-reproducing systems has taken place, including cellular automata, computer programs, kinematic machines, molecular machines, and robotic colonies. A comprehensive overview of the field is documented in [2] and [3].

Within the context of extraterrestrial colonization, current phased approaches to Martian exploration see the development of an enduring robotic presence on the Moon in the next five years. Several space agency roadmaps, of which [4] is typical, suggest that individual countries will deploy advanced robots as-needed to expand the size of an established colony. It is well known, however, that for every unit mass of payload to be launched into space, eighty additional units of mass are required to be launched as well [5] – hence the motivation to endow robots with the capacity for self-reproduction. These machines would be able to utilize on-site resources to enlarge their numbers when deemed necessary for a given task. Extraterrestrial systems with such capability are less dependent than traditional colonies on the fiscal constraints of multiple launches of robots. Self-reproduction may therefore provide a highly cost-effective solution to the problem of establishing extraterrestrial colonies.

The remainder of this section presents a rationale for investigating technology that enables an extraterrestrial self-reproducing system to be resilient to changes in the environment. In the next section we highlight theory that enables a novel solution to the resiliency problem. The following sections deal with system modeling, validation, and the simulated results of cooperative reproduction to achieve a self-reconfigurable system. The last section presents conclusions.

Motivation

In a landmark conceptual study on a self-replicating lunar factory [6], a system that included paving, mining, casting, and mobile assembly and repair robots was proposed. Inspired by this work, [7] suggested a factory system comprising self-replicating multi-functional robots that could mine and transport materials and components within a lunar manufacturing facility.

*Address all correspondence to this author.

The work also demonstrated the feasibility of a self-replicating robot with a prototype made of LEGO Mindstorms components. At the same time (and in the years since), a number of researchers have developed modular self-replicating and/or self-reconfigurable robots (see, for instance, [8–13]). A current survey of the state of the art and the challenges facing modular, self-reconfigurable robot systems is given in [14], and a vast listing of papers in the field is available online at [15]. As these references and those therein indicate, the focus has shifted to provable control of the modules of a single self-reconfigurable robot – the realization of various topologies, efficient and distributed control of a large number of modules, recovery from module failures, and even module self-repair [16, 17]. Approaches for local control include reinforcement learning [18], cellular automata [19], and hormone-inspired swarming for self-organization [20].

By virtue of the harsh environment an extraterrestrial robotic colony operates in, the robots in [6] and [7] would need to learn, adapt, and possibly evolve to be tolerant of external disturbances that can affect the collective's overall goals. Hence, this paper examines the performance of a system consisting of multiple robots, where each individual robot is capable of self-reproduction. This self-reproduction is a distinguishing feature that the modules in the robotic systems [8–13] are incapable of.

Robustness in complex systems has been previously studied, using the Highly Optimized Tolerance conceptual framework for example [21,22]. While these references document complex systems that are generally robust, an inescapable characteristic of the systems is their fragile nature, in that small disturbances can cause catastrophic cascading failures. However, there are numerous instances of autonomous robustness as well as resiliency to small and large environment fluctuations in complex natural systems. Examples include physiological regulation in multicellular organisms; group regulation in colonies of social insects; the evolution of species through adaptation and natural selection; and the rebounding of complex systems from earthquakes, tsunamis, hurricanes, asteroid strikes, etc. The apparent lack of resiliency in a robust complex system, as well as the additional capacity for self-reproduction, motivate this work.

Thus, given an extraterrestrial self-reproducing system with defined objectives (e.g., the mining of resources), the goal of this paper is to address open questions on: the adaptation to changes in the external environment; the rapid response to applied disturbances and disruptions to the internal system states; and, the operation of the collective in the presence of uncertainty.

THEORETICAL BACKGROUND

The theoretical framework of this paper is Generation Theory [23], and the Theory of Rational Behavior [24].

Generation Theory

Generation Theory [23] formalizes self-reproduction by "machines," a term describing any entity that is capable of producing an offspring regardless of its physical nature. These machines utilize resources to self-reproduce. A selected resource is manipulated by the parent machine via an embedded generation action to produce an outcome, which itself may or may not be a machine. Thus, we can state the following:

Definition 1. *A* generation system *is a quadruple* $\Gamma = (U, M, R, G)$*, where*

> U *is a* universal set *that contains machines, resources and outcomes of attempts at self-reproduction;*
> $M \subseteq U$ *is a* set of machines*;*
> $R \subseteq U$ *is a* set of resources *that can be utilized for self-reproduction; and,*
> $G : M \times R \to U$ *is a* generation function *that maps a machine and a resource into an outcome in the universal set.*

It is possible that $M \cap R \neq \emptyset$, and also $M \cup R \neq U$, as illustrated in Fig. 1. The former implies that machines can belong to the set of resources, and the latter states that outcomes of attempts at generation may be neither machines nor resources.

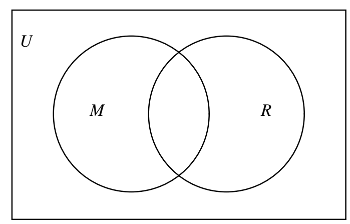

Figure 1. PICTORIAL REPRESENTATION OF DEFINITION 1.

When a machine $x \in M$ processes a resource $r \in R$ to generate an outcome $y \in U$, we write:

$$y = G(x, r). \tag{1}$$

In (1), we say that "x is capable of generating y," and call the process *reproduction*. If we have $x = G(x, r)$ then we say that "x is capable of generating itself," and call the process *replication*.

We also make use of concepts from graph theory [25] in this paper. Equation (1) may be represented by a *directed reproduction graph*, γ, as shown in Fig. 2. In this diagram, machine x and outcome y are vertices, resource r is an edge, and the direction of the edge indicates that it is machine x that uses resource r to generate outcome y.

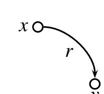

Figure 2. THE DIRECTED REPRODUCTION GRAPH OF (1).

Definition 2. *The* directed graph representation *of a generation system $\Gamma = (U, M, R, G)$ is the directed supergraph (V, E) containing all directed reproduction graphs that produce machines in M. Thus the vertex set, V, of the supergraph is equal to the machine set, M, and the edge set, E, of the supergraph is the set $\{r \in R \mid \exists\, x, y \in M : y = G(x, r)\}$.*

Hence, depictions of generation systems only show the machines that are produced rather than all possible outcomes.

Theory of Rational Behavior

The Theory of Rational Behavior [24] seeks to explain a remarkable property of the collectives that appear in nature. These collectives, which have different fractions of professions as in beehives for example, maintain an appropriate fractional distribution among the various social functions even if one of the castes is removed. The theory proposes an axiomatic behavior of individuals, examines the resulting behavior of a collective, and identifies the properties of systems of many elements.

Individual Behavior.

Rational Behavior first introduces a metric phase space X, which is the space of all possible decisions. It considers the decision-to-decision process as a dynamical system $x_{\phi,N}(x_0, t_0, t)$, $x \in X$, where x_0 is the initial decision at t_0, the initial time, t represents time, $\phi(x)$ is a scalar functional parameter, and N is an integer parameter, $N \in \{1, 2, \ldots\}$. We have $x_{\phi,N}(x_0, t_0, t_0) = x_0$. Also introduced in X is a continuous positive normalized measure μ. Let Φ be a class of positive scalar functions $\phi(x)$, $x \in X$, such that

(a) $\phi(x)$ is bounded and measurable;
(b) every nonzero Lebesgue measure neighborhood $A_0(\phi_0)$ [mes $A_0(\phi_0) > 0$] of an arbitrary value ϕ_0 of a function ϕ is the image of $B_0 \subset X$ also of nonzero measure ($\mu B_0 > 0$).

Finally, let B be an arbitrary μ-measurable set in X ($\mu B > 0$) and $B^i \subset X$ be the pre-image of a sufficiently small neighborhood of the value ϕ^i of the function $\phi(x)$.

Definition 3. *The behavior of an individual taking decisions $x_{\phi,N}(x_0, t_0, t)$, $\phi \in \Phi$, $N \in \{1, 2, \ldots\}$, in the space X is said to be* rational *if the following two properties hold:*

Ergodicity: *For almost all (in the measure μ) points $x_0 \in X$,*

$$\lim_{T\to\infty} \frac{1}{T} \int_0^T \chi_B\left(x_{\phi,N}(x_0, t_0, t)\right) dt = \Psi_{B,\phi}(N) > 0,$$

where χ_B is the characteristic function of the set B:

$$\chi_B(z) = \begin{cases} 1, & \mathit{if}\; z \in B, \\ 0, & \mathit{if}\; z \notin B. \end{cases}$$

Selectivity: *For almost all (in the measure μ) points $x_0 \in X$,*

$$\frac{\lim_{T\to\infty} \frac{1}{T}\int_0^T \chi_{B^i}\left(x_{\phi,N}(x_0, t_0, t)\right) dt}{\lim_{T\to\infty} \frac{1}{T}\int_0^T \chi_{B^j}\left(x_{\phi,N}(x_0, t_0, t)\right) dt} = \pi_{B^i,B^j,\phi}(N)$$

$$\overset{N\to\infty}{\to} \begin{cases} \infty, & \mathit{if}\; \phi^i < \phi^j, \\ \mu B^i / \mu B^j, & \mathit{if}\; \phi^i = \phi^j, \\ 0, & \mathit{if}\; \phi^i > \phi^j, \end{cases}$$

$$B^i \cap B^j = \emptyset.$$

As a result of ergodicity, the individual explores the whole decision space. From selectivity, we see that $\phi(x)$ may be regarded as a penalty function, the consequence of taking decision x. In addition, less penalized decisions are pursued as N gets larger. Thus, N may be regarded as an indication of the "level of selectivity" of an element, with more selective individuals possessing a higher value of N.

Collective Behavior.

To extend Rational Behavior to a collection of individuals, the penalty function is modified to depend on the collective's behavior in addition to an individual's behavior. As a result, the penalty for the i-th individual is $\phi(x_1, \ldots, x_i, \ldots, x_M)$, where M is the number of individuals in the collective. A brief summary of the collective's behavior modeled under *Homogeneous Fractional Interaction* follows.

Suppose that there are two options x_1 and x_2 in the decision space X, and at the initial time t_0, $m(t_0)$ individuals are in x_1. Thus, $M - m(t_0)$ individuals are in x_2. Define the fraction of individuals in x_1 as $\nu = m(t_0)/M$, $\nu \in [0, 1]$. The fraction of individuals in x_2 is then $\nu_2 = 1 - \nu$. Let the penalty of the collective be a scalar function of the state of the collective, $f(\nu) > 0$. Then the penalty of the i-th individual in the collective is:

$$\phi_i = \begin{cases} f(\nu), & \text{if in } x_1; \\ f(\nu - \frac{1}{M}), & \text{if in } x_2. \end{cases} \tag{2}$$

for all individuals in the collective. This model holds for decision spaces of arbitrary size. Thus, the penalty for each individual is the same (i.e., homogeneous), and depends on the fraction of the collective's individuals in a particular decision (i.e., fractional). While individuals may switch from one decision to another, equal fractions result in equal penalties. The goal here is to use rational individual behavior to have elements figure out what is best for them, converge to the best fraction of individuals in a particular decision, and simultaneously realize an optimal value of the collective behavior penalty function. This process is represented in Fig. 3, where ν^* is an optimal fraction for the penalty of the collective.

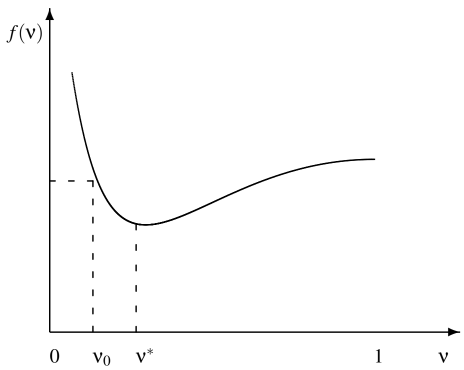

Figure 3. PENALTY VERSUS FRACTION OF THE COLLECTIVE.

SYSTEM MODELING AND VALIDATION

To facilitate the combined implementation of Generation Theory and the Theory of Rational Behavior to a model of an extraterrestrial self-reproducing robotic colony, we start with a trivial scenario under restrictive assumptions and then systematically remove these assumptions to obtain a generally-applicable collective model. The posed scenario is based on the simplified requirements of a lunar robotic colony, as outlined in the motivation. Two types of self-replication are discussed: single step self-replication and multiple step self-replication.

Single Step Self-Replication

Consider the case where one seed machine starts a robotic colony, and this machine is only capable of degenerate reproduction [23] (see Fig. 2), i.e., all offspring produced by the seed robot are themselves incapable of reproduction.

We make the following assumptions:

(1) The seed robot has to decide between two possible tasks, mining and reproduction. The seed robot has to be occupied with either one of these tasks at each time step.
(2) The environment possesses inexhaustible resources that can be mined. Each time step that the seed robot spends mining augments a collection of resources. During a time step spent in reproduction, one resource is converted directly into an offspring machine.
(3) Offspring machines are incapable of mining, and since they are degenerate, are also incapable of reproduction.

In this restrictive setting, suppose that the seed machine is subject to an arbitrarily chosen penalty function of fixed cost:

$$\phi(x) = \begin{cases} -1, & \text{if the seed robot mines;} \\ -2, & \text{if the seed robot reproduces.} \end{cases}$$

At each instant, the seed robot has an external penalty imposed as a result of taking one of the two possible decisions at that time step. Let the seed robot be capable of perfectly estimating the reward for the i-th decision x_i, and this estimate is $\hat{a}_i$. Further, suppose we allow memory so that the seed machine is capable of remembering the reward received for a particular decision. One model for the decision dwelling times T_{x_i} for this element with level of selectivity N is the following:

$$T_{x_1} = T_{x_2}\left(\frac{\hat{a}_2}{\hat{a}_1}\right)^N \quad \text{if switching to } x_1, \tag{3}$$

$$T_{x_2} = T_{x_1}\left(\frac{\hat{a}_1}{\hat{a}_2}\right)^N \quad \text{if switching to } x_2. \tag{4}$$

Because of the strong assumptions, the use of memory, and perfect estimation, a seed robot is quickly able to figure out that the best option is to spend a greater period of time mining rather than reproducing. Fig. 4(a) demonstrates this when $N = 1$, and Fig. 4(b) illustrates the effect of incrementing N by 1 once all possible decisions have been taken. In both instances, the rationality is independent of the starting decision state.

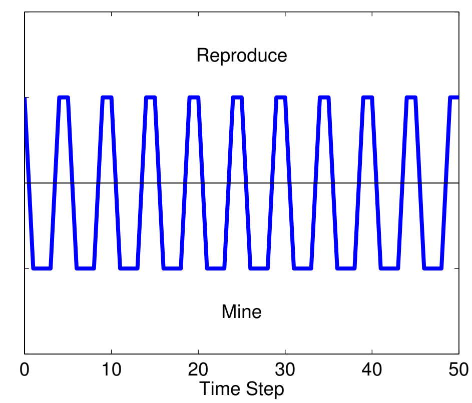

(a) $N = 1$

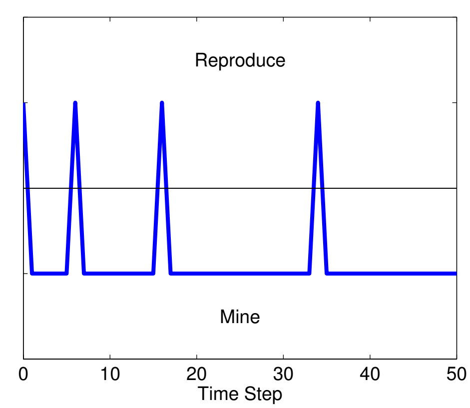

(b) N IS INCREMENTED FOR EVERY TOUR OF THE DECISION SPACE

Figure 4. PREFERENTIAL SELECTION OF DECISIONS.

Let us now relax some of the assumptions made previously. We allow offspring machines the capacity to mine, although they are still degenerate and unable to reproduce. Let M represent the total number of machines in the colony. Thus, at every time step, the $M-1$ offspring mine to each produce one unit resource. We no longer consider resources to be plentiful, and the colony is initiated with zero resources amassed. Mining is therefore required to stockpile resources, and the maximization of the total number of resources, R, is taken to be indicative of a successful colony. Accordingly, the penalty function for the seed robot becomes:

$$\phi(x) = \begin{cases} \frac{-1}{R} & \text{if the seed robot mines,} \\ \frac{-1}{R} & \text{if the seed robot reproduces.} \end{cases}$$

We also associate a cost to reproduction, R_{cost}, the value of which determines the number of resources needed to produce an offspring machine. Hence, we now have dynamic coupling between R and M through R_{cost}. These dynamics are given by the following time step updates. If the seed robot mines, then $R \leftarrow R + M \times 1 = R + M$, and if the seed robot reproduces, then $R \leftarrow R + (M-1) \times 1 - R_{cost}$ and $M \leftarrow M + 1$. Thus, reproducing depletes the current store of resources.

As expected, there is an initial delay in decision space exploration while enough resources are stockpiled for reproduction (Fig. 5(a)). This is followed by the realization that resource maximization occurs through the exclusive pursuit of a mining strategy. The greedy nature of rational elements is illustrated in Fig. 5(b) where $R_{cost} = 0$. The seed robot sees fit to continuously reproduce and have its increasing numbers of progeny maximize the resources of the collective.

To prepare for a collective that is capable of replication, we revise the goals so that a maximum of *both* resources and machines are produced. The penalty function gets modified to:

$$\phi(x) = \begin{cases} \frac{-1}{R} & \text{if the seed robot mines,} \\ \frac{-1}{M} & \text{if the seed robot reproduces.} \end{cases}$$

The directed graph representation of such a generation system is indicated in Fig. 6.

Multiple Step Self-Replication

It is unlikely that robotic self-replication can be achieved in a single step; as demonstrated in [7] however, self-replication is possible as a multiple step process. It is instructive to model the directed graph representation of this robotic system (Fig. 7). We take M to be the set of all entities that are made up of two or more LEGO Mindstorms kit components fixed together in some way. Let

$$M = \{x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8, x_9\}, \text{ and}$$
$$R = \{r_1, r_2, r_3, r_4, r_5, r_6, r_7, r_8, r_9\}.$$

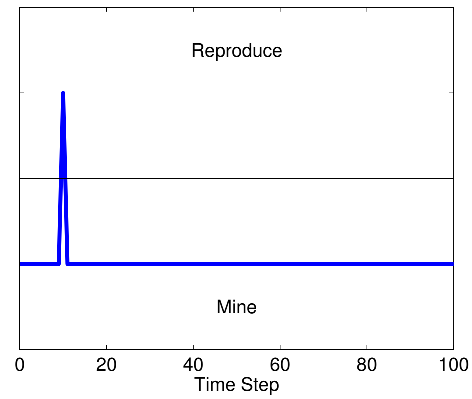

(a) $R_{cost} = 10$, $N = 1$

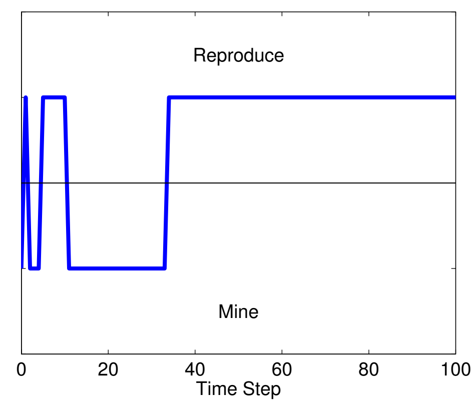

(b) $R_{cost} = 0$, $N = 1$

Figure 5. DECISION SELECTION.

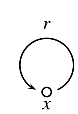

Figure 6. SELF-REPLICATION.

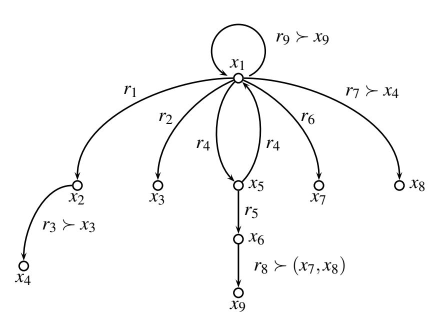

Figure 7. THE SUTHAKORN-KWON-CHIRIKJIAN ROBOT.

For notational purposes, if machine x_i belongs to an ordered list of the elements of resource r_j, then we say that x_i is *contained* in r_j, and we write $x_i \prec r_j$, where "$\prec$" is the *containment relation*. The following definitions and sequence of generation steps produce the directed graph representation of Fig. 7.

x_1 := prototype robot

r_1 := (conveyor-belt/sensor unit, docking unit, electrical connector, central controller unit (CCU), electrical cable)

x_2 := chassis assembly station

$x_2 = G(x_1, r_1)$

r_2 := chassis

x_3 := chassis aligned in assembly position

$x_3 = G(x_1, r_2)$

r_3 := (robot control system, x_3)

x_4 := RCX-chassis assembly

$x_4 = G(x_2, r_3)$

r_4 := gripper assembly/disassembly station := (CCU, electrical connector, ramp and lift system, gripper)

x_5 := prototype robot with gripper

$x_5 = G(x_1, r_4)$

$x_1 = G(x_5, r_4)$

r_5 := (left LEGO hook, right LEGO hook, CCU, electrical connector, stationary docking sensor, motorized pulley unit)

x_6 := motor and track assembly station

$x_6 = G(x_5, r_5)$

r_6 := (left LEGO track, right LEGO track)

x_7 := tracks aligned onto hooks

$x_7 = G(x_1, r_6)$

r_7 := (motor/sensor unit, x_4)

x_8 := RCX-chassis-motor assembly, moved to position

$x_8 = G(x_1, r_7)$

r_8 := (x_7, x_8)

x_9 := prototype robot on hooks

$x_9 = G(x_6, r_8)$

r_9 := x_9

$x_1 = G(x_1, r_9)$

Figs. 6 and 7 are very similar. Indeed, by aggregating resources 1 through 9 and machines 2 through 9 into one super-resource, the multiple step replication process can be converted into the single-step replication process modeled previously. This super-resource would have a higher associated reproductive cost, and so the effect of reproductive cost on self-replication needs to be investigated. Since single step self-replication captures the essence of multiple step self-replication, we limit our future discussion to single-step replication models only.

COOPERATIVE REPRODUCTION

We incorporate the theory of Homogeneous Fractional Interaction into the single step self-replication model to analyze the self-reconfigurability of the collective. In this scenario, all offspring are non-degenerate, and are hence capable of *both* mining and reproducing at any given time step. We seek to determine the appropriate fraction of mining robots for a particular reproductive cost. The penalty function for the i-th robot is:

$$\phi_i(x) = \begin{cases} \frac{-1}{R} & \text{if the robot mines,} \\ \frac{-1}{M} & \text{if the robot reproduces.} \end{cases}$$

For a reproductive cost $R_{cost} = 1$, it turns out that the optimal fraction of the colony engaged in mining activities is ≈ 0.25. This is depicted in Fig. 8 for 30 time steps, at the end of which there are 57,671 individuals in the robotic colony.

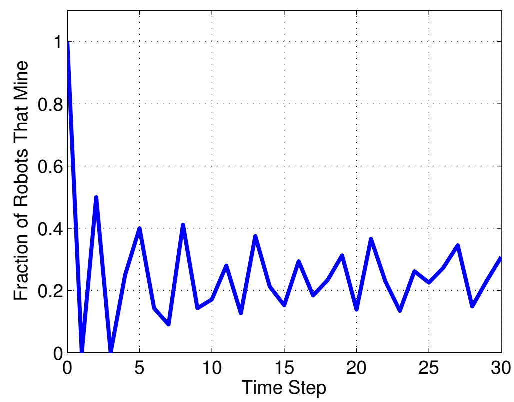

Figure 8. TIME HISTORY OF FRACTION OF MINING ROBOTS FRACTION, $R_{cost} = 1$, $N = 1$.

Adaptation of the Collective

We have previously indicated that the reproductive cost can vary if there is an increase in the number of steps required for self-replication. Alternatively, the cost of mining for resources can increase due to changes in the colony's surroundings, causing a corresponding increase in the reproductive cost. Thus, one mechanism for investigating the adaptability of a robotic colony to environmental disturbances is to determine the collective's response to changes in the reproductive cost.

Consider the reproductive cost shown in Fig. 9, where there is an increase in R_{cost} from 1 to 5 at the twentieth time step. The fraction of the colony engaged in mining activities is illustrated in Fig. 10. Once the reproductive cost changes, the collective quickly adapts and increases the fraction of individuals that are mining for resources to ≈ 0.7, although oscillations to the fraction are now magnified. After 40 time steps, there are 37,948 individuals in the robotic colony.

Response to Internal Disturbances

We require an extraterrestrial robotic colony to respond rapidly to applied disturbances and disruptions to the internal system states. Examples of disruptions include the eradication of a caste or the sudden specialization of a part of the collective (by the creation of a new caste) to take advantage of novel conditions in the environment. In our model of the colony, suppose that the ability to process resources (e.g. through casting or finishing) suddenly appears. The use of processed resources for reproduction generates the same type of robots as we had

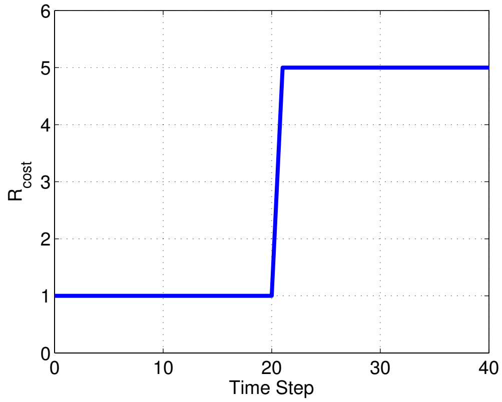

Figure 9. TIME HISTORY OF REPRODUCTIVE COST.

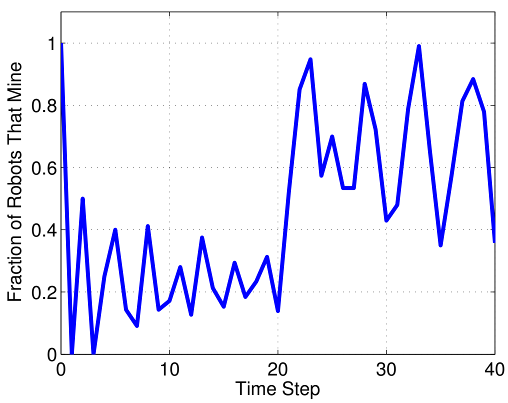

Figure 10. TIME HISTORY OF MINING CASTE, $N = 1$.

previously, and the reproductive cost now consumes the store of processed resources, P, instead of the store of raw resources. We include a typical advantage to using processed resources – that of a slower rate of depletion, P_{cost}, of raw resources. Specifically, the dynamics of R, M, and P are coupled through R_{cost} and P_{cost} as indicated by the following time step updates. If a robot mines, then $R \leftarrow R+1$; if a robot reproduces, then $P \leftarrow P - R_{cost}$ and $M \leftarrow M+1$; and if a robot processes resources, then $R \leftarrow R - P_{cost}$, and $P \leftarrow P+1$. We stipulate that $P_{cost} < R_{cost}$. The penalty function for the i-th robot now becomes:

$$\phi_i(x) = \begin{cases} \frac{-1}{R} & \text{if the robot mines,} \\ \frac{-1}{M} & \text{if the robot reproduces,} \\ \frac{-1}{R} & \text{if the robot processes resources.} \end{cases}$$

The goal is to demonstrate that our model of a collective takes advantage of an internal disturbance by reconfiguring the colony to maximize the store of raw resources. Fig. 11 shows that a new caste emerges quite rapidly at the expense of both of the old castes. The constant decline in the fraction of mining robots at the end of the simulation period may be attributed to the plentiful store of resources and the realization that fewer raw resources are required to expand the colony. After 40 time steps, there are 49,568 individuals in the robotic colony.

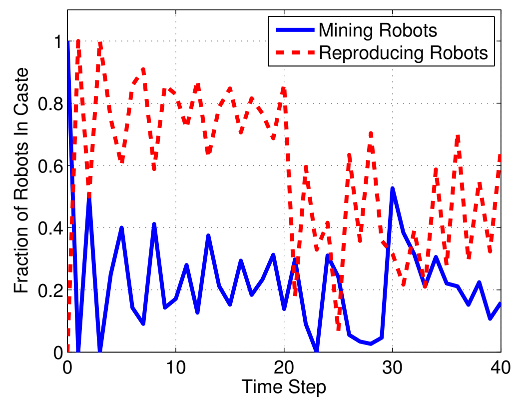

Figure 11. TIME HISTORY OF MINING AND REPRODUCING CASTES, $R_{cost} = 1$, $P_{cost} = 0.5$, $N = 1$.

Effect of Environmental Uncertainty

Here, we include the effect of noisy measurements in our model of single step self-replication. Each robot senses the penalty of the collective in a variable manner, that is, every estimate $\hat{a}$ is corrupted by a Gaussian random variable with zero mean and unit variance. The effect of this noise over the first few time steps is to delay the approach of the fraction of mining robots to ≈ 0.25 (see Fig. 12). However, the colony is able to reject the noise in their individual measurements and determine the appropriate fraction of individuals in each caste. After 30 time steps, there are 55,116 individuals in the robotic colony.

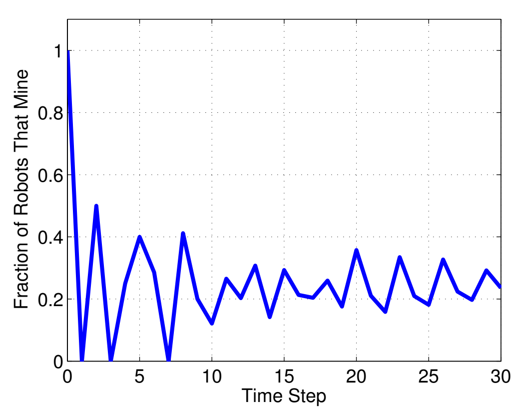

Figure 12. TIME HISTORY OF MINING CASTE WITH UNCERTAIN MEASUREMENTS, $R_{cost} = 1$, $N = 1$.

Conclusions and Future Work

This paper has combined Generation Theory and the Theory of Rational Behavior to investigate the self-reconfigurability of an extraterrestrial robotic colony. Through simulations, we have demonstrated that the colony is able to adapt to changes in the external environment, respond rapidly to applied disturbances and disruptions to the internal system states, and operate in the presence of uncertainty. Avenues for future work include incorporating the process of natural selection (in the form of probabilities and fitness functions) to the behavior of a collective. Such an analysis could demonstrate the emergence of evolutionary traits with time. The work in this paper can be extended to a decision space with a greater number of options. Finally, a corroboration of the simulations with hardware experiments is required.

REFERENCES

[1] von Neumann, J., 1966. *Theory of Self-Reproducing Automata*. University of Illinois Press.

[2] Freitas Jr., R. A., and Merkle, R. C., 2004. *Kinematic Self-Replicating Machines*. Landes Bioscience.

[3] Sipper, M., 1998. "Fifty years of research on self-replication: An overview". *Artifical Life,* **4**(3), pp. 237–257.

[4] Foing, B., 2005. "Roadmap for robotic and human exploration of the moon and beyond". In Proceedings of the 56th International Astronautical Congress, no. IAC-05-A5.1.01.

[5] Wertz, J. R., and Larson, W. J., eds., 1999. *Space Mission Analysis and Design*, 3rd ed. Microcosm Press.

[6] Freitas Jr., R. A., and Gilbreath, W. P., eds., 1982. Advanced Automation for Space Missions, no. N83-15348 (NASA Conference Publication CP-2255).

[7] Chirikjian, G. S., Zhou, Y., and Suthakorn, J., 2002. "Self-replicating robots for lunar development". *IEEE/ASME Transactions on Mechatronics,* **7**(4), Dec.

[8] Yim, M., Roufas, K., Duff, D., Zhang, Y., Eldershaw, C., and Homans, S., 2003. "Modular reconfigurable robots in space applications". *Autonomous Robots,* **14**(2-3), March, pp. 225–237.

[9] Zykov, V., Mytilinaios, E., Adams, B., and Lipson, H., 2005. "Self-reproducing machines". *Nature,* **435**, 12 May, pp. 163–164.

[10] Griffith, S., Goldwater, D., and Jacobson, J. M., 2005. "Self-replication from random parts". *Nature,* **437**, 29 September, p. 636.

[11] Salemi, B., Moll, M., and Shen, W.-M., 2006. "SUPERBOT: A deployable, multi-functional, and modular self-reconfigurable robotic system". In Proceedings of the 2006 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 3636–3641.

[12] Klavins, E., 2007. "Programmable self-assembly". *IEEE Control Systems Magazine,* **27**(4), August, pp. 43–56.

[13] Gilpin, K., Kotay, K., and Rus, D., 2007. "Miche: Modular shape formation by self-dissasembly". In Proceedings of the 2007 IEEE International Conference on Robotics and Automation, pp. 2241–2247.

[14] Yim, M., Shen, W.-M., Salemi, B., Rus, D., Moll, M., Lipson, H., Klavins, E., and Chrikjian, G. S., 2007. "Modular self-reconfigurable robot systems – challenges and opportunities for the future". *IEEE Robotics & Automation Magazine,* **14**(1), pp. 43–52.

[15] Slee, S. Self-reconfigurable robots: Papers. `http://imawhiner.com/~sgs/robots/papers.php`.

[16] Christensen, D. J., 2007. "Experiments on fault-tolerant self-reconfiguration and emergent self-repair". In Proceedings of the IEEE Symposium on Artificial Life, pp. 355–361.

[17] Stoy, K., and Nagpal, R., 2004. "Self-repair through scale-independent self-reconfiguration". In Proceedings of the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems, Vol. 2, pp. 2062–2067.

[18] Varshavskaya, P., Kaelbling, L. P., and Rus, D., 2004. "Learning distributed control for modular robots". In Proceedings of the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems, Vol. 3, pp. 2648–2653.

[19] Butler, Z., Kotay, K., Rus, D., and Tomita, K., 2002. "Cellular automata for decentralized control of self-reconfigurable robots". In Proceedings of the 2002 IEEE International Conference on Robotics and Automation, Vol. 1, pp. 809–816.

[20] Shen, W.-M., Will, P., Galstyan, A., and Chuong, C.-M., 2004. "Hormone-inspired self-organization and distributed control of robotic swarms". *Autonomous Robots,* **17**(1), July, pp. 93–105.

[21] Carlson, J. M., and Doyle, J., 2002. "Complexity and robustness". In Proceedings of the National Academy of Science, Vol. 99, pp. 2538–2545.

[22] Csete, M. E., and Doyle, J. C., 2002. "Reverse engineering of biological complexity". *Science,* **295**, 1 March, pp. 1664–1669.

[23] Kabamba, P., 2006. The von Neumann threshold of self-reproducing systems: Theory and computation. Tech. Rep. CGR-06-11, The University of Michigan.

[24] Meerkov, S. M., 1979. "Mathematical theory of behavior-individual and collective behavior of retardable elements". *Mathematical Biosciences,* **43**(1-2), pp. 41–106.

[25] Diestel, R., 2005. *Graph Theory*, 3rd ed. Springer-Verlag Heidelberg.