

- [10] J.-T. Lim and S. M. Meerkov, "On asymptotic reliable closed serial production lines," *Control Eng. Practices*, vol. 1, pp. 147–152, 1993.
- [11] J. Li, "Throughput analysis in automotive paint shops: a case study," *IEEE Trans. Autom. Sci. Eng.*, vol. 1, no. 1, pp. 90–98, Jul. 2004.
- [12] —, "Performance analysis of production systems with rework loops," *IIE Trans.*, vol. 36, pp. 755–765, 2004.
- [13] —, "Overlapping decomposition: A system-theoretic method for modeling and analysis of complex manufacturing systems," *IEEE Trans. Autom. Sci. Eng.*, vol. 2, no. 1, pp. 40–53, Jan. 2005.
- [14] S. R. Agnihotri and R. S. Kenett, "The impact of defects on a process with rework," *Eur. J. Oper. Res.*, vol. 80, pp. 308–327, 1995.
- [15] Y. Narahari and N. Viswanadham, "Transient analysis of manufacturing systems performance," *IEEE Trans. Robot. Autom.*, vol. 10, no. 2, pp. 230–244, Apr. 1994.
- [16] N. Viswanadham, K. R. Pattipati, and V. Gopalakrishna, "Performability studies of manufacturing systems with multiple part types," *IEEE Trans. Robot. Autom.*, vol. 11, no. 5, pp. 692–709, Oct. 1995.
- [17] Y. Narahari and L. M. Khan, "Modeling reentrant manufacturing system with inspection stations," *J. Manuf. Syst.*, vol. 15, pp. 367–378, 1996.
- [18] S. Yerelan and B. Tan, "Analysis of multistation production systems with limited buffer capacity part 2: The decomposition method," *Mathematics and Computer Modeling*, vol. 25, pp. 109–123, 1997.
- [19] A. S. Humphrey, G. D. Taylor, and T. L. Landers, "Stock level determination and sensitivity analysis in repair/rework operations," *Int. J. Oper. Prod. Manage.*, vol. 18, pp. 612–630, 1998.
- [20] J. Li, D. E. Blumenfeld, and S. P. Marin, "Manufacturing system design to improve quality buy rate," General Motors Res. Develop. Ctr., Warren, MI, Tech. Rep. R&D-10031, 2005.

Social Foraging Theory for Robust Multiagent System Design

Burton W. Andrews, Kevin M. Passino, and Thomas A. Waite

Abstract—An analogy between an agent (e.g., an autonomous vehicle) and a biological forager is extended to a social environment by viewing a communication network as implementing interagent sociality. We first describe engineering design within an evolutionary game-theoretic framework. We then explain why sociality may emerge in some environments and for some agent objectives. Next, we derive the evolutionarily stable design strategy for an agent manufacturer: 1) choosing whether the agent it produces should cooperate with other agents in a search problem and 2) choosing the group size of a multiagent system tasked with a cooperative search problem. We show the impact of "agent relatedness," a measure of common descent between two agents based on their underlying manufacturers, on the choices in scenarios 1) and 2). Our predictions are evaluated in an autonomous vehicle simulation testbed. The results illustrate a new methodology for manufacturers to make robust, optimal choices for multiagent system design for a given set of objectives and domain of operation.

Note to Practitioners—The design of autonomous multirobot systems with various applications, such as in parts production or search and destroy operations in a military environment, is of growing importance. Here, we integrate economic and technical issues into an unified engineering design framework for the manufacturers of robots. Our approach leads to manufacturer design decisions that are robust relative to the market for a manufacturer's products. Robot component aspects, such as sensors and communications as well as mission performance aspects, can be captured and coupled into the design process. We use the design of intervehicle cooperation and robot group size to illustrate this approach. The practical significance lies in the fact that we take a broad perspective on engineering design, one closer to the real world, due to the considerations of marketplace economics. Moreover, the approach provides a framework to study design choices that escape systematic analysis in other frameworks (e.g., group size).

Index Terms—Agent, autonomous vehicle, cooperative control, design, evolution, evolutionarily stable strategy (ESS), foraging theory, group size, manufacturer, multiagent, social.

I. INTRODUCTION

There is a considerable current interest in "cooperative control" for multiagent systems. One area is cooperative robotics where, for instance, cooperative task allocation is studied [2]–[4]. This work, however, focuses on the design of an agent's decision-making strategy, emphasizing 1) reaction to different situations in its domain of operation and 2) the design of strategies that perform well while operating in real time. Rather than study the design of similar strategies, we assume such a strategy is in place and examine the cooperative/non-cooperative and group size design choices for the manufacturer. We integrate economics into engineering design to make high-level, mission-planning-type choices (e.g., team composition), not decisions that

Manuscript received February 25, 2005. This work was supported in part by the Air Force Research Laboratory Air Vehicles Directorate (AFRL/VA) and Air Force Office of Scientific Research (AFOSR) Collaborative Center of Control Science under Grant F33615-01-2-3154 and in part by the Ohio Space Grant Consortium Fellowship. This paper was recommended for publication by Associate Editor Y. Narahari and Editor N. Viswanadham upon evaluation of the reviewers' comments.

B. W. Andrews and K. M. Passino are with the Department of Electrical and Computer Engineering, The Ohio State University, Columbus, OH 43210 USA (e-mail: andrews.191@osu.edu; passino.1@osu.edu).

T. A. Waite is with the Department of Evolution, Ecology, and Organizational Biology, The Ohio State University, Columbus, OH 43210 USA (e-mail: waite.1@osu.edu).

Digital Object Identifier 10.1109/TASE.2006.872124

are made in real time. We know of no work that analytically approaches, at the level of manufacturer design, search design choices for multi-agent systems.

The specific multiagent control problem we consider is that of cooperative search. For instance, autonomous vehicles may need to perform a search and rescue operation or search for targets in a military operation. Robots on a manufacturing floor may need to search for tasks or buffers to process, or in a distributed dynamical system (e.g., a temperature grid), a group of software controllers or agents may need to search for and eliminate regions of error [5]. In each case, the agents must cooperate to achieve their goal. As discussed before, our objective is to go beyond the existing cooperative control theory, which primarily addresses real-time control decisions of the agents, and study high-level design decisions, such as group size. Given the limitations of existing cooperative control theory and that the models required for the application of classical control methods from engineering to such a multiagent control problem would be extremely complex (and to our knowledge, none has addressed or solved this problem), a natural approach is to use foraging theory, which models similar problems in nature.

Foraging theory is used in the field of behavioral ecology to model animal decision making [6]–[8]. Evolutionary considerations often lead to optimization models of adaptive behavior, which is behavior shaped by natural selection. Foraging models predict animal behavior such as whether an animal should attack certain types of prey or how many animals should join a group under given circumstances. The applicability of foraging theory to a search problem in engineering is clear. In engineering, a decision-making agent, such as an autonomous vehicle or software module, can be viewed as a forager in a domain of operation as discussed in [1] and [9]. Foraging theory can then be used to determine an agent's optimal control strategy in a particular situation. While a solitary agent is the focus of [1] and [9], and a recent application of solitary foraging theory to robotics was studied in [10], we use social foraging theory to study manufacturer design decisions for multiple cooperative autonomous robots. Our approach is to extend the work in [11] to address the cooperative search problem. We incorporate vehicle cost and quantify sensing accuracy in terms of communication network and collaborative signal processing capabilities. Using evolutionary game theory, we predict evolutionarily stable design strategies for a manufacturer designing a group of agents or a single agent that may cooperate with other agents in the environment. Our overall approach rests on viewing design as an evolutionary process as we will discuss next.

A. Engineering as Evolutionary "Design"

Engineering designs are evaluated and purchased by consumers. Designs that have a great impact or generate huge profits for manufacturers will thrive in the marketplace, while less successful designs are likely to be discarded. Other manufacturers are likely to adopt designs similar to those that are successful. As new designs are put on the market, the process is repeated. This is an evolutionary process, where fitness corresponds to profitability. Designs that perform well within the fluctuating, unpredictable market tend to persist.

Here, we describe the evolutionary analogy in more detail with respect to multiagent systems. Our focus is on the evolution of aspects of an agent controller that determine cooperation and group size in a particular environment. The manufacturing system that produces the physical components of an agent (i.e., chips, parts, and algorithms) can be thought of as the agent's genetic makeup and is a direct byproduct of the manufacturers' design choices. In this sense, the manufacturers' design process and the apparatus responsible for producing the physical realization of the design can both be considered as the agent's genetic makeup and are both subject to change via natural selection. The genetic makeup determines the agent's "phenotype," which includes

everything from the physical components that comprise the resulting agent to how the agent maneuvers and interacts with other agents. The "lifetime" of an agent can be thought of as a single mission, an operation in which the agent strives to achieve some goal. Evolution occurs over the repetition of missions assuming an agent's existence ends with its mission.

Different components of an agent may be developed by different manufacturers, and if there exists a common manufacturer between two separate agents, the two agents are said to be similar. For instance, target sensors on two separate autonomous vehicles may be developed by the same manufacturer while all other parts of the vehicles are designed by different manufacturers. In other cases, two agents may have many different manufacturers in common. The degree to which this agent similarity exists is described by a coefficient of relatedness r , calculated as the percentage of an agent's manufacturers held in common with another agent. This definition departs from the biological definition of r used in Hamilton's rule [12], which quantifies r as the probability that two individuals carry the same allele from common descent. (A recent departure in biology from the classic coefficient quantifies relationships between individuals based on familiarity [13]). Relatedness is important because of its potential influence on the design of agents. For instance, an agent might be designed to assist a related agent. This will increase the fitness of the "relative" who has a similar genetic makeup as the original agent, thus increasing the probability that the agent's components are passed on to future generations. That is, such cooperation results in mission success and the ultimate production and sale of more of this product by the manufacturer.

An agent's success is determined by some currency related to the achievement of the mission's goal. One simple currency is a point system. A more successful agent acquires more points. If cooperation emerges and agents form a group, then the points acquired by the group are distributed among the group members. Groups may be formed for many reasons and, hence, points may be distributed among members of a group in many ways. Once an agent obtains some number of points, there is a further division of these points between the components of the agent and, hence, the manufacturers of these components, in the form of profits.

Since the goal of the manufacturer is to make money, the components of a successful agent are expected to persist in future missions. However, the creation of exact replicas in future generations is impossible due to imperfections in components or the manufacturing process (i.e., there are mutations). Also, the manufacturers may, in their design process, "explore" combinations of successful aspects of designs to try and enhance profits (analogous to sexual recombination). This could come via mergers between manufacturing companies, or via a manufacturer merging its subcontractors that make components of the design. While these mutations and combinations result in design variations that may often lead to the degradation of mission performance or even failure, sometimes they may increase successfulness. As missions are repeated, the selection will lead to increasingly robust designs that succeed in a "typical" environment. It is possible to analyze the robustness of engineering designs using game theory and the concept of an evolutionarily stable strategy (ESS) [14] and we will do so here (a loosely related general characterization of robustness of engineering designs appears in [15]).

B. Summary and Contributions

We approach the group-size aspect of the cooperative search problem by considering two design problems: 1) the design of a single agent's strategy of whether to cooperate with an already existing group of arbitrary size and 2) the design of an entire group of cooperative agents of a particular size. We use the ESS concept to predict manufacturer-designed strategies that are evolutionarily stable. An

evolutionarily stable design strategy is robust to changes or mutations in the strategies employed by other manufacturers in the environment in the sense that a population of manufacturers using the ESS in the design game cannot be invaded by another manufacturer strategy. Effects of agent relatedness as well as model parameter variations are considered. For each design scenario, an equilibrium group size is predicted by noting group size tendencies when manufacturers use the ESS for the design of agents in specific environmental situations.

One of the primary contributions of our work is to bridge fields (behavioral ecology, economics, and engineering) to provide a new perspective and methodology for approaching engineering design problems for multiagent systems. To do so, we extend foraging theory [11] to fit the cooperative search problem under consideration. We define $P_s(G)$ as the probability of mission success for a group of G vehicles, s indicating success, and use this probability as the cost function for vehicular cooperative search. The use of $P_s(G)$ embeds concepts from risk-sensitive foraging theory [1], [16], [17] by using the variance in the discovery of objects to make design decisions. Decisions that exploit high variance options (linked to smaller group sizes) are said to be “risk prone,” and those that exploit low variance options (linked to larger group sizes) are said to be “risk averse.” We include vehicle costs (not done for animals in [11]), and we incorporate the group searching advantage concept into the vehicular problem by denoting such an advantage for G vehicles as $K(G)$ and defining it via the dependence of sensing accuracy on a communication network, collaborative signal processing, and their limitations. Other contributions include expansions of the original analysis in [11] to predict group size equilibria based on the use of evolutionarily stable strategies as opposed to solely Nash equilibrium strategies. This is done by formulating the design games explicitly and allows us to specify mixed ESS cases where expected group sizes may be noninteger values (which may be common, especially for small G). This last set of contributions is likely to be useful for behavioral ecologists.

Our predictions are evaluated in a multi-autonomous vehicle simulation. As expected, the vehicular case does not perfectly fit the theory due to the existence of a finite number of objects and a finite domain within which perfect search coverage is impossible for a moving vehicle. However, the simulation provides a useful demonstration of the applicability of the theory to vehicular, cooperative search problems. In particular, we see important properties of mission success for a group of vehicles that are consistent with those of the model used in our theory. As a consequence, the theory is able to predict robust design strategies that translate to increased profits for manufacturers.

II. COOPERATION AND GROUP SIZE

Social foraging theory considers the benefits individuals accrue by foraging socially. Using concepts from social foraging theory, we present a cooperative search model and use game theory to analyze agent manufacturer design strategies under two scenarios. First, a manufacturer must decide whether to design an individual agent to cooperate with an already-existing group. We examine the evolutionary stability of each strategy and discuss how this game theoretic analysis is a means by which the manufacturer can determine, before manufacturing and deployment, strategies that are robust for a given situation in the sense that they will survive in the market amid design strategy mutations. This, in turn, provides insight into group size tendencies in the environment. The second scenario addresses the issue of group design. We suppose the manufacturer now must decide how many agents are needed in a multiagent system. Once again, this is a design decision by the manufacturer prior to deployment of the group, and the game theoretic analysis allows for the determination of evolutionarily stable design strategies. Provided that manufacturers

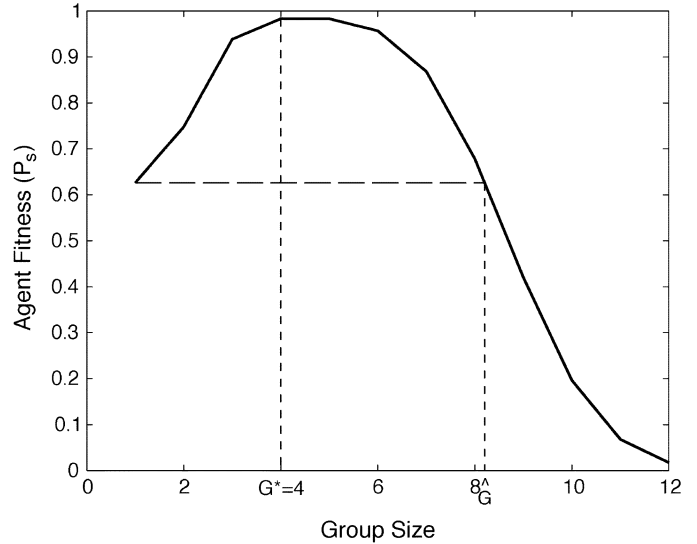


Fig. 1. Plot of $P_s(G)$. Parameter values are $\alpha = 8$, $T_c = 12$ min, $c = 1$ point, $V_r = 3$ points, $v = 1$ points, and $\lambda = 0.006$ objects/s.

design based on the ESS, we gain insight into the equilibrium size at which a designer will tend to design a group.

Before proceeding, a couple of remarks should be made. In the above discussion we ignore the explicit role of a customer. If you wish, you can think of the customer as purchasing and deploying the agent according to the manufacturer’s design guidelines (e.g., whether to cooperate with other customers who are going to deploy agents). The customer and manufacturer have an agreement that if the agent succeeds, the points it gets on a mission are converted to a proportional amount of monetary profits that are given to the manufacturer. If the agent fails, then no profits are given to the manufacturer. Here, for simplicity, we embed the notion of a customer in the design process and focus on the perspective of the self-interested manufacturer.

As will be discussed in further detail below, the design game has a population of manufacturers and, at each stage, we imagine that two manufacturer’s designed agents are drawn from the population and compete. We say “imagine” since we do not actually send them out to compete in a real mission. The two agents accrue points and corresponding profits are provided to their manufacturers. Then, a process based on selection and differential fitness acts on the proportions of the different designs in the population resulting in some manufacturers increasing the numbers of agents they produce (the relatively more successful agents) and other manufacturers decreasing the number of agents they produce (the relatively less successful agents), or going out of business altogether. In this situation, the manufacturer’s interest is solely in the success of the agent it produces, and it must decide in the design challenge of this section whether to cooperate with the agents produced by other self-interested manufacturers to improve its success and, hence, increase its marketshare. Since there is a one-to-one correspondence between the manufacturer and its product, we will sometimes speak of the agents as competing, but, of course, they are just the instruments of the true competitors—the manufacturers.

A. Agent-Based Cooperative Search Model

Here, we extend the model in [11, Ch. 2] to fit a cooperative search problem in engineering. Consider multiple agents that search a domain for objects, each with point value v . The goal of the manufacturer is to design one of these agents to acquire V_r points (r indicating required) within T_c time units (c indicating critical). For example, T_c may be mission time for military autonomous vehicles or a time constraint for

a robot in a search operation. Points obtained flow back to manufacturers in the form of profits, hence V_r may be translated to a profit requirement for manufacturers. If V_r is not obtained, the agent and, thus, its manufacturers (and its subcontractors for components) fails. We also assume a cost of c points associated with the deployment of the agent. This might be due to manufacturing, fuel, or maintenance costs. To summarize, the manufacturer wants a profit of V_r points for a payment of c points for the agent, resulting in a rate of return on the investment of V_r/c .

First, consider a solitary agent. The agent must discover $N_1 = (V_r + c)/v$ objects during T_c in order to be successful (subscript 1 indicates that this is the number of objects required for a solitary to succeed). For simplicity, we assume V_r is defined such that N_1 is an integer. If multiple agents cooperate in a group of size G , we assume they will equally share the profits from finding objects so that the group must obtain GN_1 objects for the success of every agent in the group. Note that the success of the group implies success from each member of the group. If we consider the design of a group of agents by a manufacturer, the total investment for the group is Gc points, and if the group succeeds (i.e., finds GN_1 objects), the manufacturer's net profit is $GN_1v - Gc$, giving a rate of return

$$\frac{GN_1v - Gc}{Gc} = \frac{V_r}{c}.$$

Thus, the manufacturer's objective is still to achieve a rate of return of V_r/c on its investment in each agent; however, a group may have the capability of obtaining far more points to offset costs than a solitary as will be seen below.

We assume objects are discovered by an individual agent at a constant rate λ (implying an infinite number of objects) via a Poisson process [11]. A group collectively discovers objects at a rate of $K(G)\lambda$, where $K(G) \geq 0$ is the searching advantage of the group. The probability of a group of size G succeeding is

$$P_s(G) = 1 - \sum_{i=0}^{GN_1-1} \frac{e^{-K(G)\lambda T_c} [K(G)\lambda T_c]^i}{i!}. \quad (1)$$

Due to the equal sharing of points discussed before, the probability of success of a group equals the probability of success of each member of the group. Search paths of solitary agents are assumed not to overlap so that encounter rates are independent of other vehicles' existence.

When group sizes grow large, diminishing reward portions often outweigh group benefits. This property is likely to occur, for instance, with a group searching advantage of the form

$$K(G) = \frac{\alpha}{1 + e^{\beta-G}} \quad (2)$$

where $\alpha > 1$ is a constant. Choosing $\beta = 1 + \ln(\alpha - 1)$ guarantees that $K(1) = 1$. The searching advantage increases quickly for low G , but as groups grow too large, the advantage becomes insignificant and quickly saturates. For example, in an autonomous vehicle application, the network bandwidth may be limited or quality may degrade as G increases. This concept is illustrated in Fig. 1, which shows $P_s(G)$. Note that although we have linearly interpolated between points, $P_s(G) \in [0, 1]$ is only defined on $G \in \{1, 2, \dots\}$. Using biological terms, we have defined the fitness of an agent as its probability of success. In general, we find that the fitness function peaks at $G = G^*$ and that $\lim_{G \rightarrow \infty} P_s(G) < P_s(1)$. Agents maximize their per capita fitness when the group size is G^* (although not noticeable in Fig. 1, $P_s(4)$ is indeed greater than $P_s(5)$ for the example given resulting in $G^* = 4$). As group size expands past G^* , the group benefit begins to be outweighed by the loss in members' success due to point sharing in an overpopulated group. Eventually, an individual's fitness drops below that of a solitary agent.

B. Cooperative Agent Design

We first provide a game theoretic explanation for the choice of manufacturers designing an agent that can either search for objects as an individual or by cooperating with an already-existing group. This decision must be made given that other manufacturers exist and are facing the same problem in the same environment. Each agent cooperates with other agents in the formation of a group only if the manufacturers design them to do so. The existing group corresponds to a standard group size in the environment. The problem here represents a "free entry" game from social foraging theory [11]. We begin with the case where agents are "genetically unrelated" (i.e., there are no similarities between agents with respect to their manufacturers) and then discuss the effect of relatedness on design strategies.

1) *Unrelated Agents*: Consider a simplified version of the above situation in which two agents are randomly drawn from an infinite pool of agents and deployed in an environment where a group of $G - 1$ agents already exists. Agents are assumed to be unrelated, and each agent possesses a manufacturer-designed strategy of either joining the group or searching as a solitary. We define the "payoff" to the manufacturer as the expected number of successful components designed by the manufacturer. The complete utility of this definition will become clear later when we discuss related agents. In the unrelated case we consider here, all components of the agent designed by a manufacturer yield profits for the manufacturer (since the agent was designed by that manufacturer); hence, the payoff to the manufacturer is simply the total number of components of the agent times the probability of the agent succeeding. This is the expected number of successful components of the agent. For simplicity, we assume that all agents in the environment have the same total number of components. This, in a sense, normalizes the payoff matrix for an agent playing the game so that the total number of components does not need to be included as a factor in the entries of the payoff matrix (this term would cancel when analyzing ESS conditions). Having said this, the payoff matrix for Agent 1 is

$$J_1 = \begin{bmatrix} P_s(1) & P_s(1) \\ P_s(G) & P_s(G+1) \end{bmatrix}$$

and the payoff matrix for Agent 2 is $J_2 = J_1^\top$. For agent 1 (2), row 1 (column 1) corresponds to the strategy Go Alone, and row 2 (column 2) corresponds to the strategy Join. The payoff matrix and, hence, the ESS depends on $P_s(1)$, $P_s(G)$, and $P_s(G+1)$. For instance, if $P_s(1) < P_s(G) < P_s(G+1)$ [e.g., Fig. 2(a)], or if $P_s(1) < P_s(G+1) < P_s(G)$ [e.g., Fig. 2(b)], the ESS is the pure strategy $[0, 1]^\top$ (i.e., Join) via diagonal dominance [18]. In this case, an individual can always do better as a member of a group of size G or $G+1$ than as a solitary. Also notice that if $P_s(G+1) < P_s(G) < P_s(1)$ [e.g., Fig. 2(d)], then the pure ESS is $[1, 0]^\top$ (i.e., Go Alone) by diagonal dominance.

An interesting case, one not considered in [11] or elsewhere in the literature, arises if $P_s(G+1) \leq P_s(1) \leq P_s(G)$ [e.g., Fig. 2(c)] where the ESS is mixed and given by

$$\left[\frac{P_s(1) - P_s(G+1)}{P_s(G) - P_s(G+1)}, \frac{P_s(G) - P_s(1)}{P_s(G) - P_s(G+1)} \right]^\top. \quad (3)$$

The mixed ESS is a result of a solitary having a larger probability of succeeding than a member of a $G+1$ group, yet a smaller probability of succeeding than a member of a G group.

The above analysis predicts the ESS of a manufacturer designing an agent in the presence of a standard group size. It also predicts the tendency for the existing group size to increase or decrease based on the ESS. We emphasize, though, that this is not a model of the dynamics of the engineering design process. Rather, our analysis is a means of determining, during design and prior to implementation, agent strategies that are evolutionarily stable for the current environmental situation. Such stability implies a robustness to strategy variations by other agents in




Fig. 2. Possible game scenarios and the corresponding ESS for an agent in the entry-decision game when $r = 0$. (a) $P_s(1) < P_s(G) < P_s(G+1)$ ESS: Join; (b) $P_s(1) < P_s(G+1) < P_s(G)$ ESS: Join; (c) $P_s(G+1) < P_s(1) < P_s(G)$ ESS: Mixed; (d) $P_s(G+1) < P_s(G) < P_s(1)$ ESS: Go Alone.

a given environmental situation since it is a strategy that cannot be invaded by mutant strategies (i.e., deviations from the ESS).

Given the implementation of the ESS, we can predict tendencies for group size to increase or decrease and, consequently, reach an equilibrium level. If the group size increases or decreases based on these tendencies, it will tend to level off at \hat{G} in Fig. 2. Note that although not considered in [11], \hat{G} may be a noninteger if the only integer group size with fitness $P_s(1)$ is $G = 1$. In this case, \hat{G} will be an expected value depending on the mixed strategies played by the agents. In particular, if we denote the ESS given in (3) as $[p_a, p_j]^T$, then the expected \hat{G} is $\hat{G} = p_a G + p_j (G+1)$. It can be shown then that $P_s(\hat{G}) = P_s(1)$.

2) *Related Agents*: Relatedness, as discussed in Section I-A, describes the similarities between the component manufacturers of two agents. Relatedness has the potential to influence design decisions since points and, hence, profits from relatives all flow back to the manufacturers. This concept is mathematically described by a modified version of Hamilton's Rule [11]: an agent should have a given strategy if $rE_r + E_s > 0$, where $r \in [0, 1]$ is the coefficient of relatedness (here, the proportion of common physical or software components), E_r is the effect of the strategy on other related agents, and E_s is the effect of the strategy on the agent itself. Since r is defined as the percentage of an agent's manufacturers held in common with another agent, the ease with which r may be computed clearly depends on the number of components of an agent. Clearly, however, in practice, the costs of the components are known and these can be used to specify the percentage of components of an agent attributed to each component manufacturer.

Here, we evaluate the cooperative agent-design game when agents are genetically related. As before, suppose two agents must be designed with either the strategy Go Alone or Join a group of size $G - 1$ that is already in existence. However, now all agents are genetically related. Defining Player 1 (Player 2) as Agent 1 (Agent 2), we assume that Agent 1 is related to Agent 2 with the coefficient of relatedness r_1 and to each member of the external group with coefficient r_2 . The payoff matrices are defined by incorporating ideas from Hamilton's rule. Each cell for a manufacturer (or its corresponding agent's) payoff

matrix gives, as before, the payoff to the manufacturer as the expected number of successful components of the manufacturer. This is the same payoff definition as in the unrelated-agent game. However, the manufacturer may profit from components of other agents in the environment if those agents are related to the agent produced by the manufacturer. The payoff matrix entries for Agent 1 are

$$\begin{aligned} J_1^{11} &= P_s(1) + r_1 P_s(1) + r_2 (G-1) P_s(G-1), \\ J_1^{12} &= P_s(1) + r_1 P_s(G) + r_2 (G-1) P_s(G), \\ J_1^{21} &= P_s(G) + r_1 P_s(1) + r_2 (G-1) P_s(G), \text{ and} \\ J_1^{22} &= P_s(G+1) + r_1 P_s(G+1) + r_2 (G-1) P_s(G+1) \\ J_2 &= J_1^T. \end{aligned}$$

The ESS for the game described above may be pure or mixed depending on which scenario in Fig. 2 applies to the current group size. We first examine the pure ESS cases. The pure strategy $[1, 0]^T$ (i.e., Go Alone) is an ESS if

$$r_2 (G-1) (P_s(G) - P_s(G-1)) + P_s(G) < P_s(1) \quad (4)$$

by diagonal dominance. Similarly, the pure strategy $[0, 1]^T$ (i.e., Join) is an ESS if

$$(r_1 + r_2 (G-1)) (P_s(G+1) - P_s(G)) + P_s(G+1) > P_s(1). \quad (5)$$

Setting $r_1 = r_2 = 0$ gives rise to the genetically unrelated case, and the analysis from Section II-B.1 holds.

There are three pure ESS scenarios corresponding to panels (a), (b), and (d) of Fig. 2. First, suppose $r_1 = r_2 = 0$ and $P_s(G) < P_s(1)$ so that we have the situation in Fig. 2(d) with the ESS being Go Alone. Since, at this location on the curve, $P_s(G) < P_s(G-1)$, increasing r_2 will decrease the left-hand side of (4) and will not alter the ESS. Thus, if $P_s(G) < P_s(1)$, the ESS is to Go Alone independent of r_1 or r_2 . Next, if $P_s(1) < P_s(G) < P_s(G+1)$ [panel (a)], the ESS when $r_1 = r_2 = 0$ is Join. As either r_1 or r_2 increases, the left-hand side of (5) increases since $P_s(G+1) > P_s(G)$ and, thus, the ESS is Join independent of r_1 or r_2 . The third scenario is when $P_s(1) < P_s(G+1) < P_s(G)$ [panel (b)]. As in the panel (a) case, $r_1 = r_2 = 0$ results in an ESS of Join. However, increasing r_1 or r_2 now decreases the left-hand side of (5), thus decreasing the tendency for Join to be the ESS. The design decision of whether to join involves a tradeoff between the fitness of the individual agent and the fitness of its relatives. A higher degree of relatedness leads to stronger dependence on the relatives' fitness and, hence, a stronger disinclination to join the group and thereby decrease the members' fitness. This shows that the tendency for the size of a group to increase when $P_s(1) < P_s(G+1) < P_s(G)$ decreases as genetic relatedness increases.

Now consider the case where a pure ESS does not exist, again something that has not been considered in the literature. The mixed ESS [18] is

$$\left[\frac{J_1^{12} - J_1^{22}}{J_1^{12} + J_1^{21} - J_1^{11} - J_1^{22}}, \frac{J_1^{21} - J_1^{11}}{J_1^{12} + J_1^{21} - J_1^{11} - J_1^{22}} \right]^T.$$

The "equilibrium" group size will be either the group size under which a mixed ESS exists or the smallest group size such that (4) holds. This equilibrium will lie somewhere between G^* and \hat{G} .

C. Cooperate and Group Size Design Games

Now suppose that manufacturers are to design a group of agents that compete in the market with other groups. This corresponds to the "group-controlled entry" game in social foraging theory [11]. However, contrary to the biological case, group size here is determined *a priori* as a part of the design process. We view the group size design decision as a type of incremental process that analyzes the effects of increasing

the size of an already established group by one member. The analysis is based on a game similar to that of the cooperative agent-design game. Specifically, suppose two groups, each of size G , exist in a domain in which two individual agents external to the groups also exist. Thus, there is a total of $2G + 2$ agents in the domain. The group manufacturer must make a design decision as to whether to have its group of agents remain at size G when deployed or increase its size to $G + 1$ by cooperating with one of the external agents. Thus, the design of a group begins at size $G = 1$ and increases via the cooperation with additional agents until the desired group size is reached.

We label the two groups in the domain as Group 1 and Group 2 and choose Player 1 (Player 2) of the game to be Group 1 (Group 2). We assume that Group 1 is related to Group 2 with the coefficient of relatedness r_1 and to each of the two solitary agents in the domain with r_2 . The relatedness coefficient between a group and a single agent or another group is an average measure of the similarity between all of the manufacturers of the group and the manufacturers of the other group or agent. For example, if there are G members of a group and the i th member has relatedness r_i with an agent outside the group, then the relatedness of the group to the external agent is $(1/G) \sum_{i=1}^G r_i$. The relatedness of the same group to another group with G members is $(1/G^2) \sum_{i=1}^G \sum_{j=1}^G r_{ij}$ where r_{ij} is the relatedness between the i th member of the first group and the j th member of the second group. Group size decisions are influenced by their effects on the related agents. Since the manufacturers of Group 1 are common to each member of the group, we assume all G agents of Group 1 are related with coefficient 1.

The payoff matrix entries for Group 1 are

$$\begin{aligned} J_1^{11} &= GP_s(G) + 2r_2P_s(1) + Gr_1P_s(G), \\ J_1^{12} &= GP_s(G) + r_2(P_s(1) + P_s(G+1)) + Gr_1P_s(G+1) \\ J_1^{21} &= GP_s(G+1) + r_2(P_s(G+1) + P_s(1)) + Gr_1P_s(G), \\ J_1^{22} &= GP_s(G+1) + 2r_2P_s(G+1) + Gr_1P_s(G+1) \end{aligned}$$

where row 1 corresponds to the strategy to remain at the current group size G (the Remain strategy), and row 2 corresponds to the strategy to increase the group size to $G+1$ (the Add strategy) via cooperation with one external agent. Also, $J_2 = J_1^T$. The Remain strategy is an ESS if

$$r_2(P_s(1) - P_s(G+1)) + G(P_s(G) - P_s(G+1)) > 0 \quad (6)$$

and Add is an ESS if

$$r_2(P_s(G+1) - P_s(1)) + G(P_s(G+1) - P_s(G)) > 0. \quad (7)$$

We analyze the ESS conditions as before using the panels in Fig. 2 as a guide. When the group size G is such that panel (a) of Fig. 2 holds, then Add is an ESS independent of r_2 or r_1 , via diagonal dominance. Hence, the group size will tend to increase. Due to the assumed nature of the $P_s(G)$ curve, the group size will reach a level such that either $P_s(G) = P_s(G+1)$ or $P_s(G) > P_s(G+1)$. Let us first examine the $P_s(G) > P_s(G+1)$ case (i.e., there does not exist a G such that $P_s(G) = P_s(G+1)$). If $r_2 = 0$, the group size will tend to settle at the optimal size G^* . This is because Remain will become the ESS as soon as $P_s(G) > P_s(G+1)$, and the group size will never have a tendency to move past the peak of the fitness curve and into the panel (b) of Fig. 2. If $r_2 \neq 0$, the first term in (6) (with coefficient r_2) becomes negative at the transition from panel (a) to (b). Hence, the left-hand side of (6) decreases with increasing r_2 , which decreases the tendency for the group to remain at the current group size G . In other words, the “equilibrium” group size increases past G^* with increasing r_2 . Groups will then tend to form at the minimum size such that (6) holds. If there does not exist a G such that $P_s(G) = P_s(G+1)$, then a mixed ESS

will never exist since one of either (6) or (7) will always hold. Also note that $J_1^{12} + J_1^{21} - J_1^{11} - J_1^{22} = 0$ is independent of r_2 and r_1 so that the mixed ESS does not exist.

If, on the other hand, a G exists such that $P_s(G+1) = P_s(G)$ (something not considered in the literature), then we have one of two scenarios. The first is $r_2 \neq 0$, in which case, Remain is the ESS if $P_s(1) > P_s(G+1)$ and Add is the ESS if $P_s(G+1) > P_s(1)$. The second scenario is $r_2 = 0$, in which case there is no ESS since Add and Remain are essentially the same strategy (all entries of the payoff matrix are the same). The expected equilibrium group size must lie somewhere between G and $G+1$; however, this expected group size cannot be solved for analytically since the ESS analysis does not provide any predictions on frequencies of strategies in a population when all entries of the payoff matrix are equal.

In summary, we have determined which strategy (remaining at G or increasing to $G+1$) of a manufacturer is evolutionarily stable given that all other manufacturers are facing the same problem. The fact that a given strategy is an ESS holds only under the assumption that repeated play of the game via manufacturer design iterations always occurs with a group size of G . In other words, the group size G does not change from generation to generation as a result of the last game’s outcome. Rather, this analysis determines strategies that are ESSs and, hence, robust for a given situation, and provides a prediction of the tendency for a standard group size to increase as a result of implementing the ESS.

D. Parameter Effects

The parameters of the model underlying our analysis are found in (1). The variation of each parameter changes the shape of the fitness curve $P_s(G)$ and, consequently, the ESS of a manufacturer in either the cooperative agent design or group size design game. If changing a parameter simply shifts the $P_s(G)$ curve up or down (e.g., increasing T_c or λ shifts $P_s(G)$ up), then the analysis from Sections II-B and II-C holds. However, it should be noted that some of the discussion in these sections is under the assumption of a $P_s(G)$ of the shape in Fig. 1. The same approach to the problem as in Sections II-B and II-C may be taken to study design strategies for curves of different shapes (e.g., ones with more than one local maximum), but aspects of the analysis such as the panels from Fig. 2 may not apply.

By studying such parameter changes, we find an important trend. The group size at which groups tend to form increases as the difficulty of a mission decreases by: decreasing the agent cost, decreasing the point requirement, increasing object reward, increasing object encounter rate, increasing critical time, or increasing the group searching advantage. This pattern highlights the concept of risk sensitivity. An increase in group size results in a decrease in the variance of the number of object discoveries. In other words, small groups are “risky” [11]. An underlying theme in risk-sensitive foraging theory is that when requirements are expected to be met (i.e., when critical times are long or requirements are small), safer or less variant options result in a higher probability of success. In contrast, risky or more variant options result in a higher probability of success when requirements are not expected to be met. These same concepts arise here. For easy missions, manufacturers will tend to design agents that form larger groups, and when missions are difficult, manufacturers will tend to design agents that form smaller groups.

III. APPLICATION TO AUTONOMOUS VEHICLES

Here, we demonstrate the applicability of the analysis in Section II to an autonomous vehicle cooperative search problem. Consider a domain where multiple autonomous vehicles are designed and deployed to search for objects as a group. This may be the result of a manufacturer designing a single agent to cooperate with an already existing




Fig. 3. Simulated probability of success for a group of G vehicles.

group as discussed in Section II-B or an entire group of agents as discussed in Section II-C. Each vehicle has a 500×500 -m sensor “footprint” and can sense or find an object when its footprint is positioned over it. Vehicles search the domain systematically by covering one domain length and, upon reaching the edge, turning and covering the opposite direction precisely one sensor width away from the previous search strip (i.e., via a “lawnmower pattern”). Vehicle starting positions are dispersed evenly across the domain guaranteeing nonoverlapping search paths. Objects are distributed via a Poisson process, and the expected encounter rate of a solitary is calculated as $\lambda_0 = N_o W u / A$ where N_o is the number of objects, W is sensor width, u is vehicle velocity (120 m/s), and A is the domain area.

The vehicle’s sensing equipment is intrinsically imperfect and, thus, subject to error. However, vehicles are capable of communication via a network, allowing for the sharing of information, such as shape or conspicuousness of objects. Such information sharing increases sensing effectiveness. Based on this, we denote the probability of a vehicle detecting an object in its footprint as $\gamma(G)$, an increasing function of group size G . Note that $0 \leq \gamma(G) \leq 1$. If a single agent encounters objects at a rate of λ_0 objects/s, the resulting rate of detection of objects for a solitary is $\lambda = \lambda_0 \gamma(1)$ objects/s. The collective detection rate for a group of G agents is $G\gamma(G)\lambda_0$, setting this equal to the collective detection rate of the model $K(G)\lambda$ and rearranging $K(G) = G\gamma(G)/\gamma(1)$. Therefore, a particular $K(G)$ corresponds to the probability of detection function $\gamma(G) = K(G)\gamma(1)/G$.

Parameter values are arbitrarily chosen to illustrate the relevant concepts as $N_o = 150$ (implying $\lambda_0 = 0.01$ objects/s), $\gamma(1) = 0.6$, $v = 1$ point, $c = 1$ point, and $V_r = 3$ points. Assume $K(G)$ given by (2) with $\alpha = 8$. Since $\gamma(G) = K(G)\gamma(1)/G$, information sharing via the communication network and group information processing allows agents to increase the accuracy of object detection as group sizes increase for relatively small groups. However, for the example here, as groups increase past $G = 4$, this detection probability decreases. This may be due to network bandwidth limitations and information processing overhead.

Fig. 3 shows $P_s(G)$, resulting from a simulated group of vehicles under the preceding parameters. Fig. 3 is similar to Fig. 1. The most significant difference is that the simulated curve is shifted down from the theoretical curve. This is expected since, in simulation, each vehicle’s sensor footprint swings outside the domain when turning around. This decreases λ from the estimated $\lambda = 0.006$ objects/s, shifting $P_s(G)$ down.

Now that we have established the existence of a fitness curve of the shape assumed in Section II, the rest of the analysis in that section holds with respect to the decisions of a manufacturer designing

an agent or group to operate in each scenario. For instance, the expected equilibrium group size for the cooperative agent design case when $r_1 = r_2 = 0$ is $\hat{G} = 7.4405$. The noninteger group size is a result of a mixed ESS in the design game. The standard group size has a tendency to increase until it reaches a size of 7, at which point, it is evolutionarily stable, according to (3), for the design strategy of a manufacturer to be to design an agent to search alone with a probability of 0.5595 and cooperate with the group with a probability of 0.4405. If we increase relatedness so that, for example, $r_1 = 0.3$ and $r_2 = 0.7$, the equilibrium group size becomes $G = 6$ since there is no group size under which a mixed ESS exists, and 6 is the smallest group size such that (4), the Go Alone pure ESS holds. This is a smaller size than \hat{G} since decreasing the size from \hat{G} increases $P_s(G)$, which is profitable to the external group manufacturers and, hence, to the related manufacturers of the agent being designed.

The equilibrium group size for the group size design case is $G^* = 4$ when $r_1 = r_2 = 0$. If we let $r_1 = 0.3$ and $r_2 = 0.7$, the equilibrium increases to $G = 5$ since 5 is the smallest group size such that (6), the Go Alone pure ESS holds. This increase arises because manufacturers of a group have some desire, due to relatedness, to increase the profits of the manufacturers of the solitary agents in the environment. Adding the solitaires to the group increases their probability of success and the profits of their manufacturers.

IV. CONCLUSION

We have introduced a method for producing robust designs for manufacturers that produce agents that operate in a potentially cooperative framework. The robustness stems from the merging of engineering with the field of behavioral ecology via evolutionary game theory. A manufacturer design strategy that is evolutionarily stable is robust to alterations of the strategies of other manufacturers designing agents in the environment. We apply the concepts discussed to a particular multiagent application where autonomous vehicles perform a cooperative search. In doing so, we address the problem of group size in multiagent systems, which, to our knowledge, has not been considered in other cooperative control areas, such as in cooperative robotics. The robustness of evolutionarily stable group design strategies provides success for manufacturers, and success translates to monetary profits, which is the manufacturers’ ultimate goal. Our analysis is centered around the use of the probability of mission success as the cost function which we choose to optimize. It would be straightforward to extend the theory to other cost functions, such as the number of points an agent is expected to obtain for a given mission.

REFERENCES

- [1] B. W. Andrews, K. M. Passino, and T. A. Waite, “Foraging theory for decision-making system design,” in *Proc. IEEE Conf. Decision and Control*, Atlantis, Paradise Island, Bahamas, Dec. 2004.
- [2] P. R. Chandler, M. Pachter, and S. Rasmussen, “UAV cooperative control,” in *Proc. Amer. Control Conf.*, Arlington, VA, Jun. 2001, pp. 50–55.
- [3] J. Finke, K. M. Passino, and A. Sparks, “Cooperative control via task load balancing for networked uninhabited autonomous vehicles,” in *Proc. IEEE Conf. Decision and Control*, Maui, HI, Dec. 2003, pp. 31–36.
- [4] A. E. Gil, K. M. Passino, S. Ganapathy, and A. Sparks, “Cooperative scheduling of tasks for networked uninhabited autonomous vehicles,” in *Proc. IEEE Conf. Decision and Control*, Maui, HI, Dec. 2003, pp. 522–527.
- [5] N. Quijano, A. Gil, and K. Passino, “Experiments for dynamic resource allocation, scheduling, and control,” *IEEE Contr. Syst. Mag.*, vol. 25, no. 1, pp. 63–79, Feb. 2005.
- [6] D. Stephens and J. Krebs, *Foraging Theory*. Princeton, NJ: Princeton Univ. Press, 1986.

- [7] A. Houston and J. McNamara, *Models of Adaptive Behavior*. Cambridge, U.K.: Cambridge Univ. Press, 1999.
- [8] C. Clark and M. Mangel, *Dynamic State Variable Models in Ecology*. New York: Oxford Univ. Press, 2000.
- [9] P. Pirolli and S. Card, "Information foraging," *Psychol. Rev.*, vol. 106, no. 4, pp. 643–675, 1999.
- [10] P. Ulam and T. Balch, "Niche selection for foraging tasks in multi-robot teams using reinforcement learning," in *Proc. 2nd Int. Workshop Mathematics and Algorithms of Social Insects*, Atlanta, GA, Dec. 15–17, 2003, pp. 161–167.
- [11] L. Giraldeau and T. Caraco, *Social Foraging Theory*. Princeton, NJ: Princeton Univ. Press, 2000.
- [12] W. D. Hamilton, "The genetic evolution of social behavior, i & ii," *J. Theoretical Biol.*, vol. 7, pp. 1–52, 1964.
- [13] M. A. Nowak and K. Sigmund, "Evolution of indirect reciprocity by image scoring," *Nature*, vol. 393, pp. 573–577, Jun. 1998.
- [14] J. Maynard Smith, *Evolution and the Theory of Games*. Cambridge, U.K.: Cambridge Univ. Press, 1982.
- [15] J. Carlson and J. Doyle, "Highly optimized tolerance: robustness and design in complex systems," *Phys. Rev. Lett.*, vol. 84, pp. 2529–2532, 2000.
- [16] A. I. Houston and J. M. McNamara, "The choice of two prey types that minimizes the probability of starvation," *Proc. Behavioral Ecology Sociobiology*, vol. 17, pp. 135–141, 1985.
- [17] D. W. Stephens and E. L. Charnov, "Optimal foraging: some simple stochastic models," *Behavioral Ecology and Sociobiology*, vol. 10, pp. 251–263, 1982.
- [18] J. W. Weibull, *Evolutionary Game Theory*. Cambridge, MA: MIT Press, 1998.

A Case Study on Integrated Production Planning and Scheduling in a Three-Stage Manufacturing System

Hong-Sen Yan and Xiao-Dong Zhang

Abstract—This paper presents an integrated optimization model of production planning and scheduling for a three-stage manufacturing system, which is composed of a forward chain of three kinds of workshops: a job shop, a parallel flow shop consisting of parallel production lines, and a single machine shop. As the products at the second stage are assembled from the parts produced in its upstream workshop, a complicated production process is involved. On the basis of the analysis of the batch production, a dynamic batch splitting and amalgamating algorithm is proposed. Then, a heuristic algorithm based on a genetic algorithm (known as the integrated optimization algorithm) is proposed for solving the problem.

Note to Practitioners—This paper presents a method for integrated production planning and scheduling in a three-stage manufacturing system consisting of a forward chain of three kinds of workshops, which is common in such enterprises as producers of automobiles and household electric appliances, as in the case of an autobody plant usually with the stamping workshop, the welding and assembling workshop, and the painting workshop. Herein, the production planning and scheduling problems are simultaneously addressed in the way that a feasible production plan can be obtained and the inventory reduced. A batch splitting and amalgamating algorithm is proposed for balancing the production time of the production lines. And a case study of the integrated planning and scheduling problem in a real autobody plant verifies the effectiveness of our method.

Index Terms—Batch splitting, integrated optimization, multistage manufacturing system, production planning, scheduling.

I. INTRODUCTION

In a manufacturing setting, production planning is essential for achieving efficient resource allocation over time in meeting demands for finished products [1]. Scheduling is a key factor for manufacturing productivity [2]. Effective scheduling can improve on-time delivery, reduce inventory, cut lead time, and improve machine utilization [2]. Thus, production planning and scheduling problems have been studied extensively. However, the current literature mainly focuses on the production planning and scheduling problems of a single workshop [3]. Little has been written about simultaneously optimizing the production plan and schedule of a multistage manufacturing system [4]–[6]. Most of them focus on the master production schedule (MPS) and capacity requirement planning (CRP) in a manufacturing resource planning (MRP) II environment, while the scheduling problem is seldom considered. Since these methods keep the planning separate from scheduling, they often generate an infeasible production plan, which has to be modified to obtain a feasible schedule [3]. Undoubtedly, considering these two problems simultaneously is advisable for avoiding an infeasible solution [3], [7]–[9]. However, none of

Manuscript received June 4, 2005; revised November 26, 2005. This work was supported in part by the National Natural Science Foundation of China under Grant 50475075, in part by the Special Science and Research Foundation for Doctoral Academic Subject of Colleges and Universities of China under Grant 20040286012, and in part by the Secondary Program for Cultivating Middle-Aged and Young Science Leaders of Colleges and Universities of Jiangsu Province, China. This paper was recommended for publication by Associate Editor Y. Narahari and Editor N. Viswanadham upon evaluation of the reviewers' comments.

The authors are with the Research Institute of Automation, Southeast University, Nanjing 210096, China (e-mail: hsyang@seu.edu.cn).

Digital Object Identifier 10.1109/TASE.2006.872107