ELSEVIER

Journal of Economic Behavior & Organization
Vol. 47 (2002) 179–195

JOURNAL OF
Economic Behavior
& Organization

www.elsevier.com/locate/econbase

Communication and cooperation

John H. Miller*, Carter T. Butts, David Rode

Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA 15213, USA

Received 17 June 1999; received in revised form 26 April 2000; accepted 2 June 2000

Abstract

Communication plays a vital role in the organization and operation of biological, computational, economic, and social systems. Agents often base their behavior on the signals they receive from others and also recognize the importance of the signals they send. Here we develop a framework for analyzing the emergence of communication in an adaptive system. The framework enables the study of a system composed of agents who evolve the ability to strategically send and receive communication. While the modeling framework is quite general, we focus here on a specific application, namely the analysis of cooperation in a single-shot Prisoner's Dilemma. We find that, contrary to initial expectations, communication allows the emergence of cooperation in such a system. Moreover, we find a systematic relationship between the processing and language complexity inherent in the communication system and the observed behavior. The approach developed here should open up a variety of phenomena to the systematic exploration of endogenous, strategic communication.

JEL classification: C63; C70; D80

Keywords: Cooperation; Communication; Adaptation; Cheap talk; Signalling; Tagging; Prisoner's dilemma; Genetic algorithms

1. Introduction

Communication plays a central role in the organization and operation of biological, computational, economic, and social systems. Agents often base their behavior on the signals they receive from others and also recognize the importance of the signals they send. The scope of behavior mediated by communication is enormous. In biological systems, such activities as mating (with signals ranging from the species-unique flashes of fireflies to the ritualized repartee of singles bars) and coordinated hunting rely on signals being passed

* Corresponding author. Tel.: +1-412-268-3229; fax: +1-412-268-6938; URL: http://zia.hss.cmu.edu.
E-mail address: miller@santafe.edu (J.H. Miller).

0167-2681/02/$ – see front matter © 2002 Elsevier Science B.V. All rights reserved.
PII: S0167-2681(01)00159-7

between agents. In computational systems, formal protocols are carefully developed to allow packets of information to flow across decentralized networks. In economics, behaviors as diverse as used-car buying and price collusion require careful communication (in the latter case, the impromptu use of the lower-ordered digits in auction bids has been observed recently as a means to facilitate collusion). To gain insight into the development and impact of communication on agents, we formulate an adaptive computation model (Holland and Miller, 1991) of endogenous agent communication. While the modeling framework is quite general, we focus here on a specific application, namely the analysis of cooperation in a single-shot Prisoner's dilemma.

The focus of our model is on systems in which communication emerges endogenously from the adaptive behavior of agents. To achieve this goal, we impose a number of conditions. Firstly, while all agents have access to a standard set of identifiable symbols, the meaning of each symbol must be induced by each individual agent, rather than specified in advance. Secondly, we do not allow agents to form any centralized enforcement mechanisms (for example, institutions like court systems that develop and interpret contract wording, formal ostracism procedures, etc.)—the only consequence of communication is the (decentralized) reaction it induces from the other agents. Thirdly, we do not allow agents to have any more information about their partners than that implicit in the communication itself. To insure that this latter assumption has some consequence, we restrict the number of repeated interactions among agents. Thus, agents cannot easily identify a particular partner, other than through the (potentially easily mimicked) communication that is received. Finally, we want to explore a system in which the potential exists for relatively complex communications. To achieve this last requirement, we allow agents to use simple computer programs (finite automata) to process, send, and react to communication.

The above framework is very general and applies to many different phenomena. Here, to provide an interesting context for our work, we focus on cooperation in a single-shot Prisoner's Dilemma. Along with its ubiquitous applications and analysis (see, for example, Axelrod (1984) and Axelrod and Dion, 1988), the Prisoner's Dilemma has an additional advantage: as is well known, the dominant strategy in the single-shot game is to defect. Given that communication is "cheap talk" (as it contains no enforceable commitments), that agents rarely interact repeatedly, and that the identification of particular individuals is difficult, it would appear that, *a priori*, communication should not alter the prediction of mutual defection. Clearly, this creates a scenario that is far less amenable to communicative interaction than most, and thus, this game provides a stark test for the impact of communication.

The emergence of communication has been explored by Werner and Dyer (1992) and, more recently, Steels (1996). This previous work has focused on mostly biological issues and uses very different computational and analytic techniques. Our work is also related to models of "tagging" that have been used to explore behavior in the iterated Prisoner's dilemma. Under tagging, agents are able to "recognize" one another via an observable marker and, based on this observation, decide whether or not to interact. Holland (1993) suggested that tags might allow new patterns of social interaction to develop. Stanley et al. (1994) allowed partners to recognize one another and to base refusals to play in a Prisoner's Dilemma on past experience. In Riolo (1996), agents were more likely to interact with other agents that looked alike (based on a predefined metric). Lindgren and Nordahl (1994) explored a more extended tag matching regime that allowed agents to modify their underlying matching rule. Like tagging, our model of communication allows agents to attempt to identify their partners. However, our model provides a much more sophisticated context for this identification that goes well beyond passive observation. Here, agents base their behavior on sophisticated and interactive communication schemes that are endogenously derived. Our agents can place new meanings on existing signals, invent new patterns of signals, and strategize about the signals to send in reaction to the ones received.

The adaptive model we create should serve as a productive environment from which to explore other issues and applications of agent communication. The behavior of the agents in our model is controlled by coevolving automata, and such machines embrace a broad class of behaviors rich in possibilities. Moreover, the underlying structure of the system is such that we can impose simple notions of communication processing and language complexity and link such ideas to their behavioral consequences. While we largely focus on the use of numerical experiments to derive and test hypotheses, the careful observation of the model should promote the development of other analytic tools. Along with a direct demonstration of how we can study models of adaptive communication, we also gain some insight into the emergence of cooperation via communication. To foreshadow the results, we find that with sufficient communication "complexity" epochs of cooperation can emerge in the model. A close examination of the dynamics of the system reveals that such cooperation is the result of a well-defined adaptive walk through the space of strategies that both exalts and exploits communication.

2. The computational model

2.1. The basic game

In the basic game, two agents are paired and must play a single-shot Prisoner's Dilemma. Prior to choosing a move (either to cooperate or defect), agents are allowed to communicate. Communication takes place via the exchange of communication tokens (predefined on $\{\emptyset, 1, \dots, T\}$, where T is exogenously given). The $\emptyset$ token has a special interpretation, namely, that the agent issuing the token has decided on a final move in the game.[1] At each step of the communication, each player simultaneously sends a single communication token to the other player. Communication continues until either both players issue the $\emptyset$ token, indicating that both have picked a final move in the game, or the number of steps in the communication (pairs of exchanged tokens) exceeds a preset chat limit.[2]

An example may help clarify the above ideas. Two agents meet and simultaneously issue a 1_A and 2_B (where the A and B distinguish the two agents, respectively). Agent A, when it hears 2_B, may decide to cooperate in the game, and thus, it will start to issue $\emptyset$

[1] Once such a token is communicated, the issuing agent is committed to the chosen (but undisclosed) move, and will continue to send the $\emptyset$ token until the communication ends.

[2] In the experiments reported here the chat limit was non-binding. That is, it was set high enough so that agents that actually reached the chat limit were guaranteed not to converge on a final move in the game.

tokens until either B decides on a final move or they reach the chat limit. When B hears the initial 1_A, it might decide to send a 1_B (which A receives at the same time B hears the $\emptyset$ sent by A). A will not react to the 1_B, since it has already picked a move, but B will react to the $\emptyset$ token it receives, by, say, deciding to defect in the game. Once this decision is made, both players simultaneously send $\emptyset$ tokens, and the single-shot game is played with A cooperating and B defecting. Note that the only information an agent has about its partner is the sequence of communication tokens emitted by that partner in response to the agent's own communication tokens—no other identifying information is available to the agents.

At the end of the communication the two players receive payoffs. In our game, the payoffs follow typical values: mutual cooperation pays 3.0 each, mutual defection pays 1.0 each, a sucker (a cooperator facing a defector) receives 0.0 and the corresponding defector in this latter interaction receives a temptation of 5.0. It is also possible that at the end of communication one or both of the players have not chosen a move. We give any player who has not chosen a move by the chat limit a *no act* payoff of −5.0. Moreover, a player who has picked a move, but faces an opponent who has not, receives a *no deal* payoff of 2.0. These latter payoff amounts were made to encourage agents to make final moves, while not trying to unduly penalize an agent who faces an indecisive opponent.[3]

We analyze a system composed of a population of P agents. During each generation of the system, every agent is paired with every other agent for one, single-shot Prisoner's Dilemma. Agents accumulate their payoffs from each game, and at the end of each generation their strategies undergo selection and modification via a simple adaptive algorithm the details of which are discussed later.

2.2. *Agent structure*

The behavior of each agent is controlled by a simple computer program represented as a finite automaton. Miller (1988, 1996) used this approach to study cooperation in a repeated Prisoner's Dilemma game, and showed, among other things, how such a representation can allow the exploration of a very interesting class of adaptive systems. Each automaton is composed of a fixed number of internal states on $\{1, \ldots, S\}$, where the maximum number of states, S, is given exogenously. Each state contains an action, either the sending of a communication token on $\{1, \ldots, T\}$, or an allowable move in the final game (here cooperate or defect) linked with the sending of the $\emptyset$ communication token. Note that the $\emptyset$ token does not communicate which move an agent has chosen, but only the fact that a move has been selected.

At the start of the game an agent enters state 1 of its automaton. Upon entering a state, the agent initiates whatever action is defined for that state. The communication tokens are always simultaneously passed to the agents. Associated with each state of the machine is a transition table that determines the next state that the automaton will enter contingent on the communication token (an element of $\{\emptyset, 1, \ldots, T\}$) received. If the agent has not chosen a final move, it enters a "new" state based on the transition table for its current state

[3] Of course, one open question is whether an agent might benefit by developing a strategy that makes opponents indecisive.

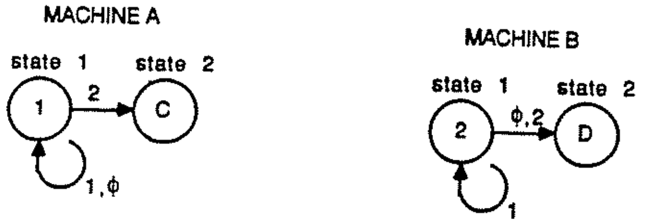

	period 1	period 2	period 3	Final Action
MACHINE A:	1	φ(C)	φ(C)	C
MACHINE B:	2	2	φ(D)	D

Fig. 1. Automata structure and play (see text for full details).

(since transitions are possible to any of the automaton's states, it may re-enter its current state). If the action in the current state is a final move in the game, then the automaton will not undergo any further transitions—thus, once an agent has picked a final move, its only behavior is to wait for the opponent (and send Ø tokens).

Fig. 1 illustrates two simple two-state automata. In this game, the automata are allowed to communicate using the tokens $\{Ø, 1, 2\}$. The actions for each state are given by the labels inside of the state nodes. The transitions are shown by the labeled arrows, where the labels indicate the communication token received from the opponent. Initially, both machines begin in state 1. In this state, machine A will issue a 1 and B will send a 2. When A receives the 2 from B, it has a transition into state 2, at which point it decides to cooperate and to send Ø tokens for the remainder of the game. Agent B, upon hearing a 1, remains in state 1, and sends another 2. Since A has already decided on a final move, the reception of a 2 from B has no impact on A. Agent B, however, when it receives the Ø sent by A, moves to state 2. In this state, B will defect in the game and begins to send a Ø. At the next step, both players issue Ø tokens, the communication phase ends, and a single-shot game is held in which player A cooperates and B defects.

Given the automaton representation, a strategy in this game is defined by giving every state an action (either send a communication token (from $\{1, \ldots, T\}$), cooperate and send Ø, or defect and send Ø), and a transition table for each state that maps any possible communication token received from the opponent ($\{Ø, 1, \ldots, T\}$) to a "new" state ($\{1, \ldots, S\}$). Initially, each agent is given a random strategy. To choose an action, with 0.5 probability we randomly pick a communication token from $\{1, \ldots, T\}$ (with equal probability on each token), otherwise a final move is assigned (either cooperate or defect, with equal probability). Each element of the transition table is randomly selected with uniform probability across all possible states ($\{1, \ldots, S\}$).

2.3. *Evolving automata*

A population of P agents is initially created using the random generation mechanism discussed above. At the beginning of each generation of the algorithm, every agent plays a single game against every other agent in the population. At the beginning of each of these games the automaton is reset to state 1 to begin play. An agent's final payoff is the average score received by its automaton across all of its games. Based on this average score, a new population of P agents is created by selecting better performing agents from the current population and randomly modifying some of these strategies.

The selection mechanism is a simple tournament selection. Two agents are randomly selected (with replacement) from the population, and the one with the better average score is placed into the new population. Note that this selection mechanism favors better performing strategies from the current generation, but it does not necessarily guarantee that the best strategy will be selected or that the worst will be culled.

Once a strategy is selected for reproduction, there is a 0.5 probability that it will be modified by mutation before being placed in to the new population. If a strategy undergoes mutation, then one of its states is randomly selected, and with a 0.5 probability the action of that state is randomly altered to an alternative action (using the same action selection mechanism described above for generating random machines), otherwise a randomly chosen transition of the state is altered (with uniform probability across the alternatives).

The above selection and modification procedures are carried out P times, resulting in a new population the same size as the original one. Once this new population is formed, payoffs are reset to zero and a new generation of the algorithm is begun—agents are again paired with each other, play the game, receive payoff, and then undergo selection and modification.

3. Results

Given the above system, we can explore the dynamics of the model. At the outset, a reasonable hypothesis is that there will be no observable cooperation in this system. Recall that agents in the game are playing (essentially) a single-shot Prisoner's Dilemma, and the dominant strategy in such a game is to defect. In our system, agents do play each other across generations, but given the difficulty of identifying particular individuals (the only identifying characteristic is the communication stream), *a priori* it would seem that such communication would be insufficient to allow cooperation to emerge in the game; Talk is cheap here and hard to back up with future promises. Thus, we would predict the following:

- *Hypothesis*: No systematic tendency toward cooperation should be observed in the system.

While the confirmation of such a hypothesis would not be too surprising, we find the following:

- *Observation*: Repeated outbreaks of mutually cooperative behavior occur in the system.

The apparent contradiction of the initial hypothesis by the observation suggests an interesting area for investigation.

The analysis proceeds in the following manner. Our initial focus concerns a model with a population of fifty agents ($P = 50$), who evolve four-state automata ($S = 4$), and can use two communication tokens ($T = 2$). The results of the model from these parameters appear

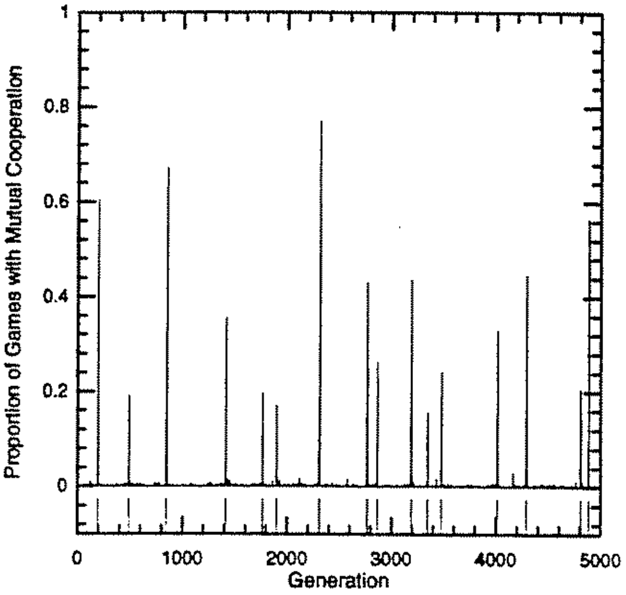

Fig. 2. Proportion of mutually cooperative plays by generation for a system composed of 50 agents ($P = 50$), four-state automata ($S = 4$), and two communication tokens ($T = 2$) (bottom ticks represent start of cooperative outbreaks).

representative of our overall findings.[4] The initial focus of the analysis is on characterizing the cooperation we observe in the system. We then turn to a more systematic analysis of how the key parameters (machine size and number of available tokens) influence cooperation. Finally, we provide an explanation for how cooperation emerges in this system, through an analysis of the strategic dynamics.

3.1. Patterns of cooperation

As a first step in the analysis, we characterize the cooperation observed in the system. We use as a simple measure of cooperation: the proportion of all single-shot games during a given generation in which both players cooperate.[5] Fig. 2 plots the rate of mutual cooperation by generation for a typical run of the model (the underlying data generated by this experiment

[4] We have done extensive robustness testing of most of the models assumption's, including separately coding the model by two different developers in two different languages, using different selection, crossover, and mutation mechanisms, strategic implementations (for example, allowing "bluffing"), and population sizes. We have also altered payoff parameters and agent matching procedures (including restricting interactions to predefined social networks and avoiding self interaction). Analyses of these alternatives indicate that while we do find some minor quantitative differences among the models, the main conclusions remain rather robust. As in all adaptive models, movements toward the extremes, for example, large mutation rates, will eventually disrupt the system, nonetheless, the large equivalence class of parameter values under reasonable assumptions is reassuring.

[5] Of course, other measures of cooperation exist (for example, the number of cooperative acts observed during play), but these alternatives are closely correlated with our measure of mutual cooperation.

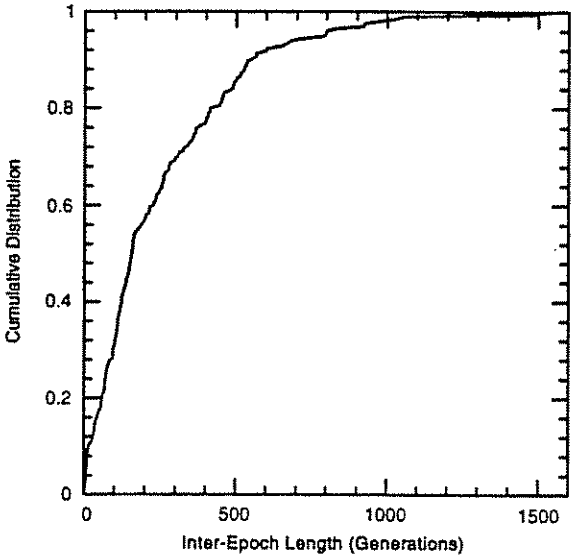

Fig. 3. Cumulative distribution of generations between cooperative epochs for a system composed of 50 agents ($P = 50$), four-state automata ($S = 4$), and two communication tokens ($T = 2$), across 50,000 generations.

are the basis for many of the subsequent graphs). As can be seen from the figure, while in general the rate of mutual cooperation is very low, there are occasional periods in which "high" levels of cooperation emerge in the system.[6] Though not directly observable from the figure, during the periods in which mutual cooperation is low, mutual defection strongly dominates play.

A close examination of the phase portrait of mutual cooperation (a plot of the cooperation rate at time t versus $t + 1$) indicates that low-levels of cooperation are typically followed by low levels of cooperation, but on occasion, cooperative outbreaks occur in which cooperation rates rise quickly to a peak and then rapidly fall off. Thus, we observe relatively short epochs of sustained cooperation. In the experiment shown in Fig. 2, 1.6% of the time the system has above threshold rates of cooperation. Over the 5000 generations of the experiment, we observe an average of 3.5 cooperative epochs per 1000 generations, each of which averages 4.5 generations in length.[7]

The picture that emerges from the above is a system dominated by mutual defection, with occasional, short-lived epochs of mutual cooperation. While systematic, the cooperative epochs are not periodic. Fig. 3 plots the cumulative distribution of the interval between

[6] For convenience, we define a high level of cooperation to be one in which at least 10% of the games have mutually cooperative plays. This threshold becomes important in defining periods of sustained (above threshold) cooperation. We have found that our results generally scale in the obvious way with different threshold levels. In many of the following figures, tick marks at the bottom of the graph are used to demarcate the start of cooperative epochs.

[7] For the analysis, we ignore the first 100 generations of the experiment so as to avoid biases caused by initial stochastic effects.

cooperative epochs for an experiment of 50,000 generations using similar parameters to those above. The median inter-epoch length was around 155 generations, with a range from 2 to 1537 generations, and the power spectrum is uncorrelated. Finally, we have found that the system exhibits little path dependence across experiments, that is, the major results are similar if we run one experiment for 20,000 generations or aggregate data across five experiments each run for 4000 generations. This observation suggests that the state of the system is somehow being recycled during the course of the experiment.

3.2. *Patterns of communication*

Fig. 4 plots the average length of communication (number of token exchanges per game) observed in the experiment shown in Fig. 2 (the tick marks at the bottom of the graph indicate the start of each cooperative epoch). While there is a lot of variance surrounding the average chat length, at the start of each cooperative epoch it tends to peak. Note also that the lower-bound of the average chat length increases at the start of each epoch, and then slowly decays (this is indicated by the "holes" at the bottom of the graph). Thus, we find that the system tends to have low levels of communication most of the time, but when cooperation breaks out, communication increases dramatically and then slowly decays back to its low background level. This decay is slow enough that non-cooperation returns as the norm even while communication remains high.

Another way to characterize communication in the system is by observing the number of unique conversations during any given generation. We consider two conversations, the same

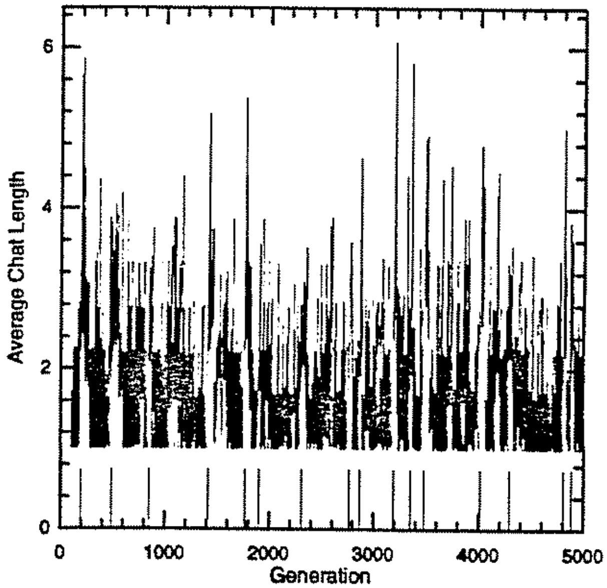

Fig. 4. Average length of communication by generation for the experiment shown in Fig. 2 (bottom ticks represent start of cooperative outbreaks).

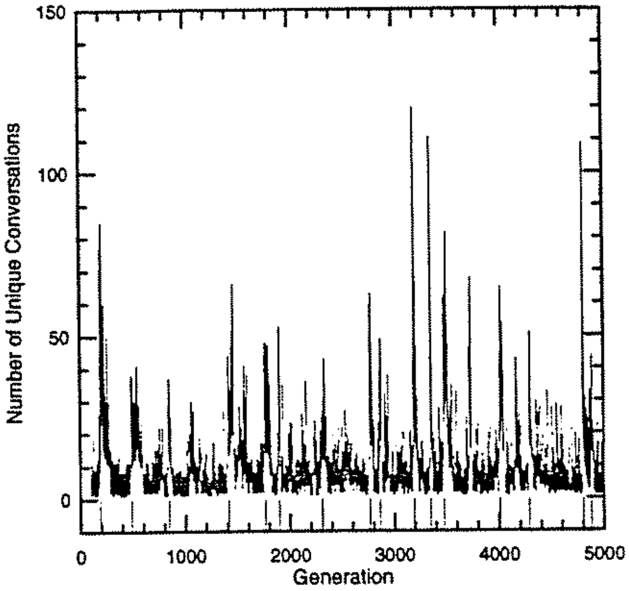

Fig. 5. Number of unique conversations by generation for the experiment shown in Fig. 2 (bottom ticks represent start of cooperative outbreaks).

if the set of communication tokens exchanged and the final moves are identical across the two games (irrespective of the players). Thus, if all players simply defected without trying to communicate, we would have only one unique conversation. In general, we observe very low levels of unique conversations during most generations. However, whenever cooperation does break out, the system experiences a rapid rise in the number of unique conversations exchanged (see Fig. 5). The number of unique conversation often remains high for a number of generations after the cooperation has subsided.

To summarize the above observations, we find that the system tends to be dominated by low levels of cooperation and communication. Nonetheless, we also observe occasional outbreaks of mutual cooperation that are sustained over a number of generations. During these cooperative epochs, both the length and number of unique conversations increases dramatically. While the cooperation quickly dissolves, the observed increases in both chat length and number of unique conversations decay at a much slower rate.

3.2.1. Communicative complexity and cooperation

One advantage of the framework we propose here is that we can link some simple measures of the complexity of the communication to key behavioral outcomes.[8] We use two notions of complexity. The first, as measured by the maximum number of states, S, available to each machine, determines the degree to which a machine can process and react to the incoming communications—we call this processing complexity. The second, as measured

[8] The word "complexity" has a variety of both formal and informal notions. We use the word here to help provide some additional context for our measures. Arguably, other terms like "capacity" might be appropriate substitutes.

by the number of communication tokens, T, available in the language, provides a notion of the richness of the basic language—we call this language complexity.

A key area of interest is how processing and language complexity influence system-wide behavior. Here, the behavioral event of interest is the emergence of cooperation. Using the model developed above, we ran numerical experiments that varied the number of states ($S = 3, 4, 6, 8, 10, 12$)[9] and tokens ($T = 1, 2, 3, 4$) in the system. For each of these conditions, we ran 20 experiments of 5000 generations each. At the end of each experiment, we collected the average number of cooperative epochs per 1000 generations, C; the average length per epoch, L; and the total number of periods per 1000 generations with above threshold mutual cooperation, P. Epochs were defined as sustained periods of above threshold (0.10) mutual cooperation, and no data were collected over the first 100 generations of the system to avoid biases caused by the initial conditions.

To provide a descriptive summary of the results of the numerical experiments, we use an ordinary least squares regression on the 480 observations and estimate:

$$C = -8.22 + 2.45S + 2.30T + 0.23ST - 0.13S^2 - 0.46T^2, \qquad R^2_{\text{adj}} = 0.91,$$

$$L = 2.48 + 0.54S - 0.20T + 0.02ST - 0.02S^2 - 0.04T^2, \qquad R^2_{\text{adj}} = 0.80,$$

$$P = -35.64 + 10.17S + 6.10T + 2.15ST - 0.48S^2 - 2.04T^2, \qquad R^2_{\text{adj}} = 0.93.$$

With the exception of the coefficients for T and T^2 in the second equation, all coefficients were significantly different from zero at the 0.005 level or better.

The regressions reveal a systematic relationship between cooperation and the underlying levels of processing and language complexity. The two measures of complexity are able to account for a large amount of the variance observed in the data. Fig. 6 shows the predicted number of cooperative epochs as a function of the number of states and tokens (using the parameter estimates from the C equation above). Using the estimated equations evaluated at the average values of S and T, we find that increasing machine size by one state results in C increasing by about 1, L increasing by 0.3, and P increasing by 8.4. Increasing the number of tokens by one results in a predicted increase of C by 1.7, L by 0.15, and P by 12. Thus, as either processing or language complexity increases, we expect to see more, and longer, epochs of mutual cooperation. We also find that these two types of complexity have a positive interaction effect with one another.

3.3. How does cooperation emerge?

As discussed above, since communication has no directly enforceable consequences and identification of individuals is difficult, we would expect only defection. Yet, we find cooperation emerging in a very systematic way.

[9] Note that, we did not run systems with less than two states because cooperation did not emerge in such worlds. With only two states, machines are incapable of both communicating and then either cooperating or defecting based on the behavior of the opponent (since each of these acts requires a separate state in the machine). Thus, there is no way for a two-state machine to both communicate (even a single token) and effectively react to opponents, and we find that the system quickly evolves to being filled with agents that only defect.

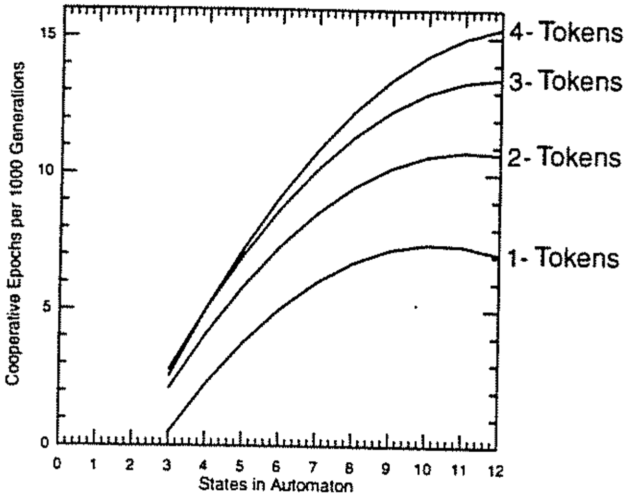

Fig. 6. Predicted number of cooperative epochs.

Our explanation for the emergence of cooperation in the model relies on a close examination of the strategic dynamics that occur in the system. To develop such an explanation we have carefully examined the strategies that arise just prior to, during, and just after a cooperative epoch. These observations across many such events, linked with some tests of their implications, give us some confidence in the soundness of the hypothesis presented below.

The first issue to be confronted is whether or not the observed cooperation is an artifact of the random nature from which strategies are derived. That is, since we randomly alter strategies as the system evolves, cooperation could just be a spurious result of this creation process. Of course, there is an important difference between systems that rely on random events (such as natural selection) and those that are randomly determined. If our observations are indeed being driven only by random noise, then a simple test is to observe the system without any selection mechanism operating. Without selection, the only changes in the strategies come about via mutation, and thus the model behaves as if we were taking an undirected random walk through the space of strategies. Observation of the cooperation rates under such conditions reveal a very different picture from that observed in Fig. 2. Clear epochs are not observable and cooperation rates oscillate wildly throughout the generations. Thus, our observed patterns of cooperation cannot be adequately explained by pure random effects.

Our explanation for the emergence and destruction of cooperation in this system is based on how the adaptive process sequentially creates and destroys key strategy types. While it may be the case that other mechanisms could be involved, we feel that we have strong evidence for the following, five-stage cycle.

3.3.1. Stage 1: domination by "no communication and defect (NCD)"

In this stage, strategies evolve so that when they meet an opponent they do not even try to communicate, but instead simply defect. In a world dominated by defection, such a strategy

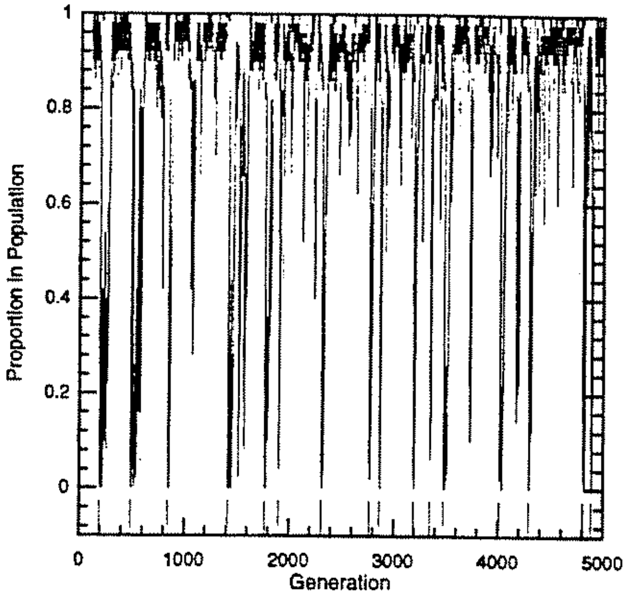

Fig. 7. Proportion of "no communication and defect" strategies for the experiment shown in Fig. 2 (bottom ticks represent start of cooperative outbreaks).

is very stable. Fig. 7 shows the proportion of strategies that implement NCD. As is apparent in the figure, this is a very common strategy throughout most generations, with occasional break downs most often associated with the start of cooperative epochs.

An interesting issue is why we do not see strategies that first try to communicate and then defect (which would lead to the same payoff). We hypothesize that such strategies do not persist for two reasons. The first is that if agents get trapped into perpetual communication they will get a much lower payoff in the game, and so there is an incentive for agents to avoid such a situation by not communicating. A second, and perhaps more important, reason is that machines that do a lot of communicating are much more vulnerable to mutations that can alter their behavior in a maladaptive way (since the additional states they use to communicate can be mutated to either inadvertently cooperate or perpetually communicate). Any such mutations will dramatically reduce the survivability of the machine, and eliminate them from the population. In general, we suspect that there is a large basin of attraction surrounding the NCD behavior that easily traps the adaptive system.

3.3.2. *Stage 2: emergence of "communicate and reciprocate communication (CRC)"*

During this stage, a few strategies emerge that have the following behavior: begin by communicating, and if the opponent communicates, cooperate, if they say nothing, defect. Both CRC and NCD machines can coexist, since both end up in mutual defection, and thus, achieve identical payoffs.

While the emergence of CRC behavior seems improbable, in fact, it can often get invoked by a single mutation. The reason for this is that during stage 1, there is very little selective

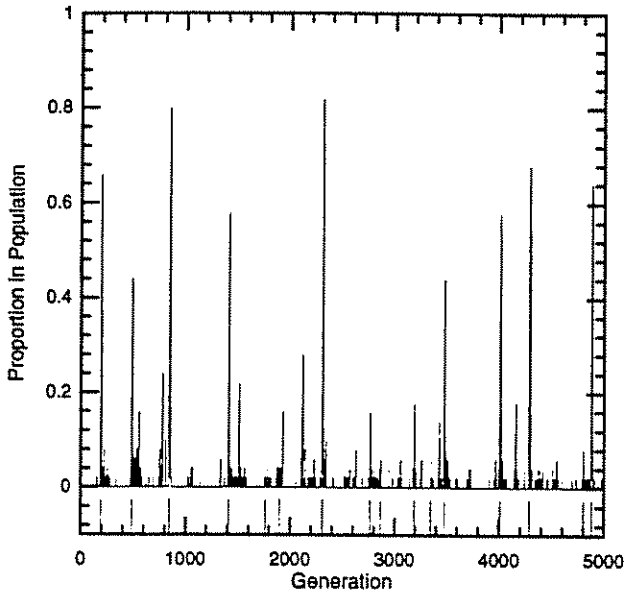

Fig. 8. Proportion of "communicate and reciprocate communication" strategies for the experiment shown in Fig. 2 (bottom ticks represent start of cooperative outbreaks).

pressure on the NCD strategies except at the first state of the machine. That is, as long as the first state of the machine has the defect action the machine will exhibit NCD behavior. Thus, mutations can freely accumulate in other parts of the machine without direct consequence. If, however, the action of the first state is mutated, these other parts of the machine can be invoked.[10] Typically, these new machines are very maladaptive, but occasionally CRCs are created. Once the CRC machines arise, the stage is set for the emergence of cooperation.

3.3.3. Stage 3: emergence of cooperation

Once more than one CRC machine arises (which, given their identical payoff to the NCD machines, will happen by chance via selection during stage 2), the CRC machines can achieve above average payoffs. When two CRC machines meet one another, they both communicate, cooperate, and receive the high mutual cooperation payoff. However, the CRC machines can also protect themselves against the NCD machines, since the lack of communication from the NCD machine leads to mutual defection. Since the NCD machines average the mutual defection payoff while the CRC ones get both these payoffs and the mutual cooperative ones, selective pressure allows the CRC machines to rapidly dominate the population, and in so doing, cooperation emerges. Fig. 8 plots the proportion of CRC strategies in the population for each generation. As is apparent from the figure, almost

[10] Human cultures often carry along a variety of rarely invoked artifacts, such as obscure rules of etiquette or words in the dictionary. On occasion mutations allow these artifacts to be promulgated.

all of the cooperative outbreaks are associated with a rapid increase in CRC. A strategy was considered CRC if it immediately cooperated on communication, and thus, strategies that develop more elaborate handshakes (for example, the exchange of two communication tokens before cooperation) are not included in the graph.

3.3.4. *Stage 4: emergence of mimics who "communicate and defect (CD)"*

At the beginning of this stage, the population is dominated by machines that communicate with one another and cooperate. Of course, this leaves the population vulnerable to the emergence of strategies that mimic the communication, and then, instead of cooperating, defect. Thus, mimics are able to get their opponents to cooperate while they defect—leading to very high payoffs for the mimic and low payoffs for the opponents. Once mimics arise, selective pressure will quickly destroy the cooperation and leave in its wake a population of defecting mimics.

Note that in our model the emergence of effective mimics is facilitated by the adaptive system. New agents arise in this system from the reproduction and possible mutation of currently existing, successful agents (in this case, the CRCs). Thus, even if the CRCs have created an elaborate communication handshake that requires multiple responses to a variety of signals before cooperation will ensue, all that is needed for a mimic to arise is that during reproduction the final action that causes cooperation in the CRC is mutated to a defection. Thus, successfully cooperating parents plant the seeds of their own destruction when they reproduce, allowing the creation of deadly mimics who do not have to expend a lot of effort "learning the secret handshake".

3.3.5. *Stage 5: "communicate and defect (CD)"*

During the final stage, the mimics have become dominant and the world now consists of agents who communicate and defect. As previously mentioned in the discussion of stage 1, there are pressures that will tend to select agents who minimize the amount of communication they do before defecting. Though these dynamics may be relatively slow, the system will eventually return to stage 1, and the cycle begins anew.

Previously we found that machines with more states (or languages with more tokens) tend to have more, and longer, outbreaks of cooperation. Analysis indicates that such systems often develop more complex communications before cooperation is established (for example, two or three tokens are exchanged before cooperation is achieved), and we suspect that the observed differences in epoch structure occur because with more states (tokens) such paths are both easier to create and harder to mimic. Mimicry is more difficult to achieve under these conditions because mutations are more likely to disrupt other key elements of the automaton, and thereby destroy the ability in the offspring to perfectly mimic the parents. We have also observed cases where more complex machines are able to subvert potentially deadly mimics by, for example, responding to communication in a way that locks the mimic into self-destructive behavior. Additional states and tokens also permit the simultaneous existence of multiple communicative pathways that lead to cooperation. Such a diverse population is able to better withstand the attacks of mimics on individual pathways.

Analysis of our experiments indicates that the five-stage cooperative cycle described above seems to account for the emergence of cooperation (and, in those cases to date, where it was apparently violated, closer analysis indicates that a sensible variation was responsible). Clearly what we are observing is a very interesting ecosystem of strategies that coevolve with one another. While our theory captures a major pathway for the emergence of cooperation and is consistent with some simple analytic arguments (see, for example, Robson, 1990), we suspect that given the rich ecological and coevolutionary dynamics involved in the system, other pathways might be possible.

4. Conclusions

Communication is a key mediator of many different types of agent interactions. Here we developed a framework from which to analyze the emergence of communication in an adaptive system. The framework provides the ability to study a system composed of agents who evolve the ability to strategically send and receive communication. We illustrated this framework by analyzing the relationship between communication and cooperation in a single-shot Prisoner's Dilemma. We found that, contrary to initial expectations, communication allows the emergence of cooperation in such a system. Moreover, the processing and language complexity inherent in the communication system can be systematically tied to the observed behavior. The modeling framework also enabled us to explore and verify a key hypothesis about the likely mechanism that allows cooperation to emerge.

The basic framework developed here has many possible extensions. A number of exogenous factors can influence communication, for example, tokens could be passed sequentially (versus simultaneously), noise could alter the fidelity of the communication channel, or communication costs could be explicitly introduced so as to limit the number of tokens that are passed. Obviously, such factors can be analyzed within the above model. As previously mentioned, the automata representation creates a class of models rich in possibilities (essentially, any system that communicates and reacts to discrete inputs and outputs to determine a final discrete action). Thus, a variety of models across many different fields should now be amenable to analysis. For example, in the social sciences, models of organizational communication, economic bargaining, and political action could be investigated. The general approach developed here should open up a variety of phenomena to the systematic exploration of endogenous communication.

Further reading

Holland, J.H., 1975. Adaptation in Natural and Artificial Systems. University of Michigan Press, Ann Arbor.

Acknowledgements

The authors wish to thank Walter Fontana for helpful comments. Research support was provided by the NSF (SBR 97–10014, SBR 95–21914, and the graduate fellowship program) and the Santa Fe Institute.

References

Axelrod, R., 1984. The Evolution of Cooperation. Basic Books, New York.

Axelrod, R., Dion, D., 1988. The further evolution of cooperation. Science 242, 1385–1390.

Holland, J.H., 1993. The effect of labels (tags) on social interactions. Santa Fe Institute Working Paper 93-10-064.

Holland, J.H., Miller, J.H., 1991. Artificial adaptive agents in economic theory. American Economic Review Papers and Proceedings 81, 365–370.

Lindgren, K., Nordahl, M.G., 1994. Artificial food webs. In: Langton, C.G. (Ed.), Artificial Life III. Addison-Wesley, Reading, MA.

Miller, J.H., 1988. The evolution of automata in the repeated Prisoner's Dilemma, in Two Essays on the Economics of Imperfect Information. Ph.D. Dissertation. University of Michigan, MI.

Miller, J.H., 1996. The coevolution of automata in the repeated Prisoner's dilemma. Journal of Economic Behavior and Organization 29 (1), 87–112.

Riolo, R.L., 1996. The Effects of Tag-Mediated Selection of Partners in Evolving Populations Playing the Iterated Prisoner's dilemma. Santa Fe Institute Working Paper 97-02-016.

Robson, A.J., 1990. Efficiency in evolutionary games: Darwin, Nash and the secret handshake. Journal of Theoretical Biology 144, 379–396.

Stanley, E.A., Ashlock, D. Tesfatsion, L., 1994. Iterated Prisoner's dilemma with choice and refusal of partner's. In: Langton, C.G. (Ed.), Artificial Life III. Addison-Wesley, Reading, MA.

Steels, L., 1996. A self-organizing spatial vocabulary, Artificial Life Journal 2 (3).

Werner, G.M., Dyer, M.G., 1992. Evolution of communication in artificial organisms. In: Langton, C.G., Taylor, C., Farmer, J.D., Rasmussen, S. (Eds.), Artificial Life II, Addison-Wesley, Reading, MA.