Collective Stochastic Versions of Playable Games as Metaphors for Complex Biosystems: Team Connect Four

IL'YA SAFRO AND LEE SEGEL
Department of Computer Science and Applied Mathematics, Weizmann Institute, Rehovot, Israel

Received November 25, 2002; revised March 13, 2003; accepted March 24, 2003

In an effort to better understand complex biological systems, the game "Connect Four" is generalized to be a stochastic contest between two teams. Members of each team typically possess "sensors" that provide some information on the nearby deployment of the pieces. Sensing something interesting increases the probability that a given team member will move. Simulations show the relative strengths of various sensor weightings and thereby cast some light on the use of sensors in more general complex autonomous systems. © 2003 Wiley Periodicals, Inc.

Key Words: goals; sensor coordination; random play

This study is motivated by the desire to illuminate complex biological activities such as those of the immune system and the metabolic system. Our immune system is pitted against an ever-changing assault by myriads of rapidly evolving pathogens. Our metabolic system is a multipurpose chemical factory that continually adjusts its output in response to shifting conditions of external and internal environments. Both systems are *distributed* in two senses (Terms being defined are shown in italic). Distributed systems are composed of many myriads of agents (a trillion cells in the immune system, untold enzyme molecules in the metabolic system), so that "responsibility" for system actions is widely distributed through a large population. Moreover, the agents of the system are distributed in space. The numerous and spatially dispersed agents act in an *autonomous* fashion, in that there is no commanding "boss."

Researchers are beginning to examine the overall organizational principles of the immune and metabolic systems. Such examinations have obvious scientific and medical motivations, but they also serve the more general purpose of revealing design principles of large-scale autonomous systems that have undergone eons of evolution and therefore presumably have become efficient in some sense or senses. These principles may well be relevant to other distributed autonomous systems, for example those of certain program architectures in artificial intelligence [1].

Correspondence to: Lee Segel, E-mail: lee@wisdom.weizmann.ac.il

FIGURE 1

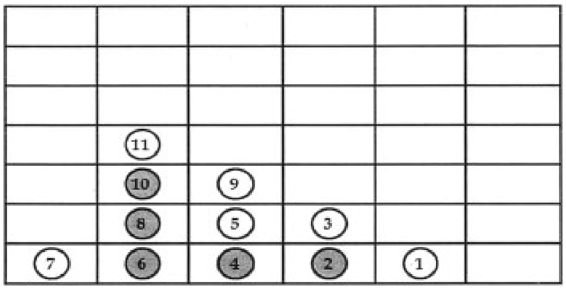

The successively numbered circles represent successive plays in a game of Connect Four. White wins on the 11th move by constructing four pieces in a (diagonal) row.

CONSTRUCTING A GAME TO GIVE INSIGHTS ON COMPLEX SYSTEMS

One standard research strategy is to construct "toy models" of complex phenomena. The terminology has significance; "toys" are fun to play with, and one can learn a lot from playing with good toys. Here we illustrate the idea that a certain kind of toy, a game, can be the focus of investigations that are interesting in themselves and also can perhaps lead to deeper understanding of the behavior of complex systems.

The game that we have chosen is "Connect Four," a generalization of tic-tac-toe. White and black checkers are dropped into a vertical frame. The checkers land on the bottom-most empty space in the column wherein they are dropped. Once a checker is dropped in place it is *dormant;* it never moves. The game ends when a player succeeds in constructing a line of checkers—vertical, horizontal, or diagonal—that is four checkers long (Figure 1). Also check a search engine for references; under "*connect four*" *game* Google lists more than 14,000 but, for example, all the first 10 references are relevant.

Ordinary games such as Connect Four and checkers pit two intelligent players against one another. To serve as biological models, such games must be generalized so that they represent confrontations between two multiagent entities, such as the struggle between the immune system and pathogens. The opposing entities should be distributed autonomous systems. Thus the individual agents should not be subject to the orders of an overarching intelligence, but rather should try to make a sensible move by gauging the state of play in their locality. Each multiagent entity (which we will call a *team*), should be much more like a partially motivated rabble than a perfectly disciplined army of fighting sages. Information is fragmentary and agents have limited cognitive abilities; teams are myopic mobs of mediocrities. Though not without primitive wisdom, an ill-informed mediocrity will not act predictably in the face of complex challenges, so that "decisions" on what move to make should be regarded as stochastic.

These thoughts lead to the following version of *Team Connect Four* (TC4). The playing field consists of N columns. There are N *team members* on each team, one of each color per column. The white team has the first move. With the aid of various sensors of local conditions, the *white team member wi* at the top of the *i*th unfilled column computes its *activity* a_{wi}. By normalizing the various a_{wi} values, a probability P_{wi} for a white team member to move is calculated as follows, where index j runs over all unfilled columns:

$$P_{wi} = a_{wi} / \left[\sum_j a_{wj} \right]. \qquad (1)$$

The white team moves according to these probabilities. Then the black and white teams continue to move alternately according to such rules. To avoid edge affects, in all simulations we wrapped the game board around a cylinder, so that the leftmost column was regarded as adjacent to the rightmost column.

This article is an initial case study of various versions of TC4, as examples of collective stochastic team games. We hope to convince the reader that such games are of interest in themselves, forming types of "artificial life" that are fun to design and to simulate; that they suggest mathematical challenges; and that they can shed light on the behavior of "real" distributed autonomous systems, such as those found in biology.

PURELY RANDOM STRATEGIES

The "control" game of TC4, the simplest version against which others should be compared, has the property that both sides move randomly. But in the present context "move randomly" can sensibly be defined in the following two ways. In *TC4-D1* (D1 means "definitely 1"), the two sides alternate in making a single move; the choice of column is made according to the probabilities of (1). In *TC4-E1,* (1) is interpreted as the probability that white moves in column *i*. One move per turn is expected (E1) but in any given turn there may be no move, or there may be more than one.

The first simulations that we shall report used a 5×6 board, 5 rows and 6 columns, when both teams used the same random strategy. It is anticipated that going first is advantageous. Indeed, in TC4-E1, simulation of 10^4 games gave 0.61 as the fraction F_1 of wins by *team one* (the team that moves first). One would expect the advantage of going first would decrease with board size. Indeed, we found that $F_1 = 0.57$ for board sizes of 10×10 (10^4 simulations), whereas $F_1 = 0.51$ for 100×100 boards (500 simulations).

Ties are possible (the board is filled and neither side has four in a row), but we found them very unlikely. For example, only 10 of 10,000 games ended in ties in the 5×6 simulation of TC4-E1. The number of ties will thus not be reported henceforth.

Upon examining our simulation results, those already obtained and those to be presented, it is natural to ask whether one can trust observed trends and numerical estimates of various probabilities. In order to estimate their accuracy consider our simulations as random processes that generate "indicator variables": 1 when the first team wins, and 0 when the first team loses. Let p be the fixed probability that the first team wins. Then in N games the variance δ^2 is given by $p(1 - p)/N$ [2, p. 325]. The team 1 win fraction, F_1, is the observed value of p. We can thus take $F_1(1 - F_1)/N$ as a simple estimate of the variance. Taking $F_1 = 1/2$ gives the maximum value of this estimate, and indeed often F_1 is close to 1/2. Thus we use the formula $\delta = 1/2\sqrt{N}$ as a rough but sufficiently accurate estimate of the standard deviation δ. When $N = 10^4$, then $\delta = 0.005$; $N = 10^3$ implies $\delta = 0.02$. These estimated standard deviations are typically small compared to the estimated means, so that indeed one can generally trust the trends that the simulations display. In particular, there is statistical significance to our simulation results (above) that F_1 is a decreasing function of board size for TC4-E1.

The reader is invited to test his/her intuition by guessing the superior strategy, TC4-D1 or TC4-E1 . . . . Our simulations showed a clear advantage to strategy E1. (In reporting our results we will use the notation [*a*, *b*] to describe a game where the first team follows procedure a and the second follows procedure b.) In 10,000 repetitions of [TC4-D1, TC4-E1] we found $F_1 = 0.39$, whereas in 10,000 repetitions of [TC4-E1, TC4-D1] the result was $F_1 = 0.72$.

To try to understand why E1 is superior to D1, consider the following conclusions from the binomial distribution. With $N = 10$ and equal probabilities $p = 0.1$ for a move in each of the 10 columns (strategy E1), one obtains $Z = 0.35$, $U = 0.39$, $T = 0.19$, $\theta = 0.06$, and $F = 0.01$ for the probabilities, respectively, of zero, one (Unity), two, three, and four moves. With probability $T = 0.19$, strategy E1 permits two moves in a row, but this is much less than the corresponding probability $Z = 0.35$ for this occurrence of two successive moves with strategy D1. Only with probability $\theta + F = 0.07$ is there an advantage to strategy E1. For intelligent players of ordinary Connect Four, three or four moves in a row is enough to yield a win under almost all circumstances, but of course this is not the case in TC4 when teams choose their moves randomly. Nonetheless, the simulations show that the rather rare possibility of executing three or four moves in a row gives decisive advantage to strategy E1.

Strategy D1 is closer to that of the standard game of TC4 (each side makes one move at a time). However, this strategy requires coordination between team members, so that only one member will move during each turn. The intricate subject of team coordination must be left to the future; accordingly, from now on we will only consider E1 as our control random strategy.

ADVANTAGES OF A SINGLE ENVIRONMENTAL SENSOR

Metabolic enzymes have *regulator sites,* often several of them, to which key metabolites bind. Such binding typically changes the shape of the enzyme and thereby alters its catalytic activity. The immune system employs numerous cellular *receptors* that bind various molecules; such binding events give information on the state of the immune system's progress toward its numerous overlapping and conflicting "goals" and thereby lead to alteration and improvement of immune performance [6]. This illustrates the key role of *sensors* in modulating system performance. The issues here include "how significantly can sensing local partial information contribute to improved performance" and "how can the system best coordinate information from a variety of different sensors." Such considerations led us to make some explorations of the role of sensors in TC4.

Let us first examine the role of a primitive *simple sensor s.* This allows a potential mover in a given column to determine whether or not there is a piece of the same color at the top of the dormant column below. If there is, then the default activity of that mover (of index i) in the absence of information, a_{*i0}, is multiplied by a *weighting factor* f_S to give activity a_{*i}:

$$a_{*i} = a_{*i0} f_S \quad (* = w \quad \text{or} \quad b). \tag{2}$$

Figure 2 shows simulations wherein each member of team one possesses a simple sensor, which weights a possible activity more heavily, according to (2), when a dormant piece of the same color is sensed. Team two follows the random strategy E1. As f_s ranges from 1 to 100, the fraction of wins by team one continually increases, albeit in a saturating fashion. Evidentally, some information is better than none.

We now let teams contend where each has a different sensor. In addition to the simple sensor *s*, we also consider a *simple-left sensor sL.* A potential mover with sensor *sL* can detect whether or not there is a dormant piece of the same color at the top of the column immediately to the left of the potential mover. It is to be expected that sensor *sL* provides less relevant information than sensor *s.*

Let us now examine the influence of the weights f_S and f_{sL} on the results of TC4-E1 (each team has one expected move). Suppose that team one possesses sensor *s* and team two possesses sensor *sL.* (We denote this game by [*s*, *sL*]). It is seen from Figure 3, graph 1, that in this situation, for the range of parameters tested and for fixed f_{sL}, team one has an advantage that becomes more and more appreciable the larger is f_S. The result makes sense: as was illustrated in

FIGURE 2

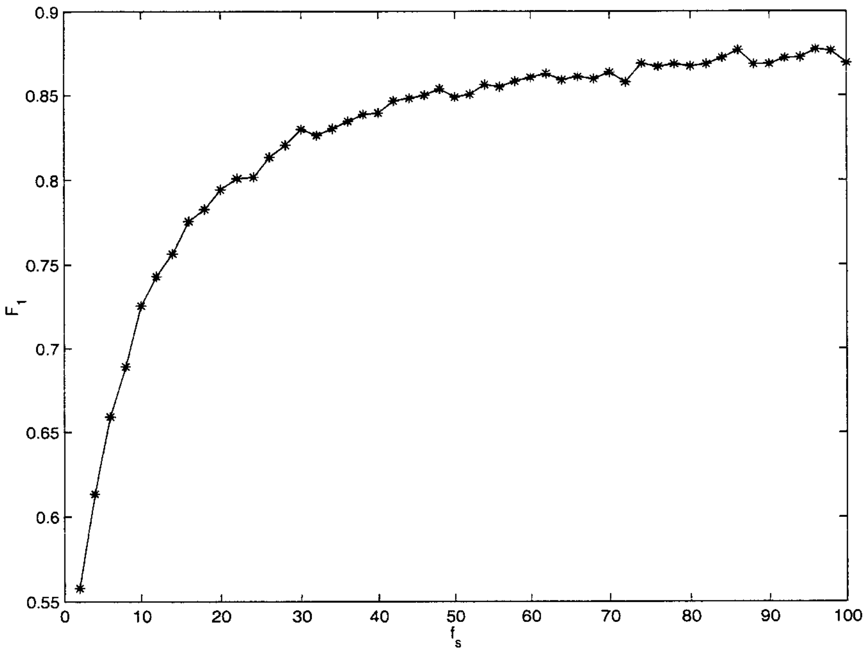

Fraction F_1 of wins in 10,000 10 × 10 games of team one, who possess a "simple sensor *s*." Such a sensor weights a possible move more heavily by a factor of f_s (horizontal axis), according to (2), when a dormant piece of the same color is sensed. Team two follows the random strategy E1.

Figure 2 for a completely ignorant opponent (who moves randomly), also here with a *comparatively* ignorant opponent the more weight that is given to a useful sensor, the better. The usefulness of sensor s is also shown in graph 3, where both teams possess this sensor; strong weighting of sensor s for team one greatly amplifies the fraction of wins by team one. Direct evidence for the superiority of sensor s compared to sensor sL is obtained by comparing graph 1 and graph 2. When team one possesses a heavily weighted sensor s, its advantage is markedly larger compared with the case when team one possesses a heavily weighted sensor sL. Additional evidence is provided by a simulation where team one has weights $f_{sL} = 20$ and $f_S = 2$ for sensors sL and s, while team two has the reversed weighting $f_S = 20$ and $f_{sL} = 2$. In spite of going first, team one only wins about half the games (not shown).

In Figure 4 it is seen that for fixed f_S in team one, the advantage of team one decreases and hence the advantage of team two increases, as f_{sL} increases. Thus it is helpful for a team member to increase the weight even of inferior information if that is the only information that is available.

Let us now consider a sensor that can provide two pieces of information, not just one as heretofore. Accordingly, suppose that a team member can detect the topmost dormant piece in two locations, the column of the member and the column to the left of the member's column. We use a subscript 1 when a member detects its own color and a 0 otherwise, with the first subscript corresponding to the column of the member. The corresponding weighting factors are taken to be $f_{00} = 1$, $f_{11} = 100$, $f_{10} = f_{01} = s$, where s is a parameter. Figure 5 provides simulation results for a contest between a randomly playing team one (strategy E1) and a second team with two-bit sensors, as above. It is seen that the fraction of wins by team 2 (F_2) is an decreasing function of s, i.e. stronger weighting of the comparatively insignificant sensory data 01 or 10 (compared to 11) leads to poorer results.

COORDINATING TWO SENSORS

An important conclusion was obtained by running a number of simulations with various pairs of weights f_s and f_{sL} when each team possesses both the inferior sensor sL and the superior sensor s (Figure 6). A representative result is the following. Suppose that for team one the weights are fixed at $f_{sL} = 7$, $f_s = 3$. This amounts to cutting the depicted surface by a plane perpendicular to the x-axis at $x = 3$. For team two, consider the weights $f_{sL} = y$, $f_s = 10 - y$, for $y = 1, 2, \ldots, 10$. The fraction of wins by team two has a maximum when y is about 7. Given the errors inherent in the stochastic simulations, the maximum could also occur for $y = 6$ or 8, but the presence of a maximum seems unmistakable for all fixed values of f_{sL} between 3 and 10 (inclusive). Thus this

FIGURE 3

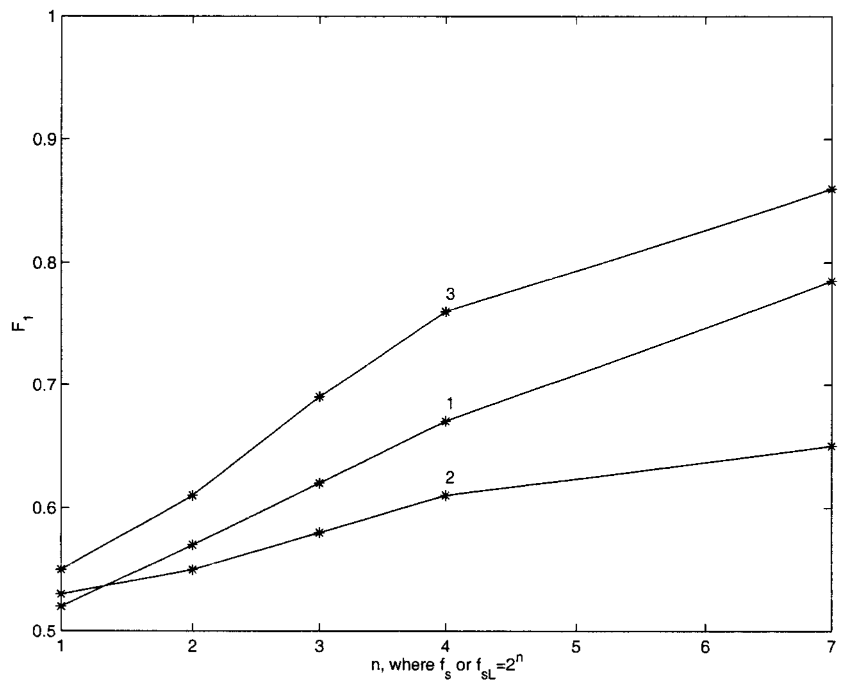

Graphs give fraction F_1 of team one wins in 5000 10 × 10 games. (1) Team one: sensor *s* with $f_s = 2^n$; team two: sensor *sL* with fixed $f_{sL} = 4$. (2) Team one: sensor *sL* with $f_{sL} = 2^n$; team two: sensor *s* with fixed $f_s = 4$. (3) Team one: sensor *s* with $f_s = 2^n$; team two: sensor *s* with fixed $f_s = 4$.

FIGURE 4

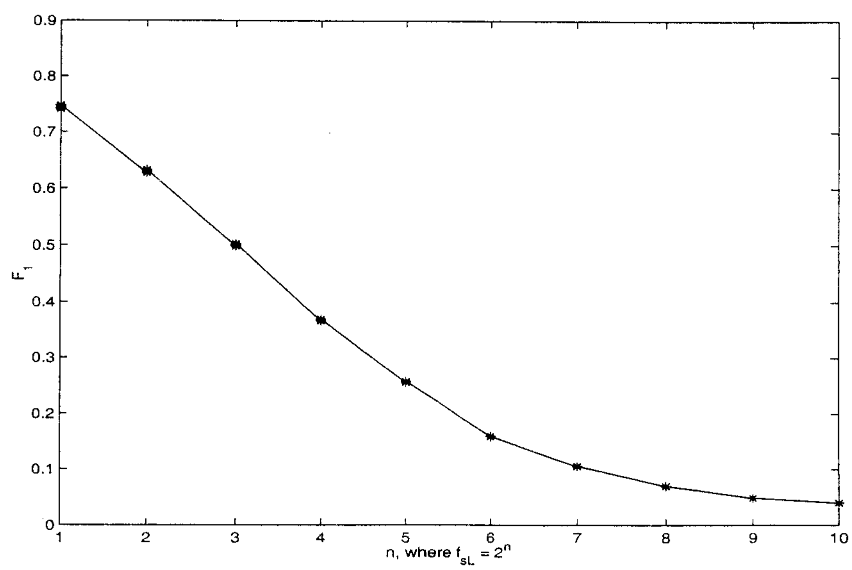

Team one: sensor *s* with fixed $f_s = 10$. Team two: sensor *sL* with $f_{sL} = 2^n$. Board is 100 × 100; 10^4 games per point.

FIGURE 5

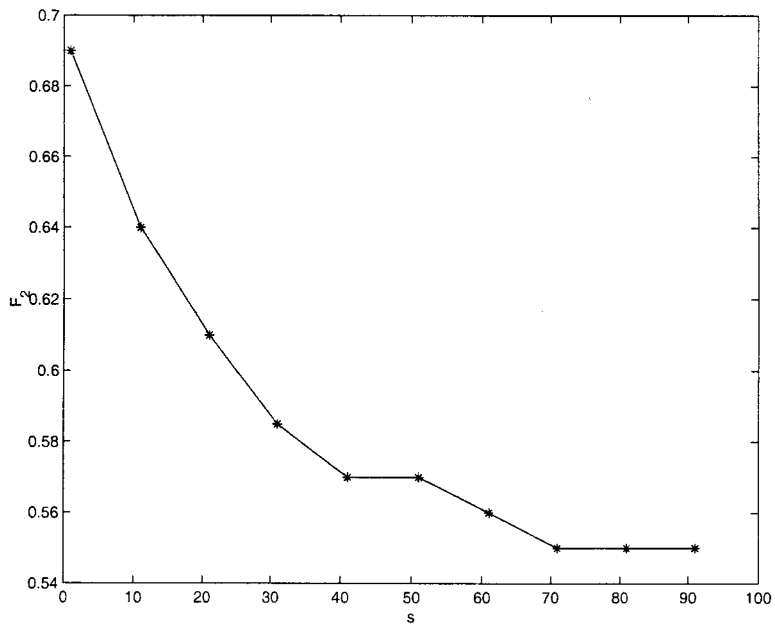

Fraction of wins for team 2, who possess sensors that can provide two bits of information. Team one: random play E1. An increase in the parameter *s* corresponds to heavier weighting of relatively insignificant information. 10×10 board; 5000 games per point. See text.

set of simulations demonstrates the intuitively reasonable property that when a team member or other agent possesses more than one sensor, there is expected to be an optimal relative weighting of the (partial) information obtained from the sensors.

The "representative result" mentioned in the previous paragraph is intuitively understandable. Figure 6 shows an optimum y in the weighting $f_s = 10 - y$ for maximizing the win fraction F_2 for the (second) team member with the superior sensor s. Too small a weight f_s means that the

FIGURE 6

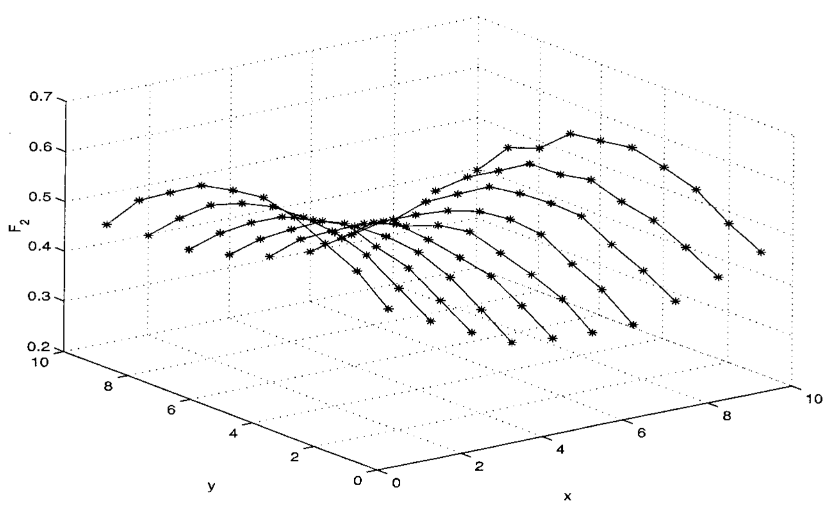

Fraction F_2 of wins by team two when both teams possess both sensors s and sL. For team 1 the weighting factors are $f_s = x$, $f_{sL} = 10 - x$; for team two, the weighting factors are $f_{sL} = y$, $f_s = 10 - y$: $x, y = 1, 2, \ldots, 10$. 10×10 board; 5000 games per point.

information from sensor s is underutilized. Too large a weight f_s means that the information from sensor sL is underutilized. (Think of the extreme situation of a relatively enormous weight for sensor s.) Consider now y fixed and x varying. There is expected to be an optimum x, which maximizes the win probability of the team with sensors sL. But this is the first team; thus Figure 6 shows as expected a minimum for the win probability F_2 of the second team.

AN ASYMMETRIC GAME

To the best of our knowledge all standard games are symmetrical, i.e., they have the property that the "pieces" and rules for both sides are the same. (There is the somewhat trivial exception of handicaps, when a superior player accepts an initial disadvantage in the number of his pieces.) A symmetrical game is appropriate as a model, say, of the competition between differently mutated organisms in a given species. But the "game" between pathogens and immune systems is not symmetrical, for the rules by which viruses and bacteria act are very different from the rules that govern lymphocytes and macrophages.

It is a challenge to try to construct interesting asymmetric games. In fact, we have already seen an example of such a game, when we pitted the two random strategies (D1 and E1) against one another. Here the procedures for the two sides differ, but the "aim of the game," construct "four in a row," is the same for both sides. One way to modify classical Connect Four so that the game's aim differs between the two sides is to regard the object of the game as constructing different sets of patterns for different sides. Player one, for example, wins when either a vertical or horizontal line of four of his pieces is constructed. Player two wins upon construction of a diagonal four-in-a-row or a square of four pieces.

A topic for future study is designing and simulating asymmetric versions of Four in a Row that can mimic various essential aspects of the conflict between the immune system and pathogens. Here is a tentative example. Team one, representing the pathogens, goes first and tries to "win" by constructing a row of four team one members. We think of the task of building this structure as analogous to the pathogens "goal" of constructing an entity, for example, a sneeze, that can propagate them to a new host. Team two, representing the immune system, can *kill,* i.e., cause the removal of, pathogens of team one, if 2 team two members eventually land on a dormant column surmounted by one or more members of team one. If this occurs, all the team one members in the column are removed. The resulting "collapse" leaves a column that is solely composed of team two members. Pathogens of team one can kill a single team two member at the top of a column if a team one member lands on that column.

Team two (the immune system) wins if all of team one are killed or if the board is filled without a win by team one (pathogens survive but do not propagate). In this game, the asymmetry in killing ability reflects the relative simplicity by which a single pathogen can do damage, compared with the fact that, at least for the adaptive immune system, coordination among different types of effectors is required for the immune system to dispose of pathogens.

In running the game just described, we found that the pathogens always won. We then made life progressively more difficult for the pathogens, by requiring construction of X-in-a-row for team one to win. Figure 7 shows the results. For $X = 14$, in a 20×20 board, the asymmetric game provides a fairly even match between pathogens and the immune system. There is thus a "proof of principle" that evenly matched asymmetric games can be designed by quantitative manipulations.

We have made an initial step toward inventing games that exemplify aspects of the conflict between pathogens and the immune system. Further steps would involve the introduction of sensors, analogously to what has been done above. An important additional step is to introduce propagation to other members of the same team of the information gained by one member's sensor—and disturbance of this propagation by the opposing team (pathogens sabotage the signals of their host's immune system).

DISCUSSION

We have examined a novel class of games, which are of interest in themselves and which also illustrate matters of concern in the study of distributed autonomous systems. Classically, there seem to be two main types of games. The first type is games in the sense of von Neumann and Morgenstern, games that feature a payoff matrix and that are normally analyzed under the assumption that the opponent will make his/her best possible move. Such games are not considered here, although there is relevant, important and interesting research concerning *interaction games* that combine incomplete information, local interaction, and randomness [3].

A second classical type of games, *combinatorial games,* have perfect information, no chance moves, and an outcome (win, lose, draw). See Fraenkel [4] for a brief survey and an extensive bibliography. Among combinatorial games Fraenkel distinguishes between *people-games,* games people play such as chess and go, and *math-games,* games mathematicians play such as Nim. He points out that people-games invariably are computationally intractable. For example, Fraenkel and Lichtenstein [13] showed (roughly speaking) that there are board positions in generalized $N \times N$ chess for which calculations of order exp N are required to determine whether these are winning positions. True, there are excellent chess playing programs that combine heuristic board evaluation schemes with massive "look ahead" computations, but this emphasizes the computational difficulty in another way. Yet, paradoxically, in contrast to math-

FIGURE 7

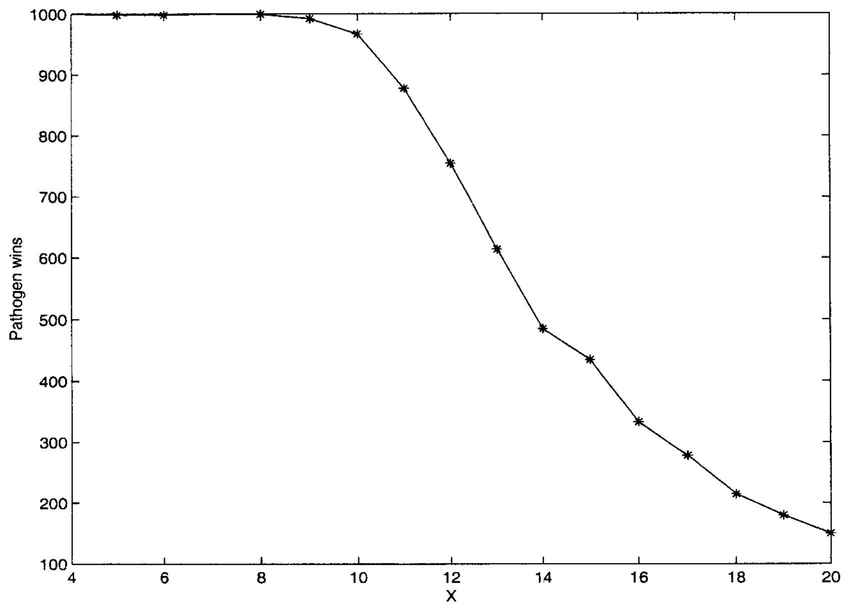

Number of “pathogen wins in 1000 simulations of 20 × 20 asymmetric games pitting “pathogens” versus “immune systems,” for various values of the parameter *X*. Each side uses the random strategy E1, where one move per side is expected. To win, pathogens must construct a line of length *X*. See text for other rules.

games, computationally difficult people-games exhibit *board feel* [4]. The presence of board feel means that even non-experts can often successfully judge which of the two sides is ahead. The possibility of exploiting such intuition supports the selection of a people-game, Connect Four, as the basis for our study.

People-games normally pit one human versus another. To serve the purpose of modeling complex systems, people-games must be generalized to become competitions between autonomous groups. Although we do not make the distinction here, one might wish to term such groups *mobs* if group members do not communicate and *teams* if they do. *Armies* are (nonautonomous) groups run by an overarching intelligence, as in standard chess.

Our first analyses of “team connect four” (TC4) examined random strategies, as controls. Two possibilities were considered, wherein each team alternately (i) moves once, or (ii) moves so that its expected number of moves is one. Our simulations showed, somewhat surprisingly, that if these two possibilities are pitted against each other then possibility (ii) is superior. Here and for all the simulation results that we obtained, it remains for the future to try to prove theorems that are suggested by these results. We did succeed in proving a result for possibility (i) in random play of tic-tac-toe, the granddaddy of TC4. By enumerating the various probabilities, we found that (0.514, 0.305, 0.181) were, respectively, the fractions of wins by the first player, wins by the second player, and draws.

Sensors of specialized information are found throughout biology, and it is of interest to develop intuition for their effectiveness. Accordingly, most of our simulations of TC4 concerned the relative advantages of various sensors of partial information concerning the board situation. This seems novel; the common way to take into account partial information, for example, in cellular automata and Checkered Life (see below), is to define once and for all a local neighborhood for all agents. Perfect knowledge of what goes on in this restricted neighborhood determines agent actions.

Biological sensors measure some variable that is important to the system; then transduction machinery converts this information into a suitable (i.e., evolutionarily selected) modification of behavior. In our games, behavioral modification is expressed by changing the relative probabilities of various “moves” in the games. In (1) we normalized activities a_{wi} and a_{bi} to obtain probabilities. This procedure requires communication among all team members. Communication can be made unnecessary if $a_{wi}\delta t$ and $a_{bi}\delta t$ are respectively regarded as the probabilities that the ith white or black player moves during a short time δt. Now each team member’s action is determined independently by each individual; the team becomes a "mob" according to the definition given above.

We found situations where there was an optimal weighting of sensory information. It must be kept in mind, however, that such optima generally exist "when all other things are equal." When conditions change, the optima change; this means that efficient use of biological sensors requires continual feedback to track the changing circumstances [6].

One precedent for using games as metaphors is Papadimitriou's examination of decision under uncertainty as a "new sort of game, in which one opponent is 'disinterested' and plays at random, while the other tries to pick a strategy that maximizes the probability of winning—a 'game against nature'" [7]. Noteworthy is the presentation of Eigen and Winkler [8] of newly invented board games (but not people-games) as a vehicle for demonstrating the essence of profound scientific concepts, particularly in statistical physics but also in biology and social science.

The present article can be regarded as a continuation of a study of a team version of checkers called Checkered Life. This game was suggested to help clarify assertions that distributed autonomous systems can profitably be regarded as simultaneously pursuing multiple overlapping and conflicting goals [10]. Examples of such goals in ordinary checkers are "jump," "control the center," "get a king," and "don't move from the back row." Checkered Life is played on a board that is 8 rows deep, as usual, but is N columns wide, where N is large. If white is to play, each member of the white team examines its neighborhood and calculates its activity based on that player's perception of how much its move can advance the various checker goals. After every standard move there is a probability that any legal square can "give birth" to a white or black checker. Segel [10] suggested that there be a tiny probability that the game is terminated by a *deus ex machina* at any given move; then the team with the larger number of surviving members is deemed the winner. Alternatively, the game can be deemed ended when one team becomes extinct. Perhaps more biological than prescribing situations for a "win," in Checkered Life and other games that somewhat mimic biology, is the possibility of tracking statistics of the evolving state of play, without paying undue attention to eventual ultimate extinction. Indeed, in such combats as that between vertebrate immune systems and attacking pathogens, each side is "interested" in its own survival, which is not necessarily promoted by the extinction of the other side.

A classical article in artificial intelligence deals with how simultaneously to pursue the multiple goals of checkers, by weighting their importance, and how to learn to do better by adjusting the weights according to experience [5]. Representative of recent research along this line is an article by Chellapilla and Fogel [11]. The present article is in accord with their assertion that "Intelligence pertains to the ability to make appropriate decisions in light of specific goals and to adapt behavior to meet these goals in a range of environments. Mathematical games provide a framework for studying intelligent behavior." Chellapilla and Fogel *evolve* neural networks for skillfully playing games including tic-tac-toe and checkers. We have not made explicit use of evolution in our studies, but our results concerning comparative performance measures can be regarded as providing fitness data for evolutionary calculations.

As an embodiment of a complex biological-like system, Checkered Life seems to suffer from the disadvantage that it is not amenable to rapid simulation, thereby making difficult the compilation of reliable statistics. As we have seen, such statistics can be compiled by running various versions of TC4. These seem to approximate the delicate balance required in a simulation game between fidelity to the essence of some important biological features and adherence to the principle that game is better if its rules are simple.

Probably the part of this article that has most relevance to biology (and other distributed complex systems) is its treatment of sensor weighting. It might seem that this matter has already been covered thoroughly in many papers and books concerned with the subject of "sensor fusion" or "data fusion" [12,9], but these studies generally treat the blending of different pieces of sensory information so as to optimally achieve a single well-defined goal. Here we are concerned with complex situations that can be regarded as characterized by multiple overlapping and contradictory goals. Note, however, that in testing sensor weightings we did not make reference to the abstraction of "goals" but rather tested the different possibilities "in the field" by actually "playing the game."

The assignment by an observer of abstract goals to teams that tend to win competitions can be accomplished by noting what sensors are heavily weighted. Thus the success of teams possessing a simple sensor suggests the conclusion that a goal for team i in TC4 is to extend any existing columns of team i pieces, $i = 1, 2$. This is not trivial, for it is not clear that extending a team's own 4-in-a-row possibilities is more important than blocking nascent 4-in-a-row possibilities for the other team. But such seems to be the case, perhaps not surprisingly, when teams move with a high degree of randomness. At all events, our game-playing has exhibited the process of inferring "goals" from evidence of competitive success.

ACKNOWLEDGEMENTS

The idea for this work developed from a conversation of LAS with J-P Boon, which revealed how much interest can be generated by carefully selected "toy" problems [14]. We thank A.S. Fraenkel for sharing some of his wide knowledge of games, and for valuable comments on a draft of this article. M. Mitchell made many suggestions, which greatly improved the article.

REFERENCES

1. Segel, L.A.; Cohen, I. The Immune System and Other Distributed Autonomous Systems; Oxford University Press: Oxford, 2001.
2. Dwass, M. Probability and Statistics; W. A. Benjamin: New York, 1970.
3. Morris, S. Interaction Games: A Unified Analysis of Incomplete Information, Local Interaction and Random Matching, University of Pennsylvania—Institute for Economic Research Report 97-02, 1997.
4. Fraenkel, A.S. In: More Games of No Chance, Proc. MSRI Workshop on Combinatorial Games, July, 2000, Berkeley, CA, MSRI Publ. Nowakowski, R.J., Ed.; Vol. 42, Cambridge University Press: Cambridge, 2002; pp 475–535.
5. Samuel, A.L. In: Computers and Thought; Feigenbaum, E.A.; Feldman, J., Eds.; McGraw-Hill: New York, 1963.
6. Segel, L.A.; Lev Bar-Or, R.J. Immunology 163, 1999, 1342–1349.
7. Papadimitriou, C.H.J. Comput Syst Sci 31, 1985, 288–301.
8. Eigen, M.; Winkler, R. Laws of the Game (Trans. by Robert and Rita Kimber); Princeton University Press: Princeton, NJ, 1993.
9. Block, I. IEEE Trans Syst Man, and Cybernetics (Part A) 1996, 52–67.
10. Segel, L.A. Complexity 5, 2000, 39–46.
11. Chellapilla, K.; Fogel, D. Proc IEEE 87, 1999, 1471–1496.
12. Dawant, B.M.; Garbay, C. Special topic section on biomedical data fusion. IEEE Trans Biomed Eng 46, 1999, 1169–1190.
13. Fraenkel, A.S.; Lichtenstein, D. J Combin Theory (Series A) 31, 1981, 199–214.
14. Boon, J.-P. J Stat Phys 2001, 102, 355–360.