Convergence of the Iterated Prisoner's Dilemma Game *

Martin Dyer[†] Leslie Ann Goldberg[‡] Catherine Greenhill[§]
Gabriel Istrate[¶] Mark Jerrum[‖]

September 21, 2002

Abstract

We consider a stochastic process based on the iterated prisoner's dilemma game. During the game, each of n players has a state, either *cooperate* or *defect.* The players are connected by an "interaction graph". During each step of the process, an edge of the graph is chosen uniformly at random and the states of the players connected by the edge are modified according to the *Pavlov* strategy. The process converges to a unique absorbing state in which all players cooperate. We prove two conjectures of Kittock: The convergence rate is exponential in n when the interaction graph is a complete graph, and it is polynomial in n when the interaction graph is a cycle. In fact, we show that the rate is $O(n \log n)$ in the latter case.

1 Introduction

In a two-player *prisoner's dilemma* game [1, 9], each player may choose to *cooperate* or to *defect.* If both players cooperate, they each receive a reward of R points. If both players defect, they each receive only P points. If exactly one player cooperates, he receives S points while his opponent receives T points. The parameters satisfy $T > R > P > S$

*This work was supported in part by the EPSRC Research Grant "Sharper Analysis of Randomised Algorithms: a Computational Approach" and by the ESPRIT Projects RAND-APX and ALCOM-FT.

†School of Computing, University of Leeds, Leeds LS2 9JT, United Kingdom, email: dyer@comp.leeds.ac.uk

‡Department of Computer Science, University of Warwick, Coventry CV4 7AL, United Kingdom, email: leslie@dcs.warwick.ac.uk

§Department of Mathematics and Statistics, University of Melbourne, Parkville VIC 3502, Australia, email: csg@ms.unimelb.edu.au Supported by an Australian Research Council Postdoctoral Fellowship

¶Center for Nonlinear Science and CIC-3 Division, Los Alamos National Laboratory, Mail Stop B258, Los Alamos, NM 87545, email: istrate@lanl.gov

‖School of Computer Science, University of Edinburgh, King's Building, Edinburgh EH9 3JZ, United Kingdom, email: mrj@dcs.ed.ac.uk

and $2R > T+S$. Thus, in a single round, it is best for each player to defect even though this is not globally optimal. In the *iterated* prisoner's dilemma game, rounds are played repeatedly and players may base their decisions on the outcomes of previous rounds. Empirical evidence [1] indicates that an effective strategy for this game is the so-called "Pavlov" strategy: If a player is "rewarded" with T or R points during a given round, then he repeats his previous move next time. If he is "punished" with P or S points then he does not repeat his previous move. A small case analysis reveals that the Pavlov strategy can also be stated as follows: A player cooperates if and only if he has made the same choice as his opponent during the previous round.

Kittock [5] studied the Pavlov strategy[2] in a distributed setting: n players are connected by an "interaction graph". During a round of the game, an edge of the graph is selected uniformly at random. The two players connected by the edge play one round, using the Pavlov strategy. That is, if a player agreed with his (previous) opponent last time he was chosen to play, then he cooperates this time. Otherwise he defects. As long as the interaction graph is connected, the game converges to a unique *absorbing state* in which every player cooperates. Kittock was interested in determining the *absorption time* — that is, the time required to reach the absorbing state. He provided empirical evidence for the conjecture that the absorption time is exponential in n if the graph is the complete graph, and polynomial in n if the graph is a cycle. In this note, we prove Kittock's conjectures.

The basic method which we use to prove both theorems is to define an appropriate potential function so that the progress of the Pavlov process may be compared to that of a one-dimensional random walk. See, for example, [6]. Similar techniques are often used to analyse the mixing time of Markov chains via *coupling arguments.* In the coupling context, the observed state space is taken to be the set of *pairs* of Markov-chain states, and the potential function is defined to be the distance between the states in the pair, with respect to some metric. See, for example, [7, 8, 2, 4].

2 Preliminaries

We are given a population of $n \geq 4$ players situated at the vertices of a connected graph $G = (V, E)$. Each player has an initial state $X(i) \in \{-1, 1\}$. The 1 values encode the decision "cooperate" and are referred to as *pluses.* The -1 values encode the decision "defect" and are referred to as *minuses.* During each step of the Pavlov process, we

[1]Axelrod [1] hosted a computer tournament in which strategies proposed by game theorists were played against each other. Surprisingly, a simple Markovian strategy won the tournament. That is, the winning strategy has the property that the decision of a given player in a given round depends only on the outcome of the previous round, and not on the rest of the history of the game. Nowak and Sigmund [9] did a computational study of all such Markovian strategies, and found the Pavlov strategy to be "best".

[2]Kittock's paper is described using AI language, but the iterated prisoner's dilemma strategy that he studies (which he calls the zero-memory HCR strategy) is precisely the Pavlov strategy. For more information about the context of Kittock's work and the HCR generalisation of the Pavlov strategy, see Shoham and Tennenholtz's paper [10].

choose a pair $\{i,j\} \in E$ uniformly at random and replace $X(i)$ and $X(j)$ by $X(i)X(j)$. The state X^* with $X^*(i) = 1$ for all $i \in V$ is an absorbing state of this process. If G contains no isolated vertices then X^* is the unique absorbing state and there exists a sequence of moves which can transform X to X^*, for every $X \in \{1,-1\}^V$.

We are interested in the *absorption time*; that is, the time required for the Pavlov process to reach the absorbing state. We investigate two families of graphs, namely cycles and complete graphs. Using a coupon-collector-like argument, Shoham and Tennenholtz [10] proved that, for a large class of strategies, the absorption time for both of these families is $\Omega(n \log n)$. We prove the following theorems, showing that the Pavlov process has optimal absorption time when G is a cycle, and exponential absorption time when G is a complete graph.

Theorem 1 *Let G be a cycle on n vertices and let $\varepsilon > 0$ be given. With probability at least $1-\varepsilon$, the Pavlov process reaches the absorbing state in*

$$\frac{49}{2} n \log\left(\frac{49n}{94\varepsilon}\right)$$

steps.

Theorem 2 *Let T_n denote the absorption time of Pavlov process on the graph K_n starting from a configuration X_0 with fewer than $3n/5$ pluses. With probability $1-o(1)$ we have $T_n \geq (1.1)^n$.*

3 Optimal absorption on cycles

Let G be a cycle on the vertex set $[n] = \{0, \ldots, n-1\}$. That is, G has n edges $\{i, i+1\}$ for $0 \leq i < n$. Here, and throughout the paper, addition and subtraction on vertices is performed modulo n.

We define a potential function $\psi : \{1,-1\}^V \to \mathbb{R}$ to measure the distance of a given state X from the absorbing state X^*. First, we must introduce some terminology. Let $X \in \{1,-1\}^V$ be given. A *run* in X is an interval $[i,j]$ where $0 \leq i,\ j < n$, such that $X(\ell) = -1$ for $\ell = i, i+1, \ldots, j-1, j$ and $X(i-1) = 1$, $X(j+1) = 1$. (It is possible to have $j < i$, since we are working modulo n.) Clearly all runs are disjoint. We can define the set $\mathcal{R}(X)$ of all runs in X. By convention, the all-minuses configuration is *not* considered a run, since it has no bordering pluses.

Suppose that $r = [i,j]$. The *length* of the run r, denoted by $\ell(r)$, equals the number of minuses in the run. We will refer to a run of length ℓ as an ℓ-run. A 1-run will also be called a *singleton* and a 2-run will also be called a *pair*. Then the potential function ψ is given by

$$\begin{aligned}\psi(X) \;=\; & |\{i : X(i) = -1\}| + \beta \cdot |\mathcal{R}(X)| + \gamma \cdot |\{r \in \mathcal{R}(X) : r \text{ is a singleton}\}| \\ & + \delta \cdot |\{r \in \mathcal{R}(X) : r \text{ is a pair}\}|.\end{aligned}$$

The parameters β, γ and δ will be set below. Note that a singleton is a barrier to absorption since a singleton minus cannot be changed to a plus in one step. The singleton must first become part of a longer run. So we set $\gamma > 0$ to penalise singletons. On the other hand, pairs give the opportunity for two minuses to be changed at one step. Thus pairs are helpful, and we reflect this by setting $\delta < 0$. We also set $\beta > 0$. Clearly $\psi(X^*) = 0$ for any values of β, γ, δ since $X^*(i) = 1$ for all i, and $\mathcal{R}(X^*) = \emptyset$. For ψ to be a well-defined potential function, we must also show that $\psi(X) > 0$ whenever $X \neq X^*$. This is achieved if $-2 < \delta < 0$, since there can be at most half as many pairs in X as there are minuses.

3.1 The analysis

We now analyse the Pavlov process using the potential function ψ. Let $X_0 \in \{1, -1\}^V$ be fixed. Clearly if $X_0 = X^*$ there is nothing to prove. So, suppose that X_0 contains at least one minus. Let X_1 be the result of performing one step of the process from starting point X_0. We will find an upper bound for $\mathbf{E}\left[\psi(X_1) - \psi(X_0)\right]$.

Note that each edge overlaps at most one run in $\mathcal{R}(X_0)$, and that there are $\ell + 1$ edges which overlap a given ℓ-run. Specifically, if $r = [i, j]$ then these $\ell(r) + 1$ edges are

$$\{i-1, i\}, \ldots, \{j, j+1\}.$$

Let $L(X_0)$ be defined by

$$L(X_0) = \sum_{r \in \mathcal{R}(X_0)} (\ell(r) + 1).$$

Then $L(X_0)$ equals the number of edges which overlap some run in X_0. Denote by $v(X_0, e)$ the value of $\psi(X_1) - \psi(X_0)$ given that the edge e has been chosen by the Pavlov process in step 1. Let r be an ℓ-run and let

$$\sigma(r) = \sum_{e \text{ overlaps } r} v(X_0, e).$$

By definition we have

$$\mathbf{E}\left[\psi(X_1) - \psi(X_0)\right] = \frac{1}{n} \sum_{e \in E} v(X_0, e),$$

since there are n edges in G. When X_0 contains both pluses and minuses we can also state that

$$\mathbf{E}\left[\psi(X_1) - \psi(X_0)\right] = \frac{1}{n} \sum_{r \in \mathcal{R}(X_0)} \sigma(r),$$

since runs are disjoint and an edge which does not overlap a run makes no change to X_0. Let $M = M(X_0)$ be defined by

$$M = \max\left\{ \frac{\sigma(r)}{\ell(r) + 1} \mid r \in \mathcal{R}(X_0) \right\}.$$

That is, M is the maximum over all runs of the average contribution of each edge in that run. The way in which M will be used is described below. We ignore two configurations: the all-pluses configuration X^*, and the all-minuses configuration. The latter is treated separately in Section 3.2.

Lemma 1 *Suppose that X_0 contains both pluses and minuses. With M, ψ and L defined as above, we have*

$$\mathbf{E}\left[\psi(X_1)\right] \leq \left(1 + \frac{ML(X_0)}{\psi(X_0)n}\right)\psi(X_0).$$

Proof. From above, we have

$$\begin{aligned}
\mathbf{E}\left[\psi(X_1) - \psi(X_0)\right] &= \frac{1}{n}\sum_{r\in\mathcal{R}(X_0)} \sigma(r) \\
&\leq \frac{1}{n}\sum_{r\in\mathcal{R}(X_0)} (\ell(r)+1)M \\
&= \frac{ML(X_0)}{n}
\end{aligned}$$

We can rearrange this inequality to give

$$\mathbf{E}\left[\psi(X_1)\right] \leq \psi(X_0) + \frac{ML(X_0)}{n} = \left(1 + \frac{ML(X_0)}{\psi(X_0)n}\right)\psi(X_0),$$

as stated. □

Suppose that the values of β, γ, δ could be set to ensure that $M < 0$. Then, by Lemma 1, the value of ψ *decreases* in expectation at every step. This will be used in Section 3.2 to calculate an upper bound for the absorption time of the Pavlov process.

Let $r = [i, j]$ be a run. Then there are two *outer rim* edges associated with r, namely $\{i-1, i\}$ and $\{j, j+1\}$. If r has length at least 3 then there are also two *inner rim* edges associated with r, namely $\{i, i+1\}$ and $\{j-1, j\}$. If r is a singleton then there are no inner rim edges, while if r is a pair $[i, i+1]$ then there is a unique inner rim edge $\{i, i+1\}$. All other edges which overlap r are strictly inside the interval $[i, j]$, and we call these edges *internal* edges.

Suppose that there are two runs in $\mathcal{R}(X_0)$ which are only separated by a single plus, i.e. $[i, j]$ and $[j+2, k]$ for some i, j, k. Then there are two edge choices $\{j, j+1\}$ and $\{j+1, j+2\}$ which cause the two runs to merge (note that these edges are both outer rim edges for the runs which they overlap). For simplicity, we will first assume that there are no edge choices which cause runs to merge. That is, in Lemma 2 we assume that all adjacent runs in $\mathcal{R}(X_0)$ are separated by at least two pluses. By carefully choosing values for β, γ and δ in this case, we show that M is negative: specifically $M = -1/14$.

In Lemma 3 we return to configurations which contain adjacent runs separated by a single plus.

Before presenting Lemma 2, we make a few general remarks. When all adjacent runs are separated by at least two pluses, choosing an outer rim edge will always cause r to increase in length by 1, introducing an extra minus. Similarly, choosing an inner rim edge will always cause r to decrease in length by 2, changing two minuses to pluses. When the length of r is small there might be additional effects from these four edges, as we shall see. When any internal edge is chosen, the run r is *split* into two runs which are separated by two pluses. If the two runs have length k and ℓ we say that this edge choice produces a (k, ℓ)*-split*.

We can now prove that M is negative for certain fixed values of β, γ and δ, when X_0 contains both pluses and minuses and all adjacent runs are separated by at least two pluses.

Lemma 2 *Let X_0 contain both pluses and minuses, and suppose that adjacent runs in X_0 are separated by at least two pluses. Then setting $\beta = 27/14$, $\gamma = 4/7$ and $\delta = -4/7$ we obtain $M = -1/14$.*

Proof. We will consider runs r of different lengths in turn, and calculate $\sigma(r)/(\ell(r)+1)$ in each case. Then M is the maximum of these values.

A 1-run. Let r be a 1-run $[i, i]$. The only edges which overlap r are the outer rim edges $\{i-1, i\}$ and $\{i, i+1\}$. When either of these edges are chosen, a vertex adjacent to the 1-run changes from a plus to a minus. This introduces an extra minus and changes a 1-run (singleton) to a 2-run (a pair), without changing the total number of runs. Therefore

$$\frac{\sigma(r)}{2} = 1 - \gamma + \delta = -\frac{1}{7}. \tag{1}$$

A 2-run. Suppose that $r = [i, i+1]$. There are 3 edges which overlap r. When either of the outer rim edges $\{i-1, i\}$ or $\{i+1, i+2\}$ are chosen the 2-run becomes a 3-run, introducing an extra minus and deleting a pair. There is only one inner rim edge, the edge $\{i, i+1\}$. When this edge is chosen, both minuses in the pair become pluses. Here we lose two minuses and delete a pair, decreasing the number of runs by 1. Adding these contributions and dividing by 3 we find that

$$\frac{\sigma(r)}{3} = \frac{2(1-\delta) - (2+\beta+\delta)}{3} = -\frac{\beta+3\delta}{3} = -\frac{1}{14}. \tag{2}$$

A 3-run. Suppose that $r = [i, i+2]$. There are 4 edges which overlap r, namely the two outer rim edges and the two inner rim edges. Choosing an outer rim edge turns the

3-run into a 4-run, introducing an extra minus. Choosing an inner rim edge turns the 3-run into a 1-run. Hence

$$\frac{\sigma(r)}{4} = \frac{2 + 2(-2 + \gamma)}{4} = \frac{-1 + \gamma}{2} = -\frac{3}{14}. \tag{3}$$

A 4-run. Suppose that $r = [i, i + 3]$ for some i. There are 5 edges which overlap r. Choosing an outer rim edge causes r to increase in length by 1, introducing a new minus. Choosing an inner rim edge causes the length of r to decrease by 2: in this case this introduces a new pair. Finally, there is one internal edge $\{i + 1, i + 2\}$. Choosing this edge produces a $(1, 1)$-split. This introduces two singletons and increases the total number of runs by 1, while removing two minuses. Adding these contributions together and dividing by 5, we obtain

$$\frac{\sigma(r)}{5} = \frac{2 + 2(-2 + \delta) + (-2 + \beta + 2\gamma)}{5} = \frac{-4 + \beta + 2\gamma + 2\delta}{5} = -\frac{29}{70}. \tag{4}$$

A 5-run. Let $r = [i, i + 4]$ for some i. There are six edges which overlap r. Choosing an outer rim edge causes r to increase in length by 1. Choosing an inner rim edge causes the length of r to decrease by 2. There are two internal edges, $\{i + 1, i + 2\}$ and $\{i + 2, i + 3\}$. Choosing either of these edges produces a $(1, 2)$-split, deleting two minuses, introducing a singleton and a pair, as well as increasing the number of runs by 1. Adding the contributions from all of these edges together, and dividing by 6, we obtain

$$\frac{\sigma(r)}{6} = \frac{2 - 4 + 2(-2 + \beta + \gamma + \delta)}{6} = \frac{-3 + \beta + \gamma + \delta}{3} = -\frac{5}{14}. \tag{5}$$

A 6-run. Let $r = [i, i + 5]$. There are 7 edges which overlap r. If an outer rim edge is chosen then r increases in length by 1. If an inner rim edge is chosen then r decreases in length by 2. There are 3 internal edges. Choosing $\{i + 1, i + 2\}$ or $\{i + 3, i + 4\}$ produces a $(1, 3)$-split, decreasing the number of minuses by 2 while increasing the number of singletons and the number of runs by 1. Finally, choosing the edge $\{i + 2, i + 3\}$ produces a $(2, 2)$-split, decreasing the number of minuses by 2, increasing the number of runs by 1 and the number of pairs by 2. Combining this information we find that

$$\frac{\sigma(r)}{7} = \frac{2 - 4 + 2(-2 + \beta + \gamma) + (-2 + \beta + 2\delta)}{7} = \frac{-8 + 3\beta + 2\gamma + 2\delta}{7} = -\frac{31}{98}. \tag{6}$$

An ℓ-run, where $\ell \geq 7$. Now suppose that $r = [i, j]$ is an ℓ-run for some $\ell \geq 7$. Choosing either of the two outer rim edges causes r to increase in length by 1. Choosing either of the two inner rim edges causes r to decrease in length by 2. There are 4 internal edges which need careful analysis. Choosing either $\{i+1, i+2\}$ or $\{j-2, j-1\}$ produces a $(1, \ell-3)$-split, introducing a singleton and increasing the number of runs by 1, while decreasing the number of minuses by 2. Similarly, choosing either $\{i+2, i+3\}$ or $\{j-3, j-2\}$ produces a $(2, \ell-4)$-split, introducing a pair and increasing the number of runs by 1, while decreasing the number of minuses by 2. There are $\ell-7$ other internal edges which split r into pairs of runs, each of length at least 3. In each case, the number of minuses decreases by 2 while the number of runs increases by 1, but the numbers of singletons and pairs are unchanged. We obtain

$$\begin{aligned}\frac{\sigma(r)}{\ell+1} &= \frac{2-4+2(-2+\beta+\gamma)+2(-2+\beta+\delta)+(\ell-7)(-2+\beta)}{\ell+1} \\ &= \beta-2-\frac{4\beta-6-2\gamma-2\delta}{\ell+1} \\ &= -\frac{1}{14}-\frac{12}{7(\ell+1)}. \qquad (7)\end{aligned}$$

Now M is equal to the maximum of the right hand sides of (1)–(7). It is easy to verify that the maximum is $-1/14$, as stated. □

For the remainder of this section, the values of $\beta = 27/14$, $\gamma = 4/7$ and $\delta = -4/7$ are fixed. These values were chosen without explanation for use in the proof of Lemma 2 above. They were originally derived by setting $\beta = 2-\eta$, $\gamma = 1/2+\eta$ and $\delta = -\gamma$, and choosing η to minimize M. The interested reader can easily verify that $\eta = 1/14$ is the optimal choice.

We now show that the value $M = -1/14$ can still be used in Lemma 1 even when the initial configuration has adjacent runs which are separated by a single plus.

Lemma 3 *Suppose that $X_0 \in \{1,-1\}$ contains both pluses and minuses. Then the conclusion of* Lemma 1 *holds with*

$$M = -\frac{1}{14}.$$

Proof. By Lemma 2, we have $M = -1/14$ whenever no two adjacent runs in X_0 are separated by a single plus. So now suppose that there are exactly s distinct values $i \in \{0, \ldots, n-1\}$ such that $X_0(i-1) = -1$, $X_0(i) = 1$ and $X_0(i+1) = -1$, where $s \geq 1$. We will call such an i a *rim vertex*. Define a new cycle $G' = (V', E')$ from G by splitting the vertex i into two new vertices, i' and i'', for each rim vertex i. Thus G' is a graph on $n+s$ vertices. Let the edges of G' be obtained from the edges of G by deleting the edges $\{i-1, i\}$, $\{i, i+1\}$ and adding the edges $\{i-1, i'\}$, $\{i', i''\}$, $\{i', i+1\}$, for each rim vertex i. Thus G' forms a cycle on $n+s$ vertices. Construct the configuration

${X_0}' \in \{1,-1\}^{V'}$ from X_0 by replacing the single plus at i by two pluses on i', i'', for each rim vertex i. That is, let

$$
{X_0}'(j) = \begin{cases} X_0(j) & \text{if } j \text{ is unprimed,} \\ 1 & \text{otherwise.} \end{cases}
$$

By definition, ${X_0}'$ has no two adjacent runs separated by a single plus. Note also that $L({X_0}') = L(X_0)$. Let ${X_1}'$ be the result of running the Pavlov process for one step from ${X_0}'$. Combining Lemma 1 and Lemma 2, we see that

$$
\mathbf{E}\left[\psi({X_1}') - \psi({X_0}')\right] \leq -\frac{L({X_0}')}{14(n+s)}.
$$

Suppose that we could show that

$$
\sum_{e \in E(G')} v({X_0}', e) \geq \sum_{e \in E(G)} v(X_0, e). \tag{8}
$$

Then we would have

$$
\begin{aligned}
\mathbf{E}\left[\psi({X_1}') - \psi({X_0}')\right] &= \frac{1}{n+s} \sum_{e \in E(G')} v({X_0}', e) \\
&\geq \frac{1}{n+s} \sum_{e \in E(G)} v(X_0, e) \\
&= \frac{n}{n+s} \mathbf{E}\left[\psi(X_1) - \psi(X_0)\right].
\end{aligned}
$$

From this we could conclude that

$$
\begin{aligned}
\frac{n}{n+s} \mathbf{E}\left[\psi(X_1) - \psi(X_0)\right] &\leq \mathbf{E}\left[\psi({X_1}') - \psi({X_0}')\right] \\
&\leq -\frac{L({X_0}')}{14(n+s)} \\
&= -\frac{L(X_0)}{14(n+s)}.
\end{aligned}
$$

Multiplying this inequality through by $(n+s)/n$ proves the lemma. Hence it suffices to establish (8).

It is not difficult to see that any edge which does not overlap a rim vertex in X_0 makes the same contribution in both the primed and unprimed settings. For these edges e belong to both $E(G)$ and $E(G')$, and

$$
v({X_0}', e) = v(X_0, e).
$$

Therefore, to prove (8) it suffices to prove that $Y' > Y$ for all rim vertices i, where

$$
Y = v(X_0, \{i-1, i\}) + v(X_0, \{i, i+1\})
$$

and

$$Y' = v(X_0', \{i-1, i'\}) + v(X_0', \{i'', i+1\}).$$

(Clearly the edge $\{i', i''\}$ makes no contribution to $E[\psi(X_1') - \psi(X_0')]$.) Let r_1 and r_2 be the two runs which are separated by i in X_0, and define a and b by

$$a = |\{j \in \{1,2\} \mid r_j \text{ is a singleton}\}| \qquad \text{and} \qquad b = |\{j \in \{1,2\} \mid r_j \text{ is a pair}\}|.$$

Then $0 \leq a + b \leq 2$. Consider choosing either $\{i-1, i'\}$ or $\{i'', i+1\}$ for X_0'. Clearly either choice will cause a minus to be introduced. For a of these choices a singleton is removed and a pair is created, while for b of these choices a pair is removed. Therefore

$$Y' = 2 - a\gamma + a\delta - b\delta.$$

Now consider choosing either $\{i-1, i\}$ or $\{i, i+1\}$ in X_0. The expected change of ψ is identical for either choice. Choosing either of these edges introduces a minus, decreases the number of runs by 1, and deletes all singletons or pairs which are present in X_0. The merged run which is created has length $\ell(r_1) + \ell(r_2) + 1 \geq 3$, so no singletons or pairs are created. Therefore

$$Y = 2(1 - \beta - a\gamma - b\delta).$$

Hence, using the values of β, γ and δ we obtain

$$Y' - Y = 2\beta + a(\gamma + \delta) + b\delta \geq 2(\beta + \delta) > 0,$$

proving the lemma. $\square$

3.2 Bounding the absorption time

We have fixed $\beta = 27/14$, $\gamma = 4/7$, $\delta = -4/7$. Recall that $L(X_0) = \sum_{r \in \mathcal{R}(X_0)} (\ell(r) + 1)$. Combining Lemmas 1, 2, 3 we obtain

$$\mathbf{E}[\psi(X_1)] \leq \left(1 - \frac{L(X_0)}{14\,\psi(X_0)n}\right) \psi(X_0), \tag{9}$$

for all X_0 which contain both pluses and minuses. We need the following result.

Lemma 4 *Let $X \in \{1, -1\}^V$ and let ψ, L be as defined above. Then*

$$\psi(X) \leq \frac{7L(X)}{4} \tag{10}$$

if X contains both pluses and minuses, while

$$\frac{47}{14} \leq \psi(X) \leq \frac{7n}{4} \tag{11}$$

for all $X \neq X^*$.

Proof. First suppose that X contains both pluses and minuses. Let $\psi(r)$ denote the potential of the run r, for all $r \in \mathcal{R}(X)$. That is,

$$\psi(r) = \begin{cases} 1 + \beta + \gamma & \text{if } r \text{ is a singleton,} \\ 2 + \beta + \delta & \text{if } r \text{ is a pair,} \\ \ell(r) + \beta & \text{otherwise.} \end{cases}$$

Clearly

$$\psi(X) = \sum_{r \in \mathcal{R}(X)} \psi(r).$$

It is not difficult to check that the inequality

$$\frac{\psi(r)}{\ell(r) + 1} \le \frac{7}{4}$$

holds, with equality if and only if r is a singleton. Hence

$$\psi(X) = \sum_{r \in \mathcal{R}(X)} \psi(r) \le \sum_{r \in \mathcal{R}(X)} \frac{7(\ell(r) + 1)}{4} = \frac{7L(X)}{4},$$

as stated. Now $L(X)$ denotes the number of edges which overlap some run in X. Since there are at exactly n edges in G, it follows that

$$\psi(X) \le \frac{7n}{4}$$

whenever X contains both pluses and minuses. Since the all-minuses configuration has potential n, this proves the upper bound in (11). Finally, note that

$$\psi(r) \ge \frac{47}{14},$$

with equality if and only if r is a pair. Therefore the lowest potential of all configurations with both minuses and pluses is obtained on any configuration which contains a unique run, this unique run being a pair. The all-minuses configuration has potential n, but we have assumed that n is at least 4. This proves the lower bound in (11). □

Proof of Theorem 1. Combining (9) and inequality (10) of Lemma 4, we can conclude that

$$\mathbf{E}\left[\psi(X_1)\right] \le \left(1 - \frac{2}{49n}\right) \psi(X_0) \tag{12}$$

for all X_0 which contain both pluses and minuses. However, (12) also holds for the all-minuses configuration, as follows. Let $\tilde{X}$ be the all-minuses configuration, defined by $\tilde{X}_0(i) = -1$ for all i. Let $\tilde{X}_1$ be the result of running the Pavlov process for one step from $\tilde{X}_0$. No matter which edge is chosen, the number of minuses decreases by 2 and the number of runs increases from 0 to 1. Therefore $\mathbf{E}\left[\psi(\tilde{X}_1) - \psi(\tilde{X}_0)\right] = \beta - 2 = -1/14$. Since $\psi(\tilde{X}_0) = n$, we conclude that

$$\mathbf{E}\left[\psi(\tilde{X}_1)\right] = \left(1 - \frac{1}{14n}\right)\psi(\tilde{X}_0) < \left(1 - \frac{2}{49n}\right)\psi(\tilde{X}_0),$$

as claimed.

So now let $X_0 \in \{1,-1\}^V$ satisfy $X_0 \neq X^*$. Starting from X_0, run the Pavlov process for t steps and let the resulting state be X_t. By applying (12) iteratively t times we obtain

$$\mathbf{E}\left[\psi(X_t)\right] \leq \left(1 - \frac{2}{49n}\right)^t \psi(X_0) \leq \left(1 - \frac{2}{49n}\right)^t \frac{7n}{4},$$

using the first statement of Lemma 4 for the last inequality. Let $\varepsilon > 0$ be given. Let $\nu = 47\varepsilon/14$. Whenever

$$t \geq \frac{49}{2}\, n \log\left(\frac{7n}{4\nu}\right)$$

we have $\mathbf{E}\left[\psi(X_t)\right] \leq \nu$. Using (11), any nonzero value of ψ must be at least $47/14$. Applying Markov's Lemma, we have

$$\mathrm{Prob}\left[\psi(X_t) \neq 0\right] = \mathrm{Prob}\left[\psi(X_t) \geq \frac{47}{14}\right] \leq \frac{14\nu}{47} = \varepsilon.$$

This completes the proof. $\square$

4 Exponential absorption on the complete graph

In this section we prove Theorem 2, showing that the absorption time of the Pavlov process is exponential on the complete graph K_n.

We will use the notation $X_t \in \{1,-1\}^n$ to refer to a configuration after t steps of the Pavlov process. Let N_t be the number of nodes in X_t with label 1. Clearly X_t is equal to the all-pluses absorbing state if and only if $N_t = n$. The basis of our proof is the observation that the process N_t is simple to analyse, even if the process X_t is not. Let p_t, q_t denote the labels of the two nodes chosen at step t. Then the transition probabilities of N_t are given by the following rule:

$$N_{t+1} = \begin{cases} N_t - 1 & \text{if } p_t q_t = -1 \text{ (probability } N_t(n-N_t)/\binom{n}{2}), \\ N_t & \text{if } p_t = q_t = 1 \text{ (probability } \binom{N_t}{2}/\binom{n}{2}), \\ N_t + 2 & \text{if } p_t = q_t = -1 \text{ (probability } \binom{n-N_t}{2}/\binom{n}{2}). \end{cases}$$

Consider the "speed-up" (M_t) of the chain (N_t), which only makes transitions that change the state:

$$M_{t+1} = \begin{cases} M_t - 1 & \text{with probability } 2M_t/(n + M_t - 1), \\ M_t + 2 & \text{with probability } (n - M_t - 1)/(n + M_t - 1). \end{cases}$$

If $M_0 = N_0$, then the time for M_t to hit n is clearly at most that for N_t. We will now show that it takes exponentially long for M_t to reach n. Consider the chain (Q_t) with transition probabilities given by

$$Q_{t+1} = \begin{cases} Q_t - 1 & \text{with probability } 3/4, \\ Q_t + 2 & \text{with probability } 1/4. \end{cases}$$

If $Q_0 = M_0 \geq 3n/5$ then (M_t) is stochastically dominated by (Q_t) as long as $M_t \geq 3n/5$.

Let $V(k) = a^k$, where $a = (\sqrt{13} - 1)/2$. Note that $a > 1$ and that a satisfies $\frac{3}{4}a^{-1} + \frac{1}{4}a^2 = 1$. Then

$$\mathbf{E}(V(Q_{t+1}) \mid V(Q_t)) = \frac{3}{4}\, a^{Q_t - 1} + \frac{1}{4}\, a^{Q_t + 2} = a^{Q_t} = V(Q_t),$$

so $V(Q_t)$ is a martingale.

Suppose that Q_0 is $\lceil 3n/5 \rceil$ or $\lceil 3n/5 \rceil + 1$. Let $T = \min\{t \mid \; Q_t < 3n/5 \text{ or } Q_t \geq n \,\}$. It is straightforward to show that $\operatorname{Prob}(Q_T \geq n) \leq a^{-2n/5}$. The analysis is in Example 4.1 of [3]. We include it here for completeness. First, since $\operatorname{Prob}(T < \infty) = 1$ and $|V(Q_{\min(T,t)})| \leq a^{n+1}$ for all t, we have $\mathbf{E}V(Q_T) = \mathbf{E}V(Q_0)$. This follows the stopping theorem for bounded martingales, see, for example, (3.6) of [3]. Now $V(Q_0) \leq a^{3n/5+2}$, so $\mathbf{E}V(Q_T) \leq a^{3n/5+2}$. Also,

$$\mathbf{E}V(Q_T) \geq a^{3n/5-1} \cdot \operatorname{Prob}(Q_T < 3n/5) + a^n \cdot \operatorname{Prob}(Q_T \geq n),$$

since $Q_T < 3n/5$ implies $Q_T = \lceil 3n/5 \rceil - 1 \geq 3n/5 - 1$. Hence

$$\operatorname{Prob}(Q_T \geq n) \leq \frac{a^{3n/5+2} - a^{3n/5-1}}{a^n - a^{3n/5-1}},$$

which is at most $a^{-2n/5}$ as long as n is sufficiently large ($n \geq 23$ suffices).

We conclude that every time (M_t) enters the interval $[3n/5, n-1]$ from below, the probability that it exits out the top of the interval (rather than the bottom) is at most $a^{-2n/5}$. Thus, the probability that the chain reaches absorption in as few as $(1.1)^n$ visits to the region is at most

$$(a^{-2/5} \times 1.1)^n = o(1).$$

This proves Theorem 2.

5 Other topics

Several issues remain for further study. A natural extension of our results would be to investigate the Pavlov process for other families of graphs. Two other cases seem particularly interesting: degree-bounded trees and random graphs. Another possible ingredient to our model is random noise (or player mistakes). The importance of this parameter has been previously recognized in [9].

Acknowledgements

We thank the referee for shortening the proof of Theorem 2.

References

[1] R. Axelrod, *The Evolution of Cooperation*, Basic Books, New York, 1984.

[2] C. Cooper and A. Frieze, Mixing properties of the Swendsen-Wang process on classes of graphs, Random Structures and Algorithms **15** (1999) 242–261.

[3] R. Durrett, *Essentials of Stochastic Processes* (Springer-Verlag 1999).

[4] M. Dyer and C. Greenhill, *On Markov chains for independent sets,* Journal of Algorithms, **35** (2000) 17–49.

[5] J. E. Kittock, *Emergent conventions and the structure of multi-agent systems*, in L. Nadel and D. Stein, eds., *1993 Lectures in Complex systems: the proceedings of the 1993 Complex systems summer school*, Santa Fe Institute Studies in the Sciences of Complexity Lecture Volume VI, Santa Fe Institute, Addison-Wesley, 1995.

[6] B. Hajek, *Hitting-time and occupation-time bounds implied by drift analysis with applications,* Adv. Appl. Prob. 14 (1982) 502–525.

[7] M. Luby, D. Randall and A. Sinclair, Markov chain algorithms for planar lattice structures, *Proceedings of the Symposium on Foundations of Computer Science (IEEE)* **36** (1995) 150–159.

[8] M. Luby and E. Vigoda, Approximately counting up to four, *Proceedings of the ACM Symposium on the Theory of Computing* **29** (1997) 682–687.

[9] M. Nowak and K. Sigmund, *A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game*, Nature 364 (1993), pp. 56-58.

[10] Y. Shoham and M. Tennenholtz, *On the emergence of social conventions: modelling, analysis and simulations*, Artificial Intelligence, 94 (1997), pp. 139–166.