Stochastic Stability in Schelling's Segregation Model with Markovian Asynchronous Update[*]

Gabriel Istrate[1,2]

[1] Department of Computer Science, West University of Timişoara, Bd. V. Pârvan 4, Timişoara, Romania. email: gabrielistrate@acm.org
[2] e-Austria Research Institute, Bd. V. Pârvan 4, cam. 045 B, Timişoara, Romania.

Abstract. We investigate the dependence of steady-state properties of Schelling's segregation model on the agents' activation order. Our basic formalism is the Pollicott-Weiss version of Schelling's segregation model. Our main result modifies this baseline scenario by incorporating *contagion* in the decision to move: (pairs of) agents are connected by a second, *agent influence* network. Pair activation is specified by a random walk on this network.

The considered schedulers choose the next pair *nonadaptively*. We can complement this result by an example of adaptive scheduler (even one that is quite fair) that is able to preclude maximal segregation. Thus scheduler nonadaptiveness seems to be required for the validity of the original result under arbitrary asynchronous scheduling. The analysis (and our result) are part of an *adversarial scheduling approach* we are advocating to evolutionary games and social simulations.

1 Introduction

Schelling's Segregation Model [1] is one of the fundamental dynamical systems of Agent-Based Computational Economics, perhaps one of the most convincing examples of Asynchronous Cellular Automata (ACA) [2] employed in the social sciences. It exhibits large-scale self-organizing neighborhoods, due to agents' desire to live close to their own kind. A remarkable feature of the model that has captured the attention of social scientists is the fact that segregation is an *emergent* phenomenon, that may appear even in the presence of just mild preferences (at the individual level) towards living with one's own kind. The model has sparked a significant interest and work, coming from various areas such as Statistical Physics [3], agent-based computational economics [4,5], game theory [6], theoretical computer science [7,8,9], or applied mathematics [10].

Schelling's segregation model is an asynchronous dynamical system on a graph (usually a finite portion of the one-dimensional or the two-dimensional lattice). It can be described, informally as follows: vertices in the graph are in one of three states: *unoccupied*, when no agent sits on the given node, or one of *red/blue* (± 1), corresponding to the color of the agent inhabiting the node. Agents have a (non-strict) preference towards living among agents of the same color. This is modeled by considering a *local neighborhood* around the agent. Depending on the density of like-colored agents in the neighborhood the agent may be in one of two states: *happy* and *unhappy*. An unhappy

---

agent may seek to trade places with another agent in order to become happy. It was originally observed via "pen-and-paper simulations", and proved rigorously in a variety of settings, that segregated states may arise even when agent only have a weak preference for its own color, and are happy to live in a mixed neighborhood, as long as it contains "enough" of its own kind. Difference in the topology, activation order, specification of the update mechanism account for the dizzying variety of variants of the model that have been investigated so far (for the intellectual context of the model and a related one, due to Sakoda, see [11]).

Qualitative properties of asynchronous cellular automata are highly dependent on activation order [12,13]. In particular, when viewed as dynamical systems, ACA may exhibit a multitude of limit cycles , and the update dynamics "chooses" one of these limit cycle in a path dependent manner. The challenge then becomes to explain the selection of one particular limit cycle among many possible ones.

One particularly interesting class of techniques, brought to evolutionary games by Foster and Young [14] (see also [6]) uses the concept of *stochastic stability* to deal with this problem. It was the fundamental insight of Peyton-Young [6] that adding continuous small perturbations to a certain dynamics might help "steer"— the system towards a particular subset of equilibria, the so-called *stochastically stable states*. Indeed, in several versions of Schelling's segregation model [15,6,5] the most segregates states are identified as precisely the stochastically stable equilibria of the dynamics.

Though such results are interesting, they are still not realistic enough enough: results about stochastic stability in models on graphs may be sensitive to the precise specification of the update order, which in realistic scenarios need not be the random one. As noted in many papers, precise specification of an asynchronous schedule in social systems can arise from many factors, including geography or agent incentives [16]. It is thus important to study validity of baseline results under different scheduling models. A dramatic example of this type is that of the related model of *logit response dynamics*, another model analyzed via stochastic stability [17]. Going in this model from a random single-node update to parallel modes (the so-called *revision process* of [18]) may lead (in general games) to the selection (via stochastic stability) of states that are not even Nash equilbria. In contrast, for local interaction games the parallel *all-logit* rule has a Gibbs limiting distribution [9], similar to the random update case.

Neither random scheduling nor parallel update can accurately model *social contagion* phenomena, i.e., agents becoming active as a result of other agents' action, via communication or imitation. Thus it is of interest to study the robustness social models to variations in the update rule. Indeed, in [19,20] we have proposed an *adversarial approach to social simulations*. Roughly speaking, this means that we consider the baseline dynamics under random scheduling, then modify the update order to arbitrary scheduling, and attempt to derive necessary and/or sufficient conditions on the scheduler that make the results from the random update case extend to the adversarial setting. We have accomplished this in [20] for Prisoners' Dilemma with Pavlov strategy, a Markov chain previously investigated in [21], and in [19] for the logit response dynamics.

The purpose of this paper is to introduce contagion in evolutionary versions of Schelling's segregation model, as studied by Pollicott and Weiss [10], and study the setting where the set of agents that becomes active is specified by a random walk on a second "communication" network. A similar model was investigated for the logit

response dynamics in [19], and is apparently consistent with some real-life contagion phenomena in power networks [22]. In the most general setting this communication network works on *pairs of vertices*. The more natural case where agents influence each other is a special case of our setting. The feature of the Pollicott-Weiss model that is of special interest to our study is that, although Schelling's model might have multiple equlibria, it is only the most segregated states that are stochastically stable. Our result shows that this extends to a scenario with social contagion: we prove a result with a similar flavor under a more general nonadaptive model of activation.

Even though we use analytic rather than experimental techniques, our results are naturally related to a long line of research that investigates the robustness of discrete models under various scheduling models [23,13]. On the other hand the notion we consider, that of stochastic stability, is highly related to the analysis of cellular automata using dynamical systems techniques [24]. In contrast to many such studies, though, that only perturb the initial system state, stochastic stability embodies the notion of stability under continuous (but vanishingly small) perturbations.

2 Preliminaries

We first review the notion of *stochastic stability* for perturbed dynamical systems described by Markov chains:

Definition 1. *Let the Markov chain P^0 be defined on a finite state set Ω. For every $\epsilon > 0$, we also define a Markov chain P^ϵ on Ω. Family $(P^\epsilon)_{\epsilon \geq 0}$ is called a* regular perturbed Markov process *if all of the following conditions hold:*

- *For every $\epsilon > 0$ Markov chain P^ϵ is irreducible and aperiodic.*
- *For each pair of states $x, y \in \Omega$, $\lim_{\epsilon > 0} P^\epsilon_{xy} = P^0_{xy}$.*
- *Whenever $P_{xy} = 0$ there exists a real number $r(m) > 0$, called the* resistance *of transition $m = (x \rightarrow y)$, such that as $\epsilon \rightarrow 0$, $P^\epsilon_{xy} = \Theta(\epsilon^{r(m)})$.*

Let μ^ϵ be the stationary distrib. of P^ϵ. State s is stochastically stable *if $\underline{\lim}_{\epsilon \rightarrow 0} \mu^\epsilon(s) > 0$.*

We use a standard tool in this area: a result due to Young (Lemma 3.2 in [6]) that allows us to recognize stochastically stable states in a Markov Chain using *spanning trees of minimal resistance*:

Definition 2. *If $j \in S(G)$ is a state, a* tree rooted at node j *is a set T of edges so that for any state $w \neq j$ there exists an unique (directed) path from w to j. The* resistance of a rooted tree T *is defined as the sum of resistances of all edges in T.*

Proposition 3. (Young) *The stochastically stable states of a regular Markov process (P_ϵ) are precisely those states $z \in \Omega$ such that there exists a tree T rooted at z of minimal resistance (among all rooted trees).*

3 The model

We consider an $N \times N$ two-dimensional lattice graph G with periodic boundary conditions (that is, a torus). Let V be the set of vertices of this graph. Each vertex of G hosts an agent, colored either red or blue. The neighborhood of a vertex v is the *four-point*

neighborhood, consisting of the cell to the left,up,right,down of the cell holding v. An agent's utility is written as $\forall i \in V$, $u_i(x) = r \cdot w(x_i) + \epsilon_i$, where r is a positive constant, assumed similarly to [5] to be the same for all agents, and $w(x)$ is defined, similarly to [10], as the difference between the number of neighbors of x having the same color and the number of neighbors of x having the opposite color. Finally, ϵ_i are (possibly different) agent-specific constants.

Next we specify our scheduling model, defined as follows:

Definition 4. [Markovian contagion]: *To each pair of vertices e we associate a probability distribution D_e on $V \times V$ such that $e \in supp(D_e)$* [3]. *We then choose the pair to be scheduled next as follows: Let p_i be the pair chosen at stage i. Select the next scheduled pair p_{i+1} by sampling from the set of pairs in D_{p_i}. We assume that for any two pairs e, e' the following condition holds:*

$$Pr[e \rightarrow e'] > 0 \Leftrightarrow Pr[e' \rightarrow e] > 0. \tag{1}$$

In other words: the next scheduled pair only depends on the last scheduled pair, succession relation $e \rightarrow e'$ specifies a *bidirected graph $H(G)$* whose vertex set is $V \times V$, and the scheduled pair can be seen as performing a random walk (possibly a non-uniform one) on $H(G)$. In particular the next chosen pair is *not* guaranteed to have different labels on endpoints. Furthermore, graph $H(G)$ should be connected, otherwise the choice of a particular initial sequence of moves could preclude a given edge from ever being scheduled sometimes in the future.

Observation 1 *A particular case of Definition 4, which justifies the name contagion is described informally as follows:* agents*, rather than pairs, are given the opportunity to switch. They randomly choose a swapping neighbor among those available to them. There exists a second, separate* influence network I*. The next scheduled agent is one of the neighbors (in I) of the previously scheduled agent. Indeed, to describe this scenario in the setting of the previous definition, define D_e to consist of pairs e' that share with e a vertex.*

Observation 2 *The random scheduler is a particular case of Definition 4, when D_e is the uniform distribution on $V \times V$.*

To complete the description of the dynamics, we only need to specify the probability that two agents inhabiting the different endpoints of a pair $e = (u, v)$ switch when pair e is scheduled. This is accomplished using the so-called *log-linear response rule* [25,6,18], specified as follows: let S be the state before the switch and T be the state obtained if the two agents at u, v switch. Then:

$$Pr[S \rightarrow T] = \frac{e^{\beta \cdot [u_1(T) + u_2(T)]}}{e^{\beta[u_1(S) + u_2(S)]} + e^{\beta[u_1(T) + u_2(T)]}}, \tag{2}$$

where u_1, u_2 are the corresponding utility functions of the two agents at the endpoints of the scheduled pair, and $\beta > 0$ is a constant. This is, of course, the noisy version of

---

[3] This translates, intuitively, to the following condition: we always give the participants in a swap the chance to immediately reevaluate their last move.

the best-response move, that would choose the move that maximizes the sum of utilities $u_1(\cdot) + u_2(\cdot)$.

The *state of the system* is defined by a vector $\overline{w} \in \{-1, +1\}^V$ is a vector encoding the labels of all vertices of the torus. To obtain a description of the dynamics as an aperiodic Markov chain we have to complete the description of the system state by a pair r of vertices, i.e., the last pair that had the opportunity to switch by being scheduled. Thus the state space of the Markov chain is $S(G) := \{\pm 1\}^V \times (V \times V)$. When $H(G)$ is strongly connected, M_β is ergodic, so it has an unique stationary distribution Π_β. It is easy to see (and similar to previous results e.g. in [6]) that the family of chains $(M_\epsilon)_{\epsilon > 0}$ is a regularly perturbed Markov process (where we define $\epsilon = e^{-\beta}$).

Definition 5. *A state* $w \in \{\pm 1\}^V$ *is called* maximally segregated *if* w *realizes the minimum value of the number of red-blue* edges *(of the torus), across all possible states on* G *with a given number of red/blue agents.*

In [10] a complete characterization of a maximally segregated state was obtained (Theorem 2 in that paper). Roughly they are horizontal or vertical "bands", possibly with a "strip" attached, or a rectangle, possibly with at most two "strips" attached. We refer the reader to [10] for details, and don't discuss it any further.

4   Main result and its interpretation.

Our main result is:

Theorem 6. *The stochastically stable states for Schelling's segregation model with Markovian contagion form a subset of the set* $Q \subseteq S(G)$,

$$Q = \{(w, e) | w \text{ is maximally segregated and } e \in V \times V\} \tag{3}$$

In other words: **the conclusion that stochastically stable states in Schelling's segregation model are maximally segregated is robust to extending the update model from a random one to those from the family from Definition 4, that incorporate Markovian contagion.**

The defining feature of the class of schedulers in the previous result seems to be that they choose the next scheduled pair *nonadaptively*: the next pair only depends on the last scheduled pair, and **not** on other particulars of the system state. Indeed, in the full version of the paper we will complement the result above by another one (very easy to state and prove), that shows that some **adaptive** scheduler (despite being quite fair) can forever preclude the system from ever reaching maximal segregation (thus "breaking the baseline stylized result").

Proof. We will employ a fundamental property, noted for models of segregation such as the one in this paper e.g. in [5]: they are *potential games* [26], i.e., they admit a function $L : S(G) \to \mathbf{R}$ such that, for any player i, any strategy profile (i.e., vector of player strategies) $x_{-i} := (x_j)_{j \in V, j \neq i}$, and any two strategies z, t for player i

$$u_i(z; x_{-i}) - u_i(t; x_{-i}) = L(z; x_{-i}) - L(t; x_{-i}). \tag{4}$$

In other words, differences in utility of the i'th player as a result of using different strategies are equal to the differences in potential among the two corresponding profiles.

The function L is defined simply as $L(s) = \sum_i u_i(s)$. Strictly speaking the potential above is defined for the original Policott-Weiss model i.e., defined on $\{\pm 1\}^V$, instead of $S(G)$. But it can be easily extended by simply applying it to any pair (s, e) (thus neglecting e). Moreover, the following property holds, which determines the resistance of moves:

Lemma 7. *Let $A \in \{\pm 1\}^V$ be a state of the system, and let $e = (i, j)$ be a pair in $V \times V$. Let B be a state obtained by making the move $m = A \to B$ (B is either A or is the state obtained from A by swapping the states A_i, A_j). Then the resistance $r(m)$ of move m is equal to:*

$$r(m) = [u_i(A) + u_j(A)] - [u_i(B) + u_j(B)] = 2(L(A) - L(B)) > 0 \qquad (5)$$

when $A \to B$ is a swap that diminishes potential, and to

$$[u_i(C) + u_j(C)] - [u_i(A) + u_j(A)] = 2(L(C) - L(A)) > 0 \qquad (6)$$

when $A = B$, but the corresponding swap $A \to C$ would be a potential improving one. In all other cases $r(m) = 0$ (that is m is a neutral move).

In other words, a move has positive resistance when one of the following two alternatives hold: (a). The move corresponds to a decrease in potential. The resistance of the move is, in this case, equal to the potential decrease. (b). The move corresponds to preserving the current state (as well as agents' utilities), but the other possible move would have led to a state of higher potential. The resistance of the move is, in this case, equal to the difference in potentials between this better state and the current one.

Proof. Follows directly from equation (2) and the definition of resistance.

We apply this result to prove the following lemma:

Lemma 8. *Consider a state $Y \in Q$ that is maximally segregated. Consider another state X, and a tree T rooted at X having minimal potential. Then there exists another tree $\overline{T}$ rooted at Y whose potential is at most that of tree T, strictly less in case when X is not a maximally segregated state.*

Proof. Note that, by the definition of utility functions, maximally segregated states are those that maximize the potential.

Since T is an oriented tree, there is an unique directed path

$$p : [Y = (s_0, e_0) \to \ldots \to (s_k, e_k) \to (s_{k+1}, e_{k+1}) \to \ldots \to (s_r, e_r) = X] \qquad (7)$$

in T from Y to X. Here $s_0, s_r \in \{\pm 1\}^V$ are states, and e_0, e_r are pairs in $V \times V$. First, we decompose T into three subsets as follows:

1. The set of edges of p (see Figure 1)
2. The set of edges of, W_Y the subtree rooted at y.
3. The edges of subtrees W_k rooted at nodes (s_k, e_k) of p, other than Y (but possibly including X).
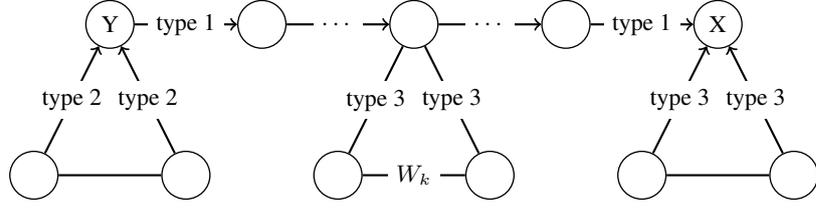
Fig. 1: Decomposition of edges of tree T. Path p is on top.

Lemma 9. *Without loss of generality we may assume that the path p contains no two consecutive vertices (s_k, e_k) and (s_r, e_r), $k < r$, with $s_k = \ldots = s_r$ and all the moves between s_k and s_r having zero resistance.*

Proof. Suppose there was such a pair k, r. Take one with maximal $k - r$. Define a tree T' by

- First connecting (s_k, e_k) directly to (s_{r+1}, e_{r+1}). This is legal since $s_k = s_r$. Indeed, since activating edge e_{r+1} move the system from s_r to s_{r+1}, activating the same edge moves the system from (s_k, e_k) to s_{r+1} as well.
- Also connecting $(s_{k+1}, e_{k+1}), \ldots, (s_r, e_r)$ directly to (s_{r+1}, e_{r+1})

It is easy to see that T' is a tree with the same resistance as T, since the only removed edges have zero resistance.

□

In particular, applying iteratively Lemma 9, we may assume that all transitions on the path from p to q either change the system state or have positive resistance.

Now, to obtain tree $\overline{T}$ we will first obtain a graph $\overline{\overline{T}}$, in which every node has *at least one path* to node Y. "Thinning out" this directed graph to a tree yields a tree $\overline{T}$ of even lower resistance. To obtain directed graph $\overline{\overline{T}}$:

1. First, add to $\overline{\overline{T}}$ the edges of T_Y.
2. Next, define path q from X to Y as follows:

$$q : [(s_r, e_r) \to (s_{r-1}, e_r) \to (s_{r-2}, e_{r-1}) \to \ldots \to (s_0, e_1) \to (s_0, e_0)]. \quad (8)$$

   In other words, q aims to "undo" the sequence of moves in path p from equation (7) from Y to X. However, since the states of the Markov chain also have as a second component a pair in $V \times V$, (corresponding to the last scheduled edge), we need to take a little extra care when defining q. Specifically q starts at X but **cannot** simply reverse the edges of p, since these do **not** correspond to legal moves. To define q, we first make a move at e_r by "undoing" the last move of p [4]. This yields state

---

[4] this is where we use a property specific to our model of Schelling segregation, as opposed to proving a result valid for general potential game: the property that we employ is that in Schelling's model any move m "can be undone". This means that there is a move n using the same pair of vertices as m that brings the system back to where it was before. Move n simply "swaps back" the two agents if they were swapped by m, and leaves them in place otherwise.

(S_{r-1}, e_r) (since pair e_r is scheduled in this as well). We then continue to "undo moves of p" until the state becomes S_0. This is possible because of condition (1): since pair e_{r+1} can be scheduled after e_r, scheduling e_r can move the system from s_r to s_{r-1}, scheduling pair e_{r-1} can move the system from s_{r-1} to s_{r-2}, and so on, until state S_0 is reached. At this moment the last activated pair was e_1. q then moves to Y by making a move (with no effect) on e_0.

Note that every such path will contain, for every $k = 1, r$, a vertex whose state is s_k.

3. For every tree component W_k obtained by removing path p from T, attached to p at (s_k, e_k) perform one of the following:

   - **Case 1: $s_k = s_{k-1}$.**
     In this case the point $(s_k, e_k) = (s_{k-1}, e_k)$ is on path q as well, therefore we also add the rooted tree W_k to $\overline{\overline{T}}$. This is possible since attaching a tree to a node depends only on the system state, but *not* on the last scheduled node. Moreover, *the resistance of W_k does not change* as a result of this attachment.

   - **Case 2: $s_k \neq s_{k-1}$ and move $(s_{k-1}, e_{k-1}) \rightarrow (s_k, e_k)$ has resistance > 0.**
     In this case, since in configuration s_{k-1} and scheduled move e_k we have a choice between moving to s_k and staying in s_{k-1}, it follows that the transition $(s_{k-1}, e_{k-1}) \rightarrow (s_{k-1}, e_k)$ has zero resistance and $L(s_k) < L(s_{k-1})$. Hence transition $(s_k, e_k) \rightarrow (s_{k-1}, e_k)$ has zero resistance.
     We now add the tree $\overline{W_k} = W_k \cup \{(s_k, e_k) \rightarrow (s_{k-1}, e_k)\}$ (rooted at node (s_{k-1}, e_k), which is on q) to $\overline{\overline{T}}$. The tree $\overline{W_k}$ has the same total resistance as W_k. All nodes from W_k, including (s_k, e_k) can now reach Y via q.

   - **Case 3: $s_{k-1} \neq s_k$, move $(s_{k-1}, e_{k-1}) \rightarrow (s_k, e_k)$ has zero resistance and all moves on p between s_k and X have zero resistance.**
     Then we add to $\overline{\overline{T}}$ this portion of p, together with W_k. This way we connect nodes in W_{k-1}, W_k to Y (via X and q). All added edges except those of one of the trees W_l have zero resistance.

   - **Case 4: $s_{k-1} \neq s_k$, the move $(s_{k-1}, e_{k-1}) \rightarrow (s_k, e_k)$ has zero resistance, but some move on p, between s_k and X has positive resistance.**
     Let $(s_{k+l}, e_l) \rightarrow (s_{k+l+1}, e_{l+1})$ be the closest move (i.e., the one that minimizes l) with positive resistance.
     If $s_{k+l} = s_{k+l+1}$ then we have already connected (s_{k+l}, e_l) to Y, as it falls under Case 1. Now just add all the (zero resistance) edges of p between s_k, e_k and (s_{k+l}, e_{k+l}), together with edges of W_k, to connect all such nodes to Y.
     If $s_{k+l} = s_{k+l+1}$ then we have already connected (s_{k+l}, e_l) to Y, as it falls under Case 2. We proceed similarly.

The previous construction has ensured that any pair (s, e) is connected by *at least* one path to Y. Thinning out $\overline{\overline{T}}$ we get a rooted tree $\overline{T}$ having resistance less or equal to the resistance of $\overline{\overline{T}}$. Since the four outlined transformations only add, in addition to trees W_k, edges of zero resistance, to compare the total resistances of T and $\overline{T}$ one should simply compare the total resistances of paths p and q. We claim that this difference in resistances of these paths is equal to the difference in potentials:

Claim. $r(p) - r(q) = 2(L(Y) - L(X)) \geq 0$.

Proving Claim (4) would validate our conclusion, since $L(Y) - L(X) \geq 0$, and $L(X) = L(Y)$ iff X is a global minimum state for the potential function. We prove this by considering the correspondence between edges of paths p and q: to each edge e of p one can associate an unique edge e' of Q that "undoes e".

By the additivity of both resistance and potential, it is enough to prove that, for every edge e of p and its associated edge of q, e', $r(e') - r(e)$ is equal to twice the difference in potentials between S_{fin}, the final state for the forward transition and S_{init}, the initial state. The first thing to note is none of the two resistances can be infinite: the transition $e \rightarrow e'$ corresponds to a move of the perturbed Markov chain (optimal or not). Its inverse corresponds to "undoing" that move, which is a legal move (eventually perturbed) in itself. We employ Lemma 7 and identify several cases:

- *The move e corresponds to **not** switching, and its resistance is zero.* Let S_1 be the common state, and let S_2 be the state corresponding to a switch. Then $S_1 = S_{init} = S_{fin}$. Also, $L(S_1) \geq L(S_2)$. So the move e' also stays in state S_1 (when it could have gone to S_2), which is the optimal action, given that $L(S_1) \geq L(S_2)$. Thus in this case both the "forward" and the "backward" transition have resistance zero, and do not count towards the sum of resistances on the path.

- *The move e corresponds to **switching**, and its resistance is zero.* Then, $S_{fin} = S_2$, $S_{init} = S_1$. By Lemma 7 $L(S_2) > L(S_1)$. The backward move has positive resistance equal to $2(L(S_2) - L(S_1))$. The result is verified.

- *The move e corresponds to **not** switching, and its resistance is nonzero.* Then $S_{init} = S_{fin} = S_1$ and $L(S_1) < L(S_2)$ (since switching would be beneficial). Therefore in the backward move the state stays S_1 (when it could have gone to S_2). The resistance is equal to $2(L(S_2) - L(S_1))$, the same as the resistance of the forward move. Therefore $r(e') - r(e) = L(S_{fin}) - L(S_{init}) = 0$.

- *The move e corresponds to **switching**, and its resistance is nonzero.* Then, $S_{fin} = S_2$, $S_{init} = S_1$, by Lemma 7 $L(S_2) < L(S_1)$ and the resistance of the forward move is equal to $2(L(S_1) - L(S_2)) = 2(L(S_{init}) - L(S_{fin}))$. The resistance of the backward move is equal to zero, so the result is verified in this case as well.
  Thus the claim is established and the proof of the theorem is complete.

5 Outlook and Further Work

Theorem 6 is only the main result in the adversarial analysis of Schelling's segregation model. It shows that stochastically segregated states are maximally segregated. Is the converse true ? Namely, is every maximally segregated state stochastically stable ? Such a result is indeed true in 1D versions of Schelling's segregation model (such as the one presented in [6]). We will discuss the 2D case with Markovian contagion in the journal version of the paper.

Other topics deserving research include studying conditions that preclude segregation, determining the convergence time of the segregation dynamics with Markovian contagion, models with Markovian contagion and concurrent updates [8,9], etc. We plan to address these and other issues in follow-up papers.

References

1. T.C. Schelling. Dynamic models of segregation. *J. Math. Sociology*, 1:143–186, 1971.

 2. N. Fates. A guided tour of asynchronous cellular automata. In *International Workshop on Cellular Automata and Discrete Complex Systems*, pages 15–30. Springer, 2013.
 3. D. Vinković and A. Kirman. A physical analogue of the Schelling model. *Proceedings of the National Academy of Sciences*, 103(51):19261–19265, 2006.
 4. R. Pancs and N. J. Vriend. Schelling's spatial proximity model of segregation revisited. *Journal of Public Economics*, 91(1-2):1 – 24, 2007.
 5. J. Zhang. Residential segregation in an all-integrationist world. *Journal of Economic Behavior and Organization*, 24(4):533–550, 2004.
 6. H. Peyton Young. *Individual Strategy and Social Structure: an Evolutionary Theory of Institutions*. Princeton University Press, 1998.
 7. C. Brandt, N. Immorlica, G. Kamath, and R. Kleinberg. An analysis of one-dimensional Schelling segregation. In *Proceedings of the 44th annual ACM STOC*, pages 789–804, 2012.
 8. A. Montanari and A. Saberi. Convergence to equilibrium in local interaction games. In *Proceedings of 50th Annual IEEE FOCS*, pages 303–312. IEEE, 2009.
 9. V. Auletta, D. Ferraioli, F. Pasquale, P. Penna, and G. Persiano. Logit dynamics with concurrent updates for local interaction games. In *European Symposium on Algorithms*, pages 73–84. Springer, 2013.
10. M. Pollicott and H. Weiss. The dynamics of Schelling-type segregation models. *Advances in Applied Mathematics*, 27:17–40, 2001.
11. R. Hegselmann. Thomas C. Schelling and James M. Sakoda: The intellectual, technical, and social history of a model. *J. Artificial Societies and Social Simulation*, 20(3), 2017.
12. N. Fates and M. Morvan. An experimental study of robustness to asynchronism for elementary cellular automata. *arXiv preprint nlin/0402016*, 2004.
13. O. Bouré, N. Fates, and V. Chevrier. Probing robustness of cellular automata through variations of asynchronous updating. *Natural Computing*, 11(4):553–564, 2012.
14. D. Foster and H. Peyton Young. Stochastic evolutionary game dynamics. *Theoretical population biology*, 38(2):219–232, 1990.
15. H. Peyton Young. The dynamics of conformity. *Social dynamics*, pages 133–153, 2001.
16. S. E. Page. On incentives and updating in agent based models. *Computational Economics*, 10(1):67–87, 1997.
17. M. Kandori, G. J. Mailath, and R. Rob. Learning, mutation, and long run equilibria in games. *Econometrica*, pages 29–56, 1993.
18. C. Alós-Ferrer and N. Netzer. The logit-response dynamics. *Games and Economic Behavior*, 68(2):413–427, 2010.
19. G. Istrate, M. V. Marathe, and S.S. Ravi. Adversarial scheduling analysis of game-theoretic models of norm diffusion. In *Proceedings of the CIE'2008*, *Lecture Notes in Computer Science*, vol. 5028, pages 273–282. Springer Verlag, 2008.
20. G. Istrate, M. V. Marathe, and S.S. Ravi. Adversarial scheduling analysis of discrete models of social dynamics. *Mathematical Structures in Computer Science*, 22(5):788–815, 2012.
21. M. Dyer, C. Greenhill, L.A Goldberg, G. Istrate, and M. Jerrum. The Convergence of Iterated Prisoner's Dilemma Game. *Combinatorics, Probability and Computing*, 11:135–147, 2002.
22. P. Hines, I. Dobson, and P. Rezaei. Cascading power outages propagate locally in an influence graph that is not the actual grid topology. *IEEE Trans. Power Systems*, 32(2):958–967, 2017.
23. N. Fates and M. Morvan. Perturbing the topology of the game of life increases its robustness to asynchrony. In *Int. Conf. Cellular Automata*, pages 111–120. Springer, 2004.
24. A. Dennunzio, E. Formenti, and Petr Kǔrka. Cellular automata dynamical systems. In *Handbook of Natural Computing*, pages 25–75. Springer, 2012.
25. L. E. Blume. Population games. In S. Durlauf and H. Peyton Young, editors, *Social Dynamics: Economic Learning and Social Evolution*. MIT Press, 2001.
26. D.Monderer and L. S. Shapley. Potential games. *Games and Economic Behaviour*, 14:124–143, 1996.