Manuscript - An Analysis of One-Dimensional Schelling Segregation

by Christina Brandt, Nicole Immorlica, Gautam Kamath, Robert Kleinberg

Author: Adrian van Schie
Advisor: Paul Dütting
Buddy: Nicolas Kirchmayr

May 19, 2016

1 Introduction

In 1969, the economist Thomas Schelling devised an agent-based model in order to understand the game theoretic aspects of racial segregation in large cities. Schelling observed that despite agents being tolerant in having neighbors of another race, over time they would still segregate themselves from agents of another race. Racial segregation can be observed for example in major cities, as depicted in Figure 1.

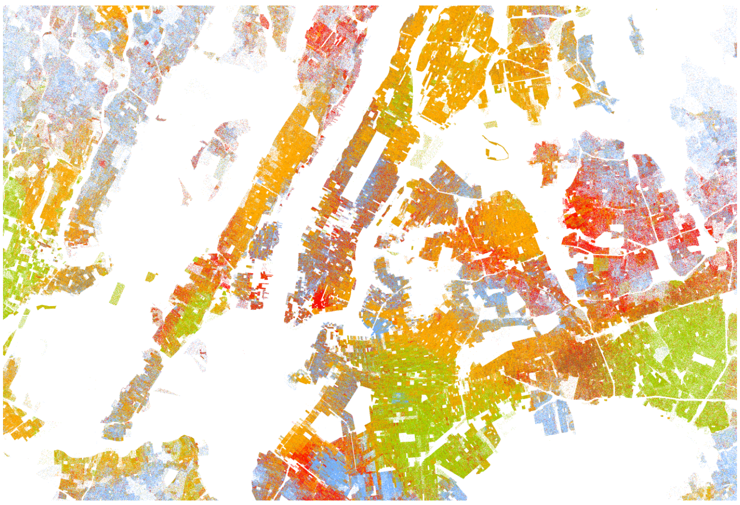

Figure 1: Racial segregation in New York city. Caucasians in blue, African Americans in green, Latinos in yellow, and Asians in green. Source: The Racial Dot Map

In their paper *An Analysis of One-Dimensional Schelling Segregation*, the four authors Christina Brandt, Nicole Immorlica, Gautam Kamath and Robert Kleinberg analyze the Schelling model of segregation in a 1-dimensional setting,

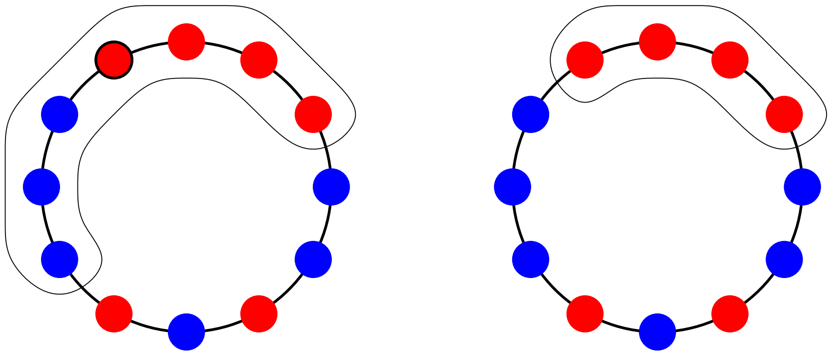

Figure 2: A ring network with n=12 nodes of two different types. With w=3. In the left figure, the nodes encircled define the neighborhood of the red node with black border. The red node is happy, since the majority of the nodes in its neighborhood are red. On the right side, the encircled nodes form a firewall.

where n individuals live in a ring network. There are two different types, and each individual is assigned a type uniformly at random. An individual is happy if the majority of the individuals in his neighborhood are of his type. If an individual is not happy, he "moves out", i.e., he switches positions with another unhappy individual of the opposite type. One question is whether the segregation process will terminate. The authors prove that for all but a tiny fraction of initial configurations the process reaches a configuration where no further swaps are possible. Another question that then arises is how long the average length of a sequence with individuals of the same type is in such a final configuration where no further swaps are possible. The authors prove in their main theorem that the length is not dependent on the ring size n, but is rather dependent on the parameter w that determines the neighborhood. To be more precise, they prove an upper bound of $O(w^2)$ for the average length of a sequence with nodes of the same type in a final configuration.

2 The Model

We consider a ring network of n nodes. Each node is inhabited by one individual whose type is chosen uniformly at random from $\{x, o\}$. The *neighborhood* of a node is determined by the *window size* w: it is defined as the set of the $2w + 1$ nearest neighbors, including the node itself. We generally assume that $w \ll n$. An individual is called *happy*, if at least a fraction τ of individuals in his neighborhood have the same type. τ is called the *tolerance parameter* and in this model we have $\tau = \frac{1}{2}$. Simply put, an individual is happy if at least w other individuals in his neighborhood have the same label.
As already mentioned, unhappy individuals may "move out". In this model, at each timestep two individuals are chosen uniformly at random as candidates for a *swap*. These two individuals swap their nodes if they are both unhappy and are oppositely labeled. After a swap, both involved individuals *are* happy. This is a crucial insight and it will be used in the proofs. The case where no further swaps are possible is called a *frozen configuration*.

A *block* is a sequence of neighboring nodes. A *run* is defined as a block where all nodes have the same label. A run of length at least $w + 1$ has a special property: each individual in this block is happy, because there are w other individuals in his neighborhood with the same label. These individuals remain happy during the whole process, i.e., no such individual will ever swap nodes. We call such a block a *firewall*, and if the type is mentioned we refer to an *x-firewall* or *o-firewall*, respectively.
Figure 2 shows an example where some of the above terms are explained pictorially.

3 Reaching a Frozen Configuration

In the first proposition we see that the probability to reach a frozen configuration converges to 1, as $n \to \infty$.

Proposition 1. *Consider the segregation process with window size w on a ring network of size n. For any fixed w, as $n \to \infty$, the probability that the process eventually reaches a frozen configuration converges to 1.*

Proof. We define a potential function $S_0(t)$ that denotes the set of individuals belonging to a firewall at time t. Furthermore, let $S_1(t)$ be the complement of $S_0(t)$, i.e., the set of individuals that do not belong to a firewall at time t. Given that both $S_0(t)$ and $S_1(t)$ are nonempty, there exists an individual $a \in S_1(t)$ neighboring an individual $b \in S_0(t)$. W.l.o.g. let b be to the right of a. These two individuals must have different labels, otherwise a would be part of the same firewall as b and belong to $S_0(t)$, which is a contradiction. Next we will see that a must be unhappy. Assume that a was happy. Its w neighbors on the right side, including b, are all of the opposite label. Then, in order for a to be happy, the w neighbors on the left side must all have the same label as a. But then, a is part of a run of length at least $w + 1$, which is a firewall by definition. This contradicts the fact that $a \in S_1(t)$.
If there are unhappy individuals of the opposite label, there is a positive probability that a will eventually swap with such an individual. After the swap, the new individual will extend b's firewall, increasing the potential function $S_0(t)$. In the other case where there are no unhappy individuals of the opposite label, the configuration is already frozen.
We have defined a potential function that takes integer values between 0 and n. There are three different cases for the initial configuration. In the case where both $S_0(t)$ and $S_1(t)$ are nonempty, the potential function $S_0(t)$ has a positive probability to increase. If $S_1(t)$ is empty, all individuals already live in a firewall and the configuration is frozen. In the last case where $S_0(t)$ is empty, the initial configuration contained no firewalls. The probability for this case is small. By partitioning the ring into $\frac{n}{w+1}$ blocks, each block has the probability $\frac{1}{2^w}$ of being a firewall in the initial configuration. The probability that none of the $\frac{n}{w+1}$ blocks is a firewall is $(1 - \frac{1}{2^w})^{\frac{n}{w+1}}$, which is $o(1)$ as $n \to \infty$.
In the first two cases we eventually reach a frozen configuration. The probability that the initial configuration is one of these two cases converges to 1 as $n \to \infty$, which concludes to proof. $\square$

4 Bounding the Average Run Length

Theorem 1 is the main theorem of the paper. The proof makes use of several lemmas and propositions, which we need to understand first.

Theorem 1. *Consider the segregation process with window size w on a ring network of size n, starting from a uniformly random initial configuration. There exists a constant $c < 1$ and a function $n_0 : \mathbb{N} \to \mathbb{N}$ such that for all w and all $n \geq n_0(w)$, with probability 1 - o(1), the process reaches a configuration after finitely many steps in which no further swaps are possible. The average run length in this final configuration is $O(w^2)$. In fact, the distribution of runlengths in the final configuration is such that for all $\lambda > 0$, the probability of a randomly selected node belonging to a run of length greater than λw^2 is bounded above by c^λ.*

In order to bound the average run length, the authors define a construct called a *firewall incubator*. They first show that such firewall incubators occur with a reasonable probability in the initial configuration, i.e., in a uniformly random *{x,o}-labeling* of the nodes. Then, they analyze the behavior of a firewall incubator. They show that if a firewall incubator satisfies a specific condition, it eventually becomes a firewall. A lower bound on the probability for this condition to occur is proved by analyzing node swaps and the application of a theorem from the field of combinatorics. The propositions and lemmas are then used in the proof of the main theorem of the paper.

The *x-bias* $\beta_t(i)$ of a node i is the number of x-labeled individuals minus the number of o-labeled individuals in the neighborhood of i at time t. If the time follows from the context, we can write $\beta(i)$. We can express this as a sum of signs, where an x-labeled individual corresponds to $+1$ and an o-labeled individual to -1, respectively. An x-labeled individual is happy if and only if $\beta(i) > 0$ and an o-labeled individual is happy if and only if $\beta(i) < 0$.

Definition 1. *A* firewall incubator *is a block F made up of three consecutive blocks D_L, I, D_R (called left defender, internal, right defender, respectively), such that:*

1. *D_L and D_R have exactly w + 1 nodes;*
2. *$\beta_0(i) > \sqrt{w}$ for all $i \in F$;*
3. *The minimum x-bias in D_L occurs at its left endpoint, and the minimum x-bias in D_R occurs at its right endpoint.*

The blocks of length w immediately to the left and right of F are denoted by A_L, A_R and are called the left and right attackers.

For the subsequent notation, signs as for the x-bias are used again. The sum of the j leftmost signs of a block B in $\{x, o\}^k$ can be written as $\chi_j(B)$, where $0 \leq j \leq k$. In the case where $j = k$ we can write $\chi(B)$ instead of $\chi_k(B)$. If we want to refer to the sum of the j rightmost signs in a block B, we write $\chi_{-j}(B)$, which is equal to $\chi(B) - \chi_{k-j}(B)$.

Definition 2. *A sequence $B \in \{x, o\}^w$ is* x-promoting *if $\chi(B) \geq 5\sqrt{w}$ and for all $j = 1, \ldots, w$, $\chi_j(B) > -2\sqrt{w}$ and $\chi_{-j}(B) > -2\sqrt{w}$.*

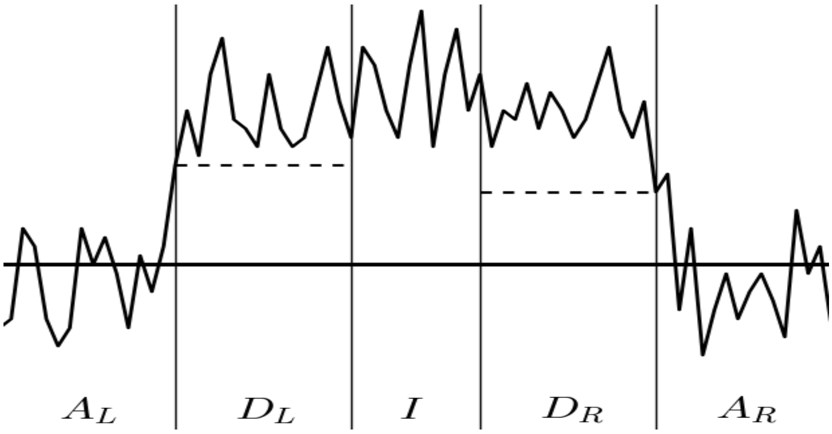

Figure 3: A firewall incubator consisting of the five blocks A_L, D_L, I, D_R, A_R. The height corresponds to the x-bias of each individual. The long horizontal line has height $\sqrt{w}$.

Lemma 1. *The probability that a uniformly random sequence* $B \in \{x, o\}^w$ *is x-promoting is* $\Omega(1)$.

Proof. First, two events are defined. For $k > 0$, let

$$\xi_k^L = \{\exists j : \chi_j(B) \leq -k\}$$
$$\xi_k^R = \{\exists j : \chi_{-j}(B) \leq -k\}.$$

The *reflection principle* applied to our case says: if B is a uniformly random element of $\{x, o\}^w$, then for all $k > 0$

$$\begin{aligned} Pr(\xi_k^L) &= Pr(\chi(B) \leq -k) + Pr(\chi(B) < -k) \\ &= Pr(\chi(B) \leq -k) + Pr(\chi(B) > k). \end{aligned} \tag{1}$$

The second equality follows by symmetry. Since $\chi(B)$ is a sum of w independent random variables that are +1 or -1,

$$\mathbb{E}[(\chi(B))^2] = w. \tag{2}$$

By using the above equalities, we get the following inequality

$$\begin{aligned} Pr(\xi_k^L) &\leq Pr((\chi(B))^2 \geq k^2) \\ &\leq \frac{\mathbb{E}[(\chi(B))^2]}{k^2} \\ &= \frac{w}{k^2}, \end{aligned}$$

for $k > 0$. The first inequality follows from equation (1), since $Pr((\chi(B))^2 \geq k^2)$ $= Pr(\chi(B) \geq k) + Pr(\chi(B) \leq -k)$. *Markov's inequality* is applied in the next inequality, as $(\chi(B))^2$ is nonnegative and $k^2 > 0$. In the last step (2) is used. Whenever $a < b$ are two numbers such that the events $\chi(B) = a$ and $\chi(B) = b$ have positive probability, the inequality

$$Pr(\xi_k^L | \chi(B) = a) \geq Pr(\xi_k^L | \chi(B) = b) \tag{3}$$

holds. From inequality (3) it follows that for any b such that the events $\chi(B) < b$ and $\chi(B) \geq b$ have positive probability,

$$Pr(\xi_k^L|\chi(B) < b) \geq Pr(\xi_k^L|\chi(B) \geq b)$$

holds. Since the unconditional probability of ξ_k^L is a weighted average of the left and right side, i.e., $Pr(\xi_k^L) = Pr(\xi_k^L|\chi(B) \geq b) \cdot Pr(\chi(B) \geq b) + Pr(\xi_k^L|\chi(B) < b) \cdot Pr(\chi(B) < b)$, we get

$$Pr(\xi_k^L|\chi(B) \geq b) \leq Pr(\xi_k^L) \leq \frac{w}{k^2}, \tag{4}$$

for all b, for $k > 0$. By symmetry, this also holds for ξ_k^R,

$$Pr(\xi_k^R|\chi(B) \geq b) = Pr(\xi_k^L|\chi(B) \geq b).$$

By the *Central Limit Theorem*,

$$\lim_{w\to\infty} Pr(\chi(B) \geq 5\sqrt{w}) = \frac{1}{\sqrt{2\pi}} \int_5^\infty e^{-\frac{x^2}{2}}\, dx.$$

Hence, there is a constant c_0 such that $Pr(\chi(B) \geq 5\sqrt{w}) > 2c_0$ for all $w \geq 25$. Note that we require $w \geq 25$, such that $w \geq 5\sqrt{w}$.
The event that B is x-promoting can be expressed in terms of the sum of signs and the events ξ_k^L, ξ_k^R. By setting $b = 5\sqrt{w}$, $k = 2\sqrt{w}$ we find a constant lower bound for the probability that B is x-promoting:

$$\begin{aligned}
Pr(\text{B is x-promoting}) &= Pr(\chi(B) \geq b \wedge \neg\xi_k^L \wedge \neg\xi_k^R) \\
&= Pr(\chi(B) \geq b) \cdot Pr(\neg\xi_k^L \wedge \neg\xi_k^R|\chi(B) \geq b) \\
&\geq Pr(\chi(B) \geq b) \cdot (1 - 2Pr(\xi_k^L|\chi(B) \geq b)) \\
&> 2c_0(1 - 2\frac{w}{k^2}) \\
&= c_0.
\end{aligned}$$

The first inequality follows from

$$\begin{aligned}
Pr(\neg\xi_k^L \wedge \neg\xi_k^R|\chi(B) \geq b) &= 1 - Pr(\xi_k^L \vee \xi_k^R|\chi(B) \geq b) \\
&\geq 1 - Pr(\xi_k^L|\chi(B) \geq b) - Pr(\xi_k^R|\chi(B) \geq b) \\
&= 1 - 2Pr(\xi_k^L|\chi(B) \geq b).
\end{aligned}$$

In the second inequality we use the constant lower bound $2c_0$ for $Pr(\chi(B) \geq b)$ and make use of inequality (4). Since $k = 2\sqrt{w}$, $(1 - 2\frac{w}{k^2})$ is equal to $\frac{1}{2}$ in the last step. $\square$

Proposition 2. *Let r be an integer such that $6 \leq r < \frac{n}{w} - 2$. For any sequence of rw consecutive nodes, the probability that a uniformly random $\{x,o\}$-labeling of the nodes contains an x-firewall incubator that starts among the leftmost w nodes and ends among the rightmost w nodes is at least c^r, where $c > 0$ is a constant independent of r, w, n.*

Proof. Given a sequence of rw consecutive nodes, partition it into r blocks $B_1, B_2, \ldots, B_r$, each of length w. Further, let B_0 denote the block preceding B_1 and B_{r+1} denote the block following B_r, each containing w nodes. Next, notation is introduced to refer to a modified labeling of the blocks $B_0, B_1, \ldots, B_{r+1}$. λ_{00} refers to the original labeling, λ_{01} refers to the labeling where the ordering of the last $4w$ labels is reversed, λ_{10} refers to the labeling where the ordering of the first $4w$ labels are reversed and λ_{11} refers to the labeling where the ordering of both the first $4w$ and last $4w$ labels are reversed. By applying Lemma 1, the probability that the $r+2$ blocks of λ_{00} are x-promoting is ${c_0}^{r+2}$. The reverse of an x-promoting sequence is also x-promoting. Therefore if λ_{00} is x-promoting, so are λ_{01}, λ_{10}, and λ_{11}. Given that all $r+2$ blocks $B_0, B_1, \ldots, B_{r+1}$ are x-promoting, by symmetry, the labelings in the set $\{\lambda_{00}, \lambda_{01}, \lambda_{10}, \lambda_{11}\}$ are equiprobable. If we can show that at least one of these four labelings has an x-firewall incubator that starts in block B_1 and ends in block B_r, then the probability that such an incubator exists is $\frac{1}{4}{c_0}^{r+2}$.
A node i at position j in the middle block B' of three consecutive x-promoting blocks B, B', B'' has an x-bias $\beta_0(i) > \sqrt{w}$, since

$$\begin{aligned}
\beta_0(i) &= \chi_{-(w-j+1)}(B) + \chi(B') + \chi_j(B'') \\
&> -2\sqrt{w} + 5\sqrt{w} - 2\sqrt{w} \\
&= \sqrt{w}
\end{aligned}$$

The $5\sqrt{w}$ and $-2\sqrt{w}$ terms follow from the definition of an x-promoting sequence.
Given that all labelings in the set $\{\lambda_{00}, \lambda_{01}, \lambda_{10}, \lambda_{11}\}$ are x-promoting, every node in $B_1, \ldots, B_r$ has an x-bias greater than $\sqrt{w}$ in all four labelings λ_{00}, λ_{01}, λ_{10}, λ_{11}.
In a firewall incubator it is required that the leftmost individual in the D_L block has minimal x-bias among all individuals of D_L. The same requirement holds for the rightmost individual in the D_R block. We can find such two individuals in at least one of the four labelings in the set $\{\lambda_{00}, \lambda_{01}, \lambda_{10}, \lambda_{11}\}$. Depending on where the minimum x-bias of $B_1 \cup B_2$ in the original labeling lies, we can take the corresponding labeling such that the minimum lies in B_1 of the new labeling. The important thing is that in the new labeling, the x-bias of the w - 1 nodes to the right of the node with minimal x-bias is greater or equal. The same can be done analogously for the right side. These two individuals define the start and end of a firewall incubator. As already seen, all the individuals in between have an x-bias greater than $\sqrt{w}$, which is a requirement for the elements in a firewall incubator. □

Definition 3. *The* satisfaction time *of a node i, denoted by $t_i^\star$, is defined to be the first time when i is selected to participate in a proposed swap with an unhappy, oppositely labeled individual. (If no such time exists, then $t_i^\star = \infty$.) A node i is called* impatient *at time t if it is unhappy and $t \leq t_i^\star$.*

Definition 4. *For a firewall incubator $F = D_L \cup I \cup D_R$ with corresponding attackers A_L, A_R, a* left attacking *x is an individual of type x who belongs to A_L in the initial configuration, and a* left defending *o is an individual of type o who belongs to D_L in the initial configuration. A* left combatant *is an individual that is either a left attacking x or a left defending o. The equivalent terms*

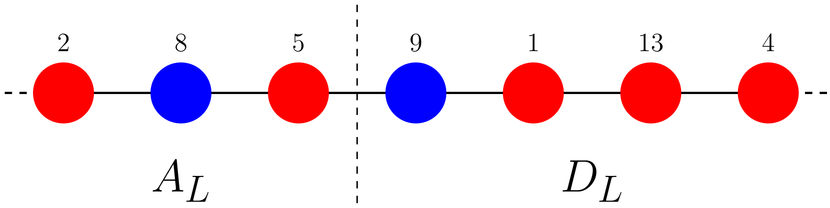

Figure 4: 7 nodes with corresponding satisfaction times. The nodes are part of a red-firewall incubator, hence, the first, third and fourth nodes are combatants. Blue labeled individuals correspond to +1 and the red ones to -1. Then, the left-transcript is [+1, -1, -1].

with "right" in place of "left" are defined similarly; henceforth when referring to combatants we will omit "left" and "right" when they can be inferred from context. The number of left attacking x's and left defending o's are denoted by a_L, d_L, and for the right combatants we define a_R, d_L similarly.

Definition 5. *The* left-transcript *(resp.* right-transcript*) is the sign sequence obtained by listing all of the left (resp. right) combatants in reverse order of satisfaction time, and translating each attacking x in this list to +1 and each defending o to -1.*
If there exists a time t_0 at which no individuals in F are impatient, any sign sequence obtained from the left-transcript (resp. right-transcript) by permuting the signs associated to individuals whose satisfaction time is after t_0, while fixing all other signs in the transcript, is called a left-pseudo-transcript *(resp.* right-pseudo-transcript*).*

Figure 4 shows a setting of nodes and the corresponding transcript.

Proposition 3. *Suppose that F is a firewall incubator and there exist left- and right-pseudo-transcripts such that all partial sums of both pseudo-transcripts are non-negative. Then F becomes an x-firewall.*

Proof. The proof is by contradiction. Let t_0 denote the earliest time at which no individual in F is impatient, i.e., each individual is either happy or has already reached his satisfaction time. This time t_0 exists, since the assumption is that left- and right-pseudo-transcripts exist. If F is not an x-firewall at time t_0, then some node $j \in F$ contains an individual labeled o. Next we see by case distinction that there exists a time $t_1, t_1 \leq t_0$, at which node j is occupied by a happy individual of label o. We already know that j is not impatient at time t_0, so either it is happy or its satisfaction time is smaller than t_0, $t_j^\star < t_0$. If it is happy, we simply set $t_1 = t_0$ and are done. In the case where the o-labeled individual has never moved, we set $t_1 = t_j^*$. The individual must have been happy at this time, otherwise it would have moved. In the last case where the o-labeled individual in node j has moved in at some time, we set t_1 to the time immediately after it moved to j. It must have been happy at this time, otherwise it would not have swapped nodes.
Let t denote the first time at which a node in F has a negative x-bias. It must hold that $t \leq t_1$, since at time t_1, the o-labeled individual at node j is happy, i.e., it has a negative x-bias. Up until time t, all nodes in F have a positive

x-bias. This means that up until time t, individuals moving out of F have label o, since all x-labeled individuals are happy.

The first node i to develop a negative x-bias is not in the internal I, but rather must be in D_L or D_R. This is because the x-bias of a node in the internal I is determined only by the labels of individuals in D_L, I, D_R and up until time t, the only individuals moving out are labeled o. W.l.o.g., let $i \in D_L$. At time 0, the x-bias of i was bounded below by the x-bias of the leftmost node in D_L, following from the definition of a firewall incubator. The neighborhood of the leftmost node in D_L is $A_L \cup D_L$. We can express the x-bias of this node in terms of the number of attacking x's and defending o's:

$$\underbrace{a_L}_{\text{x's in } A_L} - \underbrace{(w - a_L)}_{\text{o's in } A_L} + \underbrace{(w + 1 - d_L)}_{\text{x's in } D_L} - \underbrace{d_L}_{\text{o's in } D_L} = 2(a_L - d_L) + 1.$$

Since node i's x-bias is greater or equal than the x-bias of the leftmost node in D_L, we get $\beta_0(i) \geq 2(a_L - d_L) + 1$.

Up until time t, whenever the satisfaction time of an o-labeled individual in D_L is reached, it moves out because it is unhappy. Node i's x-bias increases therefore by 2 for each such swap, since an x-labeled individual moves in. When an x-labeled individual swaps out of the A_L block, it decreases the x-bias of i either by 2 or 0. In the case where the x-labeled individual is not in i's neighborhood, the x-bias remains unchanged. In the cases where the attacking x (i.e., an x-labeled individual that hasn't moved yet) moves out, either at or after its satisfaction time (but before t), the contribution to i's x-bias is -2. In the last case, where an x-individual swaps into A_L, becomes unhappy, and later swaps out, the total contribution is 0. Let a_L^t denote the number of attacking x's in A_L whose satisfaction time is before t and d_L^t the number of defending o's in the D_L block whose satisfaction time is before t. i's decrement of the x-bias is at most $2a_L^t$ until time t, and the increment is $2d_L^t$. We obtain the following inequality:

$$\begin{aligned}
\beta_t(i) &\geq \beta_0(i) + 2d_L^t - 2a_L^t \\
&\geq 2(a_L - d_L) + 1 + 2d_L^t - 2a_L^t \\
&> 2((a_L - a_L^t) - (d_L - d_L^t)).
\end{aligned}$$

$(a_L - a_L^t)$ is the number of attacking x's whose satisfaction time is greater than t and $(d_L - d_L^t)$ is the number of defending o's whose attacking time is greater than t. The above lower bound is twice the k-th partial sum of the left-transcript, where $k = (a_L - a_L^t) + (d_L - d_L^t)$ denotes the number of individuals whose satisfaction time is after t.

Since $t \leq t_0$, any left-pseudo-transcript differs from the left-transcript only by permuting a subset of the first k signs, and therefore has the same k-th partial sum. The assumption that the x-bias of i becomes negative at time t contradicts the hypothesis that there exists a left-pseudo-transcript whose partial sums are all non-negative, which concludes the proof. □

Lemma 2. *(Ballot Theorem). Consider a multiset consisting of* a *many copies of +1 and* b *many copies of -1, and let $x_1, x_2, \ldots, x_{a+b}$ be a uniformly random ordering of the elements of this multiset. The probability that all partial sums $x_1 + \cdots + x_j$ $(1 \leq j \leq a + b)$ are strictly positive is equal to max$(0, \frac{a-b}{a+b})$.*

Proposition 4. *If B is a random block of length 6w, then with probability $\Omega(\frac{1}{w})$, B contains a firewall incubator having left- and right-pseudo-transcripts whose partial sums are all non-negative.*

Here is the intuition behind the proof. We already know from Proposition 2 that B contains a firewall incubator with constant probability. If the left-transcript of B would be a uniformly random permutation of the a_L +1's and d_L -1's, we could apply the *Ballot Theorem.* Then, the probability that the partial sums are all non-negative is $\frac{a_L - d_L}{a_L + d_L}$. We have seen in the proof of Proposition 3 that at time 0 the x-bias of the leftmost node in D_L is $2(a_L - d_L) + 1$ and by definition of an x-firewall incubator must be greater than $\sqrt{w}$. It follows that $(a_L - d_L) \geq \frac{\sqrt{w}}{2}$. We also know that $(a_L + d_L) \leq 2w + 1$, since the number of nodes in the A_L and D_L blocks is w and $w + 1$, respectively. Combining these two inequalities, we see that $\frac{a_L - d_L}{a_L + d_L} = \Omega(\frac{1}{\sqrt{w}})$.
Unfortunately, the transcript is not a uniformly random permutation. A bias arises, since the number of unhappy elements of each label is not precisely equal. If for example at some point there are more unhappy o's than unhappy x's, the satisfaction time of an attacking x is more likely to happen earlier. In the paper, the authors show that the number of unhappy individuals is approximately balanced for a sufficiently long time; until some time t_0 these imbalances can be considered "small". These small imbalances are artificially corrected by introducing the concept of censored individuals and censored swaps. Suppose there are m extra unhappy individuals of one type, say x. Then choose m unhappy x-labeled individuals at random and call them *censored.* A swap is called *censored* if it involves a censored individual. Additionally, a swap involving two combatants of F is also considered censored. As long as the imbalance m is small, the probability that a swap is censored is also small. By conditioning on having no censored swaps, the transcript is indeed a uniformly random permutation. The authors then show that with probability $1 - o(1)$, no individual in F is impatient at time t_0. A pseudo-transcript can be obtained by randomly permuting the combatants whose satisfaction times are after t_0, and given that no censored swaps occurred before t_0, the pseudo-transcript is a uniformly random permutation. Then, Proposition 4 follows from Lemma 2.

We can now move on to the proof of Theorem 1, as we have seen all lemmas and propositions that will be made us of.

Proof. (of Theorem 1) In Proposition 1 we have seen that the process reaches a frozen configuration with high probability. We select a randomly sampled node a. We partition the ring in clockwise direction into blocks of length $6w$. By Proposition 4, each of these blocks has the probability $\Omega(\frac{1}{\sqrt{w}})$ of containing an x-firewall incubator having left- and right-pseudo-transcripts whose partial sums are all non-negative. Thus, for a suitable constant $c < 1$, the probability that none of the first $\frac{\lambda w}{6}$ blocks encountered on a clockwise scan of blocks with length $6w$ starting at node a contains an x-firewall in the frozen configuration is bounded above by c^{λ}. By symmetry, we can make analogous arguments for o instead of x and for the counterclockwise order. Node a cannot belong to a run of length greater than λw^2 assuming that it has individuals of both labels within this radius on both sides of itself, which completes the proof. $\square$

5 Conclusion

The authors have analyzed 1-dimensional Schelling segregation in a ring network where the neighborhood of a node is determined by a window size parameter w. They have shown that with high probability the process results in a frozen configuration where most nodes belong to a run whose size has lower and upper limits $\Omega(w)$ and $O(w^2)$. The run size is only dependent on parameter w and is not dependent on the ring size n.

In the meantime, the authors Nicole Immorlica, Gautam Kamath and Robert Kleinberg have found an improved upper bound of $O(w)$, which makes the bound on the average run length tight, i.e., $\Theta(w)$.