

Cutoff for product replacement on finite groups

Alex Zhai

Stanford University

Joint work with



Yuval Peres
Microsoft Research



Ryokichi Tanaka
Tohoku University

Product replacement walk




- Let G be a finite group, and consider an n -tuple of group elements (call this a **configuration**). For example, $G = \mathbb{Z}/11$ and $n = 6$.

Product replacement walk




- Let G be a finite group, and consider an n -tuple of group elements (call this a **configuration**). For example, $G = \mathbb{Z}/11$ and $n = 6$.
- We consider the **product replacement walk**: randomly pick two of the elements a and b , and replace a with $a + b$ or $a - b$.

Product replacement walk



- Let G be a finite group, and consider an n -tuple of group elements (call this a **configuration**). For example, $G = \mathbb{Z}/11$ and $n = 6$.
- We consider the **product replacement walk**: randomly pick two of the elements a and b , and replace a with $a + b$ or $a - b$.

Stationary distribution



- The process can be viewed as random walk on an undirected regular graph: each configuration can transition to/from $2n(n-1)$ other configurations (possibly with self-loops).

Product replacement walk



- Let G be a finite group, and consider an n -tuple of group elements (call this a **configuration**). For example, $G = \mathbb{Z}/11$ and $n = 6$.
- We consider the **product replacement walk**: randomly pick two of the elements a and b , and replace a with $a + b$ or $a - b$.

Stationary distribution



- The process can be viewed as random walk on an undirected regular graph: each configuration can transition to/from $2n(n-1)$ other configurations (possibly with self-loops).

Stationary distribution



- The process can be viewed as random walk on an undirected regular graph: each configuration can transition to/from $2n(n-1)$ other configurations (possibly with self-loops).
- The graph is not connected, e.g. the subgroup generated by the elements remains invariant throughout the process.

Stationary distribution



- The process can be viewed as random walk on an undirected regular graph: each configuration can transition to/from $2n(n-1)$ other configurations (possibly with self-loops).
- The graph is not connected, e.g. the subgroup generated by the elements remains invariant throughout the process.
- We focus on “generating n -tuples”, i.e. n -tuples that generate the whole group.

Product replacement walk




- Let G be a finite group, and consider an n -tuple of group elements (call this a **configuration**). For example, $G = \mathbb{Z}/11$ and $n = 6$.
- We consider the **product replacement walk**: randomly pick two of the elements a and b , and replace a with $a + b$ or $a - b$.

Stationary distribution



- The process can be viewed as random walk on an undirected regular graph: each configuration can transition to/from $2n(n-1)$ other configurations (possibly with self-loops).
- The graph is not connected, e.g. the subgroup generated by the elements remains invariant throughout the process.
- We focus on “generating n -tuples”, i.e. n -tuples that generate the whole group.
- As long as n is big enough (say $n \gg \log |G|$), any two generating n -tuples can reach each other by product replacement steps. Also, the vast majority of configurations are generating n -tuples.

Where does product replacement arise?



Product replacement walk is related to:

Where does product replacement arise?



Product replacement walk is related to:

- Sampling random group elements (**product replacement algorithm**)

Where does product replacement arise?



Product replacement walk is related to:

- Sampling random group elements (**product replacement algorithm**)
- Random walk on $SL_n(\mathbb{Z}/q)$


Where does product replacement arise?



Product replacement walk is related to:

- Sampling random group elements (**product replacement algorithm**)
- Random walk on $SL_n(\mathbb{Z}/q)$
- Graph-restricted variants: instead of picking any two elements, only allow certain pairs


Where does product replacement arise?



Product replacement walk is related to:

- Sampling random group elements (**product replacement algorithm**)
- Random walk on $SL_n(\mathbb{Z}/q)$
- Graph-restricted variants: instead of picking any two elements, only allow certain pairs
 - East model (graph = path)

Where does product replacement arise?



Product replacement walk is related to:

- Sampling random group elements (**product replacement algorithm**)
- Random walk on $SL_n(\mathbb{Z}/q)$
- Graph-restricted variants: instead of picking any two elements, only allow certain pairs
 - East model (graph = path)
 - Higher-dimensional versions?

- Let $\sigma_0, \sigma_1, \sigma_2, \dots$ be the sequence of configurations of the chain. (So σ_t is configuration at time t .)

- Let $\sigma_0, \sigma_1, \sigma_2, \dots$ be the sequence of configurations of the chain. (So σ_t is configuration at time t .)
- Recall for $\epsilon \in (0, 1)$ that the **mixing time** $t_{\text{mix}}(\epsilon)$ is the earliest time t such that

$$\|\mathbf{P}(\sigma_t \in \cdot) - \pi\|_{\text{TV}} \leq \epsilon,$$

where π is the stationary distribution.

- **Cutoff phenomenon:** refers to when the TV-distance from stationary rapidly drops from near 1 to near 0 within a short timescale.

Theorem (Peres-Tanaka-Z.)

Let G be a fixed finite group, and consider the product replacement walk $(\sigma_t)_{t \geq 0}$. Then, for any fixed $\epsilon > 0$,

$$\left(\frac{3}{2} - o(1)\right) n \log n \leq t_{\text{mix}}(1 - \epsilon) \leq t_{\text{mix}}(\epsilon) \leq \left(\frac{3}{2} + o(1)\right) n \log n$$

as $n \rightarrow \infty$.

Theorem (Peres-Tanaka-Z.)

Let G be a fixed finite group, and consider the product replacement walk $(\sigma_t)_{t \geq 0}$. Then, for any fixed $\epsilon > 0$,

$$\left(\frac{3}{2} - o(1)\right) n \log n \leq t_{\text{mix}}(1 - \epsilon) \leq t_{\text{mix}}(\epsilon) \leq \left(\frac{3}{2} + o(1)\right) n \log n$$

as $n \rightarrow \infty$.

- Extends work of Ben-Hamou and Peres who showed this for $G = \mathbb{Z}/2$.

Theorem (Peres-Tanaka-Z.)

Let G be a fixed finite group, and consider the product replacement walk $(\sigma_t)_{t \geq 0}$. Then, for any fixed $\epsilon > 0$,

$$\left(\frac{3}{2} - o(1)\right) n \log n \leq t_{\text{mix}}(1 - \epsilon) \leq t_{\text{mix}}(\epsilon) \leq \left(\frac{3}{2} + o(1)\right) n \log n$$

as $n \rightarrow \infty$.

- Extends work of Ben-Hamou and Peres who showed this for $G = \mathbb{Z}/2$.
- Diaconis and Saloff-Coste proved $t_{\text{mix}} = O(n^2 \log n)$ and conjectured $t_{\text{mix}} = O(n \log n)$.

Analysis of the mixing can be decomposed into three stages:

Analysis of the mixing can be decomposed into three stages:

- Initially, the elements in the configuration might be mostly constrained to some subgroup (“subgroup-confined” regime). It takes about $n \log n$ steps to escape the subgroup-confined regime.

Analysis of the mixing can be decomposed into three stages:

- Initially, the elements in the configuration might be mostly constrained to some subgroup (“subgroup-confined” regime). It takes about $n \log n$ steps to escape the subgroup-confined regime.
- Once outside the subgroup-confined regime, counts of elements can be approximated by a differential equation. After about $\frac{1}{2} n \log n$ steps, each group element appears with approximately equal frequency.

Analysis of the mixing can be decomposed into three stages:

- Initially, the elements in the configuration might be mostly constrained to some subgroup (“subgroup-confined” regime). It takes about $n \log n$ steps to escape the subgroup-confined regime.
- Once outside the subgroup-confined regime, counts of elements can be approximated by a differential equation. After about $\frac{1}{2}n \log n$ steps, each group element appears with approximately equal frequency.
- Even if group elements appear with roughly equal frequency, mixing is not guaranteed. A coupling argument for $O(n)$ steps is needed to finish the proof.

Analysis of the mixing can be decomposed into three stages:

- Initially, the elements in the configuration might be mostly constrained to some subgroup (“subgroup-confined” regime). It takes about $n \log n$ steps to escape the subgroup-confined regime.
- Once outside the subgroup-confined regime, counts of elements can be approximated by a differential equation. After about $\frac{1}{2}n \log n$ steps, each group element appears with approximately equal frequency.
- Even if group elements appear with roughly equal frequency, mixing is not guaranteed. A coupling argument for $O(n)$ steps is needed to finish the proof.

For simplicity, we’ll focus on the case $G = \mathbb{Z}/q$ for a prime q .

Analysis of the mixing can be decomposed into three stages:

- Initially, the elements in the configuration might be mostly constrained to some subgroup (“subgroup-confined” regime). It takes about $n \log n$ steps to escape the subgroup-confined regime.
- Once outside the subgroup-confined regime, counts of elements can be approximated by a differential equation. After about $\frac{1}{2}n \log n$ steps, each group element appears with approximately equal frequency.
- Even if group elements appear with roughly equal frequency, mixing is not guaranteed. A coupling argument for $O(n)$ steps is needed to finish the proof.

For simplicity, we’ll focus on the case $G = \mathbb{Z}/q$ for a prime q . Also, we’ll focus on the upper bound for mixing.

Stage 1: Getting out of the subgroup-confined regime

- Let's show we don't get stuck too long with more than, say, $\frac{2}{3}$ of our elements being zero (other subgroups can be handled similarly).

Stage 1: Getting out of the subgroup-confined regime

- Let's show we don't get stuck too long with more than, say, $\frac{2}{3}$ of our elements being zero (other subgroups can be handled similarly).
- Suppose k out of n elements are currently **not** zero.

Stage 1: Getting out of the subgroup-confined regime

- Let's show we don't get stuck too long with more than, say, $\frac{2}{3}$ of our elements being zero (other subgroups can be handled similarly).
- Suppose k out of n elements are currently **not** zero.
- Consider elements a and b picked for product replacement (so a becomes $a + b$ or $a - b$):

Stage 1: Getting out of the subgroup-confined regime

- Let's show we don't get stuck too long with more than, say, $\frac{2}{3}$ of our elements being zero (other subgroups can be handled similarly).
- Suppose k out of n elements are currently **not** zero.
- Consider elements a and b picked for product replacement (so a becomes $a + b$ or $a - b$):

$$\mathbf{P}(b = 0) \approx 1 - \frac{k}{n} \quad \implies \# \text{ non-zero stays same}$$

$$\mathbf{P}(b \neq 0, a = 0) \approx \frac{k}{n} \left(1 - \frac{k}{n}\right) \implies \# \text{ non-zero increases by 1}$$

$$\mathbf{P}(b \neq 0, a \neq 0) \approx \frac{k^2}{n^2} \quad \implies \# \text{ non-zero stays same or decreases by 1}$$

Stage 1: Getting out of the subgroup-confined regime

- Let's show we don't get stuck too long with more than, say, $\frac{2}{3}$ of our elements being zero (other subgroups can be handled similarly).
- Suppose k out of n elements are currently **not** zero.
- Consider elements a and b picked for product replacement (so a becomes $a + b$ or $a - b$):

$$\mathbf{P}(b = 0) \approx 1 - \frac{k}{n} \implies \# \text{ non-zero stays same}$$

$$\mathbf{P}(b \neq 0, a = 0) \approx \frac{k}{n} \left(1 - \frac{k}{n}\right) \implies \# \text{ non-zero increases by 1}$$

$$\mathbf{P}(b \neq 0, a \neq 0) \approx \frac{k^2}{n^2} \implies \# \text{ non-zero stays same or decreases by 1}$$

- Can analyze k as a birth-or-death process, reaches at least $\frac{n}{3}$ w.h.p. in about $n \log n$ steps.

Stage 2: Differential equation

- For each $a \in G$, let $N_a(\sigma) =$ number of a 's that appear in σ .

Stage 2: Differential equation

- For each $a \in G$, let $N_a(\sigma) =$ number of a 's that appear in σ . We'll first analyze how the $N_a(\sigma)$ behave over time.

Stage 2: Differential equation

- For each $a \in G$, let $N_a(\sigma) =$ number of a 's that appear in σ . We'll first analyze how the $N_a(\sigma)$ behave over time.
- We have the difference equation

$$\mathbf{E}[N_a(\sigma_{t+1}) - N_a(\sigma_t) \mid \sigma_t] = -\frac{N_a(\sigma_t)}{n} + \sum_{b \in G} \frac{N_{a-b}(\sigma_t)N_b(\sigma_t)}{2n(n-1)} + \sum_{b \in G} \frac{N_{a+b}(\sigma_t)N_b(\sigma_t)}{2n(n-1)}.$$

Stage 2: Differential equation

- For each $a \in G$, let $N_a(\sigma) =$ number of a 's that appear in σ . We'll first analyze how the $N_a(\sigma)$ behave over time.
- We have the difference equation

$$\mathbf{E}[N_a(\sigma_{t+1}) - N_a(\sigma_t) \mid \sigma_t] = -\frac{N_a(\sigma_t)}{n} + \sum_{b \in G} \frac{N_{a-b}(\sigma_t)N_b(\sigma_t)}{2n(n-1)} + \sum_{b \in G} \frac{N_{a+b}(\sigma_t)N_b(\sigma_t)}{2n(n-1)}.$$

- Think of $N_a(\sigma_t)$ as a function of a and t .

Stage 2: Differential equation

- For each $a \in G$, let $N_a(\sigma) =$ number of a 's that appear in σ . We'll first analyze how the $N_a(\sigma)$ behave over time.
- We have the difference equation

$$\mathbf{E}[N_a(\sigma_{t+1}) - N_a(\sigma_t) \mid \sigma_t] = -\frac{N_a(\sigma_t)}{n} + \sum_{b \in G} \frac{N_{a-b}(\sigma_t)N_b(\sigma_t)}{2n(n-1)} + \sum_{b \in G} \frac{N_{a+b}(\sigma_t)N_b(\sigma_t)}{2n(n-1)}.$$

- Think of $N_a(\sigma_t)$ as a function of a and t . Our equation approximates the differential equation (after appropriate rescaling)

$$\frac{\partial}{\partial t} N = -N + N * \left(\frac{N + N^-}{2} \right),$$

where $N^-(a, t) = N(-a, t)$.

- After taking the **Fourier transform**, convolution becomes multiplication.

- After taking the **Fourier transform**, convolution becomes multiplication.
- With (complex-valued) Fourier coefficients

$$x_k(t) = \frac{1}{n} \sum_{a \in \mathbb{Z}/q} N_a(\sigma_t) \cdot \omega^{-ka}, \quad \omega = e^{\frac{2\pi i}{q}},$$

- After taking the **Fourier transform**, convolution becomes multiplication.
- With (complex-valued) Fourier coefficients

$$x_k(t) = \frac{1}{n} \sum_{a \in \mathbb{Z}/q} N_a(\sigma_t) \cdot \omega^{-ka}, \quad \omega = e^{\frac{2\pi i}{q}},$$

we end up with the differential equation

$$\frac{\partial}{\partial t} x_k = -x_k + x_k \cdot \frac{x_k + \bar{x}_k}{2}.$$

- After taking the **Fourier transform**, convolution becomes multiplication.
- With (complex-valued) Fourier coefficients

$$x_k(t) = \frac{1}{n} \sum_{a \in \mathbb{Z}/q} N_a(\sigma_t) \cdot \omega^{-ka}, \quad \omega = e^{\frac{2\pi i}{q}},$$

we end up with the differential equation

$$\frac{\partial}{\partial t} x_k = -x_k + x_k \cdot \frac{x_k + \bar{x}_k}{2}.$$

- Note that x_0 is always identically 1.

- After taking the **Fourier transform**, convolution becomes multiplication.
- With (complex-valued) Fourier coefficients

$$x_k(t) = \frac{1}{n} \sum_{a \in \mathbb{Z}/q} N_a(\sigma_t) \cdot \omega^{-ka}, \quad \omega = e^{\frac{2\pi i}{q}},$$

we end up with the differential equation

$$\frac{\partial}{\partial t} x_k = -x_k + x_k \cdot \frac{x_k + \bar{x}_k}{2}.$$

- Note that x_0 is always identically 1. If $x_k \rightarrow 0$ for all $k \neq 0$, this means $N_a \rightarrow \frac{1}{|G|} n$.

$$\frac{\partial}{\partial t} x_k = -x_k + x_k \cdot \frac{x_k + \bar{x}_k}{2}$$

$$\frac{\partial}{\partial t} x_k = -x_k + x_k \cdot \frac{x_k + \bar{x}_k}{2}$$

- If we're outside the subgroup-confined regime, we can ensure that $\operatorname{Re} x_k < 1 - \delta$, where δ depends only on G

- After taking the **Fourier transform**, convolution becomes multiplication.
- With (complex-valued) Fourier coefficients

$$x_k(t) = \frac{1}{n} \sum_{a \in \mathbb{Z}/q} N_a(\sigma_t) \cdot \omega^{-ka}, \quad \omega = e^{\frac{2\pi i}{q}},$$

we end up with the differential equation

$$\frac{\partial}{\partial t} x_k = -x_k + x_k \cdot \frac{x_k + \bar{x}_k}{2}.$$

- Note that x_0 is always identically 1. If $x_k \rightarrow 0$ for all $k \neq 0$, this means $N_a \rightarrow \frac{1}{|G|} n$.

$$\frac{\partial}{\partial t} x_k = -x_k + x_k \cdot \frac{x_k + \bar{x}_k}{2}$$

- If we're outside the subgroup-confined regime, we can ensure that $\operatorname{Re} x_k < 1 - \delta$, where δ depends only on G

$$\frac{\partial}{\partial t} x_k = -x_k + x_k \cdot \frac{x_k + \bar{x}_k}{2}$$

- If we're outside the subgroup-confined regime, we can ensure that $\operatorname{Re} x_k < 1 - \delta$, where δ depends only on G

$$\implies \frac{\partial}{\partial t} x_k \leq -\delta x_k$$

$\implies x_k$ decreases exponentially.

- When x_k is small, decay is exponential with constant approximately 1:
 $x_k(s+t) \approx e^{-t} x_k(s)$.

$$\frac{\partial}{\partial t} x_k = -x_k + x_k \cdot \frac{x_k + \bar{x}_k}{2}$$

- If we're outside the subgroup-confined regime, we can ensure that $\operatorname{Re} x_k < 1 - \delta$, where δ depends only on G

$$\implies \frac{\partial}{\partial t} x_k \leq -\delta x_k$$

$\implies x_k$ decreases exponentially.

- When x_k is small, decay is exponential with constant approximately 1: $x_k(s+t) \approx e^{-t} x_k(s)$.
- After about $\frac{1}{2} n \log n$ steps, the x_k are about $O(1/\sqrt{n})$, which corresponds to $N_a = \frac{n}{|G|} + O(\sqrt{n})$.

Stage 3: Bounding mixing time via coupling

- We've analyzed the convergence of the frequency counts; now let's return to mixing.

Stage 3: Bounding mixing time via coupling

- We've analyzed the convergence of the frequency counts; now let's return to mixing.
- Consider a stationary product replacement process σ'_t (i.e. the initial state is drawn uniformly among generating n -tuples).

Stage 3: Bounding mixing time via coupling

- We've analyzed the convergence of the frequency counts; now let's return to mixing.
- Consider a stationary product replacement process σ'_t (i.e. the initial state is drawn uniformly among generating n -tuples).
- Suffices to couple our process σ_t to σ'_t (with probability $1 - \epsilon$) within $\frac{3}{2}n \log n + O(n)$ steps.

Reduction to frequency counts

- Let $N_{a,b}(t)$ denote the number of positions that were originally a in σ_0 and are now b in σ_t .

Reduction to frequency counts

- Let $N_{a,b}(t)$ denote the number of positions that were originally a in σ_0 and are now b in σ_t .
- Similarly, let $N'_{a,b}(t)$ be the number of positions that were a in σ_0 and are now b in σ'_t . (Note that the initial value a is still considered with respect to σ_0 .)

Reduction to frequency counts

- Let $N_{a,b}(t)$ denote the number of positions that were originally a in σ_0 and are now b in σ_t .
- Similarly, let $N'_{a,b}(t)$ be the number of positions that were a in σ_0 and are now b in σ'_t . (Note that the initial value a is still considered with respect to σ_0 .)
- By symmetry, it suffices to find a coupling where $N_{a,b} = N'_{a,b}$ for all $a, b \in G$.

Reduction to frequency counts

- Let $N_{a,b}(t)$ denote the number of positions that were originally a in σ_0 and are now b in σ_t .
- Similarly, let $N'_{a,b}(t)$ be the number of positions that were a in σ_0 and are now b in σ'_t . (Note that the initial value a is still considered with respect to σ_0 .)
- By symmetry, it suffices to find a coupling where $N_{a,b} = N'_{a,b}$ for all $a, b \in G$.
- By similar argument to before, after $\frac{3}{2}n \log n$ steps we'll have something like

$$N_{a,b} \approx \frac{n}{|G|^2} + O(\sqrt{n})$$

for all $a, b \in G$.

Coupling (rough sketch)

- Consider two configurations σ and σ' . Let

$$D = \sum_{a,b \in G} |N_{a,b} - N'_{a,b}|.$$

Coupling (rough sketch)

- Consider two configurations σ and σ' . Let

$$D = \sum_{a,b \in G} |N_{a,b} - N'_{a,b}|.$$

- Assume that we already have $N_{a,b} \approx \frac{n}{|G|^2} + O(\sqrt{n})$, and so $D = O(\sqrt{n})$. Goal is to couple until $D = 0$.

Coupling (rough sketch)

- Consider two configurations σ and σ' . Let

$$D = \sum_{a,b \in G} |N_{a,b} - N'_{a,b}|.$$

- Assume that we already have $N_{a,b} \approx \frac{n}{|G|^2} + O(\sqrt{n})$, and so $D = O(\sqrt{n})$. Goal is to couple until $D = 0$.
- Rough idea: suppose σ has more occurrences of g , while σ' has more occurrences of g' .

Coupling (rough sketch)

- Consider two configurations σ and σ' . Let

$$D = \sum_{a,b \in G} |N_{a,b} - N'_{a,b}|.$$

- Assume that we already have $N_{a,b} \approx \frac{n}{|G|^2} + O(\sqrt{n})$, and so $D = O(\sqrt{n})$. Goal is to couple until $D = 0$.
- Rough idea: suppose σ has more occurrences of g , while σ' has more occurrences of g' .
 - For product replacement, suppose we choose $(a, b) \rightarrow (a + b, a)$ for updating σ and $(a', b') \rightarrow (a' + b', b')$ for updating σ' .

Coupling (rough sketch)

- Consider two configurations σ and σ' . Let

$$D = \sum_{a,b \in G} |N_{a,b} - N'_{a,b}|.$$

- Assume that we already have $N_{a,b} \approx \frac{n}{|G|^2} + O(\sqrt{n})$, and so $D = O(\sqrt{n})$. Goal is to couple until $D = 0$.
- Rough idea: suppose σ has more occurrences of g , while σ' has more occurrences of g' .
 - For product replacement, suppose we choose $(a, b) \rightarrow (a + b, a)$ for updating σ and $(a', b') \rightarrow (a' + b', b')$ for updating σ' .
 - If it happens that $a = g$ and $a' = g'$, then we try to couple b and b' so that we always have $b - b' = g' - g$.

Coupling (rough sketch)

- Consider two configurations σ and σ' . Let

$$D = \sum_{a,b \in G} |N_{a,b} - N'_{a,b}|.$$

- Assume that we already have $N_{a,b} \approx \frac{n}{|G|^2} + O(\sqrt{n})$, and so $D = O(\sqrt{n})$. Goal is to couple until $D = 0$.
- Rough idea: suppose σ has more occurrences of g , while σ' has more occurrences of g' .
 - For product replacement, suppose we choose $(a, b) \rightarrow (a + b, a)$ for updating σ and $(a', b') \rightarrow (a' + b', b')$ for updating σ' .
 - If it happens that $a = g$ and $a' = g'$, then we try to couple b and b' so that we always have $b - b' = g' - g$.
 - This can be done (approximately) because b and b' are nearly uniformly distributed over G .

Coupling (rough sketch)

- Consider two configurations σ and σ' . Let

$$D = \sum_{a,b \in G} |N_{a,b} - N'_{a,b}|.$$

- Assume that we already have $N_{a,b} \approx \frac{n}{|G|^2} + O(\sqrt{n})$, and so $D = O(\sqrt{n})$. Goal is to couple until $D = 0$.
- Rough idea: suppose σ has more occurrences of g , while σ' has more occurrences of g' .
 - For product replacement, suppose we choose $(a, b) \rightarrow (a + b, a)$ for updating σ and $(a', b') \rightarrow (a' + b', b')$ for updating σ' .
 - If it happens that $a = g$ and $a' = g'$, then we try to couple b and b' so that we always have $b - b' = g' - g$.
 - This can be done (approximately) because b and b' are nearly uniformly distributed over G .
- If done carefully, can ensure that D has $\Omega(1)$ probability of either increasing or decreasing, and in expectation, it decreases. Simple random walk started at $O(\sqrt{n})$ is likely to hit 0 within $O(n)$ steps.

- For $G = \mathbb{Z}/q$, what is optimal dependence on q ?

- For $G = \mathbb{Z}/q$, what is optimal dependence on q ?
- What is the mixing time of graph-restricted product replacement for other graphs (e.g. trees)?

- For $G = \mathbb{Z}/q$, what is optimal dependence on q ?
- What is the mixing time of graph-restricted product replacement for other graphs (e.g. trees)?
- Can mixing time of graph-restricted product replacement be tightly bounded in terms of the structure of the graph?

- For $G = \mathbb{Z}/q$, what is optimal dependence on q ?
- What is the mixing time of graph-restricted product replacement for other graphs (e.g. trees)?
- Can mixing time of graph-restricted product replacement be tightly bounded in terms of the structure of the graph?
- Thank you!


Theorem (Peres-Tanaka-Z.)

Let G be a fixed finite group, and consider the product replacement walk $(\sigma_t)_{t \geq 0}$. Then, for any fixed $\epsilon > 0$,

$$\left(\frac{3}{2} - o(1)\right) n \log n \leq t_{\text{mix}}(1 - \epsilon) \leq t_{\text{mix}}(\epsilon) \leq \left(\frac{3}{2} + o(1)\right) n \log n$$

as $n \rightarrow \infty$.

Product replacement walk



- Let G be a finite group, and consider an n -tuple of group elements (call this a **configuration**). For example, $G = \mathbb{Z}/11$ and $n = 6$.
- We consider the **product replacement walk**: randomly pick two of the elements a and b , and replace a with $a + b$ or $a - b$.

Product replacement walk



- Let G be a finite group, and consider an n -tuple of group elements (call this a **configuration**). For example, $G = \mathbb{Z}/11$ and $n = 6$.
- We consider the **product replacement walk**: randomly pick two of the elements a and b , and replace a with $a + b$ or $a - b$.