Plurality Consensus in the Gossip Model

Emanuele Natale[†]

joint work with
L. Becchetti[†], A. Clementi[*],
F. Pasquale[*] and R. Silvestri[†]

[*] UNIVERSITA' DEGLI STUDI DI ROMA TorVergata

[†] SAPIENZA UNIVERSITÀ DI ROMA

Efficient Algorithms Group, CS Department, Salzburg University,
February 20th, 2015

UNIVERSITÄT SALZBURG

The Plurality Consensus Problem

We have a set of nodes each having one color out of $\{1, \ldots, k\}$.

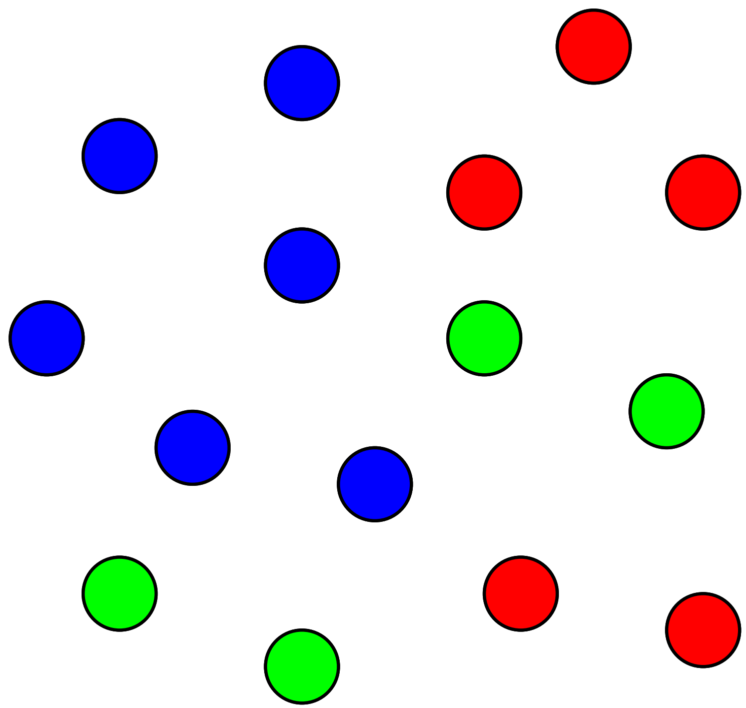

The Plurality Consensus Problem

There is a plurality of nodes having the same color.

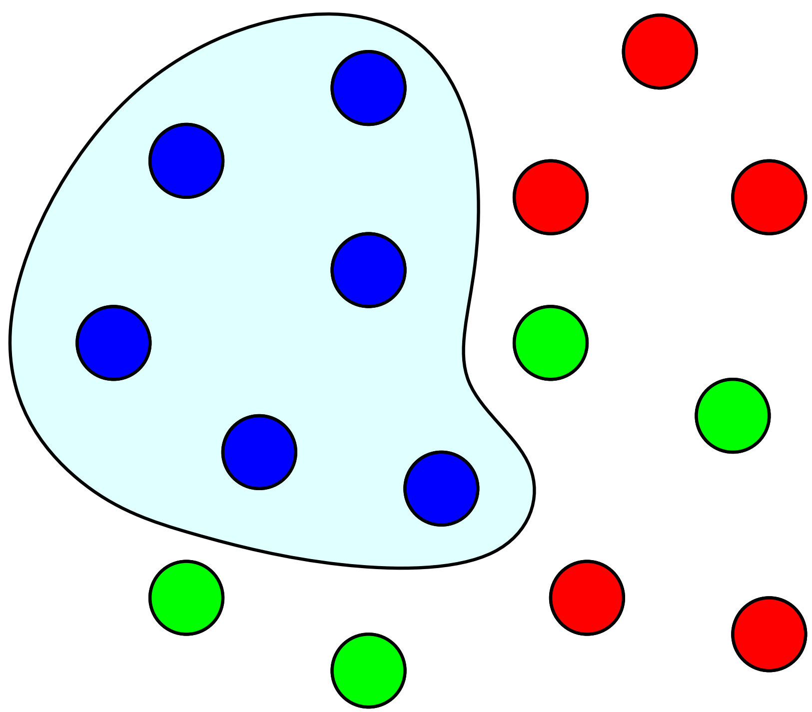

The Plurality Consensus Problem

We want to reach consensus on the plurality color.

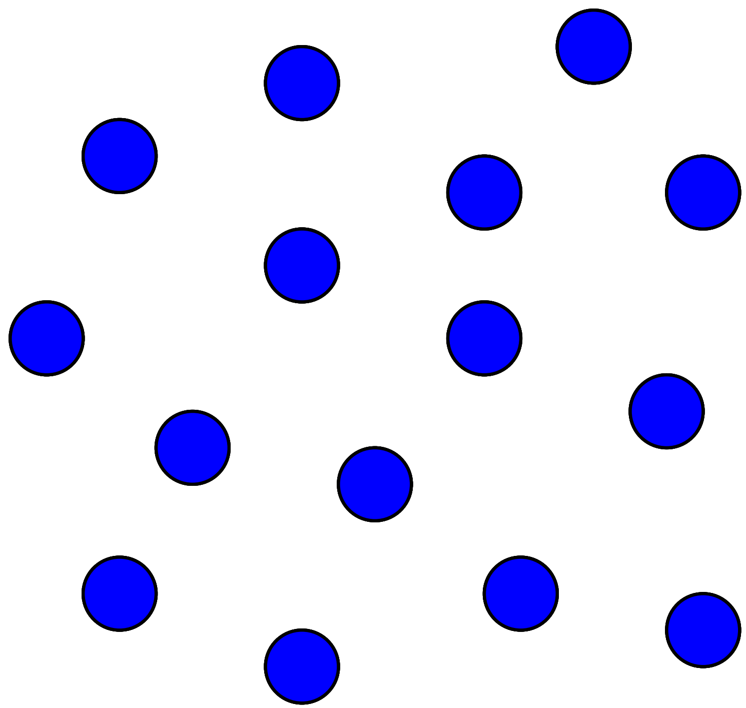

Motivations and Applications

- **Computer science**: distributed databases and sensor networks (Bénézit et al. ’09).

Motivations and Applications

- **Computer science**: distributed databases and sensor networks (Bénézit et al. '09).

- **Social networks**: opinion dynamics (Mossel et al. '14).

Motivations and Applications

- **Computer science**: distributed databases and sensor networks (Bénézit et al. '09).

- **Social networks**: opinion dynamics (Mossel et al. '14).

- **Biology**: cell cycle (Cardelli et al. '12).

Motivations and Applications

- **Computer science**: distributed databases and sensor networks (Bénézit et al. '09).
- **Social networks**: opinion dynamics (Mossel et al. '14).
- **Biology**: cell cycle (Cardelli et al. '12).
- **Chemestry**: chemical reaction networks/population protocols (Angluin et al. '07).

Pre-CS History

Voter Model ('70). Each node with a Poisson clock. When rings, takes the opinion of a random neighbor.

Pre-CS History

Voter Model ('70). Each node with a Poisson clock. When rings, takes the opinion of a random neighbor.

→

Continuos time (sequential/asynchronous) process. Well studied in statistical physics (constant number of particle types).

Pre-CS History

Voter Model ('70). Each node with a Poisson clock. When rings, takes the opinion of a random neighbor.

→ Continuos time (sequential/asynchronous) process. Well studied in statistical physics (constant number of particle types).

Probabilistic Polling (Peleg '99). Time divided in discrete rounds. All nodes *simultaneously* take the opinion of a random neighbor.

Pre-CS History

Voter Model ('70). Each node with a Poisson clock. When rings, takes the opinion of a random neighbor.

→ Continuos time (sequential/asynchronous) process. Well studied in statistical physics (constant number of particle types).

Probabilistic Polling (Peleg '99). Time divided in discrete rounds. All nodes *simultaneously* take the opinion of a random neighbor.

→ Discrete time (parallel/synchronous) process. Initiated the study of Plurality Consensus in Computer Science.

Asynchronous vs Synchronous

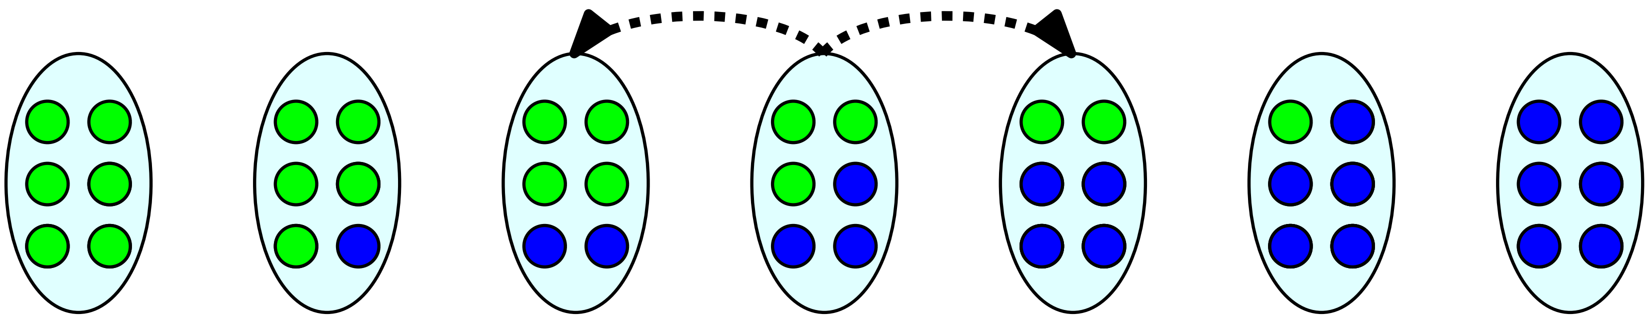

Asynchronous Case

Asynchronous vs Synchronous

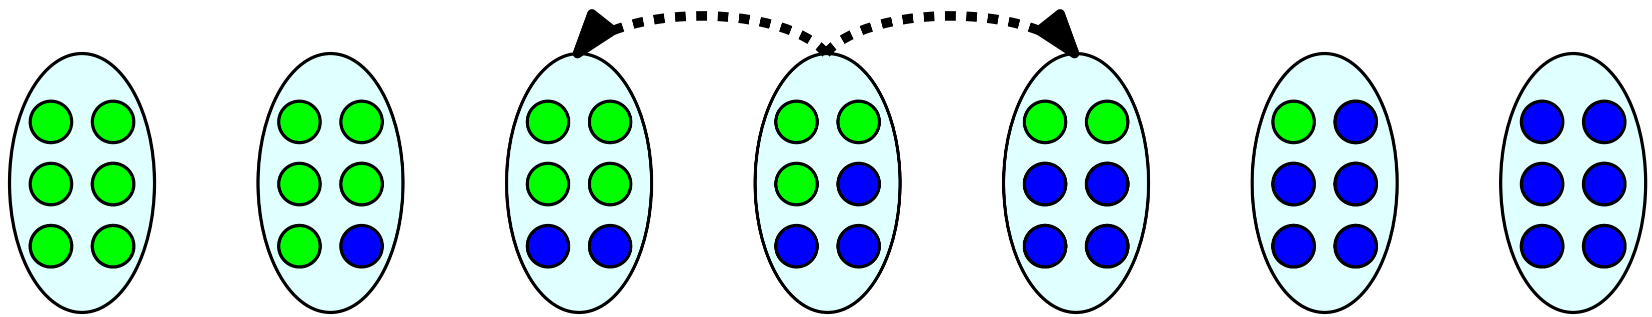

Asynchronous Case

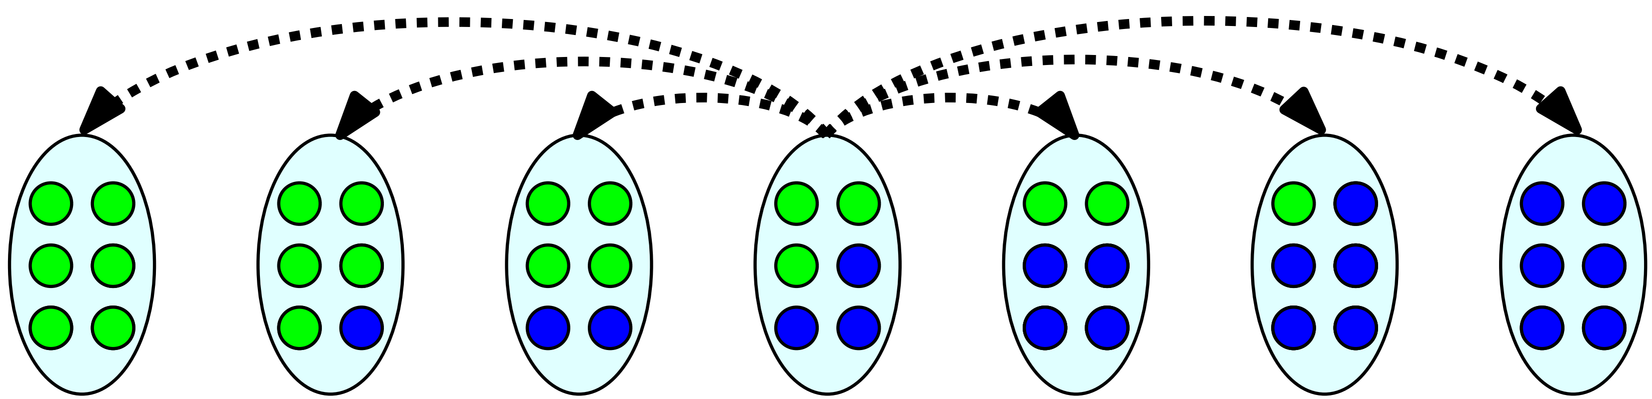

Synchronous Case

Our Setting

- **Initial bias**: the plurality is at least $(1+\epsilon)$ times any other color.

Our Setting

- **Initial bias**: the plurality is at least $(1 + \epsilon)$ times any other color.

- **Topology**: complete graph (and regular expanders).

Our Setting

- **Initial bias**: the plurality is at least $(1+\epsilon)$ times any other color.

- **Topology**: complete graph (and regular expanders).

- **Communication model**: $\mathcal{GOSSIP}$ model [Censor-Hillel et al., STOC '12]. Each node in one round can exchange messages with only one neighbor.

Our Setting

- **Initial bias**: the plurality is at least $(1+\epsilon)$ times any other color.

- **Topology**: complete graph (and regular expanders).

- **Communication model**: $\mathcal{GOSSIP}$ model [Censor-Hillel et al., STOC '12]. Each node in one round can exchange messages with only one neighbor.

- **Local memory and message size**: $O(\log n)$.

Relationships to Other Communication Models

$\mathcal{GOSSIP}$ model with neighbors chosen randomly: Telephone Call, Push&Pull, Uniform Gossip…

Relationships to Other Communication Models

$\mathcal{GOSSIP}$ model with neighbors chosen randomly: Telephone Call, Push&Pull, Uniform Gossip…

$\mathcal{LOCAL}$ model [Peleg, SIAM '00]: each node in one round can exchange messages with all its neighbors.

Relationships to Other Communication Models

$\mathcal{GOSSIP}$ model with neighbors chosen randomly: Telephone Call, Push&Pull, Uniform Gossip…

$\mathcal{LOCAL}$ model [Peleg, SIAM '00]: each node in one round can exchange messages with all its neighbors.
…on the complete graph, plurality consensus can be achieved in one round.

Relationships to Other Communication Models

$\mathcal{GOSSIP}$ model with neighbors chosen randomly: Telephone Call, Push&Pull, Uniform Gossip...

$\mathcal{LOCAL}$ model [Peleg, SIAM '00]: each node in one round can exchange messages with all its neighbors.
...on the complete graph, plurality consensus can be achieved in one round.

Censor-Hillel et al. (STOC '12):
Every task that can be solved in the $\mathcal{LOCAL}$ model in T rounds, can be solved in $O(T + \text{polylog} n)$ rounds in the $\mathcal{GOSSIP}$ model.
But...

Relationships to Other Communication Models

$\mathcal{GOSSIP}$ model with neighbors chosen randomly: Telephone Call, Push&Pull, Uniform Gossip…

$\mathcal{LOCAL}$ model [Peleg, SIAM '00]: each node in one round can exchange messages with all its neighbors.
…on the complete graph, plurality consensus can be achieved in one round.

Censor-Hillel et al. (STOC '12):
Every task that can be solved in the $\mathcal{LOCAL}$ model in T rounds, can be solved in $O(T + \text{polylog} n)$ rounds in the $\mathcal{GOSSIP}$ model.
But…using the preceding theorem, message size grows dramatically!

(Some) Related Works

	Mem. & mess. size	**# of colors**	**Time efficiency**	**Comm. Model**
Kempe et al. FOCS '03	$O(k \log n)$	any	$O(\log n)$	$\mathcal{GOSSIP}$
Angluin et al. DISC '07 Perron et al. INFOCOM '09	$\Theta(1)$	2	$O(\log n)$	Sequential
Doerr et al. SPAA '11	$\Theta(1)$	2	$O(\log n)$	$\mathcal{GOSSIP}$
Babaee et al. Comp. J. '12 Jung et al. ISIT '12	$O(\log k)$	Constant	$O(\log n)$	Sequential

(Some) Related Works

	Mem. & mess. size	**# of colors**	**Time efficiency**	**Comm. Model**
Kempe et al. FOCS '03	$O(k \log n)$ ✗	any ✓	$O(\log n)$ ✓	$\mathcal{GOSSIP}$ ✓
Angluin et al. DISC '07 Perron et al. INFOCOM '09	$\Theta(1)$	2	$O(\log n)$	Sequential
Doerr et al. SPAA '11	$\Theta(1)$	2	$O(\log n)$	$\mathcal{GOSSIP}$
Babaee et al. Comp. J. '12 Jung et al. ISIT '12	$O(\log k)$	Constant	$O(\log n)$	Sequential

(Some) Related Works

	Mem. & mess. size	**# of colors**	**Time efficiency**	**Comm. Model**
Kempe et al. FOCS '03	$O(k \log n)$	any	$O(\log n)$	$\mathcal{GOSSIP}$
Angluin et al. DISC '07 Perron et al. INFOCOM '09	$\Theta(1)$	2	$O(\log n)$	Sequential
Doerr et al. SPAA '11	$\Theta(1)$	2	$O(\log n)$	$\mathcal{GOSSIP}$
Babaee et al. Comp. J. '12 Jung et al. ISIT '12	$O(\log k)$	Constant	$O(\log n)$	Sequential

(Some) Related Works

	Mem. & mess. size	**# of colors**	**Time efficiency**	**Comm. Model**
Kempe et al. FOCS '03	$O(k \log n)$	any	$O(\log n)$	$\mathcal{GOSSIP}$
Angluin et al. DISC '07 Perron et al. INFOCOM '09	$\Theta(1)$	2	$O(\log n)$	Sequential
Doerr et al. SPAA '11	$\Theta(1)$	2	$O(\log n)$	$\mathcal{GOSSIP}$
Babaee et al. Comp. J. '12 Jung et al. ISIT '12	$O(\log k)$	Constant	$O(\log n)$	Sequential

(Some) Related Works

	Mem. & mess. size	**# of colors**	**Time efficiency**	**Comm. Model**
Kempe et al. FOCS '03	$O(k \log n)$ ✗	any ✓	$O(\log n)$ ✓	$\mathcal{GOSSIP}$ ✓
Angluin et al. DISC '07 Perron et al. INFOCOM '09	$\Theta(1)$ ✓	2 ✗	$O(\log n)$ ✓	Sequential ✗
Doerr et al. SPAA '11	$\Theta(1)$ ✓	2 ✗	$O(\log n)$ ✓	$\mathcal{GOSSIP}$ ✓
Babaee et al. Comp. J. '12 Jung et al. ISIT '12	$O(\log k)$ ✓	Constant ✗	$O(\log n)$ ✓	Sequential ✗

Characterizing the Initial Bias

$c_i^{(t)} := |\{i\text{-colored nodes}\}|$ color 1 is the plurality

Characterizing the Initial Bias

$c_i^{(t)} := |\{i\text{-colored nodes}\}|$

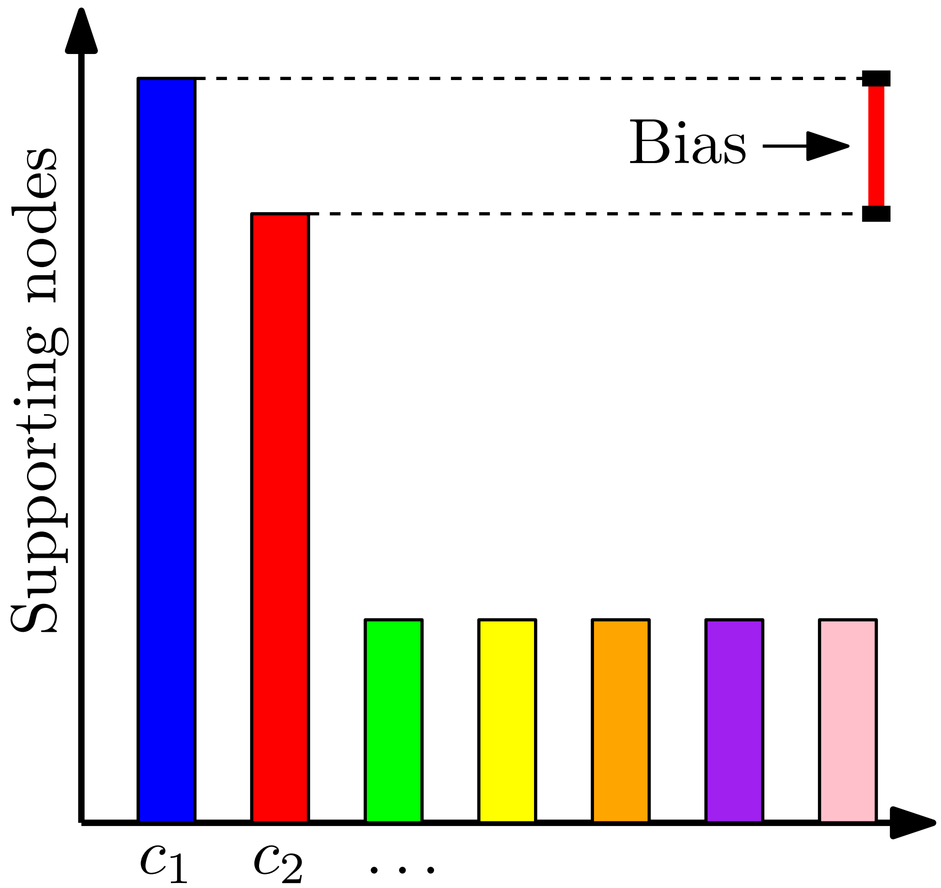

color 1 is the plurality

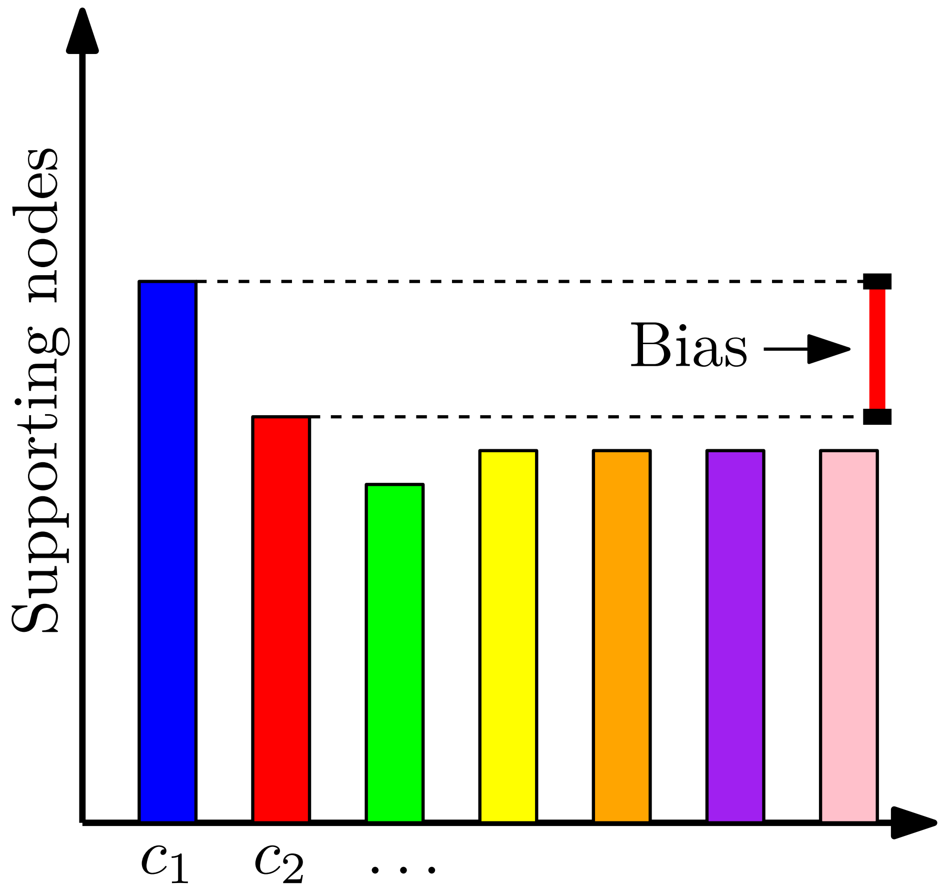

The 3-Majority Dynamics

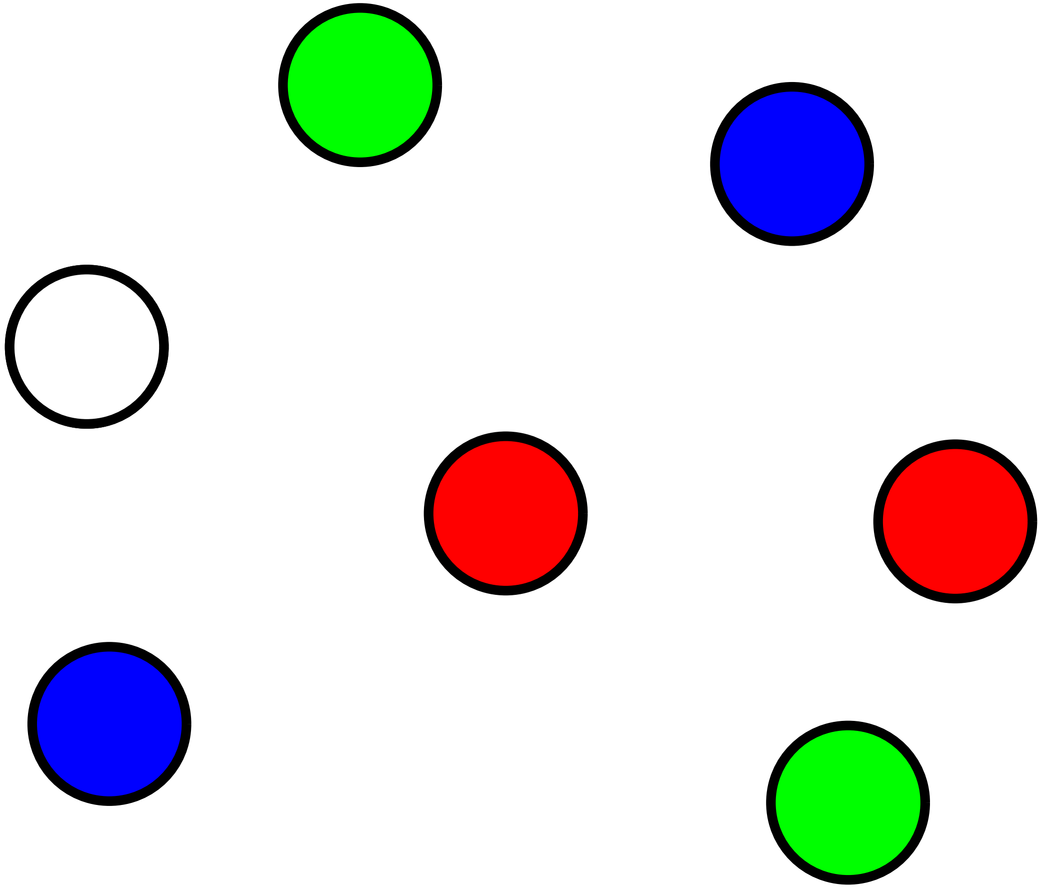

The 3-Majority Dynamics

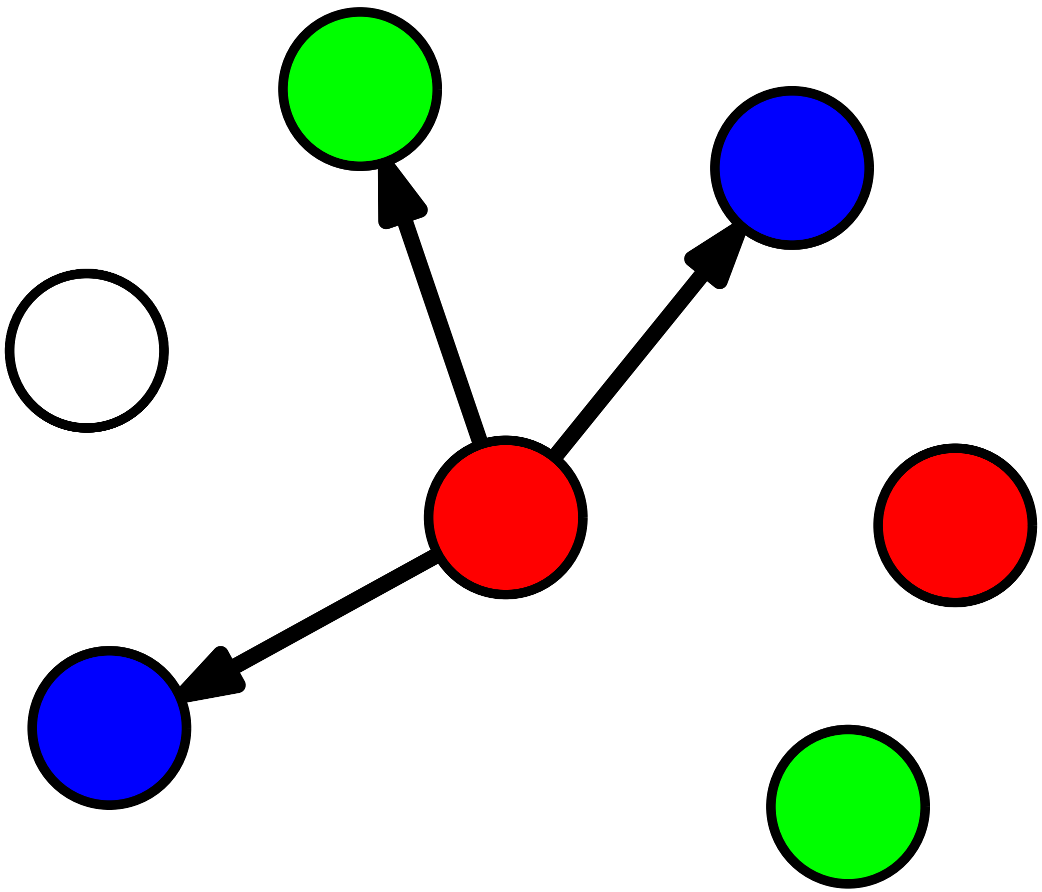

Each node observes the color of three other nodes chosen u.a.r....

The 3-Majority Dynamics

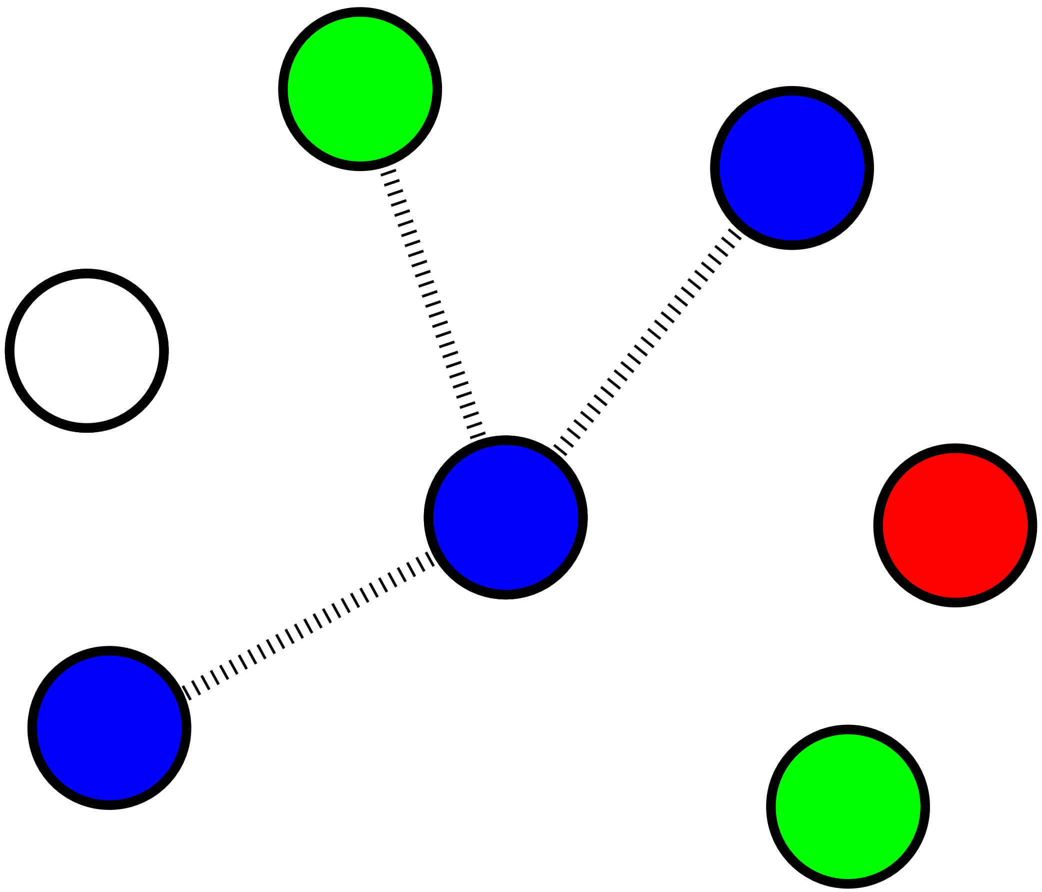

...and changes its color according to the majority of these three (breaking ties u.a.r.).

The 3-Majority Dynamics

	Mem. & mess. size	**# of colors**	**Time efficiency**	**Comm. Model**
Us+Trevisan SPAA '14	$O(\log k)$	$n^{\Theta(1)}$	$O(k \log n)$	$\mathcal{GOSSIP}$

The 3-Majority Dynamics

	Mem. & mess. size	**# of colors**	**Time efficiency**	**Comm. Model**
Us+Trevisan SPAA '14	$O(\log k)$	$n^{\Theta(1)}$	$O(k \log n)$	$\mathcal{GOSSIP}$

The 3-Majority Dynamics

	Mem. & mess. size	**# of colors**	**Time efficiency**	**Comm. Model**
Us+Trevisan SPAA '14	$O(\log k)$	$n^{\Theta(1)}$	$O(k \log n)$	$\mathcal{GOSSIP}$

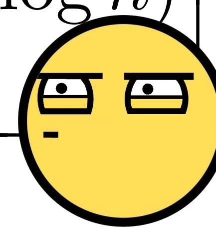

The Monochromatic Distance

$$\mathrm{md}(\mathbf{c}^{(\mathbf{0})}) := \sum_{i=1}^{k} \left(\frac{c_i^{(0)}}{c_1^{(0)}} \right)^2 = 1 + \mathcal{D}\Bigg($$

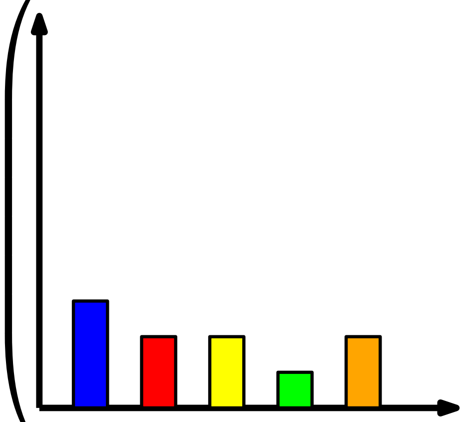

,

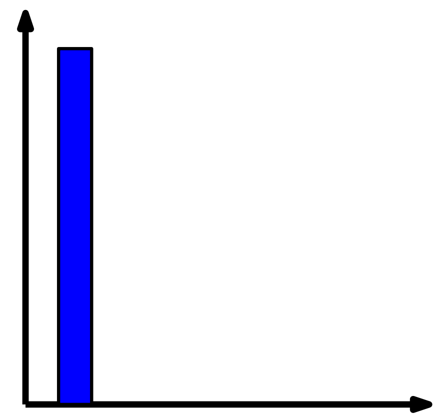

$$\Bigg)$$

The Monochromatic Distance

$$\mathrm{md}(\mathbf{c}^{(\mathbf{0})}) := \sum_{i=1}^{k} \left(\frac{c_i^{(0)}}{c_1^{(0)}} \right)^2 = 1 + \mathcal{D}\Bigg($$ 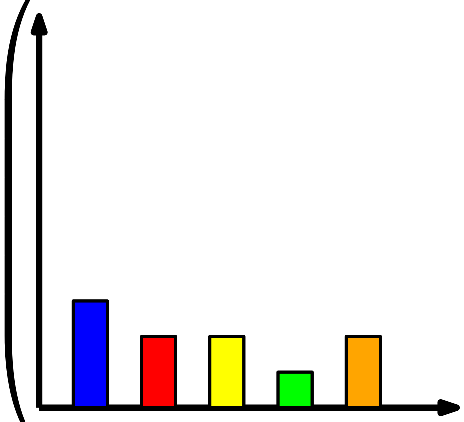 $,$ 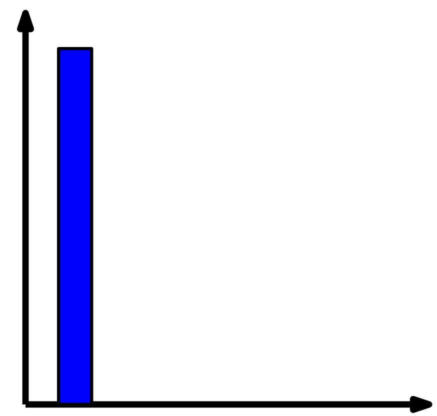 $\Bigg)$

$1 \leq \mathrm{md}\Bigg($ 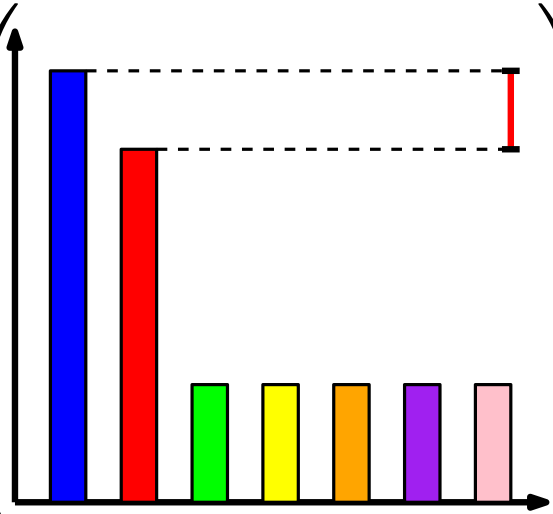 $\Bigg) \ll \mathrm{md}\Bigg($ 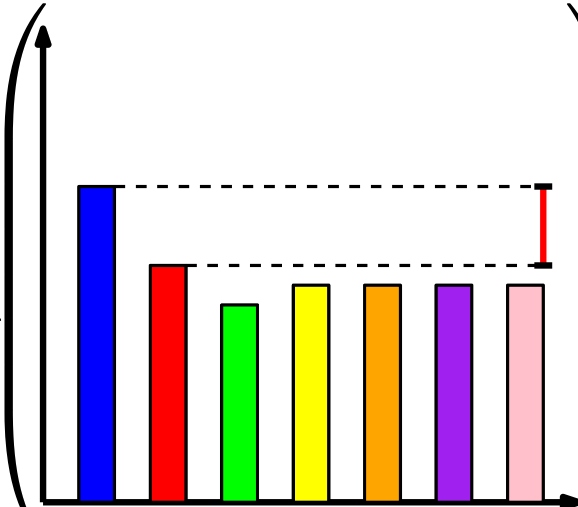 $\Bigg) \leq k$

Our Results

First analysis for $k = \omega(1)$ of the Undecided-State Dynamics:

(Angluin et al., Perron et al.,
Babaee et al., Jung et al.)

Upper Bound

If $k = O\left((n/\log n)^{1/3}\right)$ and $c_1 \geq (1+\epsilon) \cdot c_2$ with $\epsilon > 0$, then w.h.p. the Undecided-State Dynamics reaches plurality consensus in $O\left(\mathrm{md}(\mathbf{c}^{(0)}) \cdot \log n\right)$ rounds.

Our Results

First analysis for $k = \omega(1)$ of the Undecided-State Dynamics:

(Angluin et al., Perron et al.,
Babaee et al., Jung et al.)

Upper Bound

If $k = O\left((n/\log n)^{1/3}\right)$ and $c_1 \geq (1+\epsilon) \cdot c_2$ with $\epsilon > 0$, then w.h.p. the Undecided-State Dynamics reaches plurality consensus in $O\left(\mathrm{md}(\mathbf{c}^{(0)}) \cdot \log n\right)$ rounds.

Lower Bound

If $k = O\left((n/\log n)^{1/6}\right)$ then w.h.p. the Undecided-State Dynamics converges after at least $\Omega(\mathrm{md}(\mathbf{c}^{(0)}))$ rounds.

The Undecided-State Dynamics

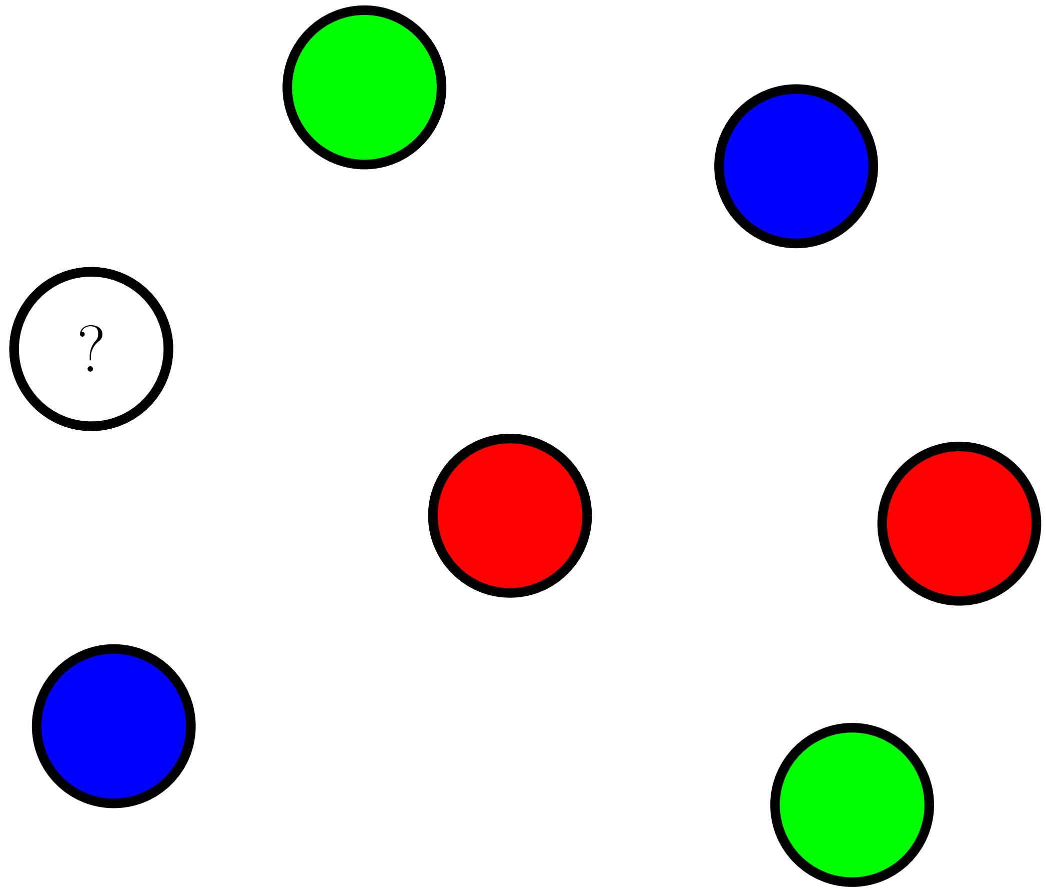

Some nodes can be "undecided".

The Undecided-State Dynamics

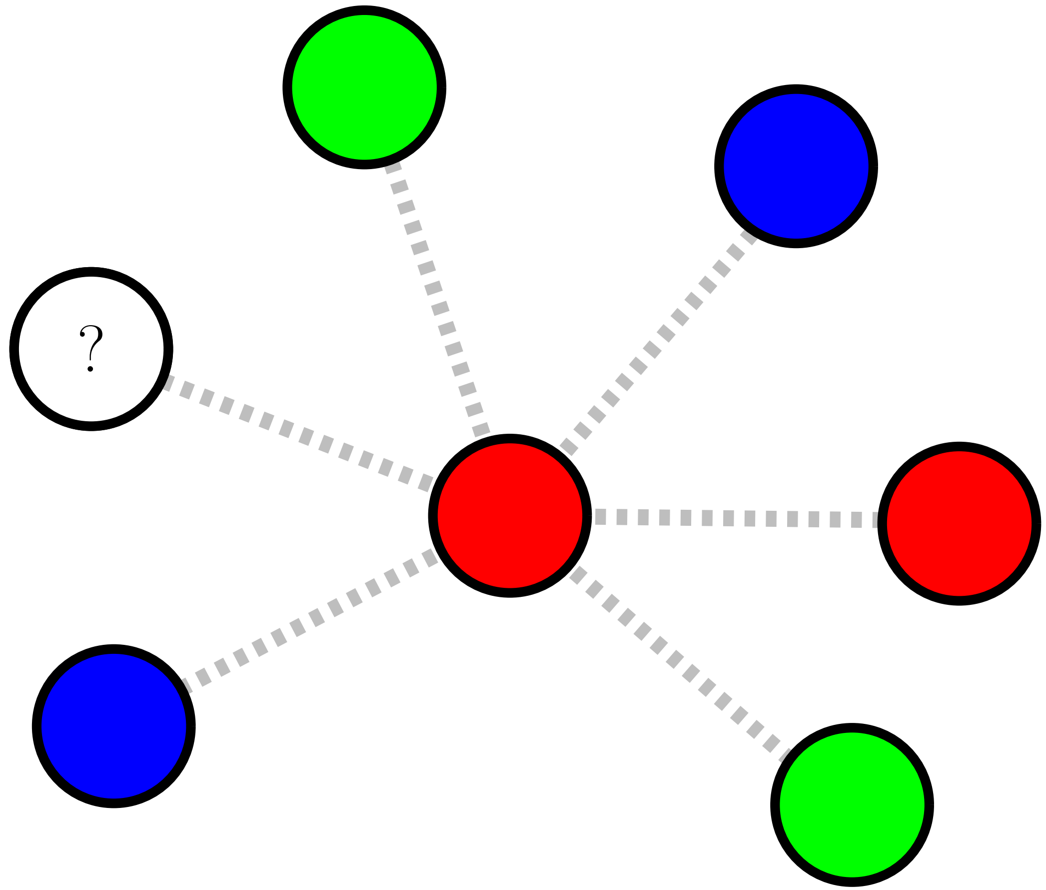

At the beginning of each round, each node observes a neighbor picked uniformly at random.

The Undecided-State Dynamics

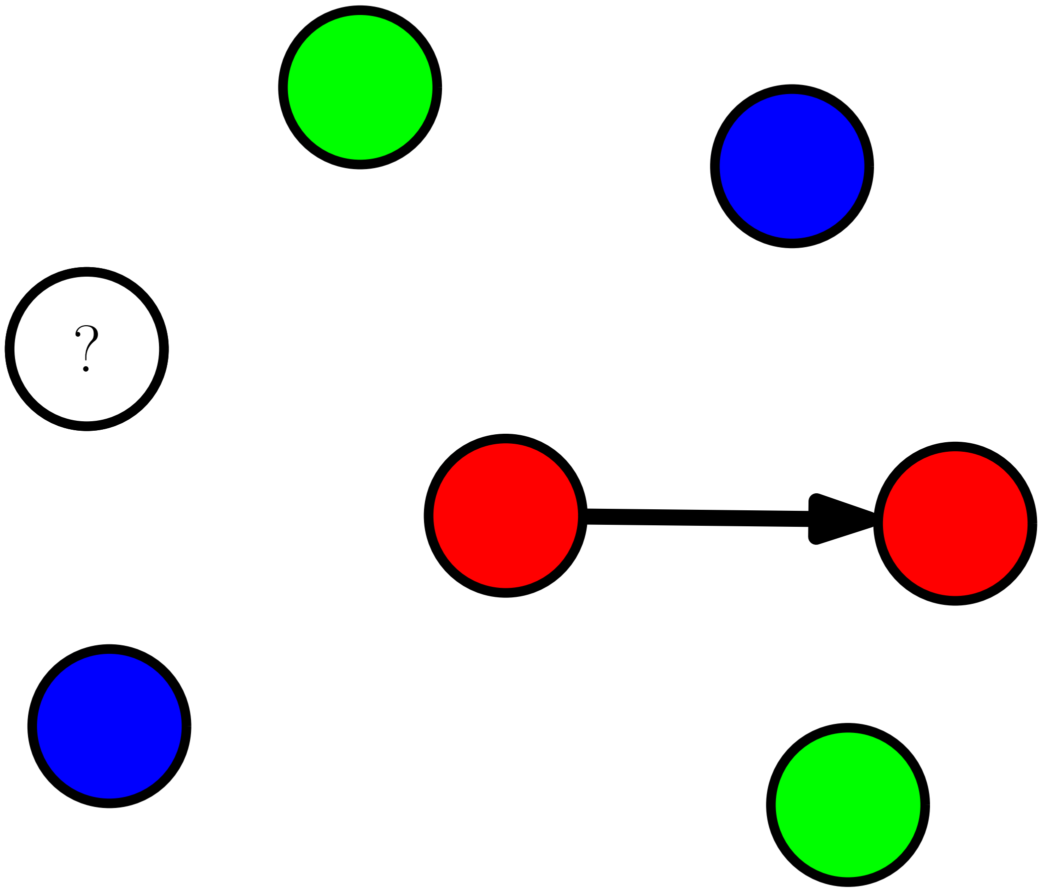

If the observed node shares the same color...

The Undecided-State Dynamics

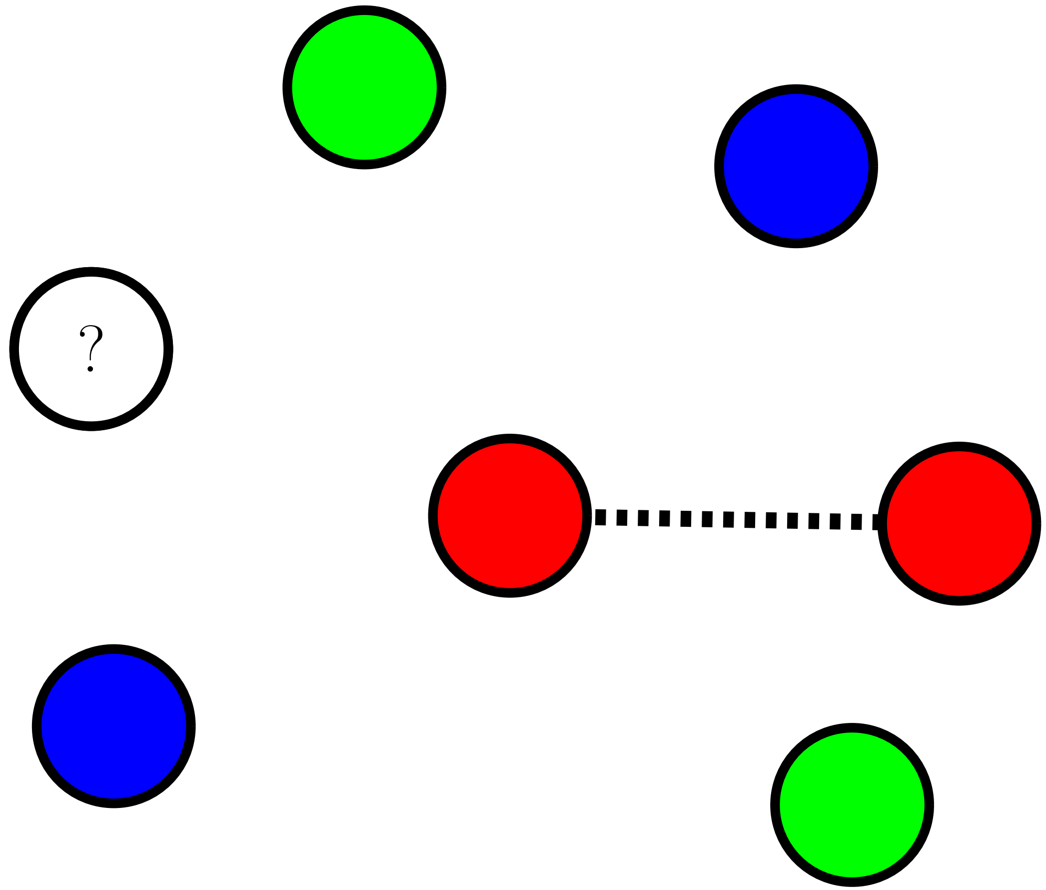

. . . nothing happens;

The Undecided-State Dynamics

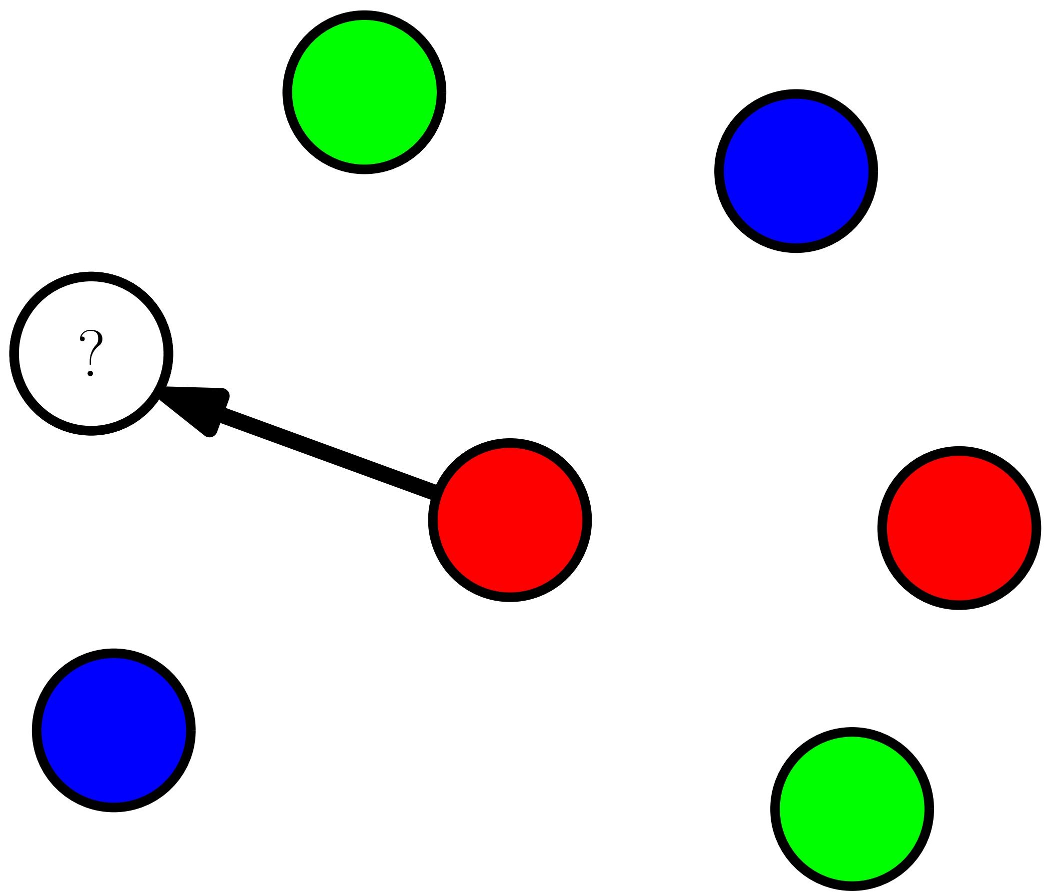

if the node observes an undecided one...

The Undecided-State Dynamics

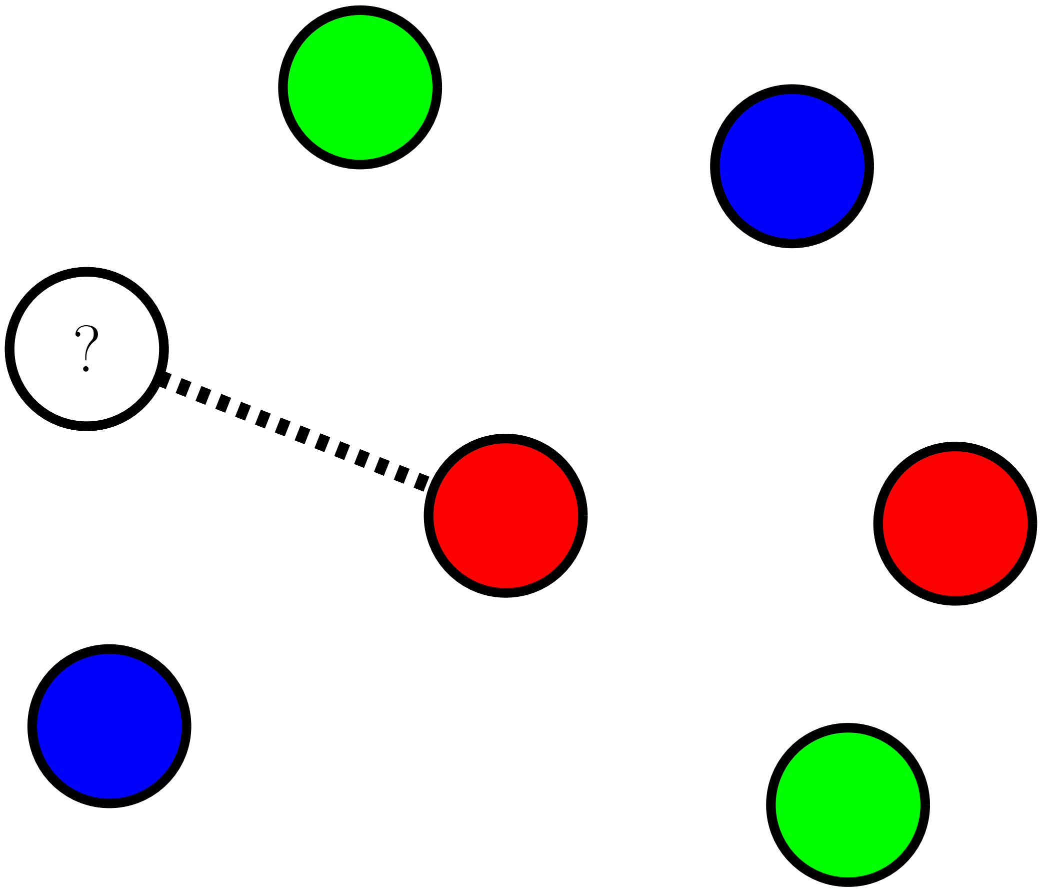

...nothing happens too;

The Undecided-State Dynamics

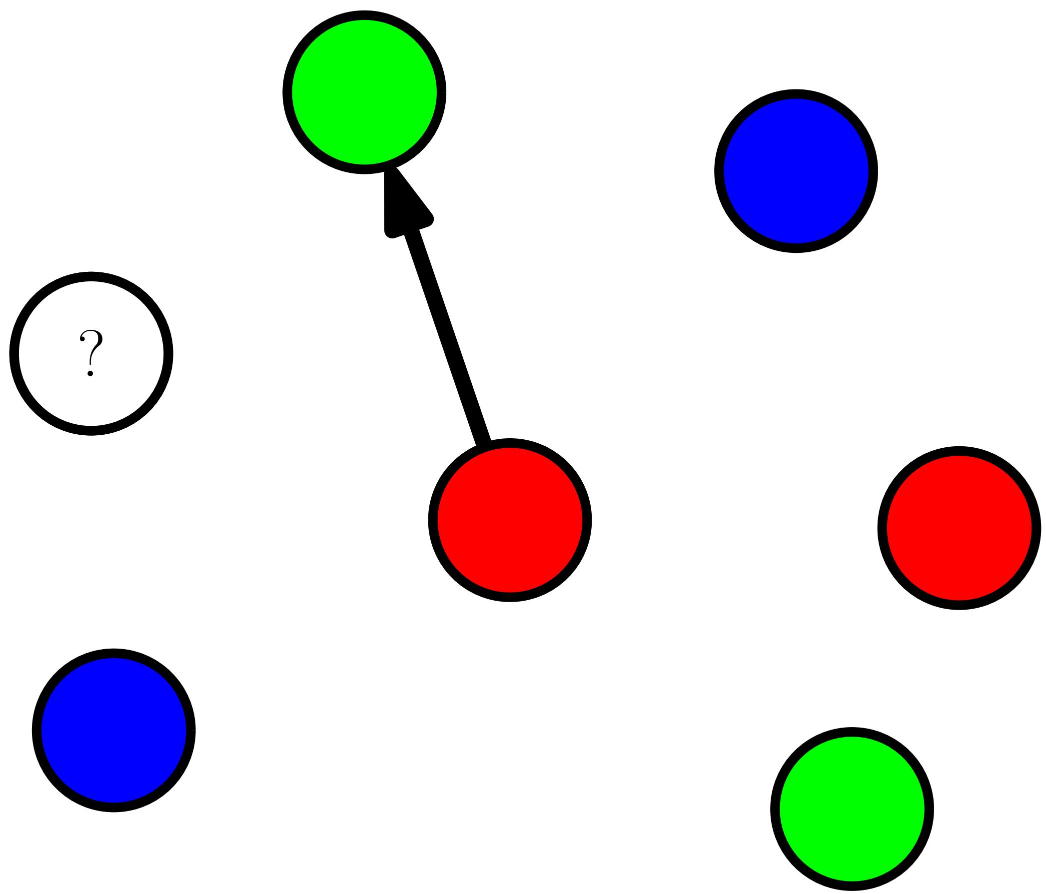

but, if the observed node has a different color. . .

The Undecided-State Dynamics

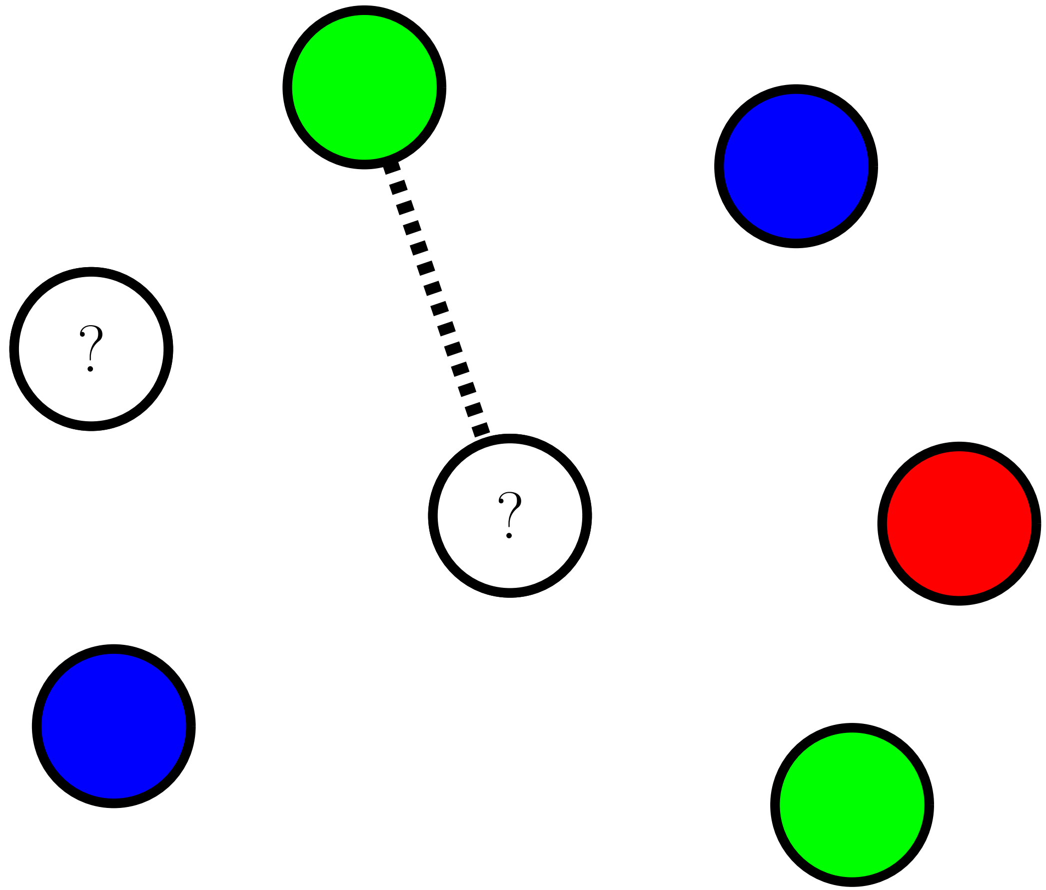

. . . then the node becomes undecided.

The Undecided-State Dynamics

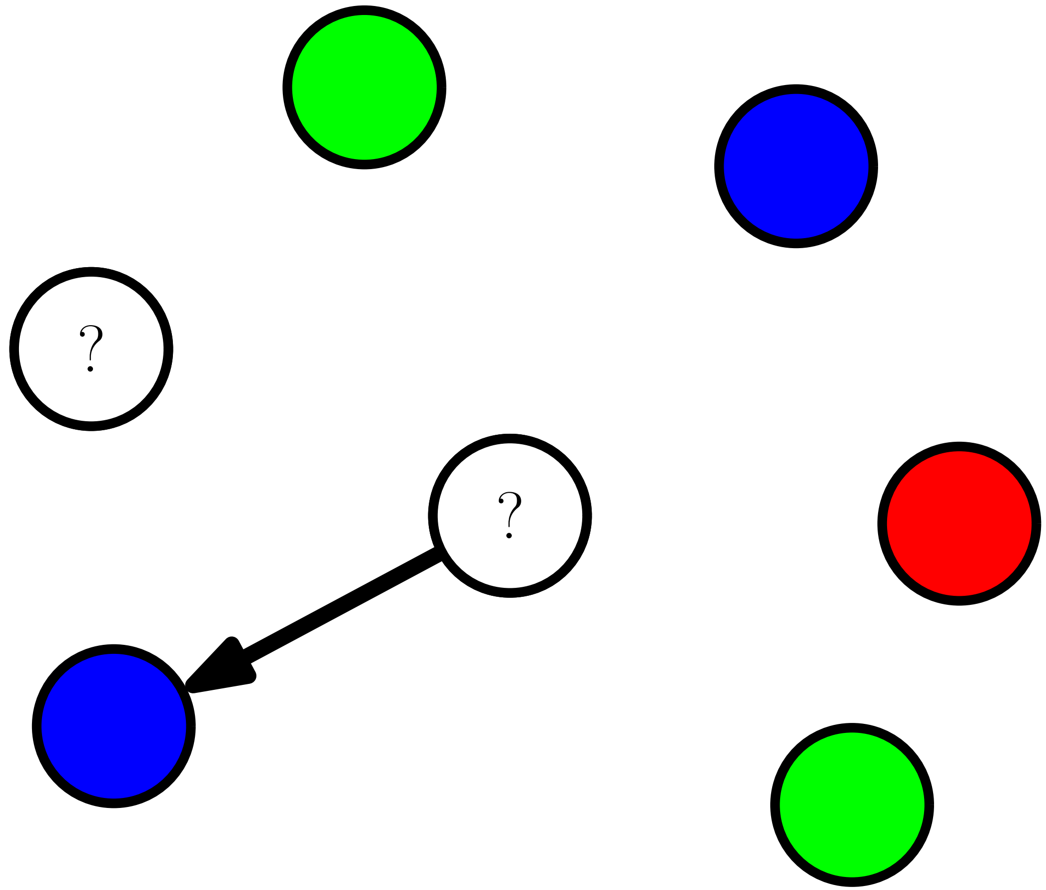

Once undecided...

The Undecided-State Dynamics

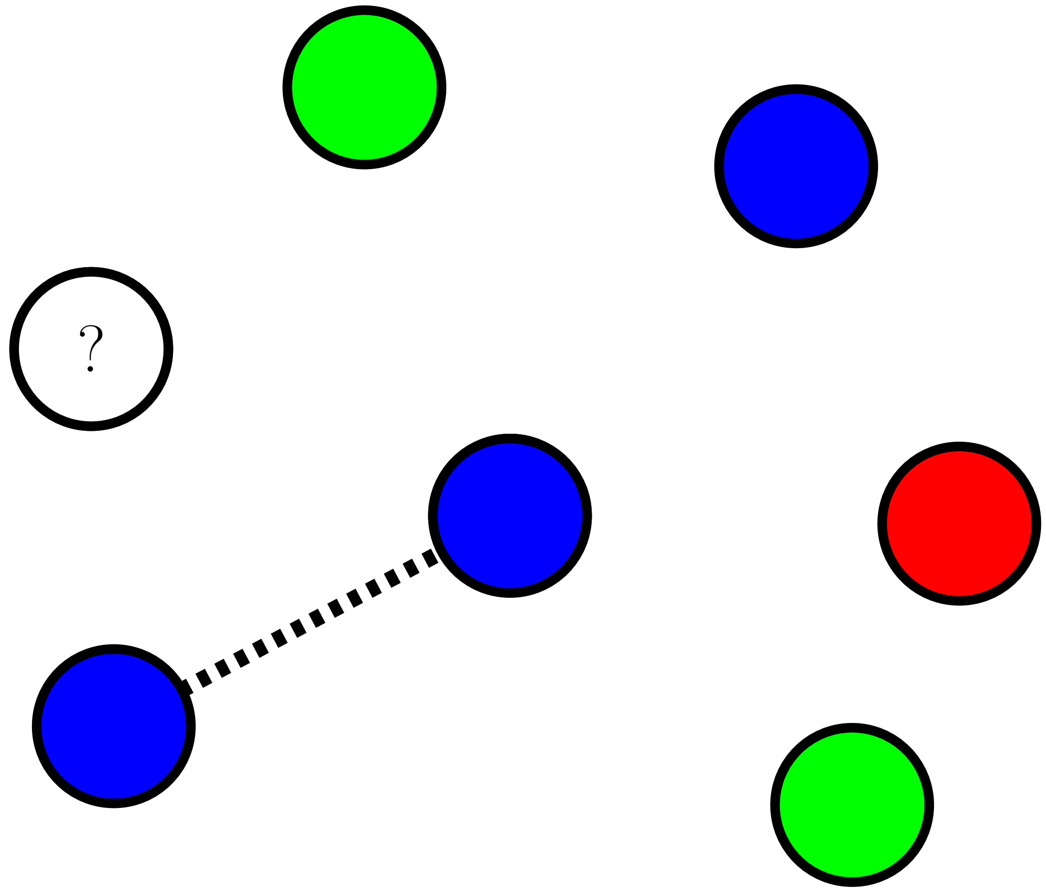

. . . the node copies the first color it sees.

Overview of the Process

$$c_i^{(t)} := |\{i\text{-colored nodes}\}|, \qquad \text{color 1 is the plurality,}$$

$$q^{(t)} := |\{\text{undecided nodes}\}|, \qquad \mathbf{c}^{(t)} := \left(c_1^{(t)}, \ldots, c_k^{(t)}, q^{(t)}\right).$$

$$\mathbf{E}\left[c_i^{(t+1)} \,\middle|\, \mathbf{c}^{(t)}\right] = c_i^{(t)} \cdot \underbrace{\frac{c_i^{(t)} + 2q^{(t)}}{n}}_{\text{Growth factor}}$$

Overview of the Process

$$\mathbf{E}\left[c_i^{(t+1)} \,\middle|\, \mathbf{c}^{(t)}\right] = $$
$$= c_i^{(t)} \cdot \underbrace{\frac{c_i^{(t)} + 2q^{(t)}}{n}}_{\text{Growth factor}}$$

Simulation of the growth factor:

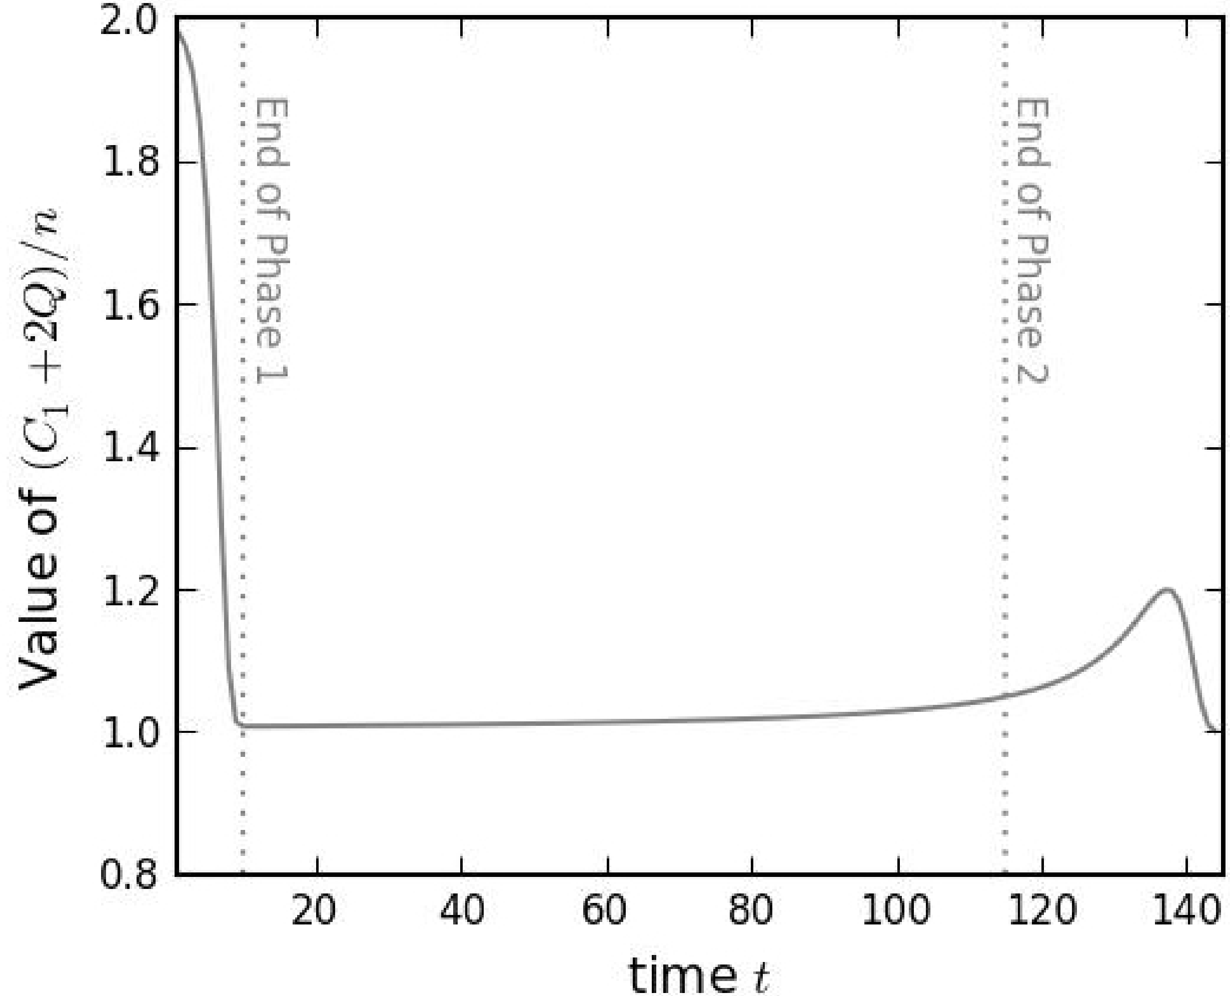

Remarks

W.h.p.

- Plurality does not change.
- Growth factor of plurality is > 1.

Overview of the Process

$$\mathbf{E}\left[c_i^{(t+1)} \,\middle|\, \mathbf{c}^{(t)}\right] = c_i^{(t)} \cdot \underbrace{\frac{c_i^{(t)} + 2q^{(t)}}{n}}_{\text{Growth factor}}$$

Simulation of the growth factor:

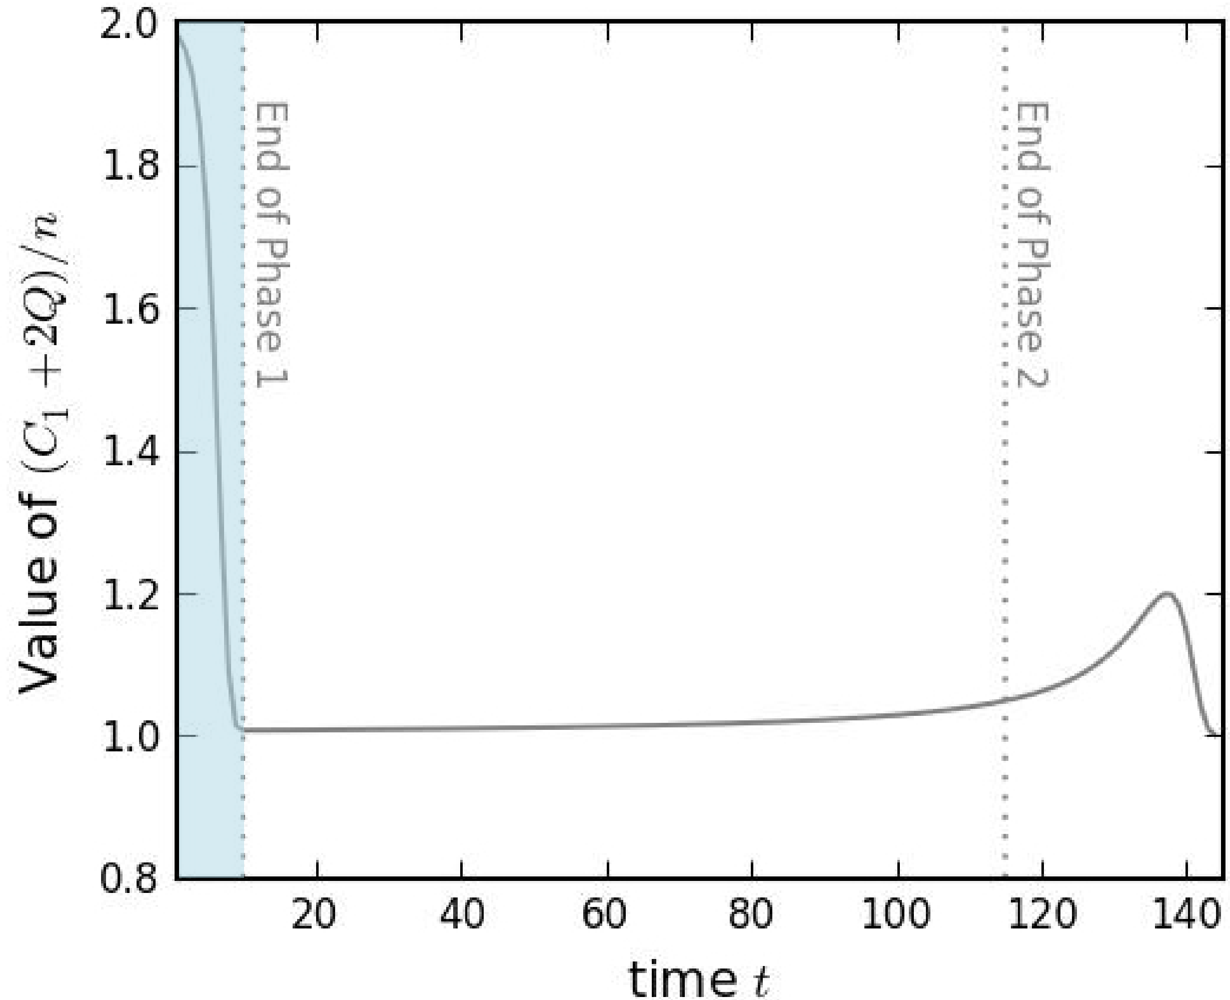

Remarks

W.h.p.

- Plurality does not change.
- Growth factor of plurality is > 1.

Overview of the Process

$$\mathbf{E}\left[c_i^{(t+1)} \,\middle|\, \mathbf{c}^{(t)}\right] = = c_i^{(t)} \cdot \underbrace{\frac{c_i^{(t)} + 2q^{(t)}}{n}}_{\text{Growth factor}}$$

Simulation of the growth factor:

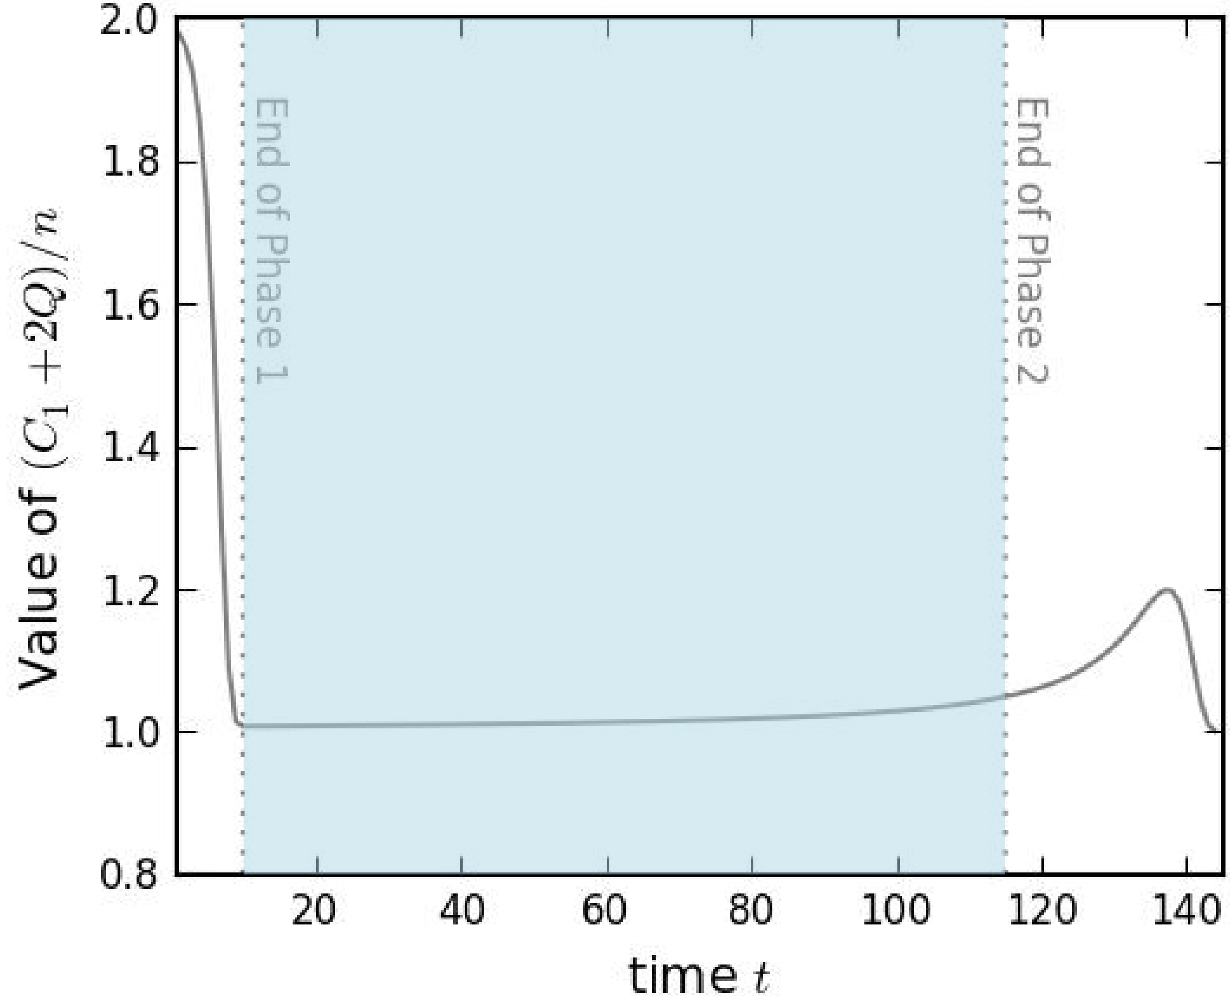

Remarks

W.h.p.

- Plurality does not change.
- Growth factor of plurality is > 1.

Overview of the Process

$$\mathbf{E}\left[c_i^{(t+1)} \,\middle|\, \mathbf{c}^{(t)}\right] = = c_i^{(t)} \cdot \underbrace{\frac{c_i^{(t)} + 2q^{(t)}}{n}}_{\text{Growth factor}}$$

Remarks

W.h.p.

- Plurality does not change.
- Growth factor of plurality is > 1.

Simulation of the growth factor:

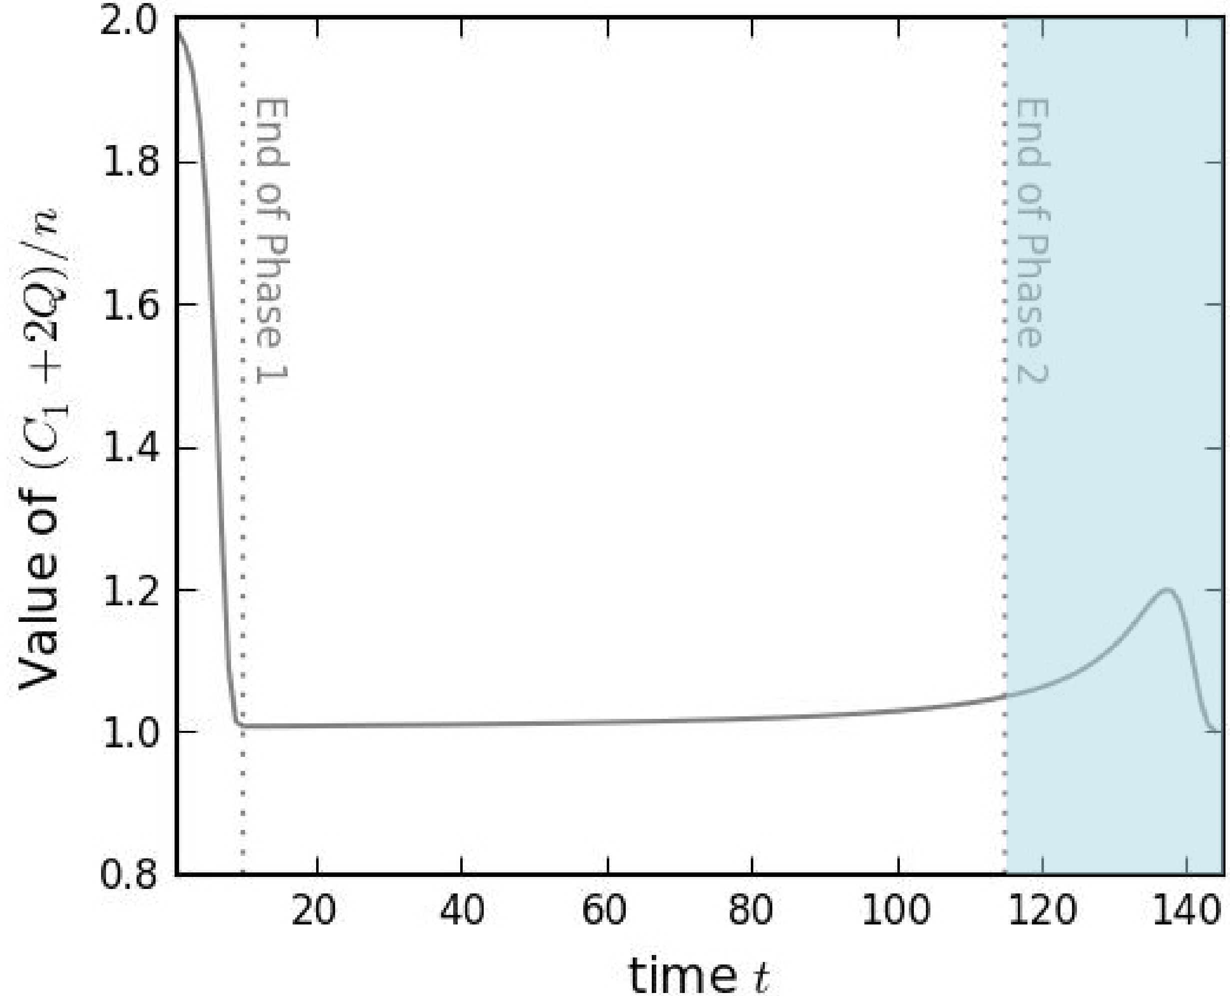

Expected Behaviour of the Process

$$
\begin{cases}
\mathbf{E}\left[q^{(t+1)} \;\middle|\mathbf{c}^{(t)}\right] = \frac{1}{n}\left[\left(q^{(t)}\right)^2 + \left(n - q^{(t)}\right)^2 - \sum_i \left(c_i^{(t)}\right)^2\right] \\
\mathbf{E}\left[c_1^{(t+1)} \;\middle|\mathbf{c}^{(t)}\right] = c_1^{(t)} \cdot \frac{c_1^{(t)} + 2q^{(t)}}{n} \\
\quad\vdots \\
\mathbf{E}\left[c_k^{(t+1)} \;\middle|\mathbf{c}^{(t)}\right] = c_k^{(t)} \cdot \frac{c_k^{(t)} + 2q^{(t)}}{n}
\end{cases}
$$

Key Idea

Tip: Look for $md(\mathbf{c}^{(t)})$ and $R(\mathbf{c}^{(t)}) := \sum_{i=1}^{k} \frac{c_i^{(t)}}{c_1^{(t)}}$.

Key Idea

Tip: Look for $md(\mathbf{c}^{(t)})$ and $R(\mathbf{c}^{(t)}) := \sum_{i=1}^{k} \frac{c_i^{(t)}}{c_1^{(t)}}$.

$$\mathbf{E}\left[\frac{c_1^{(t+1)} + 2q^{(t+1)}}{n} \,\middle|\, \mathbf{c}^{(t)} \right] =$$

$$= 1 + \frac{\left(n - 2q^{(t)} - c_1^{(t)}\right)^2}{n^2} + \frac{2\left(R(\mathbf{c}^{(t)}) - \mathrm{md}(\mathbf{c}^{(t)})\right) \cdot (c_1)^2}{n^2}$$

First Round

Round 1:
Each node observes another random one. The larger the number of colors and the more uniform the initial distribution, the higher the expected number of undecided nodes.

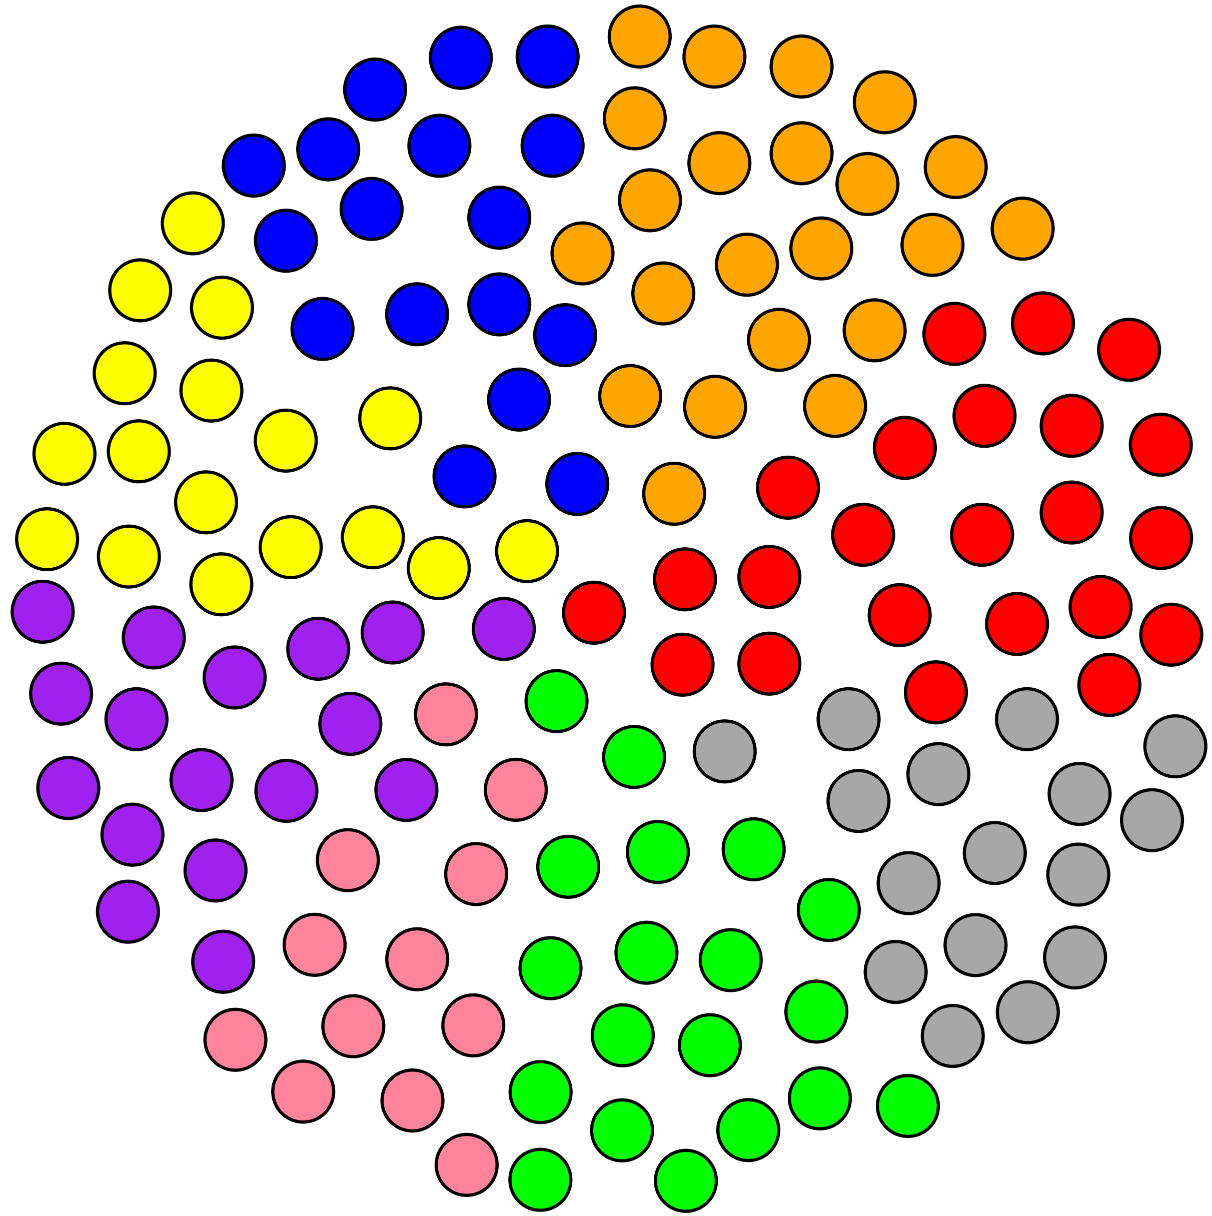

First Round

The size of each color is reduced to $\frac{(c_i^{(0)})^2}{n}$. Colors with $c_i^{(0)} = O(\sqrt{n})$ nodes are likely to disappear.

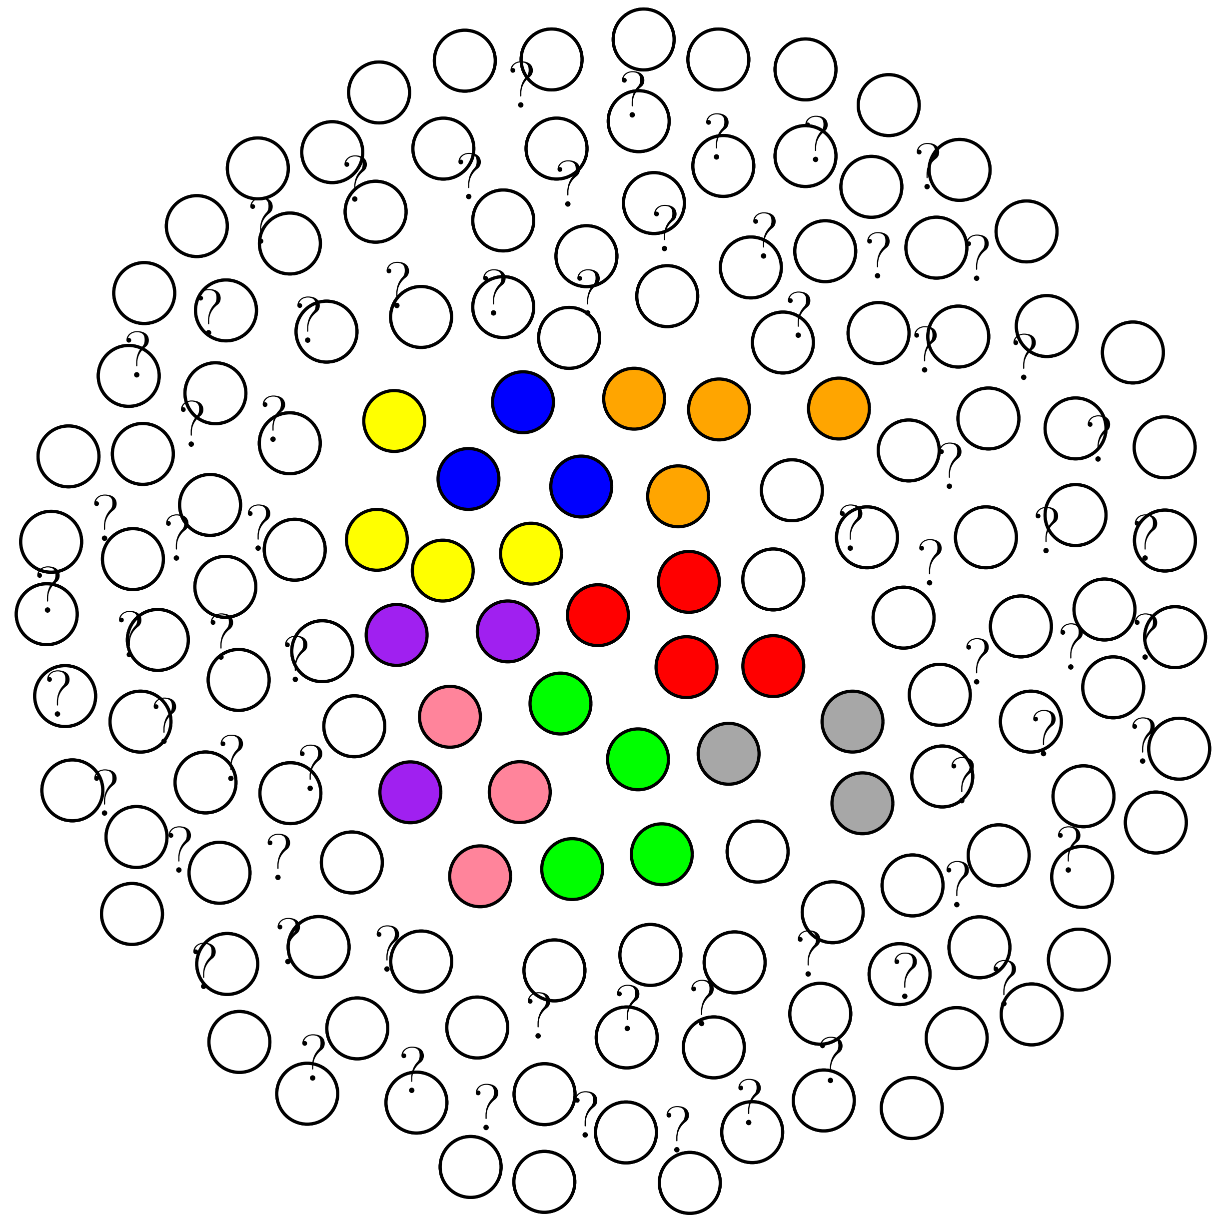

Phase 1

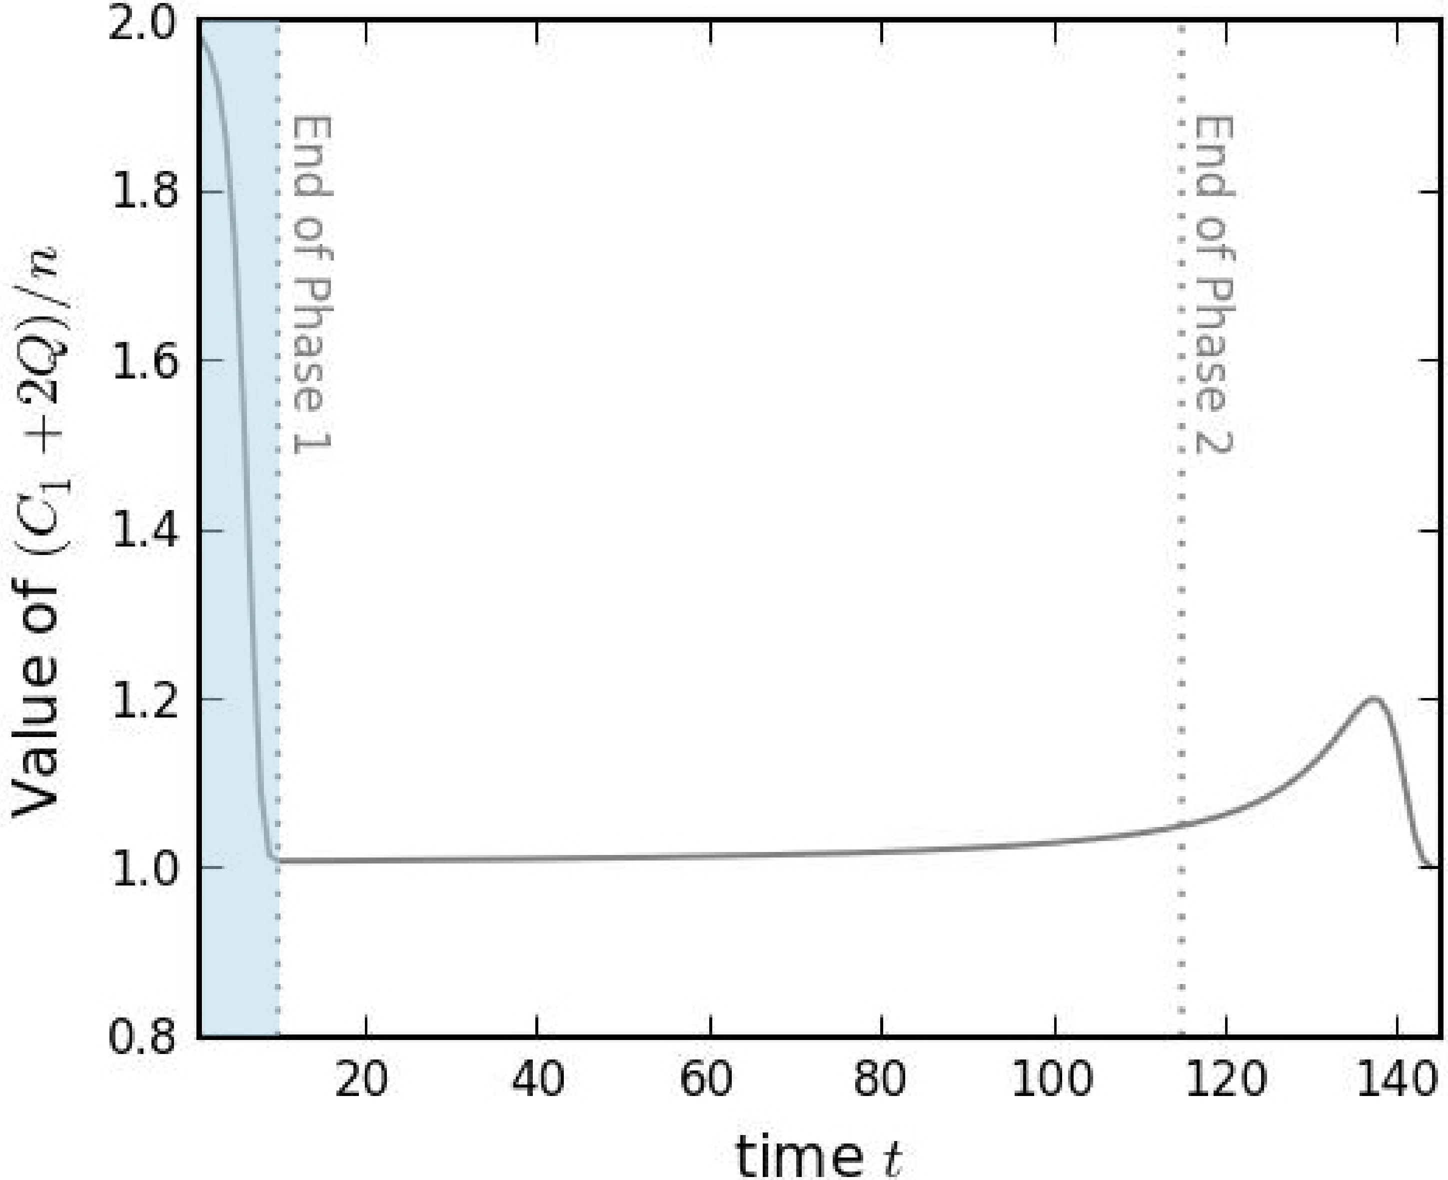

End of Phase 1
End of Phase 2
Value of $(C_1 + 2Q)/n$
time t
2.0
1.8
1.6
1.4
1.2
1.0
0.8
20
40
60
80
100
120
140

Phase 1

If the initial distribution is quite uniform there are $\Omega(n)$ undecided nodes. Undecided nodes take the first color they pull, causing colors to spread very fast.

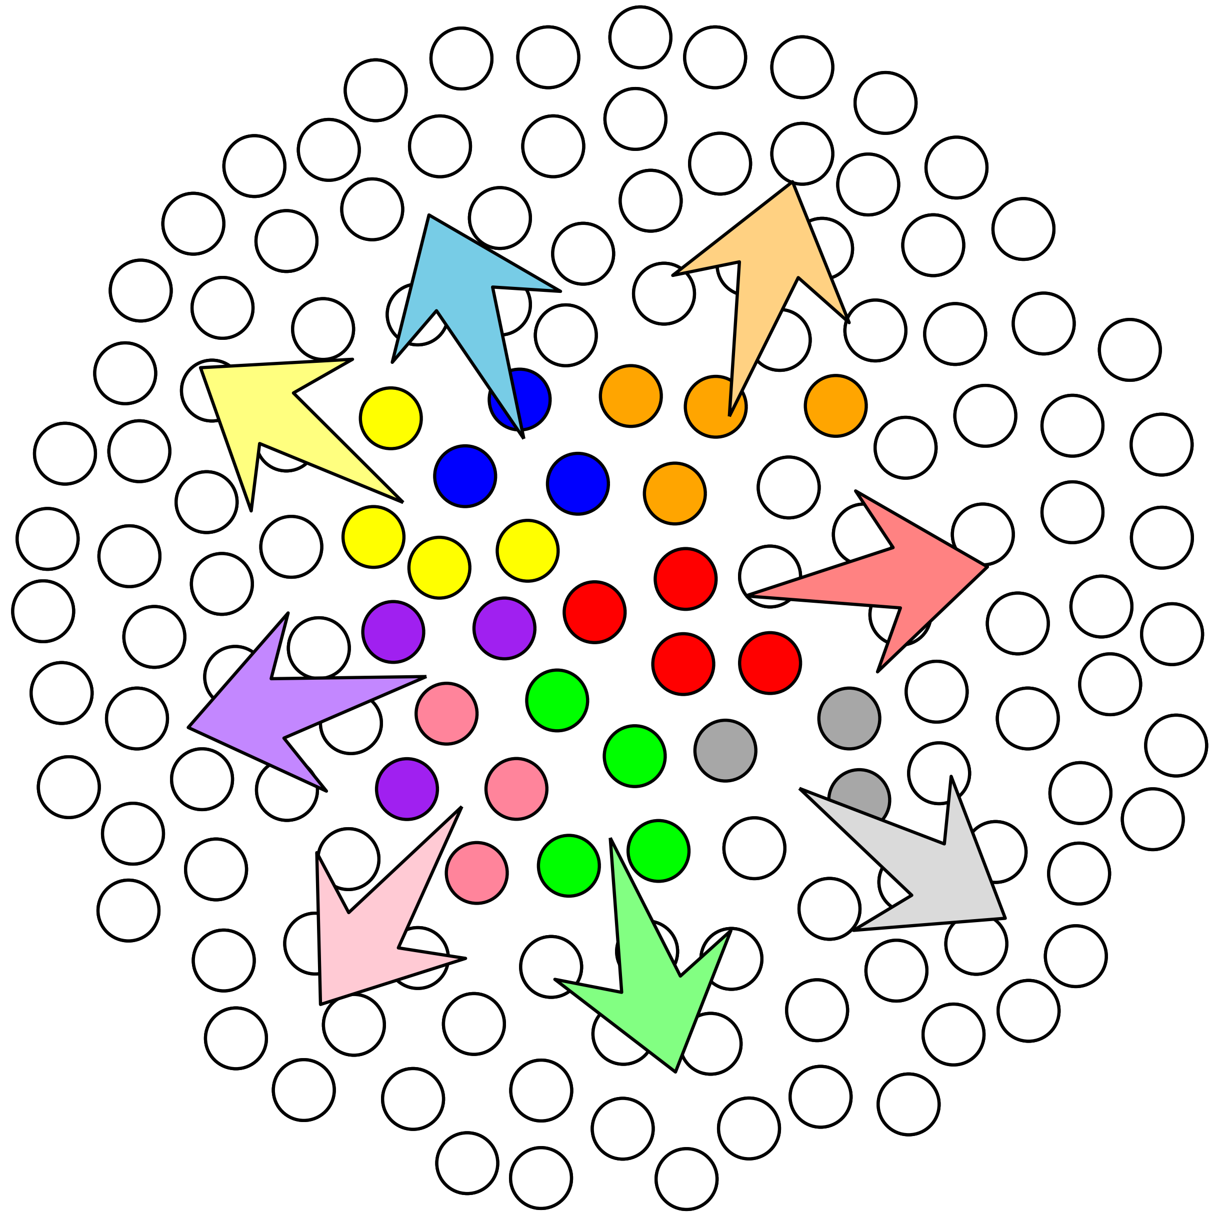

Phase 1

Lemma
Within $T = O\left(\log \frac{R(\mathbf{c})^2}{\mathrm{md}(\mathbf{c})}\right)$ rounds the system reaches a configuration such that w.h.p.

$$c_1^{(T)} = \Theta\left(\frac{n}{\mathrm{md}(\mathbf{c})}\right)$$

$$q^{(T)} = \frac{n}{2}\left(1 \pm \Theta\left(\frac{1}{\mathrm{md}(\mathbf{c})}\right)\right)$$

and, for every i, $c_1^{(0)}/c_i^{(0)}$ is approximately preserved.

Phase 2

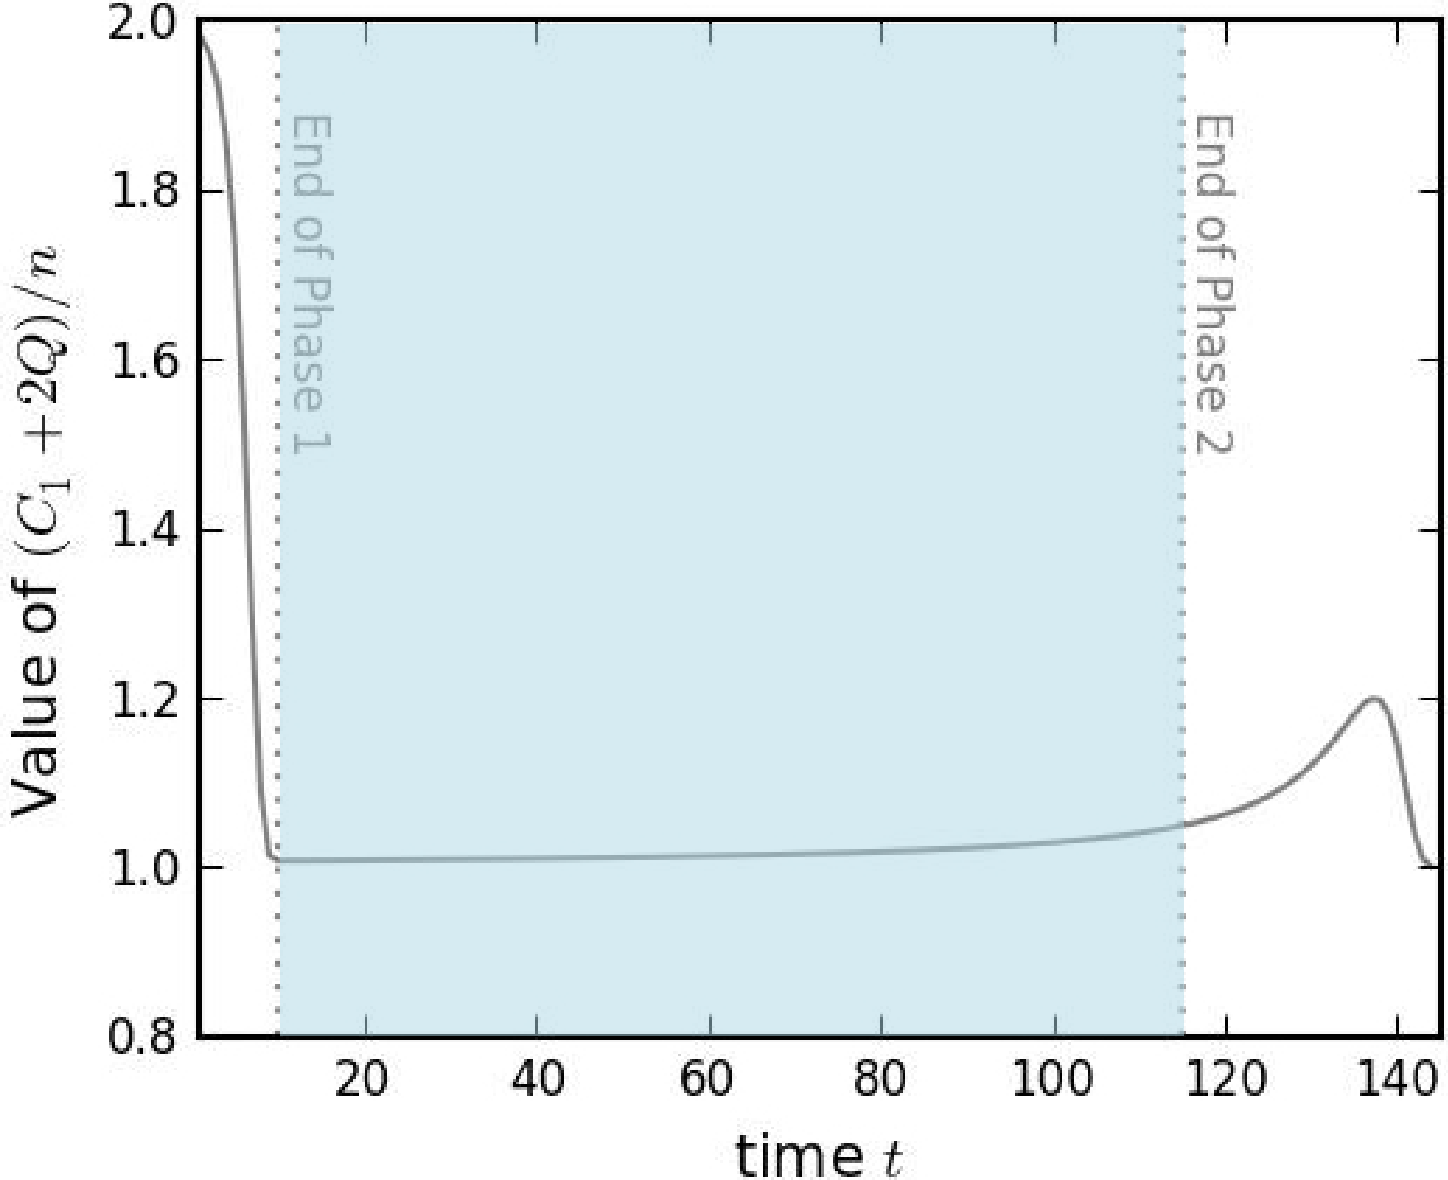

2.0
1.8
1.6
1.4
1.2
1.0
0.8
Value of $(C_1 + 2Q)/n$
20
40
60
80
100
120
140
time t
End of Phase 1
End of Phase 2

Phase 2

\# new colored

$\approx$

\# new undecided.

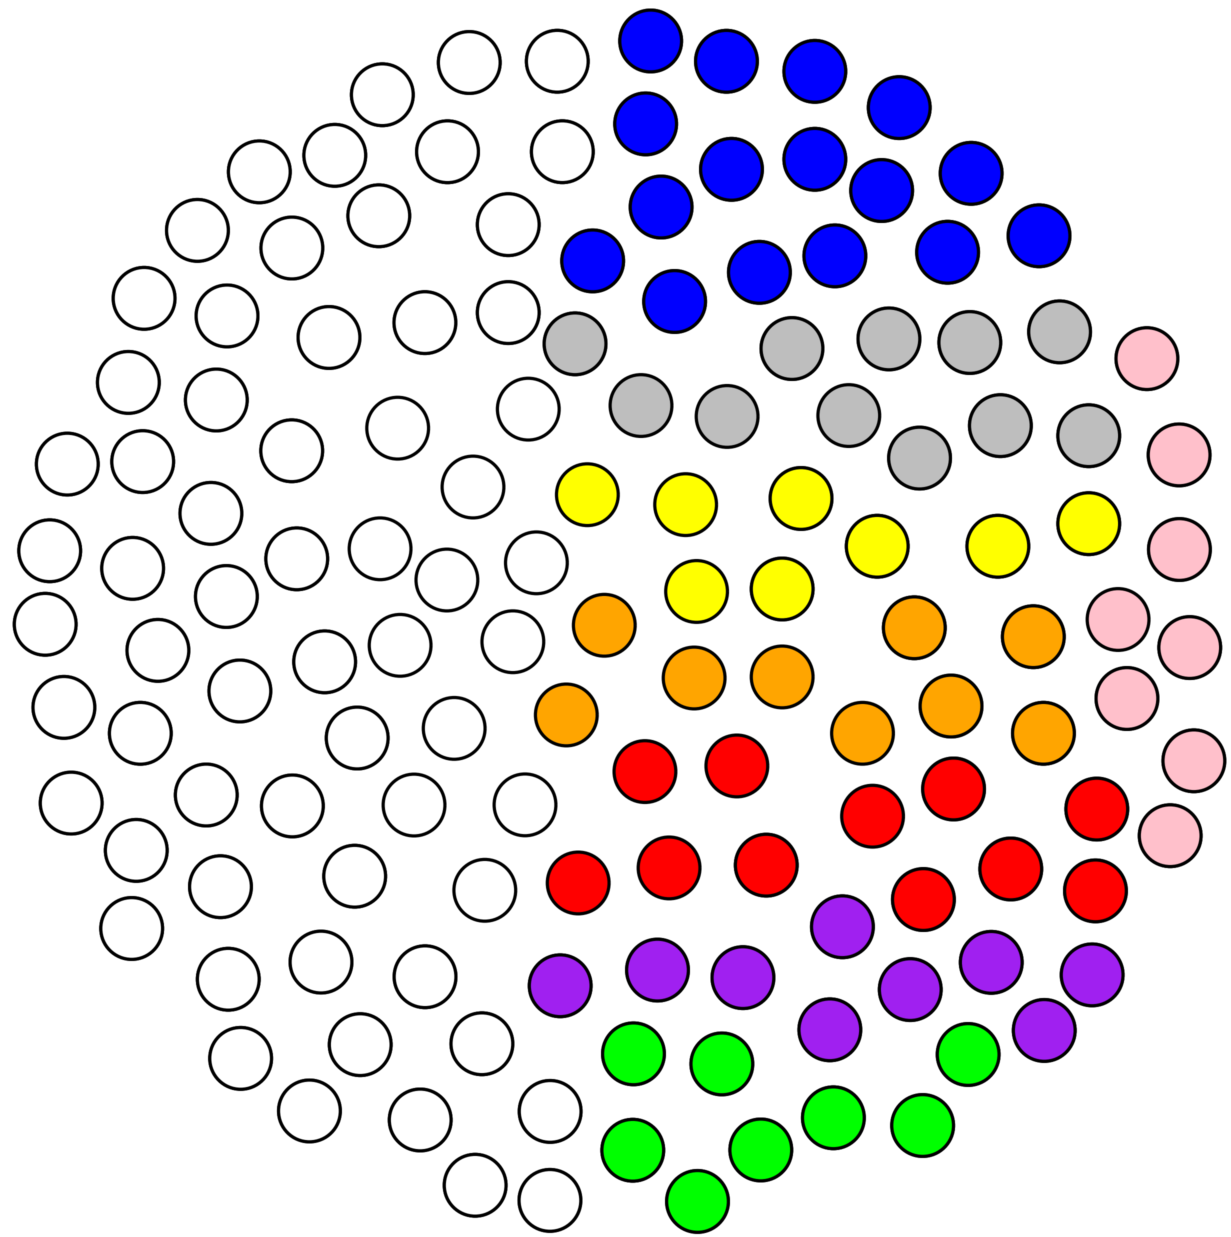

Phase 2

\# new colored

$\approx$

\# new undecided.

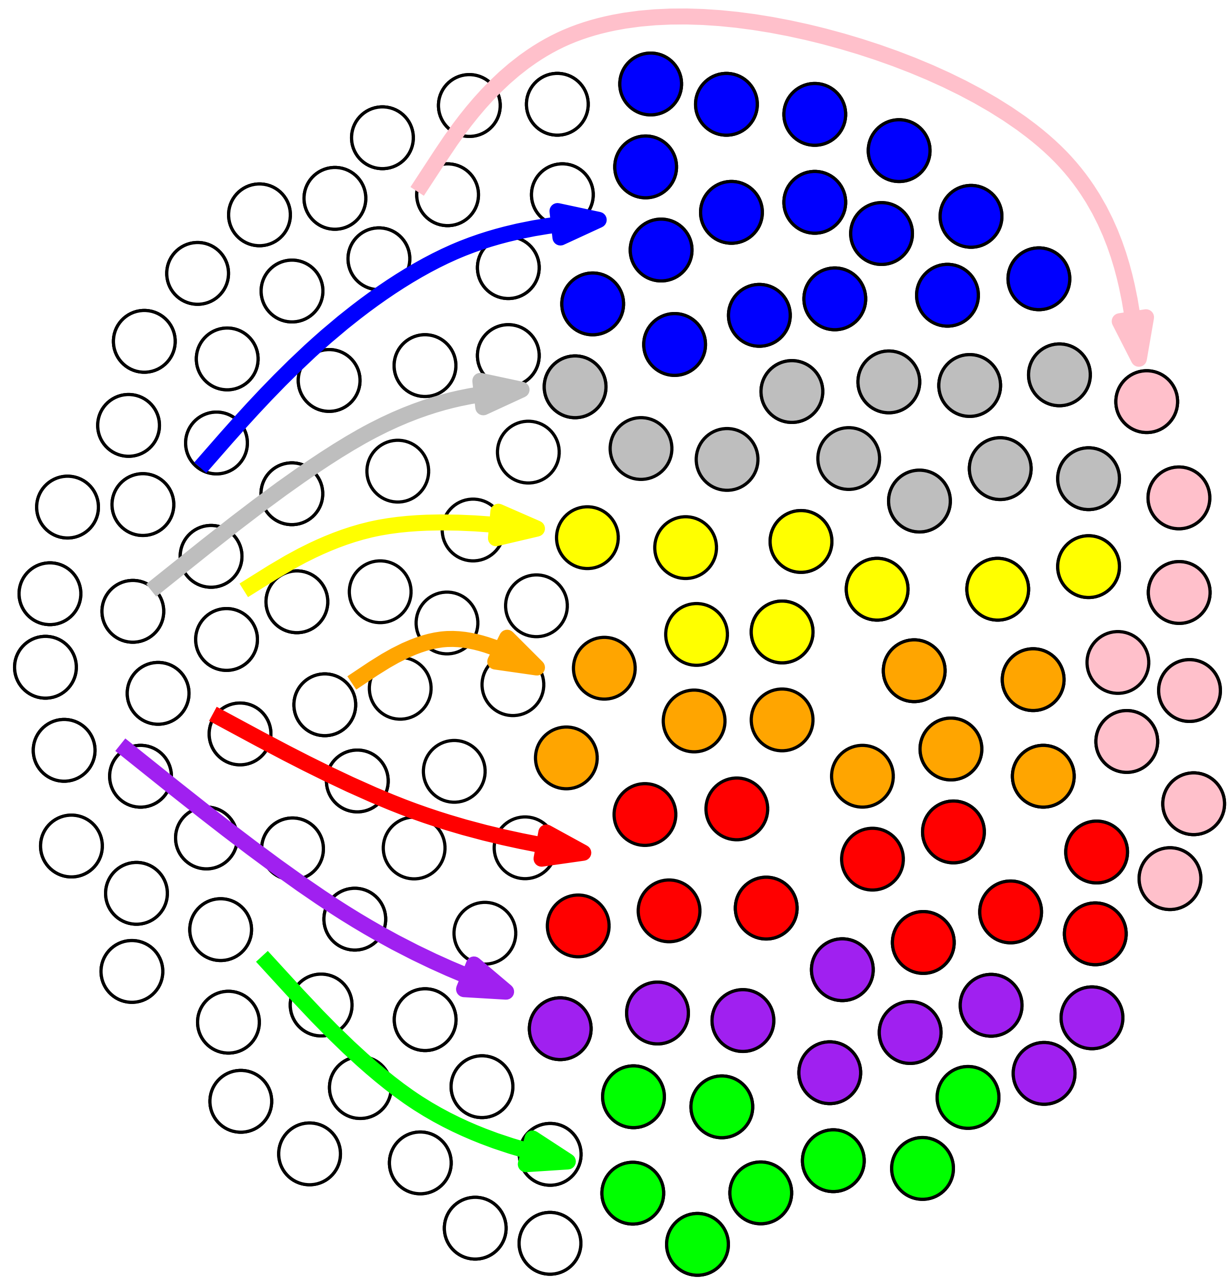

Phase 2

new colored

$\approx$

new undecided.

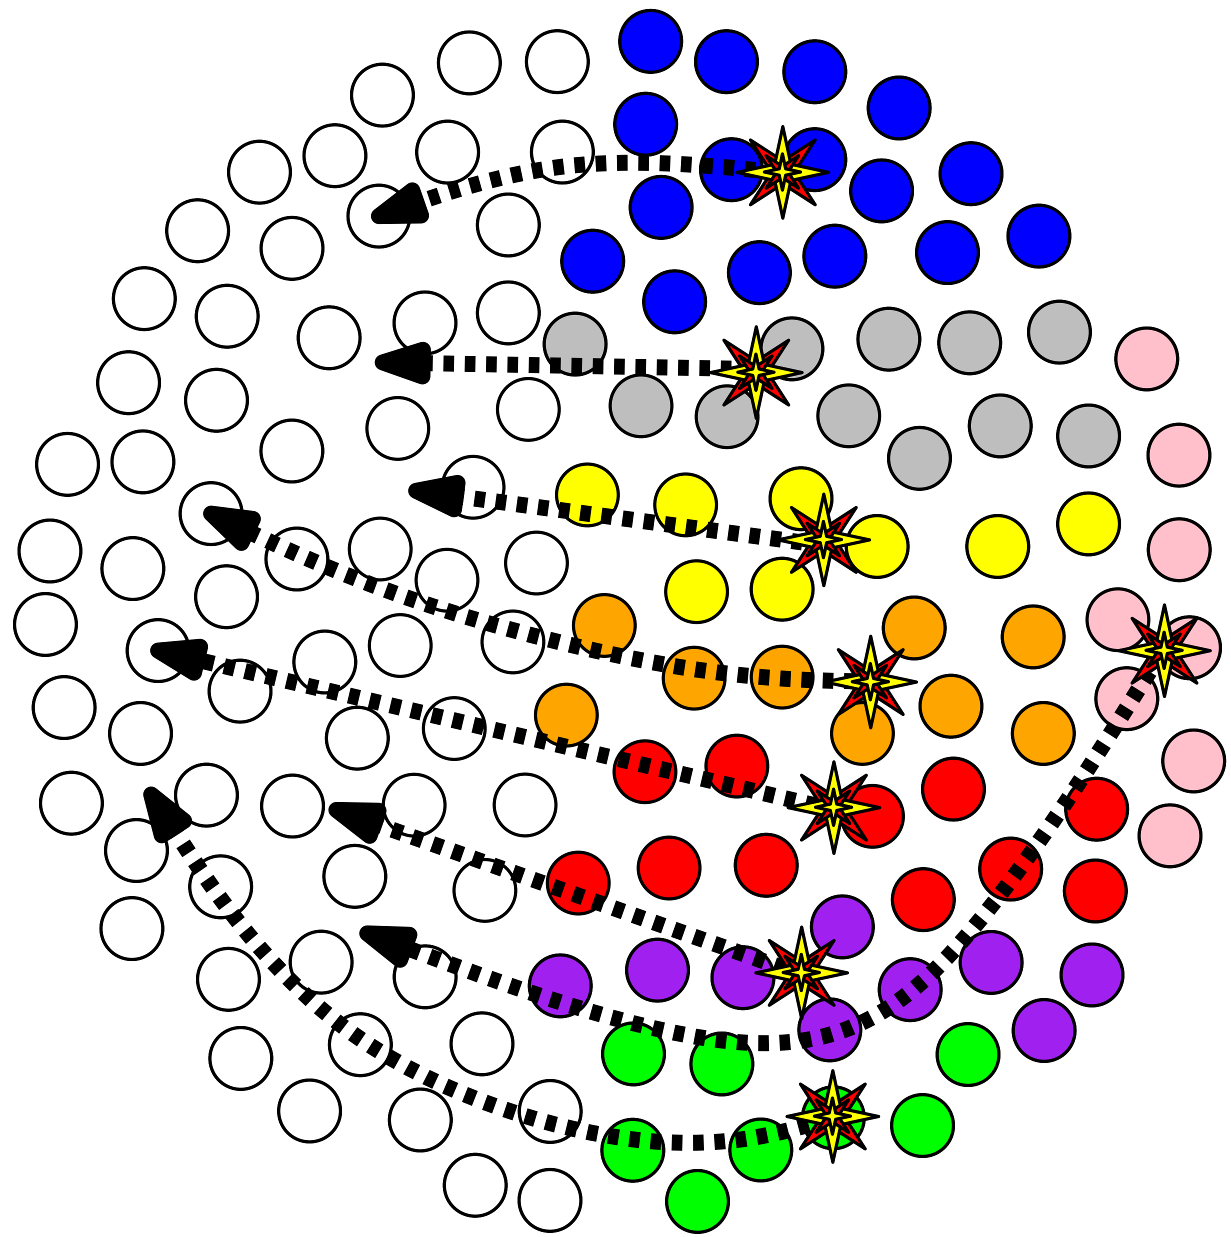

Phase 2

\# new colored
$\approx$
\# new undecided.

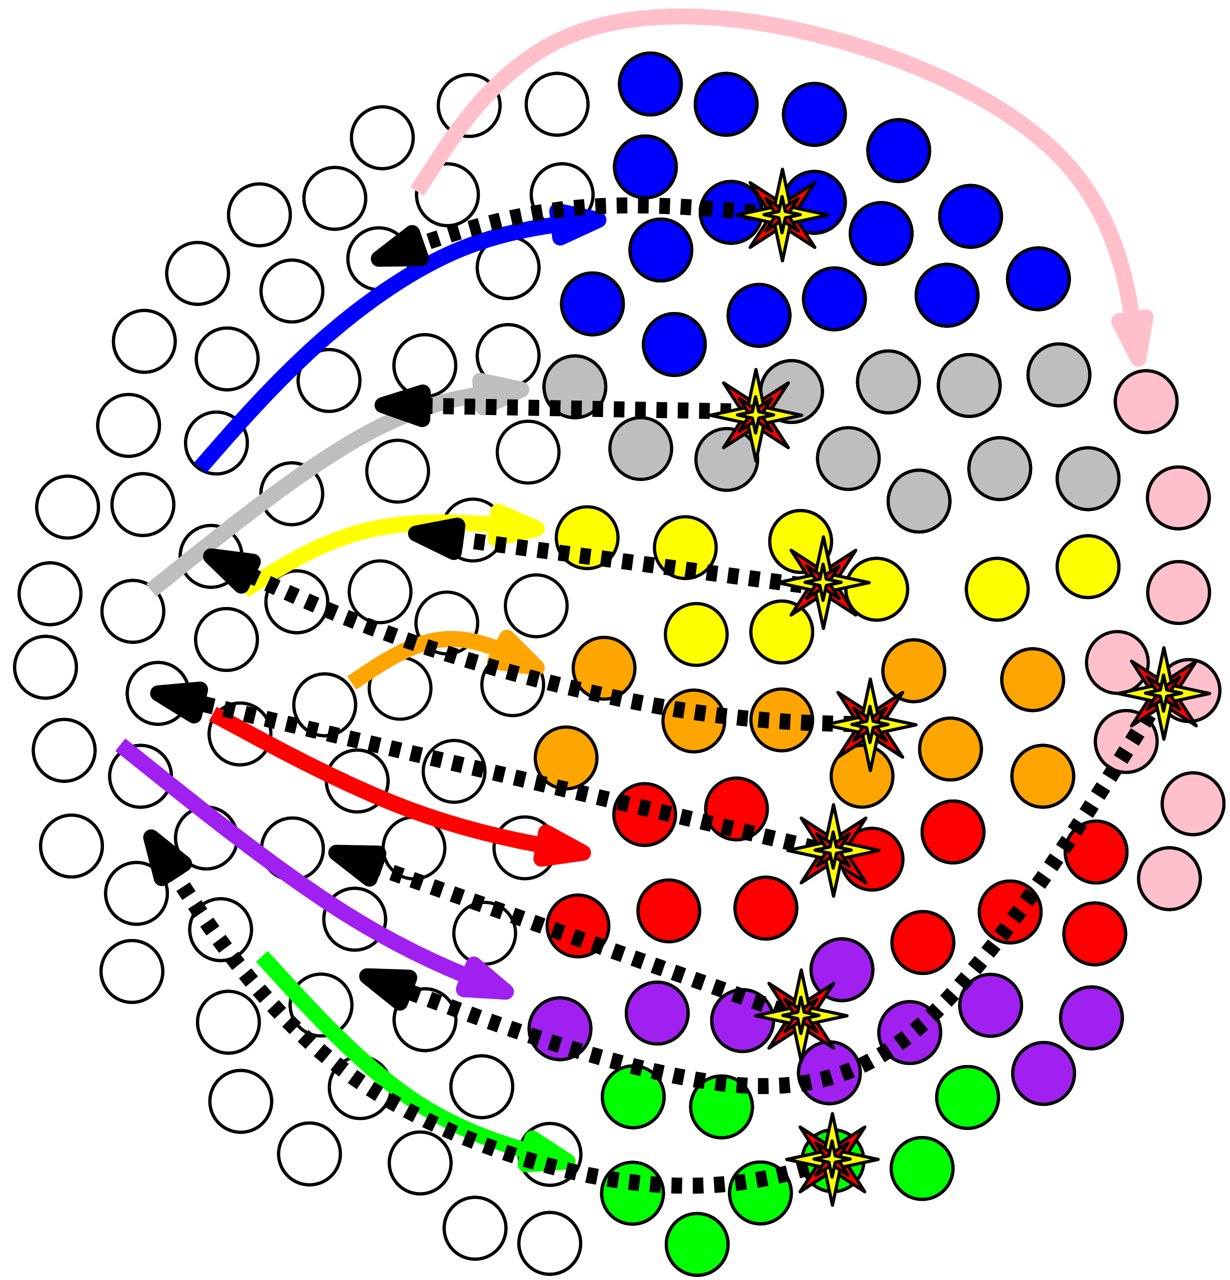

Phase 2

Average Growth:

$$\mathbf{E}\left[c_1^{(t+1)} \,\middle|\, \mathbf{c}^{(t)}\right] \approx c_1^{(t)}\left(1+\Theta\left(\frac{1}{\mathrm{md}(\mathbf{c})}\right)\right)$$

$$\mathbf{E}\left[q^{(t+1)} \,\middle|\, \mathbf{c}^{(t)}\right] \approx \frac{n}{2}\left(1-\Theta\left(\frac{1}{\mathrm{md}(\mathbf{c})}\right)\right)$$

Phase 2

Average Growth:

$$\mathbf{E}\left[c_1^{(t+1)} \;\middle|\; \mathbf{c}^{(t)}\right] \approx c_1^{(t)}\left(1+\Theta\left(\frac{1}{\mathrm{md}(\mathbf{c})}\right)\right)$$

$$\mathbf{E}\left[q^{(t+1)} \;\middle|\; \mathbf{c}^{(t)}\right] \approx \frac{n}{2}\left(1-\Theta\left(\frac{1}{\mathrm{md}(\mathbf{c})}\right)\right)$$

$\Longrightarrow$ Lower bound of $\Omega\left(\mathrm{md}(\mathbf{c})\right)$.

Phase 3

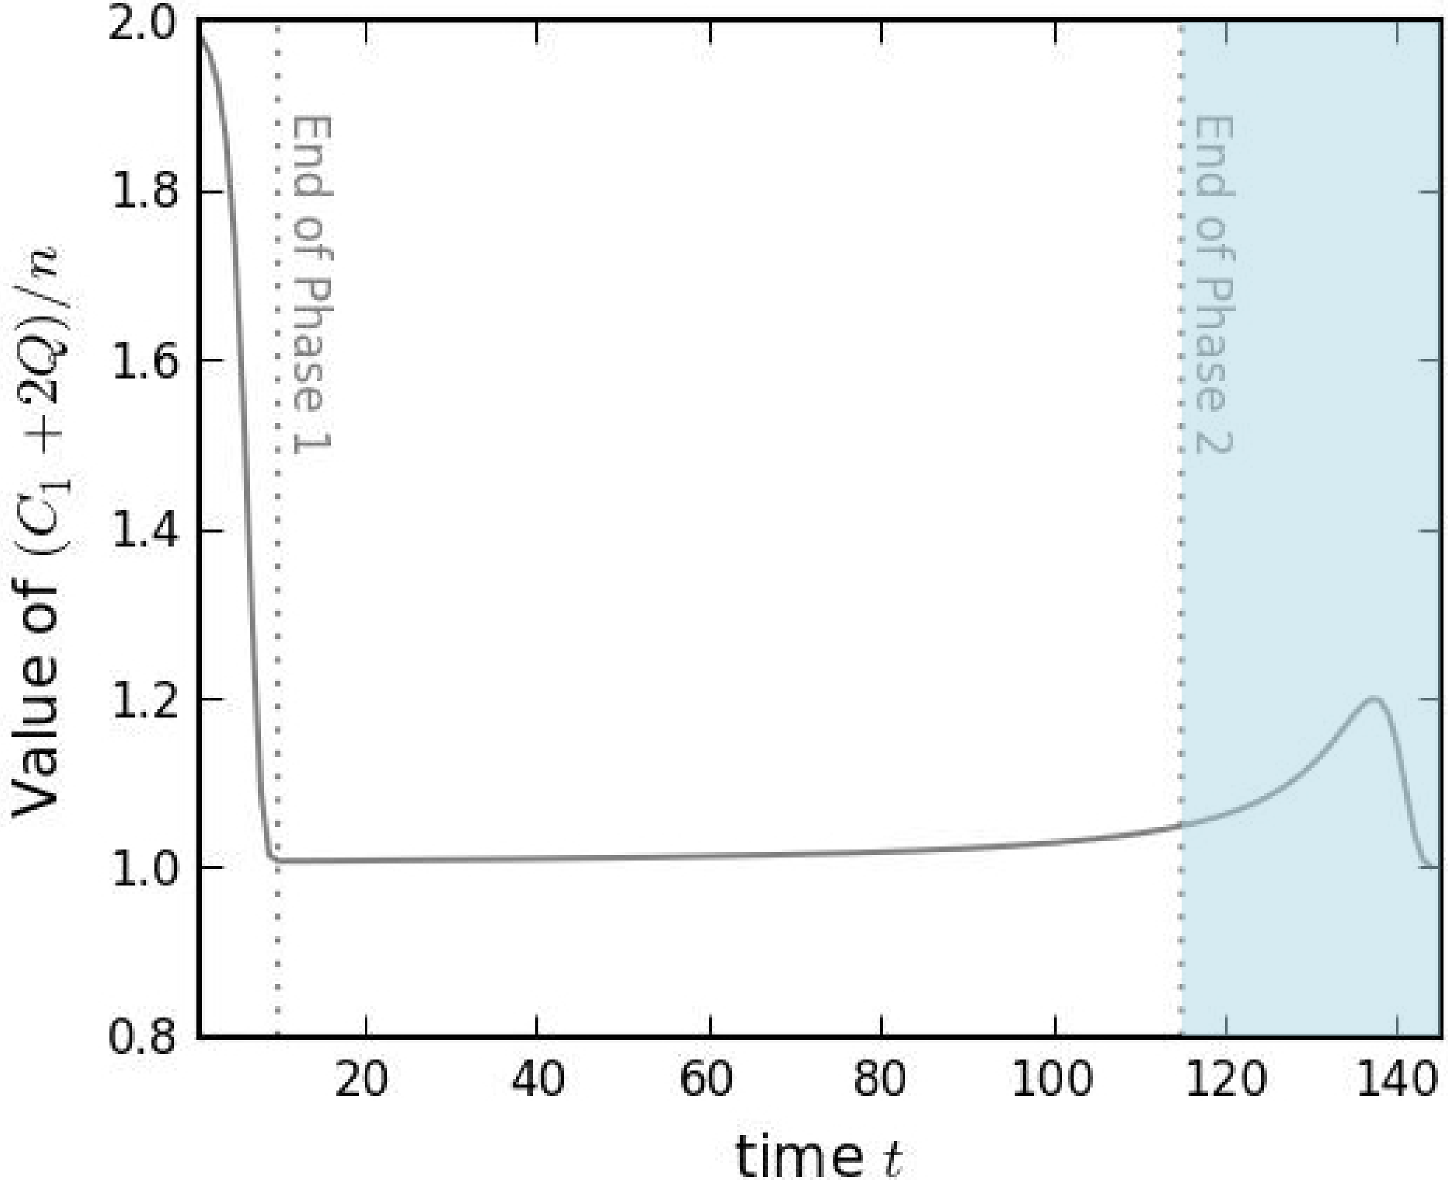

2.0
1.8
1.6
1.4
1.2
1.0
0.8
20
40
60
80
100
120
140
time t
Value of $(C_1 + 2Q)/n$
End of Phase 1
End of Phase 2

Phase 3

The plurality has a small advantage $\implies$ after long time the equilibrium breaks down.

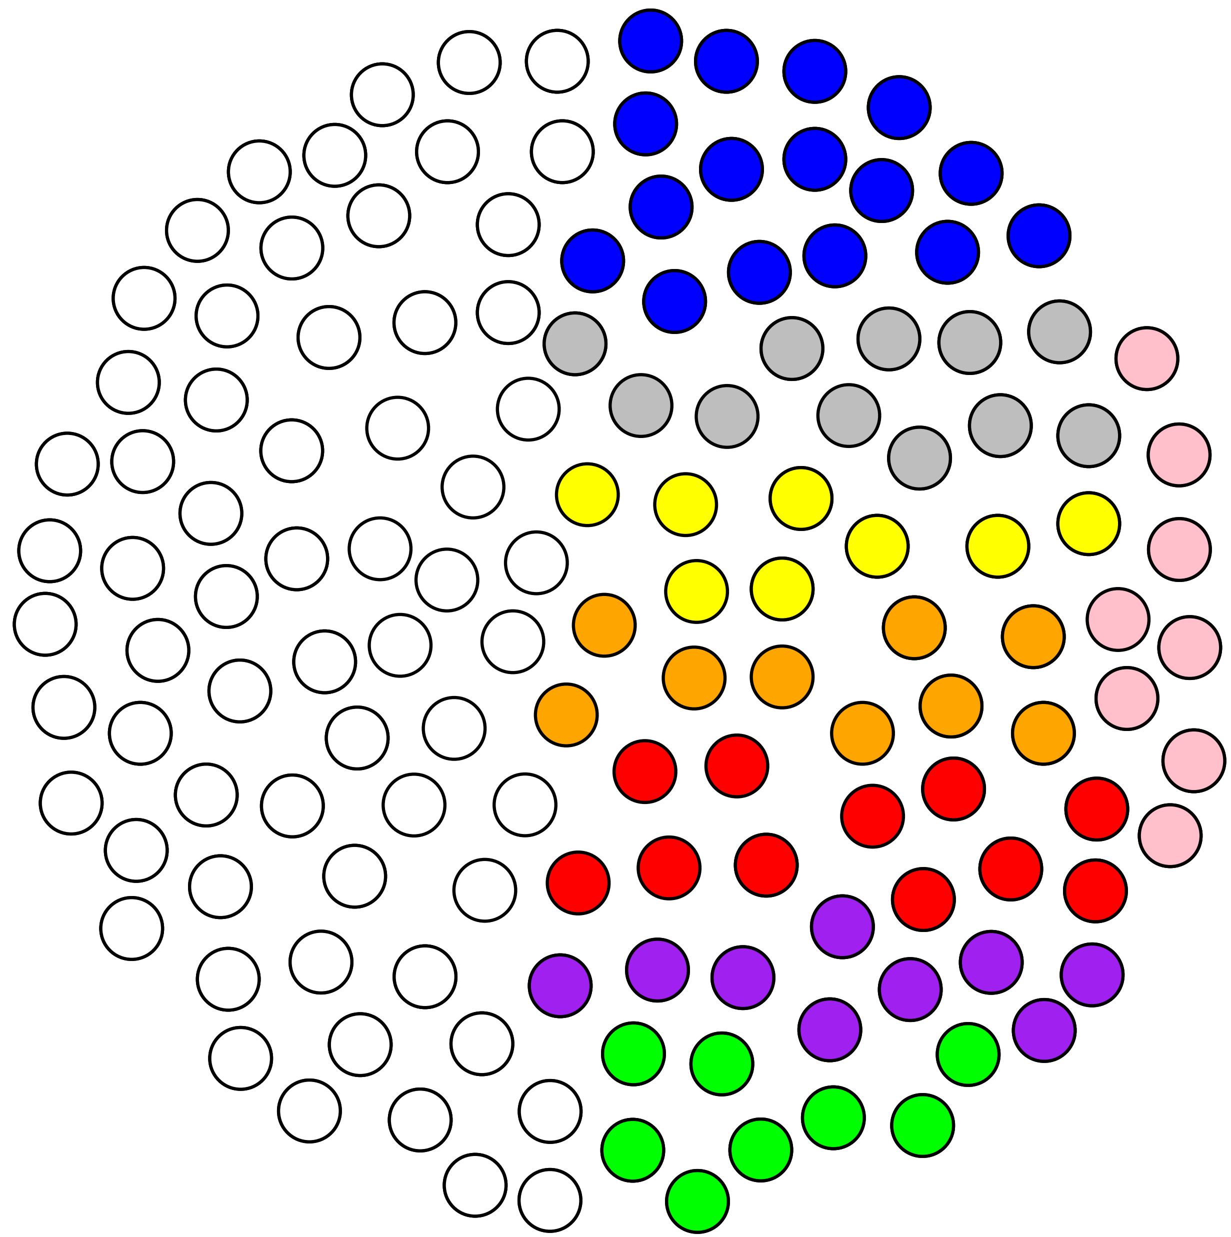

Phase 3

The plurality has a small advantage $\implies$ after long time the equilibrium breaks down.

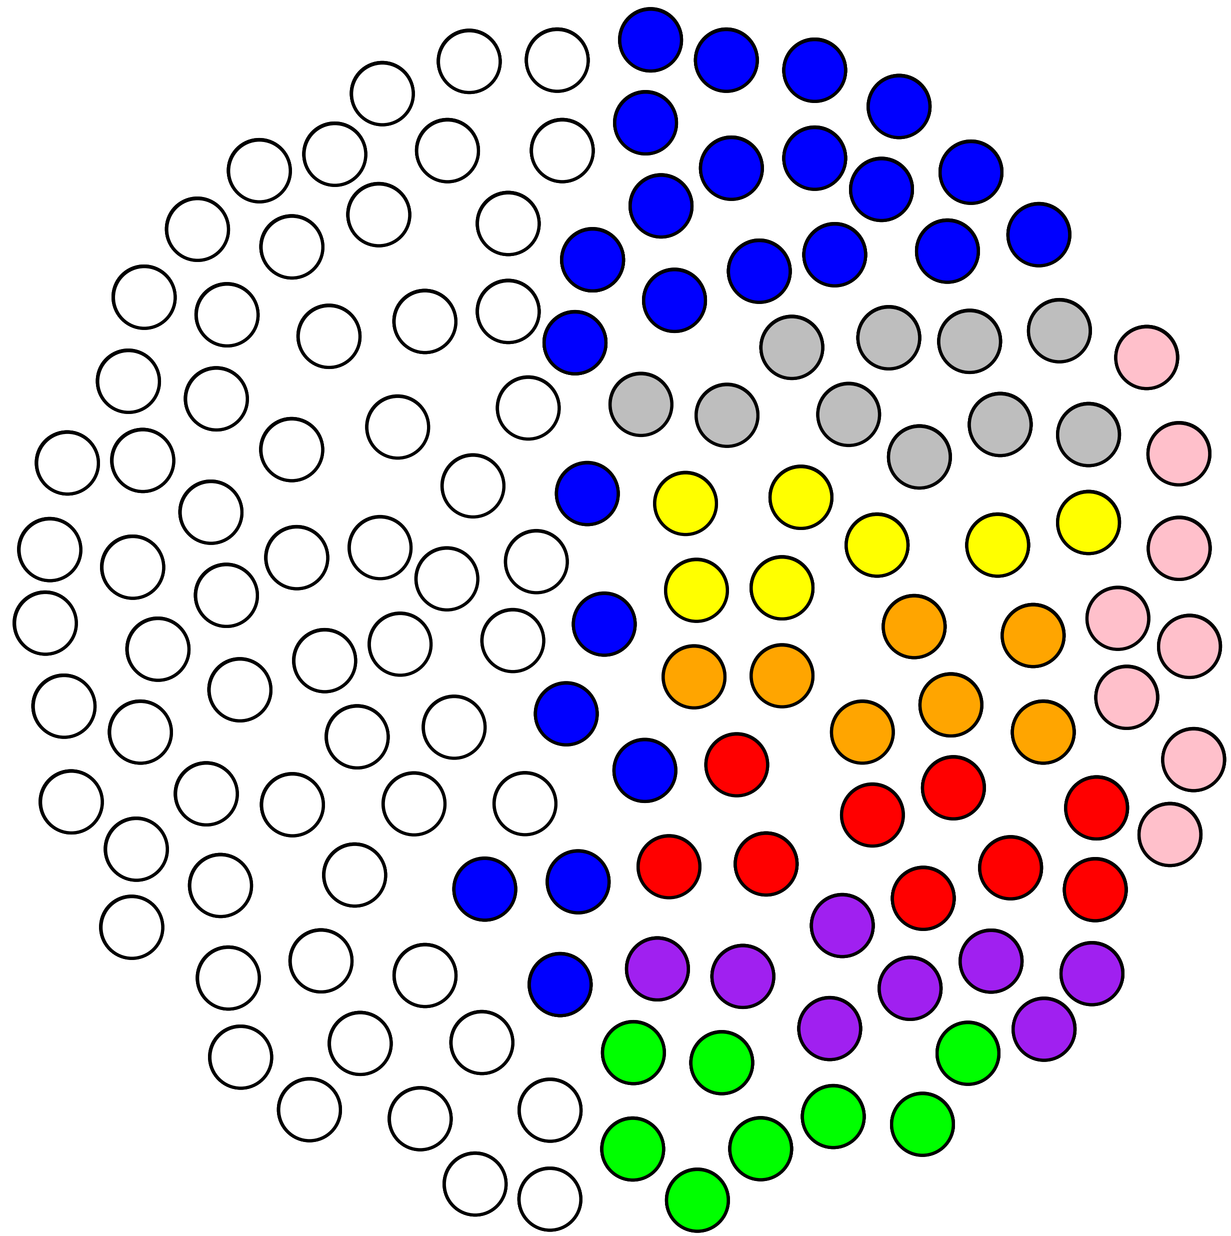

Phase 3

The plurality has a small advantage $\Longrightarrow$ after long time the equilibrium breaks down.

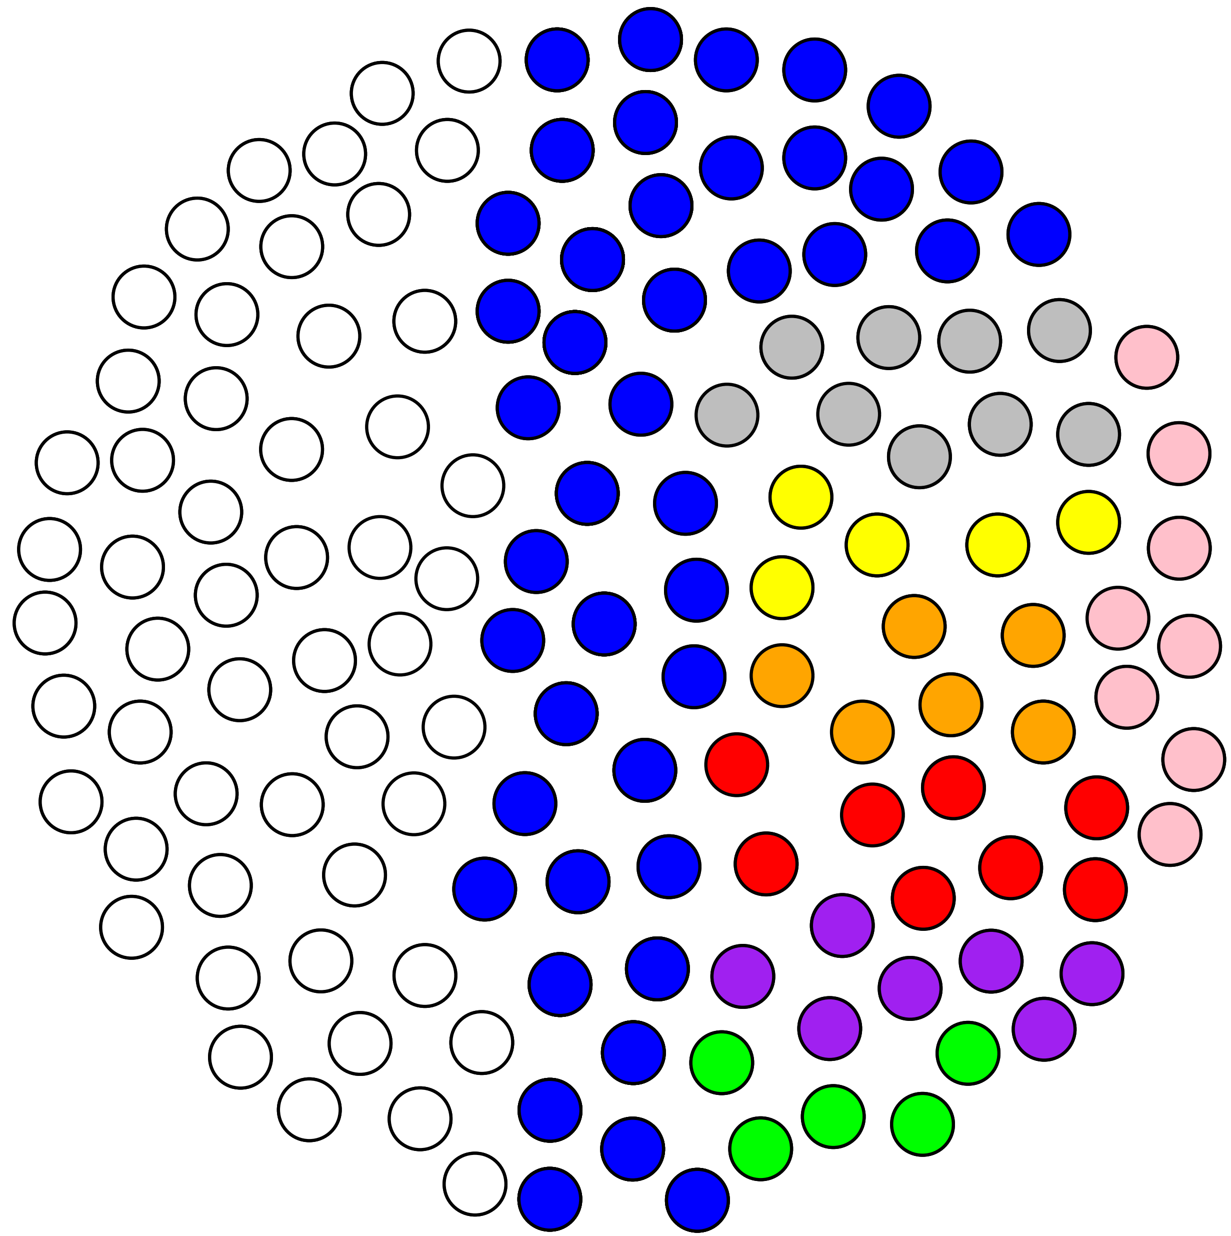

Phase 3

Average Growth:

$$\mathbf{E}\left[c_1^{(t+\mathrm{md}(\mathbf{c}))}\;\middle|\mathbf{c}^{(t)}\right] \approx c_1^{(t)}\left(1+\Theta\left(\frac{1}{\mathrm{md}(\mathbf{c})}\right)\right)^{\mathrm{md}(\mathbf{c})}$$

$$\mathbf{E}\left[q^{(t+\mathrm{md}(\mathbf{c}))}\;\middle|\mathbf{c}^{(t)}\right] \approx \frac{n}{2}\left(1-\Theta\left(\frac{1}{\mathrm{md}(\mathbf{c})}\right)\right)^{\mathrm{md}(\mathbf{c})}$$

$\implies$ After $O\left(\mathrm{md}(\mathbf{c})\log n\right)$ rounds, $R(\mathbf{c}^{(t)}) = 1 + o(1)$.

Phase 3

$$R(\mathbf{c}^{(t)}) = 1 + o(1) \implies c_1^{(t)} = \frac{n - q^{(t)}}{R(\mathbf{c}^{(t)})} \approx n - q^{(t)}$$

Phase 3

$$R(\mathbf{c}^{(t)}) = 1 + o(1) \implies c_1^{(t)} = \frac{n - q^{(t)}}{R(\mathbf{c}^{(t)})} \approx n - q^{(t)}$$

$$\implies \mathbf{E}\left[\frac{c_1^{(t+1)} + 2q^{(t+1)}}{n} \,\middle|\, \mathbf{c}^{(t)} \right] \geq 1 + \left(\frac{q^{(t)}}{n} \right)^2$$

Phase 3

$$R(\mathbf{c}^{(t)}) = 1 + o(1) \implies c_1^{(t)} = \frac{n-q^{(t)}}{R(\mathbf{c}^{(t)})} \approx n - q^{(t)}$$

$$\implies \mathbf{E}\left[\frac{c_1^{(t+1)}+2q^{(t+1)}}{n} \middle| \mathbf{c}^{(t)}\right] \geq 1 + \left(\frac{q^{(t)}}{n}\right)^2$$

$\implies$ Plurality Consensus is reached within $O(\log n)$ rounds.

Extension to d-Regular Expanders

Theorem

Given a d-regular expander graph, $k = O\left((n/\log n)^{1/3}\right)$ and $c_1 \geq (1+\epsilon) \cdot c_2$ with $\epsilon > 0$, using polylogarithmic memory and message size the plurality consensus problem can be solved in w.h.p. $O(\mathrm{md}(\mathbf{c})\mathrm{polylog}(n))$ rounds.

Extension to d-Regular Expanders

Theorem
Given a d-regular expander graph, $k = O\left((n/\log n)^{1/3}\right)$ and $c_1 \geq (1+\epsilon) \cdot c_2$ with $\epsilon > 0$, using polylogarithmic memory and message size the plurality consensus problem can be solved in w.h.p. $O(\mathrm{md}(\mathbf{c})\mathrm{polylog}(n))$ rounds.

Idea
Simulate Undecided-State Dynamics on complete graph by sampling via n parallel random walks.
(Rapidly mixing property)

Extension to d-Regular Expanders

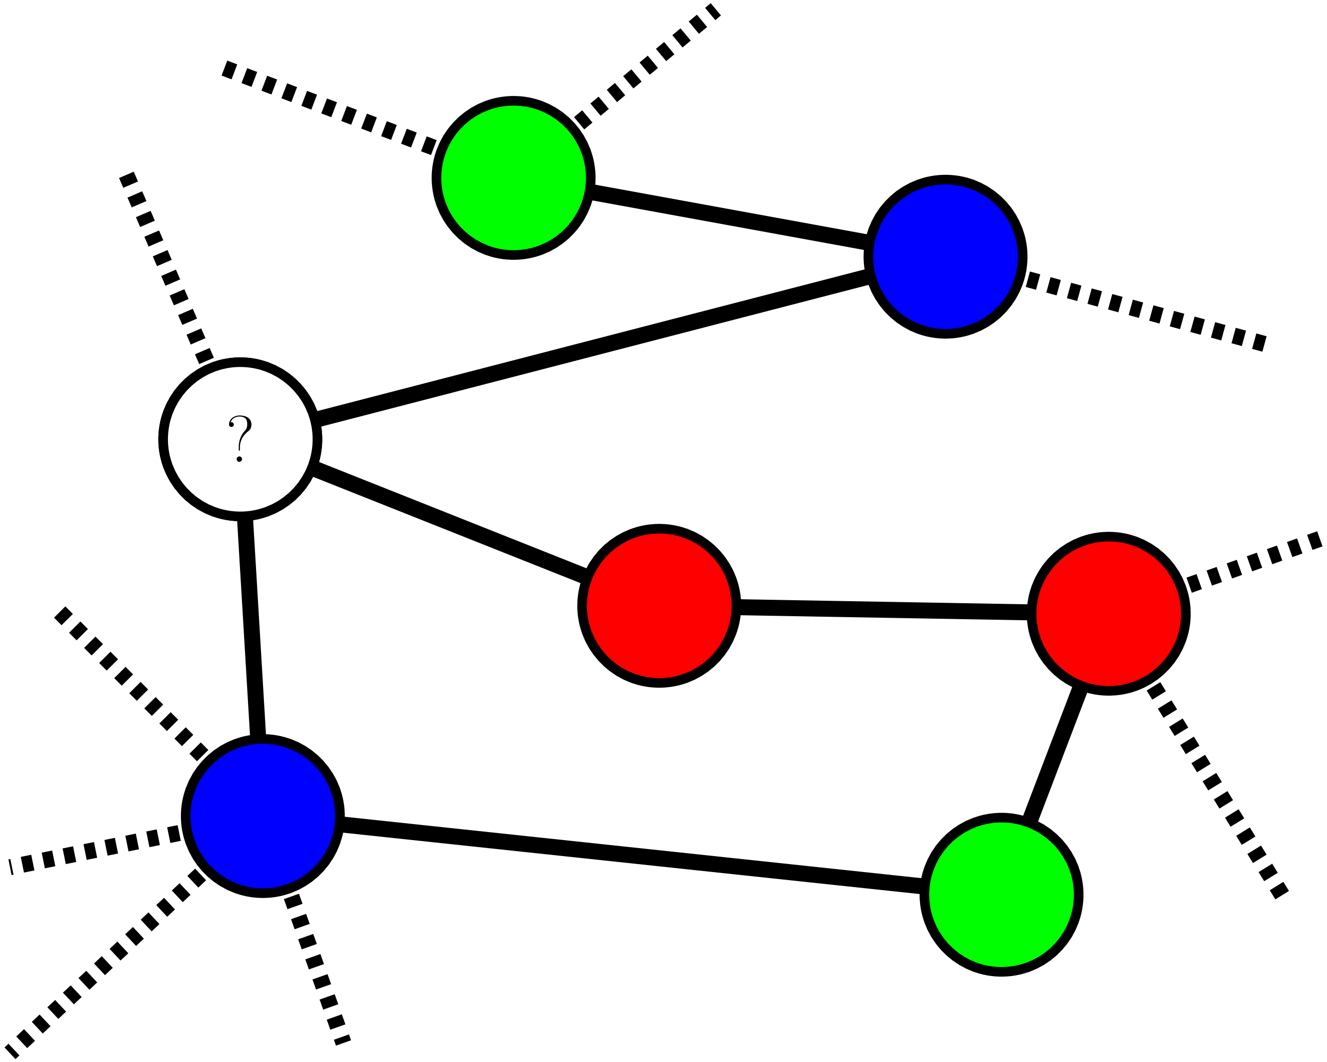

Extension to d-Regular Expanders

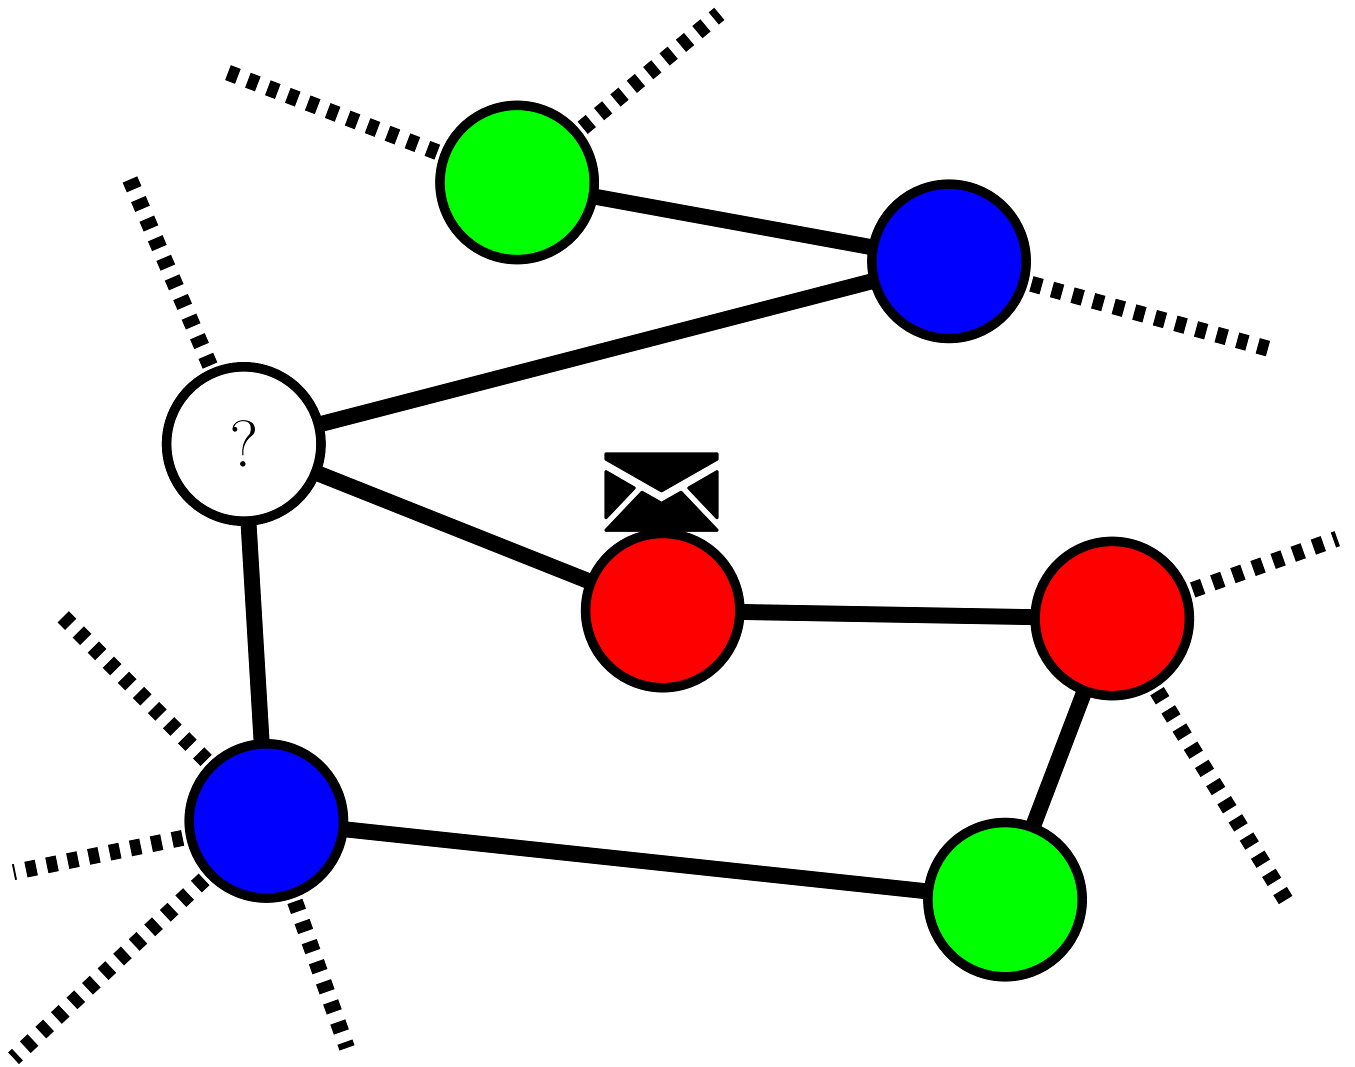

Extension to d-Regular Expanders

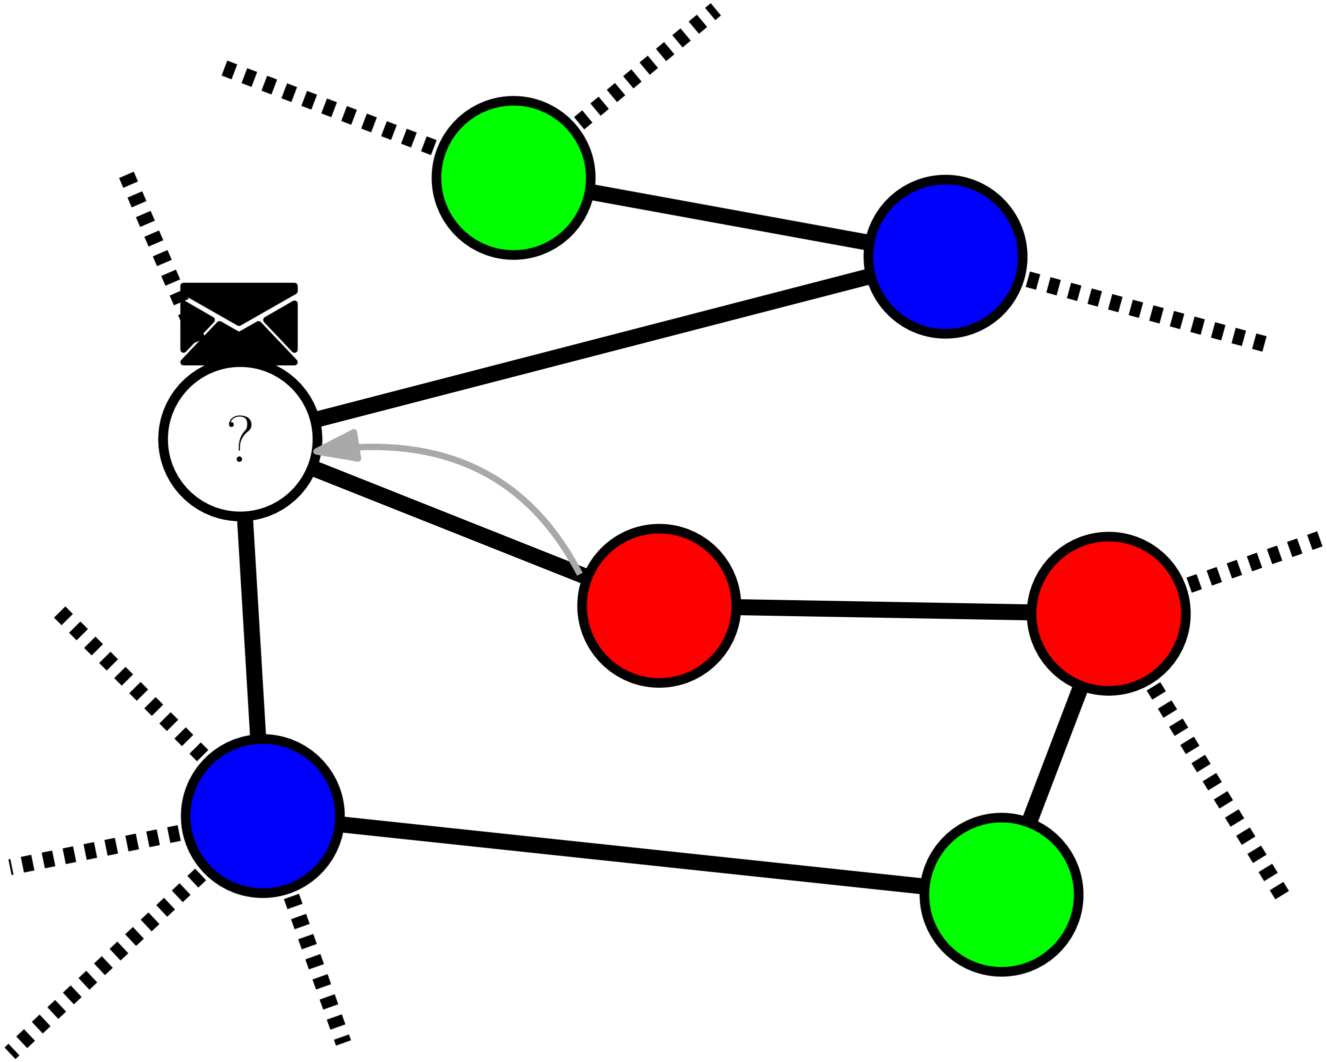

Extension to d-Regular Expanders

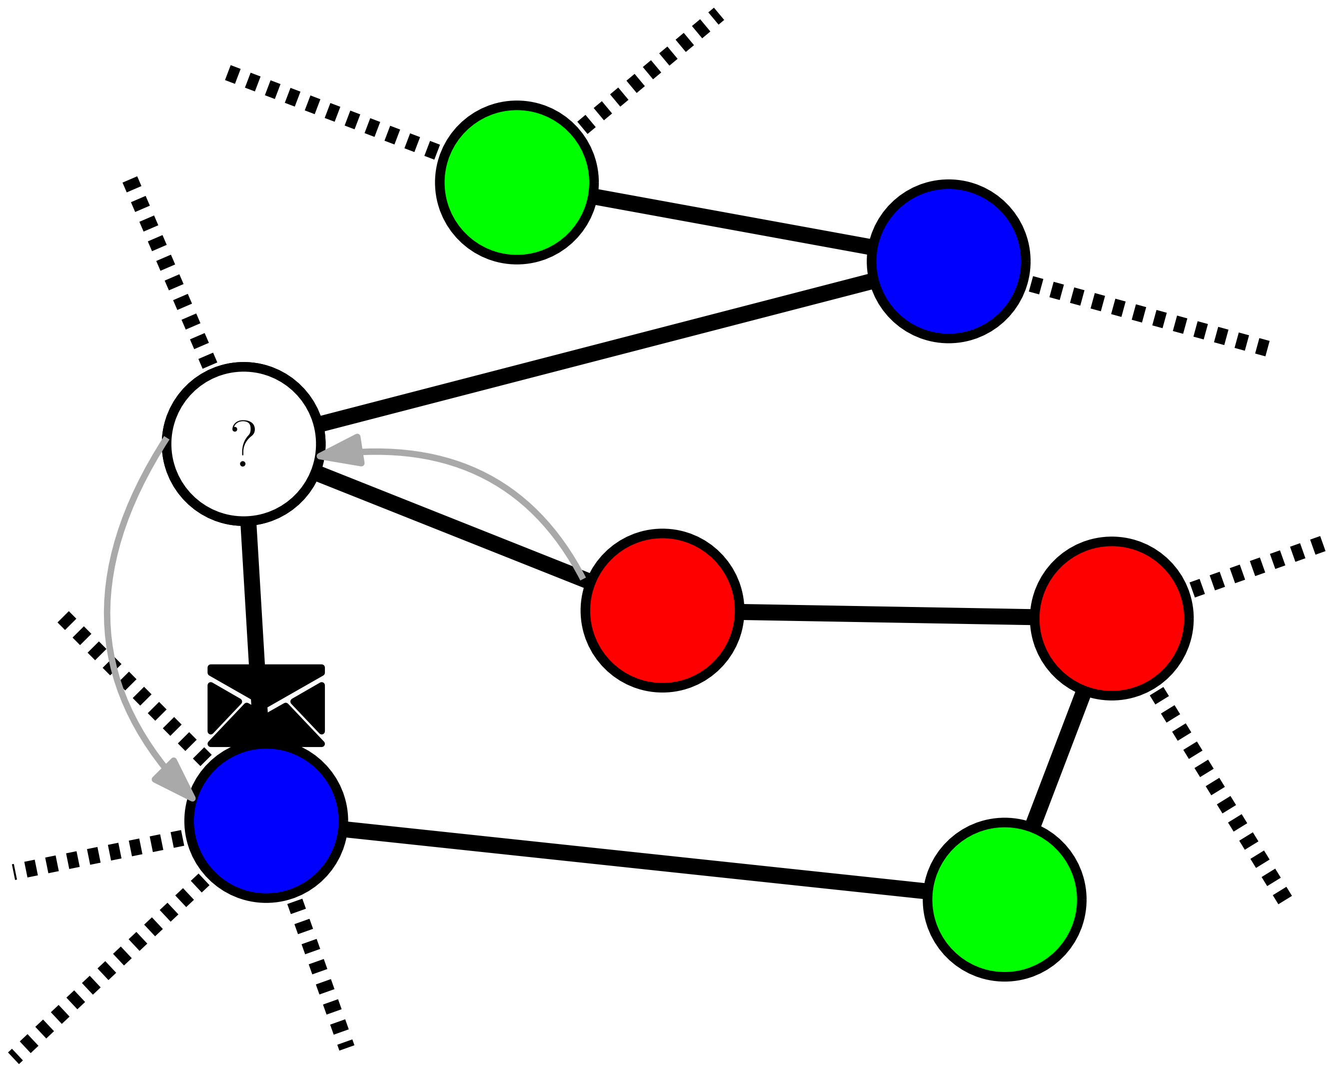

Extension to d-Regular Expanders

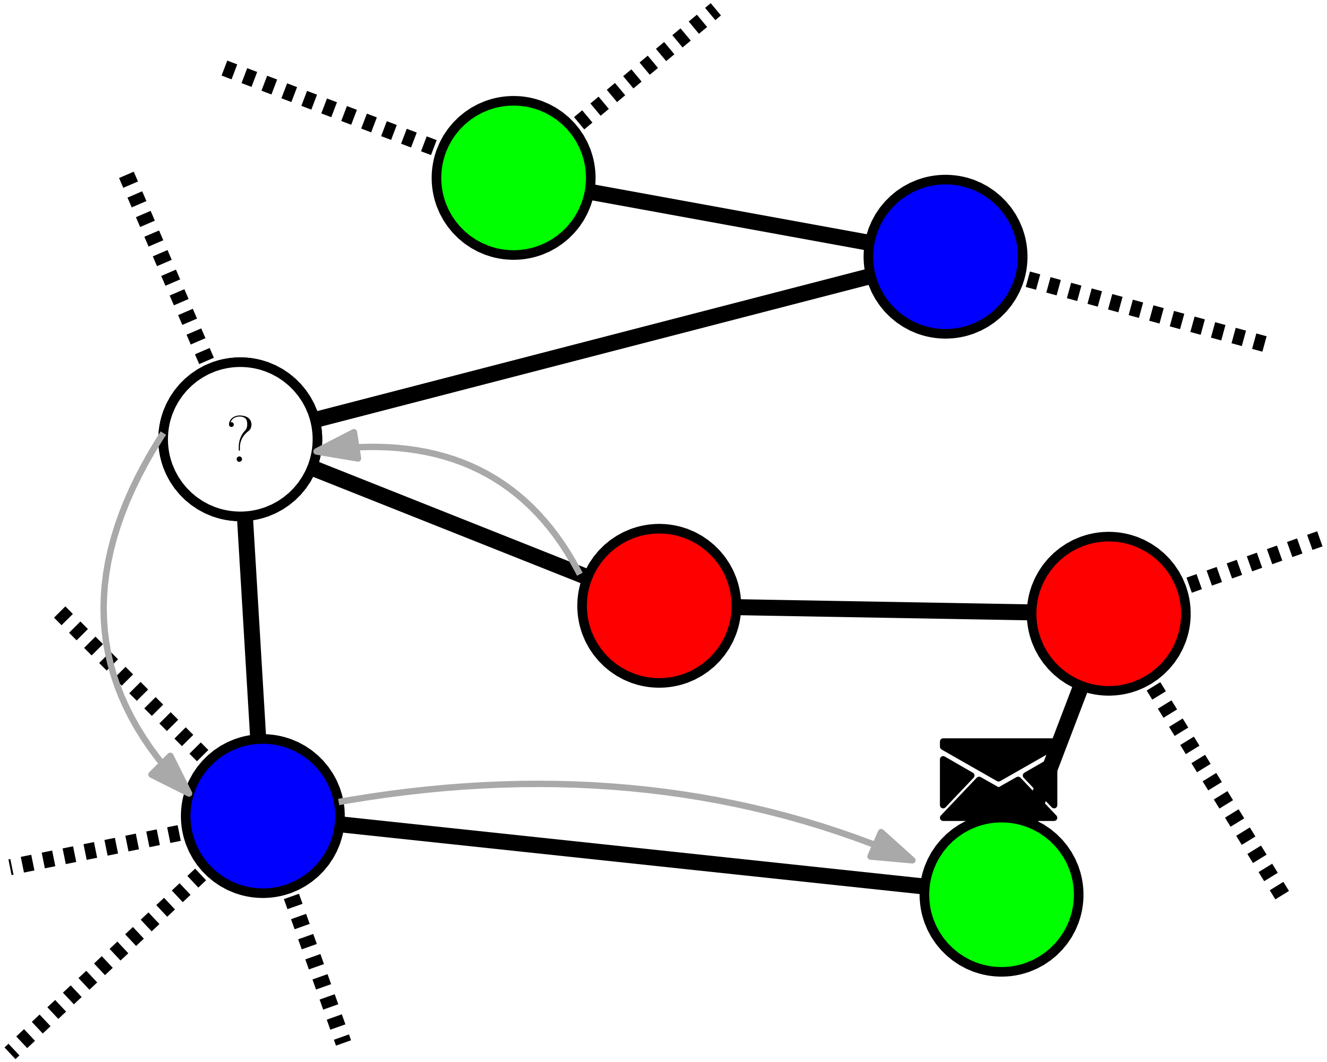

Extension to d-Regular Expanders

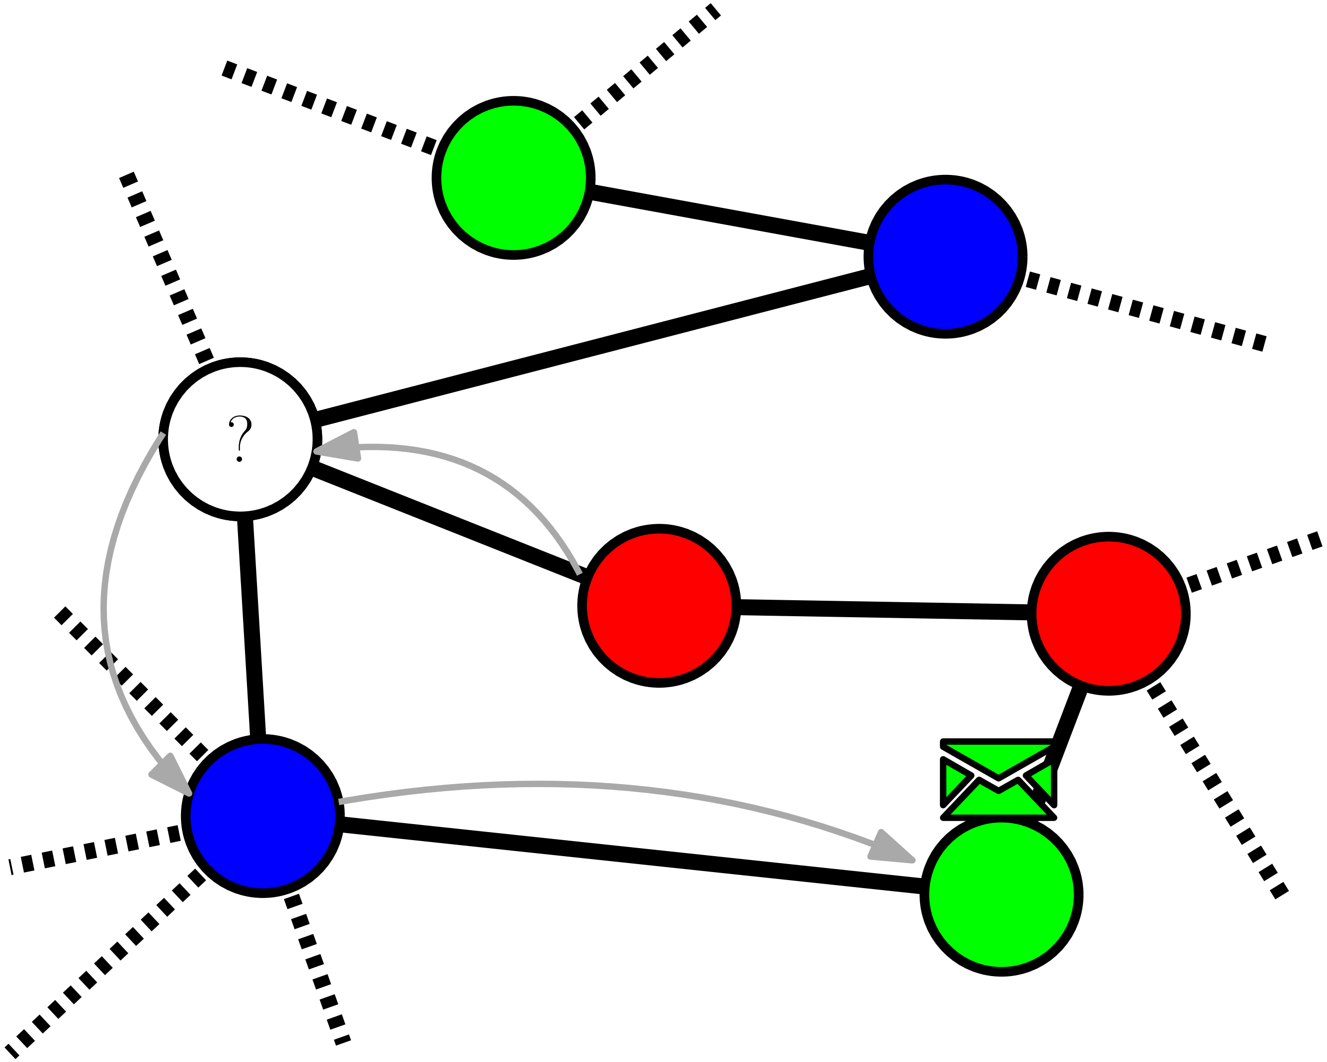

Extension to d-Regular Expanders

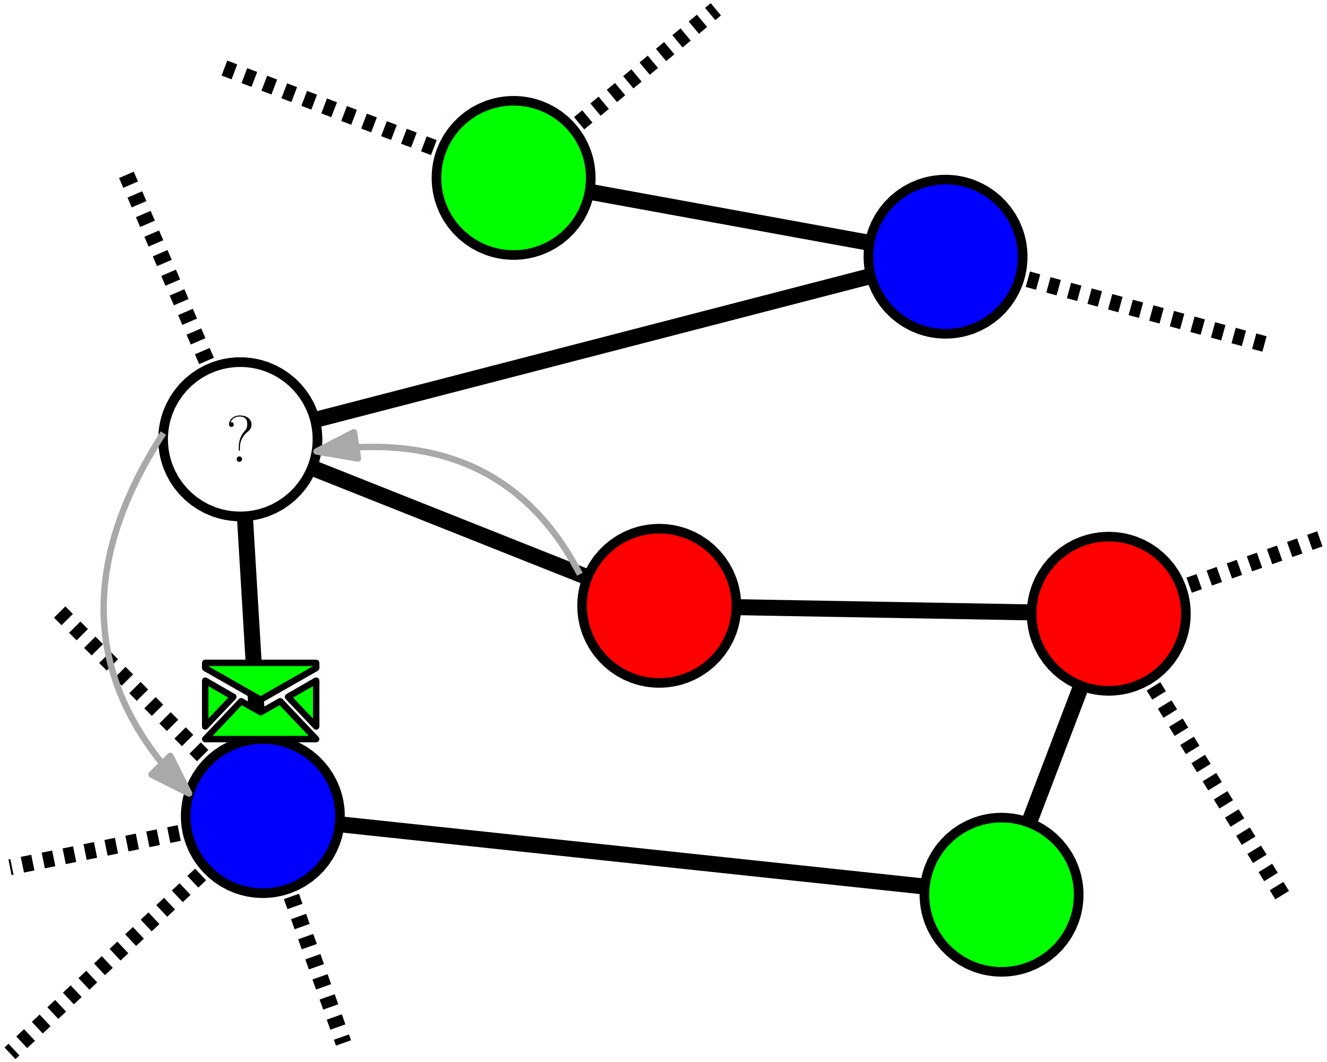

Extension to d-Regular Expanders

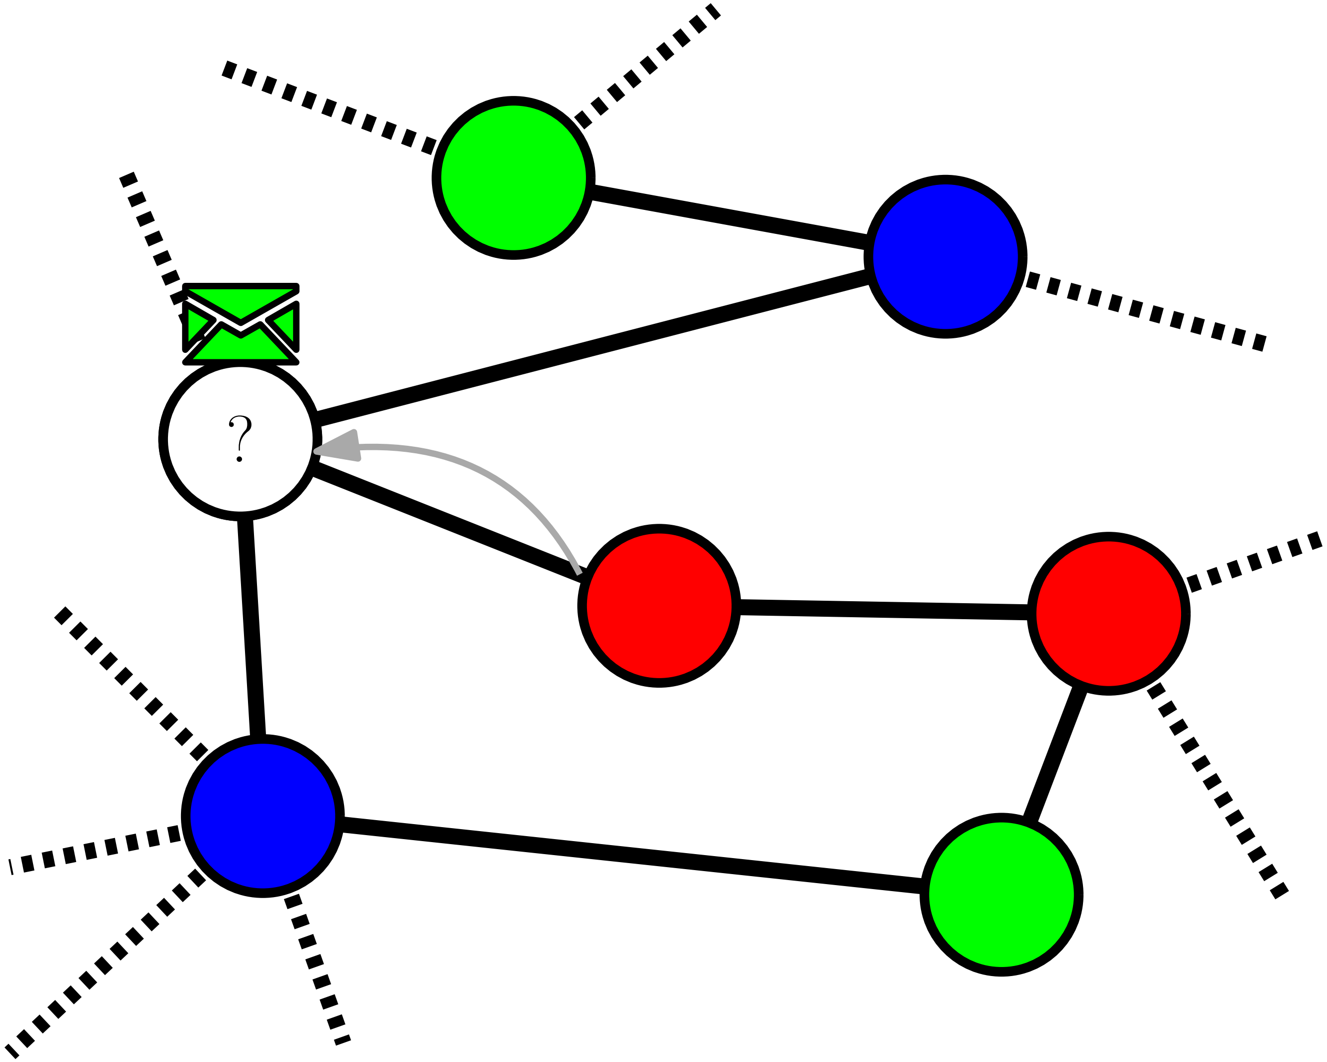

Extension to d-Regular Expanders

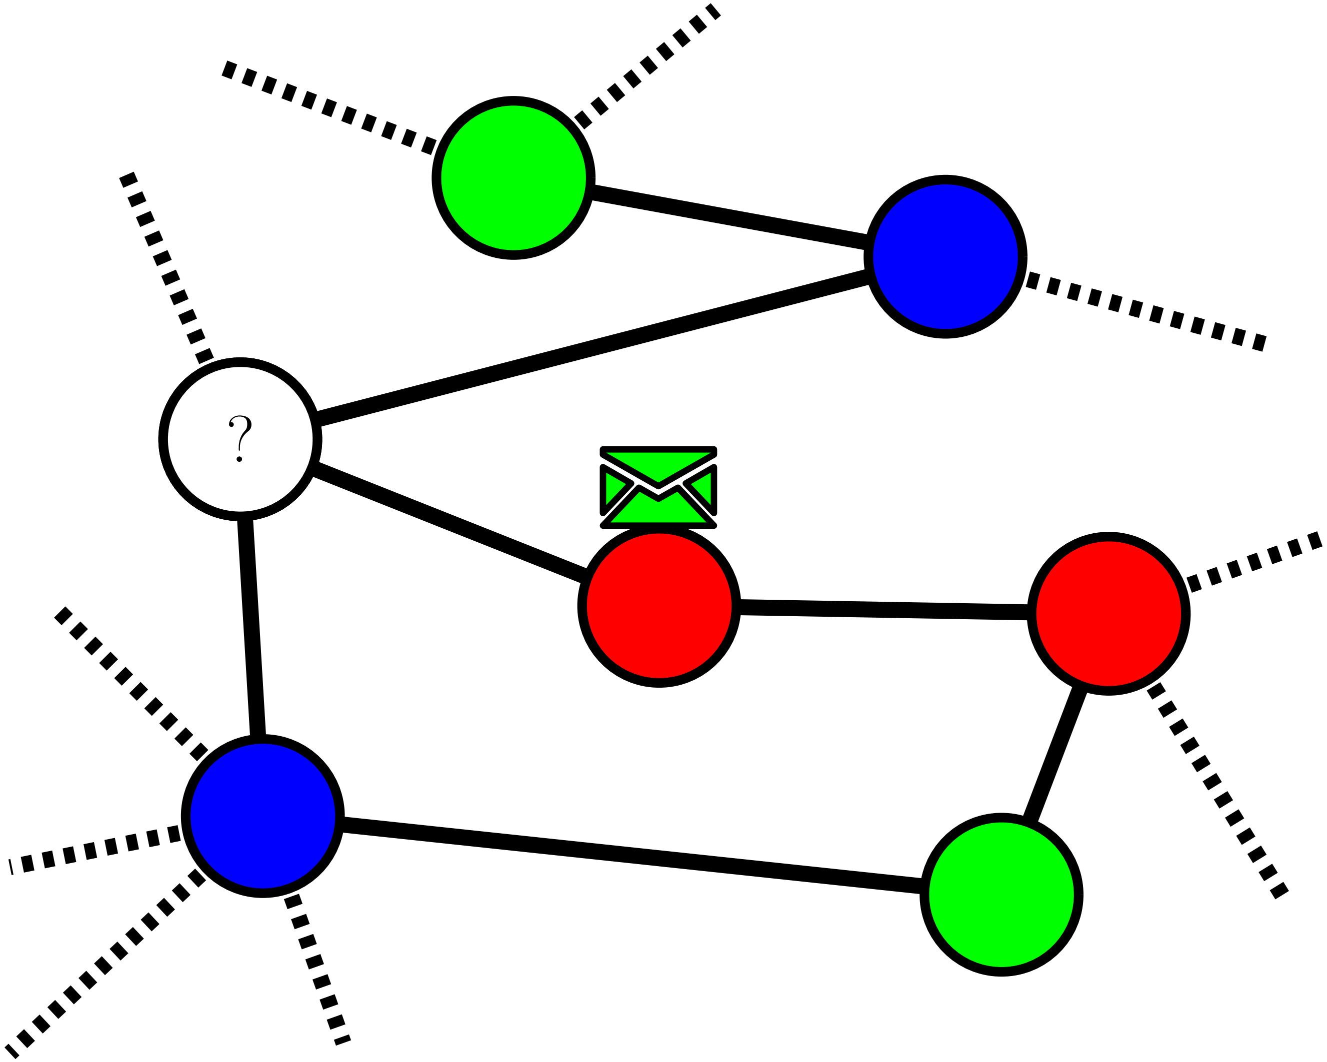

Extension to d-Regular Expanders

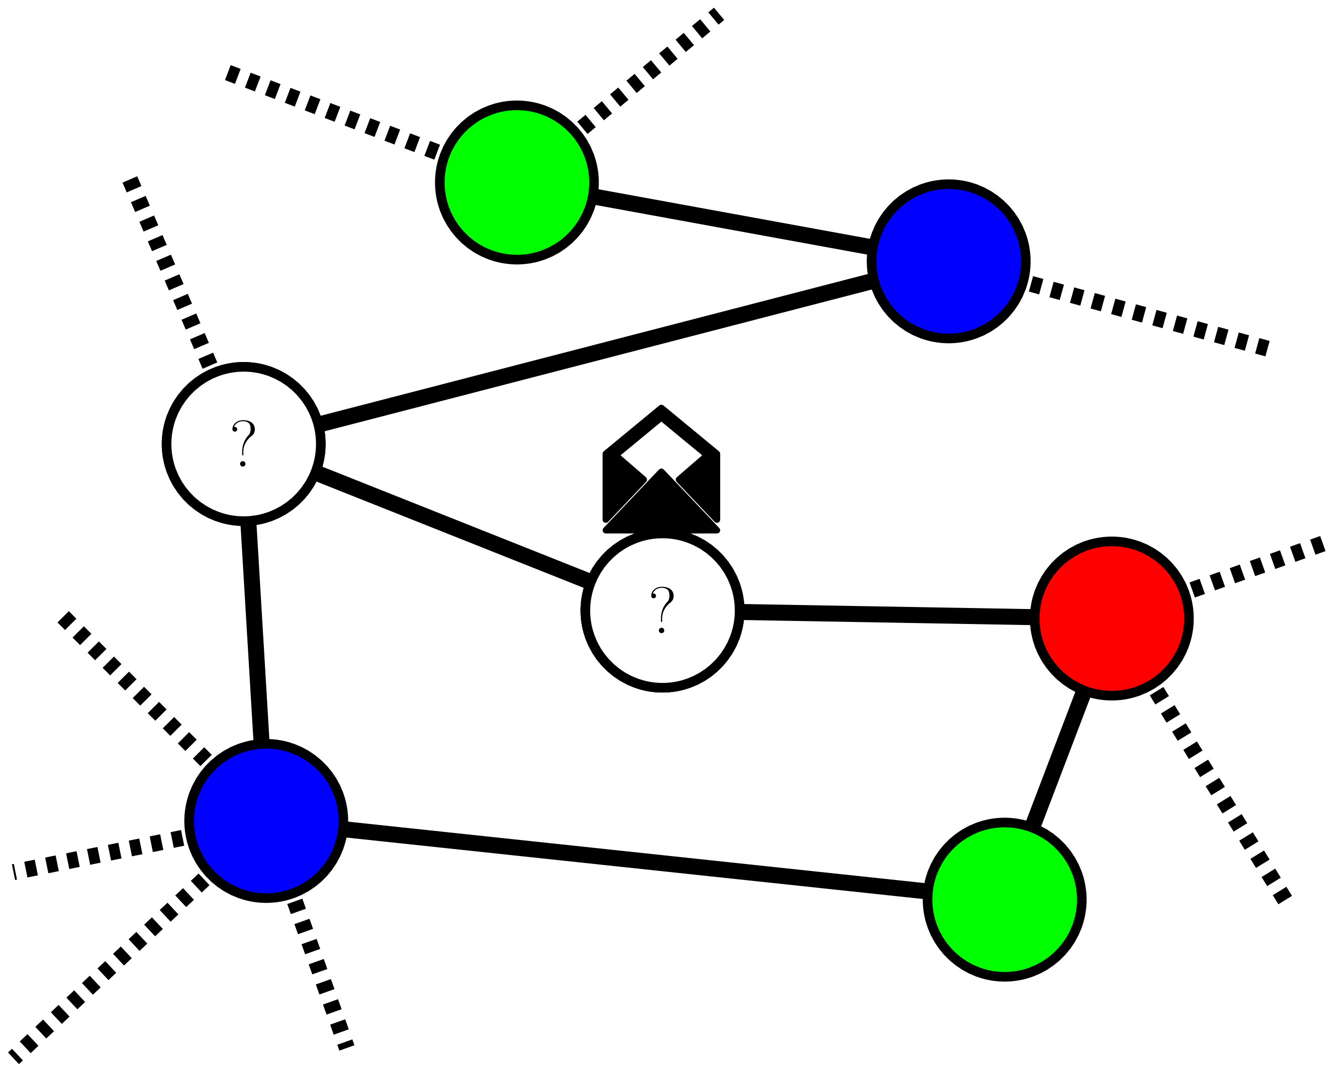

Random Walks in the $\mathcal{GOSSIP}$ Model

Issue. The $\mathcal{GOSSIP}$ model with $O(\text{polylog} n)$ limit on message size: congestion when random walks meet.

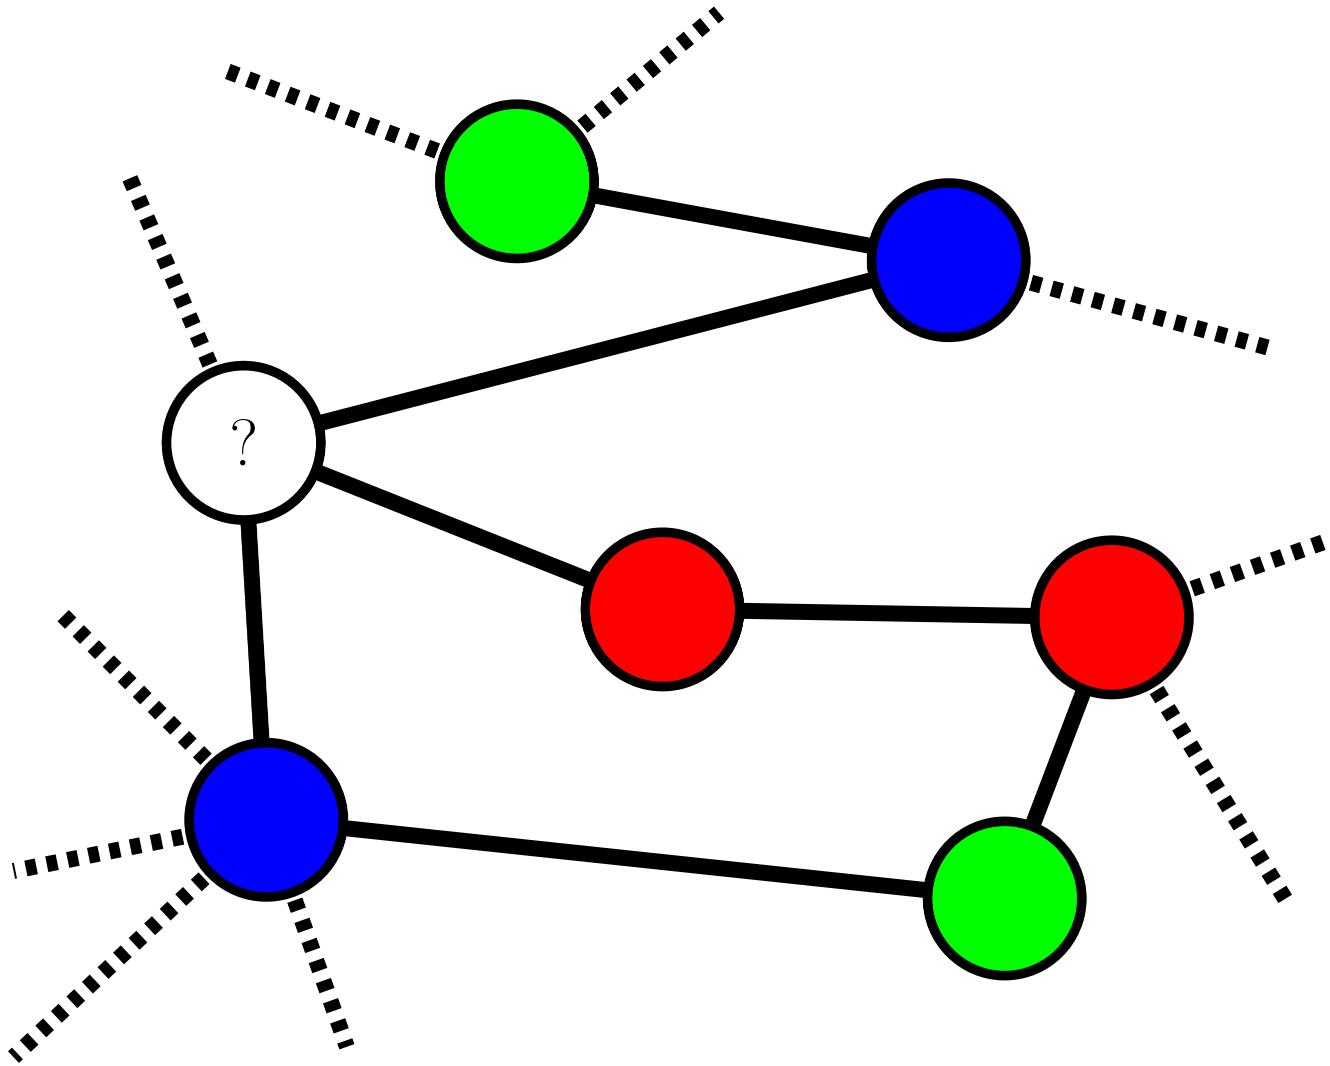

Random Walks in the $\mathcal{GOSSIP}$ Model

Issue. The $\mathcal{GOSSIP}$ model with $O(\text{polylog} n)$ limit on message size: congestion when random walks meet.

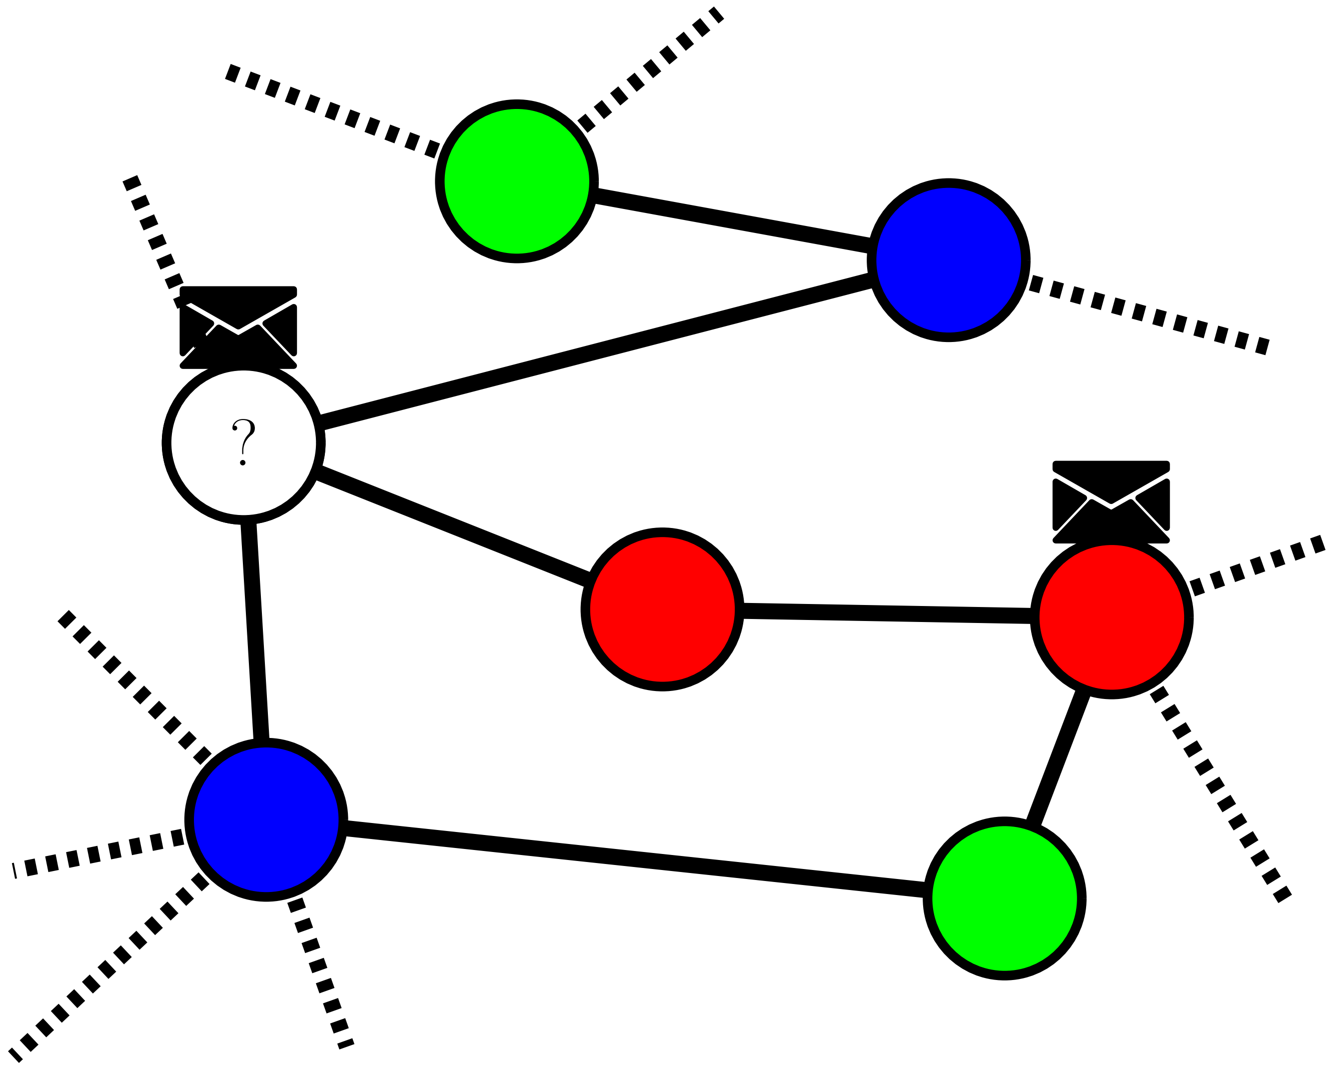

Random Walks in the $\mathcal{GOSSIP}$ Model

Issue. The $\mathcal{GOSSIP}$ model with $O(\text{polylog}\,n)$ limit on message size: congestion when random walks meet.

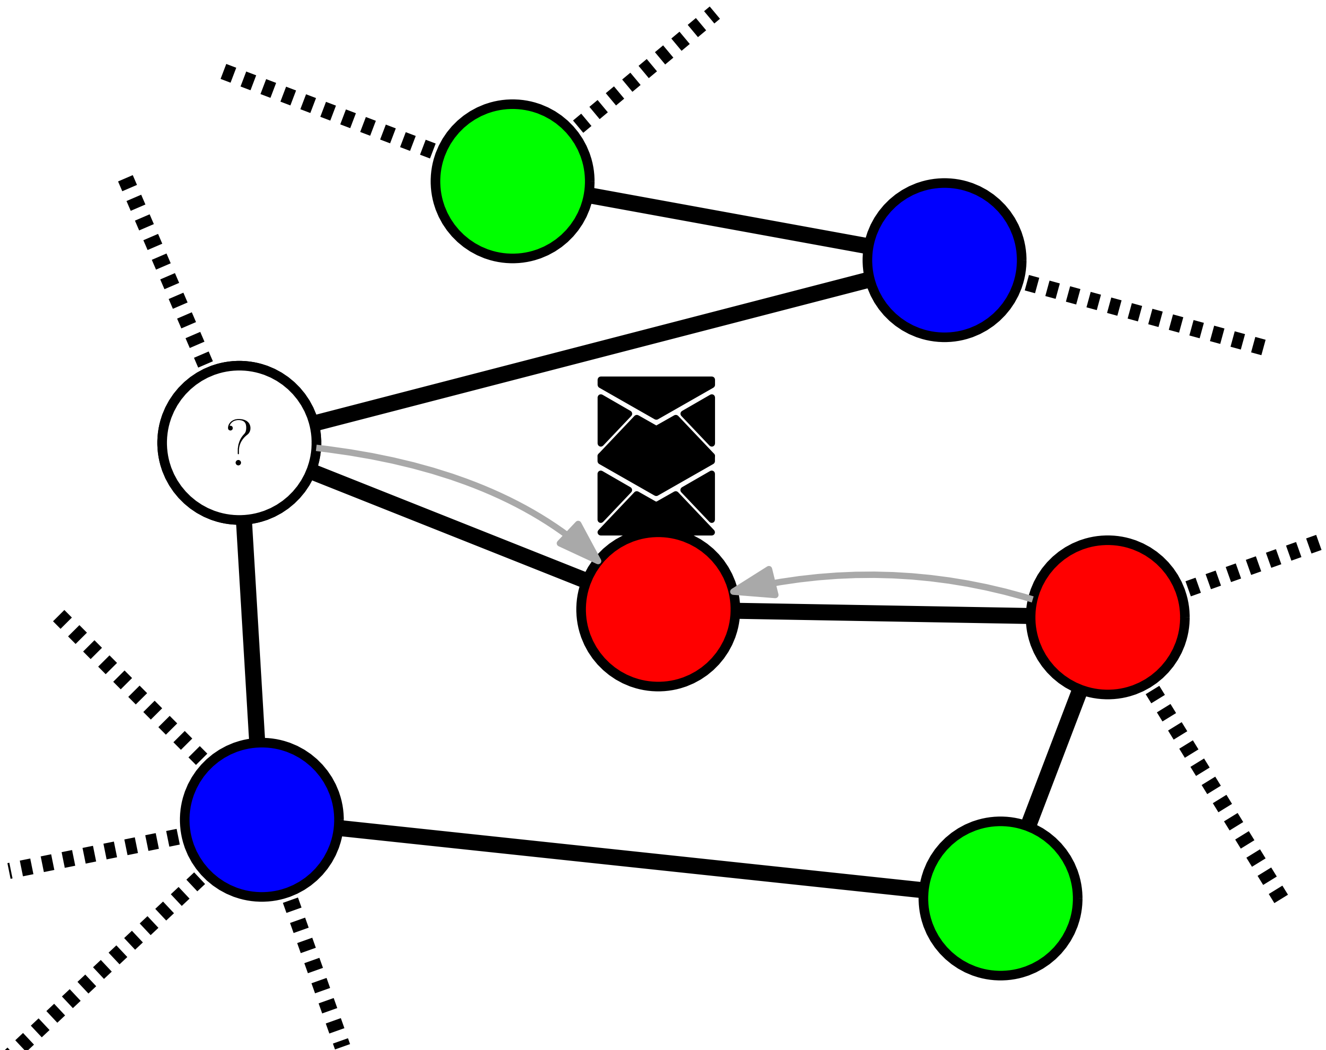

Random Walks in the $\mathcal{GOSSIP}$ Model

Issue. The $\mathcal{GOSSIP}$ model with $O(\text{polylog} n)$ limit on message size: congestion when random walks meet.

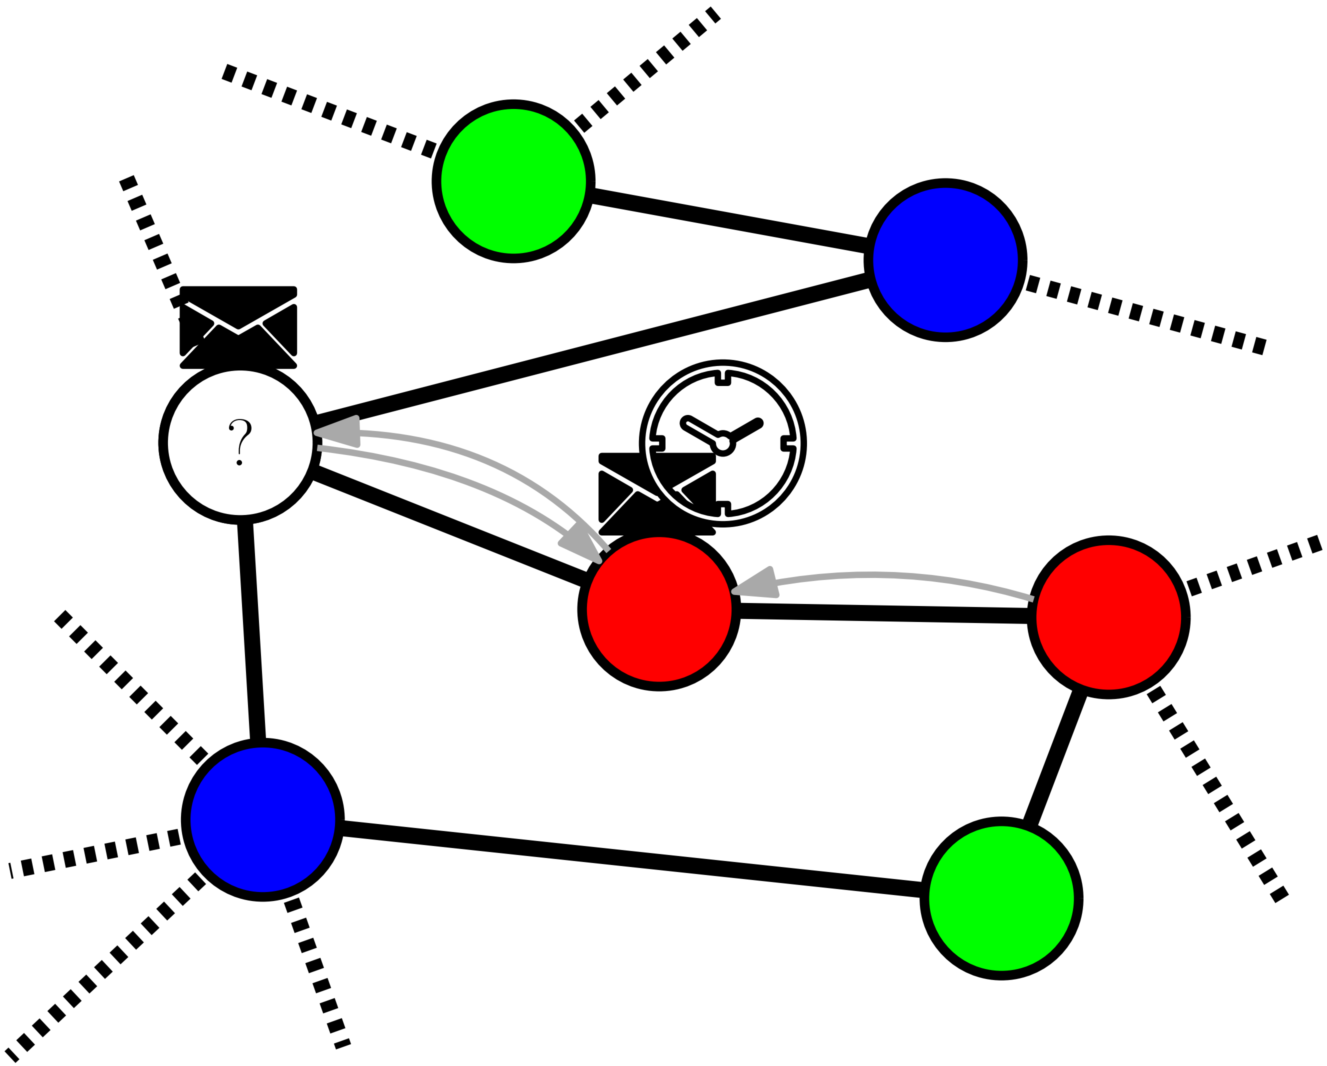

Random Walks in the $\mathcal{GOSSIP}$ Model

Case $n/\log n$ **Tokens**

- Berenbrink, Czyzowicz, Elsässer, Gasieniec (ICALP '10),
- Elsässer, Kaaser (IPDPS '15).

Random Walks in the $\mathcal{GOSSIP}$ Model

Case $n/\log n$ **Tokens**

- Berenbrink, Czyzowicz, Elsässer, Gasieniec (ICALP ’10),
- Elsässer, Kaaser (IPDPS ’15).

Our Case: n **Tokens**

Stochastic dependence: positions of tokens.

Random Walks in the $\mathcal{GOSSIP}$ Model

Case $n/\log n$ Tokens

- Berenbrink, Czyzowicz, Elsässer, Gasieniec (ICALP '10),
- Elsässer, Kaaser (IPDPS '15).

Our Case: n Tokens

Stochastic dependence: positions of tokens.

↓

Coupling: every node always sends a token
(when empty, creates a new one).

Random Walks in the $\mathcal{GOSSIP}$ Model

Case $n/\log n$ **Tokens**

- Berenbrink, Czyzowicz, Elsässer, Gasieniec (ICALP '10),
- Elsässer, Kaaser (IPDPS '15).

Our Case: n **Tokens**

Stochastic dependence: positions of tokens.

↓

(Brutal) **Coupling:** every node always sends a token (when empty, creates a new one).

Random Walks in the $\mathcal{GOSSIP}$ Model

Case $n/\log n$ **Tokens**

- Berenbrink, Czyzowicz, Elsässer, Gasieniec (ICALP '10),
- Elsässer, Kaaser (IPDPS '15).

Our Case: n **Tokens**

Stochastic dependence: positions of tokens.

↓

(Brutal) **Coupling:** every node always sends a token (when empty, creates a new one).

↓

Congestion: in t rounds, at most $\sqrt{t}$ w.h.p.

Random Walks in the $\mathcal{GOSSIP}$ Model

In t rounds, congestion at most $\sqrt{t}$ w.h.p.

Random Walks in the $\mathcal{GOSSIP}$ Model

In t rounds, congestion at most $\sqrt{t}$ w.h.p.

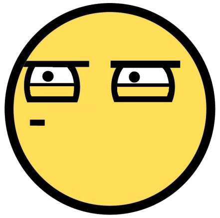

Random Walks in the $\mathcal{GOSSIP}$ Model

Becchetti, Clementi, Natale, Pasquale, Posta.
Self-Stabilizing Repeated Balls-into-Bins. (Submitted).

Random Walks in the $\mathcal{GOSSIP}$ Model

Becchetti, Clementi, Natale, Pasquale, Posta.
Self-Stabilizing Repeated Balls-into-Bins. (Submitted).

↓

Theorem
Consider the "$\mathcal{GOSSIP}$" random walks process on a complete graph with n nodes, and n tokens initially distributed in an arbitrary way.
After $O(n)$ rounds, w.h.p. the congestion is at most $O(\log n)$. Moreover, w.h.p. it keeps below $O(\log n)$ for $t = O(n^c)$ rounds (for any $c > 0$).

Random Walks in the $\mathcal{GOSSIP}$ Model

Becchetti, Clementi, Natale, Pasquale, Posta.
Self-Stabilizing Repeated Balls-into-Bins. (Submitted).

↓

Theorem
Consider the "$\mathcal{GOSSIP}$" random walks process on a complete graph with n nodes, and n tokens initially distributed in an arbitrary way.
After $O(n)$ rounds, w.h.p. the congestion is at most $O(\log n)$. Moreover, w.h.p. it keeps below $O(\log n)$ for $t = O(n^c)$ rounds (for any $c > 0$).

Open Problem: bound the congestion on graphs other than the complete one.

Summary

- md($\mathbf{c}$): global measure of bias, key of the Undecided-State Dynamic.
  $\implies$ Plurality consensus problem with **many** colors.

Summary

- md($\mathbf{c}$): global measure of bias, key of the Undecided-State Dynamic.
  $\implies$ Plurality consensus problem with **many** colors.
- Extension to regular expanders: random walks in the $\mathcal{GOSSIP}$ model.

Summary

- md($\mathbf{c}$): global measure of bias, key of the Undecided-State Dynamic.
  $\implies$ Plurality consensus problem with **many** colors.
- Extension to regular expanders: random walks in the $\mathcal{GOSSIP}$ model.

Open Problems

- md($\mathbf{c}$) $\stackrel{?}{=}$ general time lower bound on the plurality consensus problem for *any* dynamics which uses only $\log k + \Theta(1)$ bits of local memory?

Summary

- md($\mathbf{c}$): global measure of bias, key of the Undecided-State Dynamic.
  $\Longrightarrow$ Plurality consensus problem with **many** colors.
- Extension to regular expanders: random walks in the $\mathcal{GOSSIP}$ model.

Open Problems

- md($\mathbf{c}$) $\stackrel{?}{=}$ general time lower bound on the plurality consensus problem for *any* dynamics which uses only $\log k + \Theta(1)$ bits of local memory?
- Undecided-State Dynamics + sampling via random walks = efficient protocol for regular expander graphs. Similar protocols for other classes of graphs...?

Vielen Dank!