

Supplementary Information

Degree-Corrected Energy Relaxation Bounds for Competitive Recoloring on Complex Networks

Ranjan Veerabhadraswamy

Ajith Jubilson Emerson

Supplementary Note S1: exhaustive finite-instance check

The degree-corrected instability inequality and equilibrium ceiling were checked independently by exhaustive enumeration on every connected non-isomorphic simple graph with two to five vertices. The graph collection was obtained from the NetworkX graph atlas and contained 1, 2, 6, and 21 graphs for orders 2, 3, 4, and 5, respectively. For every graph, every coloring was enumerated for $k \in \{2, 3\}$.

For each coloring c , the check computed the improving set $I(c)$, the conflict energy $\Phi(c)$, and the degree-corrected quantities

$$b_i = \left\lfloor \frac{d_i}{k} \right\rfloor, \quad w_i = d_i - b_i, \quad S = \sum_i b_i, \quad H = \left\lfloor \frac{S}{2} \right\rfloor.$$

It then verified

$$\sum_{i \in I(c)} w_i \geq 2\Phi(c) - S$$

for every configuration and, whenever $I(c) = \emptyset$, verified $\Phi(c) \leq H$. In total, 60 graph-palette cases, 6,440 configurations, and 1,102 equilibrium configurations were checked. No violation was found. This finite enumeration is a software-level sanity check and is not used as a substitute for the proof in the main manuscript.

Supplementary Note S2: scheduler extensions and finite-instance checks

Let activations be independent draws from a fixed positive probability vector p , and let

$$\lambda = \min_{i: d_i > 0} \frac{p_i}{w_i}.$$

For $k \geq 2$, any improving vertex has positive degree. Moreover, $d_i \geq 1$ implies $\lfloor d_i/k \rfloor \leq d_i - 1$ and hence $w_i \geq 1$, so the ratio is defined on every vertex that can belong to $I(c)$. At level $\Phi(c) = H + r$, the structural inequality gives

$$\sum_{i \in I(c)} p_i \geq \lambda \sum_{i \in I(c)} w_i \geq \lambda(2r - \varepsilon).$$

This lower bound cannot exceed one because its leftmost term is an activation probability. Rejected activations do not change the configuration; consequently, conditional on visiting a level, the waiting time to leave it is stochastically dominated by a geometric random variable with success parameter $q_r = \lambda(2r - \varepsilon)$.

For $X_0 \geq 1$ and $\delta \in (0, 1)$, define

$$t_r = \left\lceil \frac{1}{q_r} \log \frac{X_0}{\delta} \right\rceil.$$

Then

$$\Pr(T_r > t_r \mid H + r \text{ is visited}) \leq (1 - q_r)^{t_r} \leq e^{-q_r t_r} \leq \frac{\delta}{X_0}.$$

The residence times need not be independent. Applying the union bound to the events that a visited positive excess level exceeds its allocated residence time gives total failure probability at most δ . Finally,

$$\sum_{r=1}^{X_0} t_r \leq \frac{1}{\lambda} \left(\sum_{r=1}^{X_0} \frac{1}{2r - \varepsilon} \right) \log \frac{X_0}{\delta} + X_0,$$

which is the finite-confidence certificate in the main manuscript.

The exhaustive checker was extended to four fixed positive scheduler distributions: uniform, degree-proportional after adding one, inverse-degree-proportional after adding one, and an index-weighted vector. Across the same 6,440 configurations, it performed 17,200 statewise scheduler-probability checks and 51,600 geometric-tail allocation checks at $\delta \in \{0.1, 0.05, 0.01\}$. No violation was found. These calculations validate the implementation and edge-case handling; the analytical argument remains the basis of the corollaries.

Supplementary Table S1: controlled-cell summaries

The separately supplied file `Supplementary_Table_S1.csv` contains one row for each of the 240 controlled cells in the balanced experiment. Each cell contains 40 independent seeded trajectories. The table reports topology, network size, palette size, target and realized mean degree, sample size, mean initial energy, degree-corrected ceiling H , maximum weight W , initial excess X_0 , threshold time τ_H , its sample standard deviation and descriptive normal interval, certificate B_H , normalized time τ_H/B_H , degree variance, maximum degree, final energy, accepted moves, and activations.

Supplementary Table S2: bootstrap sensitivity analysis

The separately supplied file `Supplementary_Table_S2.csv` gives normal and nonparametric percentile-bootstrap 95% confidence intervals for the global summaries, every topology-palette mean in Fig. 1b, and every marginal mean displayed in Fig. 2a–c. Bootstrap intervals used 10,000 resamples and seed 20260827. The principal global results are summarized below.

Measure	Mean	Normal 95% CI	Bootstrap 95% CI
τ_H	94.5951	[90.6950, 98.4952]	[90.8288, 98.5661]
B_H	31,755.1802	[30,635.6536, 32,874.7069]	[30,665.4151, 32,883.8989]
τ_H/B_H	0.009574	[0.009112, 0.010037]	[0.009109, 0.010048]

Across all reported summaries, the largest absolute relative change in interval width was 2.39%; within Fig. 1b it was 1.54%. No qualitative interpretation in the manuscript changed. Excluding runs for which both B_H and τ_H were zero did not materially alter the normalized-time results. The intervals are descriptive uncertainty summaries; no null-hypothesis significance tests were performed.

Reproducibility

The simulation and replay data, derived tables, and figure-generation scripts are publicly available at <https://github.com/Ranjan-V/Competitive-Graph-Recoloring/tree/v1.1-scientific-reports-upgrade>. For enquiries about data access, contact the corresponding author, Ajith Jubilson Emerson (ajith.jubilson@vitap.ac.in).