EXTINCTION SETS OF THE WHEEL GRAPHS IN A GRAPH-BASED MODEL OF HIV INFECTION

KOREA SUPERINTELLIGENCE LABS

Abstract. Mukwembi's graph-based model of HIV infection is a deterministic three-state cellular automaton on a graph G, driven by one integer parameter R, the rate at which depleted cells are replenished. Its *extinction set* $\mathcal{E}(G)$ is the set of those R for which the infection dies out from every admissible initial state. Espinosa-García, Figueroa, Fresán-Figueroa, Maldonado and Sánchez-Solís computed $\mathcal{E}(W_n)$ exhaustively for the wheels $W_n = K_1 \vee C_{n-1}$ with $4 \le n \le 10$, sampled it for $11 \le n \le 26$, and conjectured that $\mathcal{E}(W_n) = \{3\} \cup \{R : R \ge n-1\}$ for every even $n \ge 12$ and $\mathcal{E}(W_n) = \{4\} \cup \{R : R \ge n-1\}$ for every odd $n \ge 17$. We determine $\mathcal{E}(W_n)$ exhaustively for $4 \le n \le 31$ and for *every* positive R, and refute the conjecture in both of its halves: the least counterexamples are $n = 18$, where $4 \in \mathcal{E}(W_{18})$, and $n = 23$, where $5 \in \mathcal{E}(W_{23})$, the conjectured description is exactly right at every smaller n that it covers, and from $n = 27$ on it fails at *every* n we can reach. Seven of the sixteen sampled rows of the published table are incomplete, although the two extinction parameters read off that table are unaffected. On the whole verified range the data obey one linear law: for $4 \le R \le n-2$ one has $R \in \mathcal{E}(W_n)$ if and only if $n \ge 5R-3$. Five thresholds sit on that line, namely $R = 3, \dots, 7$ at $n = 12, 17, 22, 27, 32$, and the last two were predicted before they were computed. Where the slope 5 comes from we do not know, and we pose that as the question this leaves open. Every theorem below is machine-checked in Lean 4. What carries the search past $n = 26$ is an orbit quotient that is itself part of the formal development: the dihedral group D_{n-1} acts on W_n by automorphisms, the transition map is equivariant for it, and a descent argument — using neither a canonical-form theory nor the order of the group — restricts each sweep to the orbit minima, worth a measured factor 32 to 38 here.

1. Introduction

1.1. **The model.** In 2008 Mukwembi [2] proposed a graph-theoretic caricature of HIV infection in which the cells of the immune system are the vertices of a graph, infection spreads along edges, an infected cell dies after one time unit, and a dead cell is replenished — unless it is under sufficiently heavy infectious pressure, in which case the replacement cell is itself born infected. The whole model is one local rule with one integer parameter.

Let G be a finite simple graph. A *state* is a map $f\colon V(G) \to \{0,1,2\}$, where $0,1,2$ are read as *healthy*, *infected* and *dead*. For a state f and a vertex v write

$$d_f(v) \;=\; \bigl|\{\, u \in N(v) \;:\; f(u) = 1 \,\}\bigr|$$

for the number of infected neighbours of v. Given a positive integer R, the *replacement parameter*, one time step replaces f by $\Phi_R(f)$, where

$$\Phi_R(f)(v) \;=\; \begin{cases} 0, & f(v) = 0 \text{ and } d_f(v) = 0, \\ 1, & f(v) = 0 \text{ and } d_f(v) \ge 1, \\ 2, & f(v) = 1, \\ 0, & f(v) = 2 \text{ and } d_f(v) < R, \\ 1, & f(v) = 2 \text{ and } d_f(v) \ge R. \end{cases} \tag{1}$$

A healthy cell with an infected neighbour is infected; an infected cell dies; a dead cell is replaced by a healthy one, except that under pressure $d_f(v) \ge R$ the replacement is born infected. Large R therefore means *safe* replenishment, and the interesting question is how small R may be before the infection can persist forever.

An *admissible initial state* is a state with no dead cell, that is, a map $f_0 \colon V(G) \to \{0,1\}$; a graph of order n has exactly 2^n of them. Write $\mathbf{0}$ for the all-healthy state, which (1) fixes.

1.2. **Extinction sets.** Espinosa-García, Figueroa, Fresán-Figueroa, Maldonado and Sánchez-Solís [1] made the parameter the object of study. Their *extinction set* of G is

$$\mathcal{E}(G) \;=\; \bigl\{\, R \in \mathbb{Z}^+ \;:\; \text{for every admissible } f_0 \text{ there is } t \geq 0 \text{ with } \Phi_R^t(f_0) = \mathbf{0} \,\bigr\}, \tag{2}$$

its complement in $\mathbb{Z}^+$ is the *latency set* $\mathcal{L}(G)$, and the two extinction parameters are

$$\mathrm{hiv}(G) \;=\; \min \mathcal{E}(G), \qquad \mathrm{HIV}(G) \;=\; 1 + \max \mathcal{L}(G).$$

Two published facts bracket $\mathcal{E}(G)$ for a graph of order n. Mukwembi's theorem [2], restated as [1, Theorem 2.1], says that if G has order $n \geq 3$ then every admissible initial state reaches $\mathbf{0}$ at $R = n-1$, that is, $n-1 \in \mathcal{E}(G)$; and [1, Theorem 2.8] gives $2 \leq \mathrm{hiv}(G) \leq \mathrm{HIV}(G) \leq \Delta(G)+1$ for every graph with at least one edge, so, since $\Delta(G) \leq n-1$, every $R \geq n$ lies in $\mathcal{E}(G)$ as well. (All numbered references to [1] are to its v1, which is the only version as of 2026-08-30.) Hence

$$\{\, R : R \geq n-1 \,\} \;\subseteq\; \mathcal{E}(G) \tag{3}$$

always, and all the content of $\mathcal{E}(G)$ sits in the window $1 \leq R \leq n-2$. What makes that window hard is an observation of [1]: *extinction is not monotone in R*. One might expect that raising the replacement parameter can only help, but it does not, and isolated small values of R can lie in $\mathcal{E}(G)$ with every value between them and $n-1$ latent. So $\mathcal{E}(G)$ is genuinely a set, not an interval, and $\mathrm{hiv}(G)$ does not determine it.

1.3. **The wheels.** Section 4 of [1] derives, in its own words, "closed formulas for cycles, complete graphs, and complete bipartite graphs"; these are [1, Theorems 4.1, 4.2 and 4.3], giving, for every $n \geq 3$, $\mathrm{HIV}(C_n) = \mathrm{hiv}(C_n)$, which is 3 for even n and 2 for odd n, and $\mathrm{HIV}(K_n) = \mathrm{hiv}(K_n) = \lceil (n+1)/2 \rceil$; and, for all positive $m \leq n$, $\mathrm{HIV}(K_{m,n}) = \mathrm{hiv}(K_{m,n}) = \max\{m+1, \lfloor n/2 \rfloor + 1\}$. In each of the three the two parameters coincide, so each result determines the extinction set outright, as the single interval $\{R : R \geq \mathrm{hiv}(G)\}$. For the wheels

$$W_n \;=\; K_1 \vee C_{n-1} \qquad (n \geq 4),$$

it reports a computation instead. Here W_n is a rim cycle on $n-1$ vertices together with a hub joined to all of it; it has order n, its rim vertices have degree 3 and its hub has degree $n-1$. In Section 4 of [1], for $4 \leq n \leq 10$ the authors examined all 2^n admissible initial states for each relevant R, and recorded the result as exhaustive. For $11 \leq n \leq 26$ they carried out, in their words, "additional computational experiments", and state plainly that the values reported in their Table 3 on that range "should therefore be regarded as experimental evidence rather than as an exhaustive determination of the extinction sets". On the strength of that table they pose their only conjecture.

Conjecture 1.1 ([1, Conjecture 1])**.** The extinction sets of the wheel graphs satisfy

$$\mathcal{E}(W_n) = \{3\} \cup \{R \in \mathbb{Z}^+ : R \geq n-1\} \quad \text{for every even } n \geq 12,$$
$$\mathcal{E}(W_n) = \{4\} \cup \{R \in \mathbb{Z}^+ : R \geq n-1\} \quad \text{for every odd } n \geq 17.$$

Consequently $\mathrm{hiv}(W_n) = 3$ for even $n \geq 12$, $\mathrm{hiv}(W_n) = 4$ for odd $n \geq 17$, and $\mathrm{HIV}(W_n) = n-1$.

1.4. **Results.** This paper determines $\mathcal{E}(W_n)$ exhaustively for $4 \leq n \leq 31$, refutes Conjecture 1.1 in both halves and shows that from $n = 27$ on it fails at every n, corrects seven rows of the sampled table, and replaces the conjectured description by one that fits all of the data.

An exhaustive determination for every R (Section 5). For each n with $4 \leq n \leq 31$ we decide membership in $\mathcal{E}(W_n)$ for *every* positive integer R, not merely for R in a finite window. The two halves of the problem are finite for different reasons: the sweep over initial states is finite because there are 2^n of them, and the sweep over the parameter is finite because of an R-independence reduction (Proposition 3.3) — on a graph of order n the infected-neighbour count never reaches n, so the clause $d_f(v) < R$ of (1) is vacuously true once $R \geq n$ and the transition map stops depending on R altogether. The entire infinite tail is therefore decided by the single value $R = n$. This is the technical device that turns each row of the table from a statement about $\mathcal{E}(W_n) \cap [1, n]$ into a statement about $\mathcal{E}(W_n)$.

An orbit quotient, proved (Section 4). What makes a sweep over 2^{31} initial states affordable is the symmetry of the wheel: the dihedral group D_{n-1} acts on W_n by automorphisms, the transition map commutes with that action and fixes $\mathbf{0}$, so extinction is constant on orbits and one initial state per orbit suffices. Turning that sentence into a verified reduction is the content of Section 4, and the interesting part is that it needs no theory of canonical forms: soundness of the filter "keep only the least element of each orbit" is a descent (Proposition 4.2), so no completeness lemma is required and the order of the group is never used. Measured on this problem the reduction is worth a factor 32 to 38 (Proposition 4.4), against a theoretical ceiling of $2(n-1) \in \{46, 48, 50\}$.

The conjecture is false (Section 6). Its even half fails first at $n = 18$, where $4 \in \mathcal{E}(W_{18})$ although $4 \neq 3$ and $4 < 17$; its odd half fails first at $n = 23$, where $5 \in \mathcal{E}(W_{23})$ although $5 \neq 4$ and $5 < 22$. Both counterexamples are minimal in their parity class, and minimality is proved rather than assumed: the conjectured description is exactly correct at $n = 12, 14, 16$ and at $n = 17, 19, 21$. Beyond that the failure is not sporadic: at every n with $27 \leq n \leq 31$ one has $5 \in \mathcal{E}(W_n)$ while $5 \notin \{3, 4\}$ and $5 < n-1$, so both halves of Conjecture 1.1 fail at every n in that range (Theorem 6.4).

Seven table rows are incomplete (Section 7). Of the sixteen rows of Table 3 of [1] that the source itself flags as sampled, nine are confirmed exactly and seven — $n = 18, 20, 22, 23, 24, 25, 26$ — are missing one or two isolated values of R. The missing values are never the least element of $\mathcal{E}(W_n)$ and never exceed $n-2$, so the hiv and HIV columns of that table survive untouched (Corollary 7.2); it is the extinction *set* that is under-reported. We also explain, in Remark 7.3, why the source's exhaustive range stops where it does.

A linear law (Section 8). All the exhaustive data on $11 \leq n \leq 31$ are described by

$$\mathcal{E}(W_n) \ = \ \{\, 3 : n \text{ even} \,\} \ \cup \ \{\, R : 4 \leq R \leq \lfloor (n+3)/5 \rfloor \,\} \ \cup \ \{\, R : R \geq n-1 \,\},$$

equivalently: for $4 \leq R \leq n-2$ — that is, for every $R \geq 4$ below the tail (3) — $R \in \mathcal{E}(W_n)$ if and only if $n \geq 5R - 3$ (Theorem 8.1). Conjecture 1.1 keeps only the *first* element of the middle block and drops the rest, which is exactly why it is right while that block is small and wrong as soon as it is not. Five thresholds lie on the line $5R - 3$, namely $R = 3, 4, 5, 6, 7$ at $n = 12, 17, 22, 27, 32$ (Theorems 8.1 and 8.2); the last two were computed only after the pattern had been read off the first three, so they are predictions the search confirmed. Where the slope 5 comes from we cannot say, and we pose that as the open question (Question 8.3).

Every membership and non-membership claim stated as a theorem below rests on an exhaustive finite search, and every such search has been formally verified; the residue of the $n = 32$ row in Remark 8.5 is the one computation here that is not, and it carries no theorem. Section 10 says exactly which searches were run, which of them are replayed inside the proof kernel, and which are delegated to compiled evaluation.

1.5. **What is new in this version.** The first version of this paper determined $\mathcal{E}(W_n)$ for $4 \leq n \leq 26$ by brute-force sweeps, and reported the rows $27 \leq n \leq 32$ as unverified data. Section 4 is new, and with it the rows $27 \leq n \leq 31$ (Theorems 5.5 and 5.6) and the membership $7 \in \mathcal{E}(W_{32})$ carrying the fifth threshold (Theorem 8.2); only the remaining six entries of the $n = 32$ row are still data (Remark 8.5). Consequently Theorem 8.1 is verified on $11 \leq n \leq 31$ rather than on $11 \leq n \leq 26$, the systematic failure of Conjecture 1.1 past $n = 26$ (Theorem 6.4) can be stated at all, and both predicted thresholds are theorems. No verdict of the first version is withdrawn, and its range was re-derived through the new machinery as a check (Remark 5.8). One statement is corrected rather than extended: the first version wrote the linear law (5) as an equivalence for all $R \geq 4$, which the tail (3) falsifies at $R = n-1$ and $R = n$; it holds, and was only ever verified, for $4 \leq R \leq n-2$, and is stated that way here. The extinction sets themselves are unchanged.

2. Preliminaries

2.1. **States, orbits, extinction.** Throughout, G is a finite simple graph of order $n \geq 4$ with $V(G) = \{0, 1, \ldots, n-1\}$, and Φ_R is the map (1) on states of G. We abbreviate

$$\mathrm{Ext}_R(f_0) \quad :\Longleftrightarrow \quad \exists\, t \geq 0 :\ \Phi_R^t(f_0) = \mathbf{0},$$

so that (2) reads $R \in \mathcal{E}(G) \iff \mathrm{Ext}_R(f_0)$ for all 2^n admissible f_0. Since the state space is finite and Φ_R is deterministic, every orbit is eventually periodic, and because $\mathbf{0}$ is a fixed point, $\mathrm{Ext}_R(f_0)$ holds exactly when the eventual cycle of f_0 is the fixed point $\{\mathbf{0}\}$. We use both descriptions: the first gives the shape of the positive certificate, the second the shape of the negative one.

We encode a state of a graph of order n by the pair (I, D) of n-bit masks of its infected and its dead vertices, healthy meaning that neither bit is set; an admissible initial state is then a pair $(I, 0)$ with $I < 2^n$, and the 2^n of them are indexed by I. This encoding is what makes the searches below cheap: a rim vertex of W_n has exactly three neighbours, so all its threshold tests in (1) are word operations, and the hub's d_f is a population count. A mask is one machine word — of 32 bits in the searches for $n \leq 26$ and of 64 bits in the symmetry-reduced searches of Section 4 — so the encoding is faithful as long as n does not exceed the word length, which on the range treated here it does not.

2.2. **The wheels.** We fix the labelling $V(W_n) = \{0, \ldots, n-1\}$ with $0, \ldots, n-2$ the rim in cyclic order and $n-1$ the hub, so that the neighbour lists are

$$N(i) = \{\, i-1 \bmod (n-1),\ i+1 \bmod (n-1),\ n-1 \,\}\ (0 \leq i \leq n-2), \qquad N(n-1) = \{0, \ldots, n-2\}.$$

For $n \geq 4$ the three entries of a rim list are pairwise distinct. The dihedral group D_{n-1}, acting on the rim by rotations and reflections and fixing the hub, is a subgroup of $\mathrm{Aut}(W_n)$; Section 4 makes that the working instrument of this paper.

2.3. **The dynamics on a wheel, written out.** It is worth having (1) in the form in which it is actually iterated. Identify a state f of W_n with the triple (I, D, h) in which $I, D \subseteq \{0, \ldots, n-2\}$ are the infected and the dead rim vertices and $h = f(n-1) \in \{0, 1, 2\}$ is the value at the hub, and write σ for the rotation $i \mapsto i + 1 \bmod (n-1)$ of the rim.

Lemma 2.1. *Let $f = (I, D, h)$ be a state of W_n and let $H = \{0, \ldots, n-2\} \setminus (I \cup D)$ be the healthy part of the rim. Then $d_f(v) = |\{v-1, v+1\} \cap I| + [\, h = 1\,]$ for a rim vertex v, and $d_f(n-1) = |I|$. Consequently $\Phi_R(f) = (I', D', h')$ with*

$$I' \;=\; \big(H \cap T_1\big) \;\cup\; \big(D \cap T_R\big), \qquad D' \;=\; I, \qquad h' \;=\; \begin{cases} 2, & h = 1, \\ [\,|I| \geq 1\,], & h = 0, \\ [\,|I| \geq R\,], & h = 2, \end{cases}$$

where $T_k = \{v \text{ on the rim} : d_f(v) \geq k\}$; explicitly $T_1 = \sigma(I) \cup \sigma^{-1}(I)$ if $h \neq 1$ and T_1 is the whole rim if $h = 1$, and $T_k = \emptyset$ for $k \geq 4$.

Proof. The neighbours of a rim vertex v are $v \pm 1$ on the rim and the hub, and the neighbours of the hub are the whole rim; substituting these into (1) and collecting the five clauses by the value of f gives the display. The description of T_k is the observation that a rim vertex has three neighbours, so $d_f(v) \leq 3$. □

Two things follow from Lemma 2.1. One is computational: each of T_1, T_2, T_3 is a Boolean combination of I, its two rotations and the bit $[h = 1]$, so a step of (1) on W_n costs a constant number of machine-word operations plus one population count for the hub. This is why sweeps over 2^{31} initial states are feasible at all. The other is structural, and it organises the whole answer.

Lemma 2.2. *Let $n \geq 4$ and $R \geq 4$. Then no dead rim vertex of W_n is ever replaced by an infected cell: a rim vertex has degree 3, so $d_f(v) \leq 3 < R$, the set $D \cap T_R$ of Lemma 2.1 is empty, and the fourth clause of* (1) *applies at every dead rim vertex at every time.*

Proof. Immediate from $\deg_{W_n}(v) = 3$ for v on the rim. □

2.4. **What is known before any computation.** By (3), $\{R : R \geq n-1\} \subseteq \mathcal{E}(W_n)$, so the small values of R are the whole difficulty; for every wheel treated below, $R = 1$ is latent. For $R \geq 4$, Lemma 2.2 says that the only vertex at which reinfection-on-replacement can occur is the hub, so the rim cannot sustain itself and the dynamics is a front driven by the hub. This is why the value $R = 3$ behaves differently from every larger one — it is the only value at which a dead rim vertex can be born infected — and it is why $R = 3$ is the one value in the answer below that carries a parity condition. We do not use Lemmas 2.1 and 2.2 in any proof; we record them because they are the only part of the corrected law of Section 8 that is currently understood.

3. From an infinite parameter range to a finite search

A statement of the form "$R \in \mathcal{E}(W_n)$ for every R in an infinite set" is not a finite check, and neither is its negation for a single R, since Ext_R quantifies over time. This section records the three reductions that make each row of the table a finite computation, all three of which are verified.

3.1. **The positive certificate.** For a bound K let $\mathrm{run}_K(f_0)$ be the Boolean "some $\Phi_R^t(f_0)$ with $t \leq K$ equals $\mathbf{0}$", evaluated by iterating (1) and stopping early at $\mathbf{0}$, and let

$$\mathrm{sweep}(n, R, K) \;=\; \bigwedge_{I<2^n} \mathrm{run}_K\big((I, 0)\big).$$

Proposition 3.1 (Sweep certificate). *If* $\mathrm{sweep}(n, R, K)$ *is true then* $R \in \mathcal{E}(W_n)$.

Proof. By induction on K: if $\mathrm{run}_K(f_0)$ holds then either $f_0 = \mathbf{0}$, and $t = 0$ works, or $\mathrm{run}_{K-1}(\Phi_R(f_0))$ holds and a witness t for $\Phi_R(f_0)$ yields the witness $t+1$ for f_0. Quantifying over the 2^n admissible states gives (2). □

The conjunction is evaluated as one tail-recursive walk over the window $[0, 2^n)$, and it splits: the conjunctions over $[s, s+a)$ and over $[s+a, s+a+b)$ combine to the conjunction over $[s, s+a+b)$, so a single long sweep can be cut into pieces that are checked separately. None of the sweeps reported here needed to be cut.

3.2. **The negative certificate.** For the other direction we exhibit one initial state that never dies out, and we certify that by a finite forward-invariant set.

Proposition 3.2 (Latency certificate). *Let $w < 2^n$ and $K \geq 1$, and let O be the list of the first K states of the orbit of $f_0 = (w, 0)$ under Φ_R. If $f_0 \in O$, if $\Phi_R(x) \in O$ for every $x \in O$, and if $\mathbf{0} \notin O$, then $R \notin \mathcal{E}(W_n)$.*

Proof. The three hypotheses say that O is a Φ_R-closed set containing f_0 and missing $\mathbf{0}$. By induction on t, $\Phi_R^t(f_0) \in O$ for every t, hence $\Phi_R^t(f_0) \neq \mathbf{0}$ for every t, so $\mathrm{Ext}_R(f_0)$ fails and f_0 witnesses $R \notin \mathcal{E}(W_n)$. □

Taking $K = \mu + \lambda$, the transient plus the period of the orbit of f_0, always suffices. Of the 394 latency certificates used below, every one has $\mu + \lambda \leq 9$: over $4 \leq n \leq 26$ the distribution of certificate lengths is 24 of length 3, 22 of length 4, 91 of length 5, 82 of length 7, 50 of length 8 and 7 of length 9, and over $27 \leq n \leq 31$ it is $5, 5, 44, 36, 25, 3$ of the same six lengths. Latency, when it happens on a wheel, happens on a very short cycle.

3.3. **The parameter tail.** The remaining infinity is the parameter itself. It collapses.

Proposition 3.3 (R-independence). *Let G have order $n \geq 4$ and let $R, R' \geq n$. Then $\Phi_R = \Phi_{R'}$ as maps on states of G. Consequently, if $n \in \mathcal{E}(G)$ then $R \in \mathcal{E}(G)$ for every $R \geq n$.*

Proof. Every vertex of G has at most $n-1$ neighbours, so $d_f(v) \leq n-1 < n \leq R$ and likewise $d_f(v) < R'$, for every state f and every vertex v. The five clauses of (1) therefore select the same case for R as for R' at every vertex — the only clause in which R appears is the pair $d_f(v) < R$ versus $d_f(v) \geq R$, and the first alternative holds in both cases. Hence $\Phi_R(f) = \Phi_{R'}(f)$ for all f, and membership in $\mathcal{E}(G)$, being a property of the orbits of Φ_R, transfers. □

Proposition 3.3 is what makes each theorem of Section 5 a statement about all of $\mathbb{Z}^+$: the interval $[1, n-1]$ is settled value by value, and the whole tail $[n, \infty)$ is settled by the single sweep at $R = n$. Note that the tail inclusion (3) already gives $R \geq n-1$ from the published literature; we do not use it, preferring the values $R = n-1$ and $R = n$ to come out of the same search as the rest of the row, so that the reproduction of the source's own exhaustive rows in Theorem 5.1 is a control on our reading of (1) and not a restatement of a theorem we assumed.

3.4. **What is actually searched.** For a given n, the row of the table is produced by n exploratory searches, one for each $R \in \{1, \dots, n\}$, each a walk over all 2^n admissible initial states. A walk that finds no survivor returns the maximum transient it saw, which becomes the bound K of Proposition 3.1; a walk that finds one stops there and returns that state, which becomes the witness w of Proposition 3.2 together with the transient and period of its orbit. The verified proof then replays only what the certificate needs: the full sweep in the first case — or, from Section 4 on, the sweep restricted to orbit representatives — and one short orbit in the second.

4. The orbit quotient

The cost of Proposition 3.1 doubles with each unit of n, and around $n = 26$ that ends the table. The wheel is a highly symmetric graph, and this section turns its symmetry into a verified reduction of the sweep. Everything here is proved; nothing about the quotient is assumed.

4.1. **The action.** Let $m = n-1$ and let $D_m \leq \mathrm{Aut}(W_n)$ be the dihedral group generated by the rim rotation $\rho\colon i \mapsto i+1 \bmod m$ and the rim reflection $\tau\colon i \mapsto m-1-i$, both fixing the hub m. On masks these act by permuting bits: writing rot and rev for the induced maps on n-bit words, one has, for $i \leq m$,

$$(\mathrm{rot}\, x)_i = x_{\rho(i)}, \qquad (\mathrm{rev}\, x)_i = x_{\tau(i)}. \tag{4}$$

Both maps are implemented as short machine-word formulas — rot as two shifts and a mask, rev as a reversal loop — and (4) is proved about the formulas. That is the point at which a symmetry reduction usually goes wrong: a "canonical labelling" that is not induced by a permutation of the vertices reduces the wrong search. Here there is no labelling heuristic at all, and being a bit permutation is a theorem about the word operation rather than a hypothesis about a search. Everything in this section carries the standing hypothesis $4 \leq n \leq 63$, which is only the room the mask needs inside one machine word.

Proposition 4.1 (Equivariance)**.** *Let $4 \leq n \leq 63$, let σ be a permutation of $V(W_n)$ with $\sigma(N(i)) = N(\sigma i)$ for every vertex i, and let A be the induced map on masks, $(Ax)_i = x_{\sigma i}$. Then A commutes with the transition map: writing $A(I, D) = (AI, AD)$,*

$$\Phi_R(Af) \;=\; A\,\Phi_R(f) \qquad \textit{for every state } f \textit{ and every } R.$$

Proof. Bit by bit. Bit i of each of the two masks of $\Phi_R(f)$ is a function of the current value $f(i) \in \{0, 1, 2\}$ and of the infected-neighbour count $d_f(i)$ alone, namely the right-hand side of (1). The current value transports because A is a bit permutation: the pair of bits of Af at i is the pair of bits of f at σi. The count transports in two steps: $d_{Af}(i)$ counts the bits of Af over the list $N(i)$, which are the bits of f over the list $\sigma(N(i))$, and a count does not depend on the order of the list, so it equals $d_f(\sigma i)$ by $\sigma(N(i)) = N(\sigma i)$. Both sides of the display therefore have the same bit at every position. □

The hypothesis on σ is checked, not assumed: for each concrete n a decidable predicate is evaluated, which asserts that σ maps $\{0, \dots, n-1\}$ into itself and onto it, that every neighbour list of W_n has all entries below n, and that $\sigma(N(i))$ and $N(\sigma i)$ agree as multisets for every $i < n$. For the two generators this check takes between half a second and a little over a second at the orders used here, and it is run inside the proof kernel. It is not vacuous: the map $u \mapsto u+1$ on *all* of $\{0, \dots, n-1\}$, which moves the hub, fails it.

4.2. **The descent.** Call a mask *canonical* if it is the least element, in the numerical order of the underlying word, of its D_m-orbit; the test is a walk over the $2m$ group elements that exits at the first one beating x. What one wants is that a sweep restricted to canonical masks suffices. The usual route to that is a completeness theorem for a canonical form, and it is the part that is easy to get wrong. It can be avoided altogether.

Proposition 4.2 (Covering)**.** *Let* $4 \le n \le 63$ *and suppose* $\mathrm{Ext}_R\big((y,0)\big)$ *holds for every canonical* $y < 2^n$*. Then* $\mathrm{Ext}_R\big((x,0)\big)$ *holds for every* $x < 2^n$*; that is,* $R \in \mathcal{E}(W_n)$*.*

Proof. By strong induction on the integer x. If x is canonical the hypothesis applies. If not, the canonicity test, being a walk over the group, exhibits a group element g with $g \cdot x < x$. By Proposition 4.1 the map induced by g commutes with Φ_R, and it sends $\mathbf{0}$ to $\mathbf{0}$ and only $\mathbf{0}$ to $\mathbf{0}$, since it is a bit permutation; hence $\Phi_R^t(g \cdot x, 0) = g \cdot \Phi_R^t(x,0)$ for all t, and $\mathrm{Ext}_R\big((g\cdot x,0)\big)$ implies $\mathrm{Ext}_R\big((x,0)\big)$. The induction hypothesis applies to $g \cdot x < x$, which closes the induction. ☐

Two features of this argument are worth naming, because they are what make it short. It never uses that the orbit of x has $2m$ elements, or indeed the order of the group at all; and it never needs that two masks in the same orbit have the same canonical form. A failure of the canonicity test is used only as a producer of a strictly smaller mask in the same orbit, and a descent needs nothing more.

Proposition 4.3 (Quotient sweep certificate)**.** *Let* $4 \le n \le 63$ *and let* $\mathrm{sweep}^q(n,R,K)$ *be the conjunction of* $\mathrm{run}_K\big((I,0)\big)$ *over the canonical* $I < 2^n$ *only. If the two generators of* D_{n-1} *pass the automorphism check and* $\mathrm{sweep}^q(n,R,K)$ *is true, then* $R \in \mathcal{E}(W_n)$*.*

Proof. Combine Propositions 3.1 and 4.2: the sweep proves Ext_R at every canonical mask, and the covering proposition propagates it to all of them. ☐

4.3. **What it buys.** The filter still has to look at every mask below 2^n — it is the orbit *simulation*, not the enumeration, that is skipped — but the early exit makes the test constant-time on average, so the expected saving is the orbit size $2(n-1)$ up to the cost of the filter and the number of orbits with nontrivial stabiliser. It was measured by proving the same proposition twice.

Proposition 4.4 (Leverage, measured)**.** *On one machine, the elaboration time of the compiled evaluation alone, for the statement* $5 \in \mathcal{E}(W_n)$ *proved by the full sweep of Proposition 3.1 and by the quotient sweep of Proposition 4.3, is*

n	*full*	*quotient*	*ratio*
24	105 *s*	3.26 *s*	32.2
25	250 *s*	6.63 *s*	37.7
26	594 *s*	15.7 *s*	37.8

against the theoretical $2(n-1) = 46, 48, 50$*; and for a* refusal*, the false sweep at* $n = 24$*,* $R = 6$*, it is* 41.1 *s against* 1.92 *s, a ratio of* 21.4*.*

Two honest qualifications. These ratios compare the two sweeps within one and the same implementation of (1); that implementation is about 2.2 times slower than the one used for $4 \le n \le 26$ in Section 5, because its counters were moved from machine words to unbounded integers to keep the reasoning above free of machine-arithmetic side conditions. Against the older implementation the net gain is therefore about 17, not about 38. And the shortfall from $2(n-1)$ is the filter: at $n = 12$ there are 252 canonical masks against $2^{12}/22 \approx 186$, the excess being the masks with a nontrivial stabiliser. That count is itself a check on the enumeration, since the canonical masks are in bijection with the pairs (binary bracelet of length $n-1$, hub bit): there are $(2^{11} + 10 \cdot 2)/11 = 188$ binary necklaces of length 11, hence $(188 + 2^6)/2 = 126$ bracelets, and $2 \cdot 126 = 252$.

Remark 4.5 (the quotient does not over-accept). A reduction that lost states would prove too much, and the visible symptom would be a quotient sweep returning true where the full sweep returns false.

At $n = 24$, $R = 6$ — a cell genuinely outside $\mathcal{E}(W_{24})$, since $\lfloor 27/5 \rfloor = 5$ — both sweeps return false. Together with the three agreeing pairs of Proposition 4.4 and with the rejection of the hub-moving map by the automorphism check, this is the negative control for the whole of Section 4.

5. The exhaustive determination

We now state the answer. In each theorem, R ranges over all positive integers.

Theorem 5.1. *For $4 \le n \le 10$ and every $R \ge 1$,*

$$R \in \mathcal{E}(W_n) \iff R \ge n - 1.$$

Equivalently $\mathcal{E}(W_n) = \{R : R \ge n-1\}$, so $\mathrm{hiv}(W_n) = \mathrm{HIV}(W_n) = n - 1$ on this range.

This is the range that [1] analysed exhaustively itself, and the theorem agrees with its Table 3 in every cell. We state it first because it is the control on everything that follows: the same formalised reading of (1) that produces the disagreements below reproduces the source's own exhaustive numbers exactly.

Theorem 5.2. *For $11 \le n \le 16$ and every $R \ge 1$,*

$$\mathcal{E}(W_{11}) = \{R : R \ge 10\}, \quad \mathcal{E}(W_{12}) = \{3\} \cup \{R : R \ge 11\},$$
$$\mathcal{E}(W_{13}) = \{R : R \ge 12\}, \quad \mathcal{E}(W_{14}) = \{3\} \cup \{R : R \ge 13\},$$
$$\mathcal{E}(W_{15}) = \{R : R \ge 14\}, \quad \mathcal{E}(W_{16}) = \{3\} \cup \{R : R \ge 15\}.$$

The values here are those printed in Table 3 of [1]; what Theorem 5.2 adds is that they are complete, over all 2^n admissible initial states and all R, on a range the source describes as sampled.

Theorem 5.3. *For $17 \le n \le 21$ and every $R \ge 1$,*

$$\mathcal{E}(W_{17}) = \{4\} \cup \{R : R \ge 16\}, \quad \mathcal{E}(W_{18}) = \{3, 4\} \cup \{R : R \ge 17\},$$
$$\mathcal{E}(W_{19}) = \{4\} \cup \{R : R \ge 18\}, \quad \mathcal{E}(W_{20}) = \{3, 4\} \cup \{R : R \ge 19\},$$
$$\mathcal{E}(W_{21}) = \{4\} \cup \{R : R \ge 20\}.$$

Theorem 5.4. *For $22 \le n \le 26$ and every $R \ge 1$,*

$$\mathcal{E}(W_{22}) = \{3, 4, 5\} \cup \{R : R \ge 21\}, \quad \mathcal{E}(W_{23}) = \{4, 5\} \cup \{R : R \ge 22\},$$
$$\mathcal{E}(W_{24}) = \{3, 4, 5\} \cup \{R : R \ge 23\}, \quad \mathcal{E}(W_{25}) = \{4, 5\} \cup \{R : R \ge 24\},$$
$$\mathcal{E}(W_{26}) = \{3, 4, 5\} \cup \{R : R \ge 25\}.$$

The next two theorems are past the last row of Table 3 of [1], so every value in them is new.

Theorem 5.5. *For $27 \le n \le 29$ and every $R \ge 1$,*

$$\mathcal{E}(W_{27}) = \{4, 5, 6\} \cup \{R : R \ge 26\}, \qquad \mathcal{E}(W_{28}) = \{3, 4, 5, 6\} \cup \{R : R \ge 27\},$$
$$\mathcal{E}(W_{29}) = \{4, 5, 6\} \cup \{R : R \ge 28\}.$$

Theorem 5.6. *For $n = 30, 31$ and every $R \ge 1$,*

$$\mathcal{E}(W_{30}) = \{3, 4, 5, 6\} \cup \{R : R \ge 29\}, \qquad \mathcal{E}(W_{31}) = \{4, 5, 6\} \cup \{R : R \ge 30\}.$$

Proof sketch for Theorems 5.1–5.6. Fix n. For each R with $1 \le R \le n$ exactly one of the two certificates of Section 3 is supplied. For each R in the asserted set, a sweep gives $R \in \mathcal{E}(W_n)$: for $4 \le n \le 26$ the full sweep $\mathrm{sweep}(n, R, K)$ over all 2^n admissible initial states and Proposition 3.1, and for $27 \le n \le 31$ the quotient sweep $\mathrm{sweep}^q(n, R, K)$ over the dihedral-canonical masks and Proposition 4.3. The bound K used is in each case the exact maximum transient observed, so that a bound one smaller would make the search fail. For each R outside the asserted set, a single initial state w together with the first $K = \mu + \lambda$ states of its orbit gives $R \notin \mathcal{E}(W_n)$ by Proposition 3.2. The tail $R > n$ follows from the sweep at $R = n$ by Proposition 3.3, and the case distinction over $R \in \{1, \dots, n\}$ closes the iff. Altogether the six theorems consume 69 full sweeps and 27 quotient sweeps for the membership cells and 394 latency certificates for the rest.

Everything here rests on exhaustive computation: the assertion that a given R lies in $\mathcal{E}(W_n)$ is exactly the assertion that a walk over all 2^n admissible initial states, or over one representative of each of their $\approx 2^n/\big(2(n-1)\big)$ symmetry classes, found no survivor, and there is at present no proof of any single one of these cells that does not run that walk. Summed over the 96 membership cells the verified searches range over $\sum 2^n \approx 2.3 \cdot 10^{10}$ admissible initial states; of the $2.2 \cdot 10^{10}$ of those that belong to the quotiented cells, only about $3.8 \cdot 10^8$ are simulated, the rest being discarded by the canonicity filter. The 394 non-membership cells cost one short orbit each. □

Remark 5.7 (the fuel constants). The bounds K used in the sweeps are measured, not guessed, and they carry information the verdicts do not. For $R = 3$ and even n with $12 \le n \le 30$ the maximum transient found is $38, 46, 54, 62, 70, 78, 86, 94, 102, 110$ at $n = 12, 14, \ldots, 30$, and each of these ten numbers equals $4n - 10$; the slowest initial states at the smallest interesting parameter therefore take time linear in the order of the graph to die. For $R = 4$ the corresponding maxima are $25, 25, 29, 29, 32, 32, 36, 36, 40, 40, 44, 44, 48, 48, 52$ at $n = 17, \ldots, 31$, constant on consecutive pairs but not following a single linear law on the range computed. Both of these are observations about measured constants on the stated finite ranges: neither is proved, and neither is asserted outside the range on which it was measured. The verdicts of Theorems 5.1–5.6 do not depend on them — a bound K enters a proof only as a fuel that the sweep is asked to meet.

Remark 5.8 (the earlier range, recomputed through the quotient). Theorems 5.1–5.4 were obtained before Section 4 existed, by full sweeps. Independently of them, the machinery of Section 4 was used to re-derive $\mathcal{E}(W_n)$ for all $4 \le n \le 20$ — the source's own exhaustive range and the first ten rows of its sampled one — and it returns the same seventeen rows value for value, including the two corrections $4 \in \mathcal{E}(W_{18})$ and $4 \in \mathcal{E}(W_{20})$ that Table 3 of [1] misses. This is the control that a quotient which dropped an orbit would have to survive: it would have to reproduce seventeen whole rows, and two known corrections, while being wrong.

Table 1 collects Theorems 5.1–5.6 beside the published table.

6. Conjecture 1 is false

Theorem 6.1. *The even half of Conjecture 1.1 is false. There is no description of the form*

$$\mathcal{E}(W_n) = \{3\} \cup \{R : R \ge n-1\} \qquad (n \text{ even},\ n \ge 12):$$

at $n = 18$ *one has* $4 \in \mathcal{E}(W_{18})$*, while* $4 \ne 3$ *and* $4 < 17 = n - 1$.

Theorem 6.2. *The odd half of Conjecture 1.1 is false. There is no description of the form*

$$\mathcal{E}(W_n) = \{4\} \cup \{R : R \ge n-1\} \qquad (n \text{ odd},\ n \ge 17):$$

at $n = 23$ *one has* $5 \in \mathcal{E}(W_{23})$*, while* $5 \ne 4$ *and* $5 < 22 = n - 1$.

Proof sketch. In each case a single cell of Theorem 5.3 respectively Theorem 5.4 contradicts the conjectured description. The even witness is the sweep sweep$(18, 4, 25)$: all $2^{18} = 262\,144$ admissible initial states of W_{18} reach the all-healthy state within 25 steps at $R = 4$. The odd witness is sweep$(23, 5, 34)$: all $2^{23} = 8\,388\,608$ admissible initial states of W_{23} reach it within 34 steps at $R = 5$. Each is one exhaustive search and nothing else; each bound is exact, in the sense that the maximum transient over the sweep attains it. □

The two counterexamples are the least ones available, and this too is a theorem rather than a report of where a search happened to stop.

Theorem 6.3. $n = 18$ *and* $n = 23$ *are the least counterexamples in their parity classes. The conjectured description is exactly correct at* $n = 12, 14, 16$ *— the only even* n *with* $12 \le n < 18$ *— and at* $n = 17, 19, 21$ *— the only odd* n *with* $17 \le n < 23$*:*

$$\mathcal{E}(W_{12}) = \{3\} \cup \{R \ge 11\}, \quad \mathcal{E}(W_{14}) = \{3\} \cup \{R \ge 13\}, \quad \mathcal{E}(W_{16}) = \{3\} \cup \{R \ge 15\},$$
$$\mathcal{E}(W_{17}) = \{4\} \cup \{R \ge 16\}, \quad \mathcal{E}(W_{19}) = \{4\} \cup \{R \ge 18\}, \quad \mathcal{E}(W_{21}) = \{4\} \cup \{R \ge 20\}.$$

TABLE 1. $\mathcal{E}(W_n)$ determined exhaustively, against Table 3 of [1]. Rows $4 \le n \le 10$ are the source's own exhaustive range; rows $11 \le n \le 26$ are the range the source reports as sampled; rows $27 \le n \le 31$ lie past the published table.

n	$\mathcal{E}(W_n)$ (exhaustive)	Table 3 of [1]	agree
$4 \le n \le 11$	$\{R \ge n-1\}$	$\{R \ge n-1\}$	yes
12	$\{3\} \cup \{R \ge 11\}$	$\{3\} \cup \{R \ge 11\}$	yes
13	$\{R \ge 12\}$	$\{R \ge 12\}$	yes
14	$\{3\} \cup \{R \ge 13\}$	same	yes
15	$\{R \ge 14\}$	same	yes
16	$\{3\} \cup \{R \ge 15\}$	same	yes
17	$\{4\} \cup \{R \ge 16\}$	same	yes
18	$\{3, 4\} \cup \{R \ge 17\}$	$\{3\} \cup \{R \ge 17\}$	**no**: 4 missing
19	$\{4\} \cup \{R \ge 18\}$	same	yes
20	$\{3, 4\} \cup \{R \ge 19\}$	$\{3\} \cup \{R \ge 19\}$	**no**: 4 missing
21	$\{4\} \cup \{R \ge 20\}$	same	yes
22	$\{3, 4, 5\} \cup \{R \ge 21\}$	$\{3\} \cup \{R \ge 21\}$	**no**: $4, 5$ missing
23	$\{4, 5\} \cup \{R \ge 22\}$	$\{4\} \cup \{R \ge 22\}$	**no**: 5 missing
24	$\{3, 4, 5\} \cup \{R \ge 23\}$	$\{3\} \cup \{R \ge 23\}$	**no**: $4, 5$ missing
25	$\{4, 5\} \cup \{R \ge 24\}$	$\{4\} \cup \{R \ge 24\}$	**no**: 5 missing
26	$\{3, 4, 5\} \cup \{R \ge 25\}$	$\{3\} \cup \{R \ge 25\}$	**no**: $4, 5$ missing
27	$\{4, 5, 6\} \cup \{R \ge 26\}$	—	past the table
28	$\{3, 4, 5, 6\} \cup \{R \ge 27\}$	—	
29	$\{4, 5, 6\} \cup \{R \ge 28\}$	—	
30	$\{3, 4, 5, 6\} \cup \{R \ge 29\}$	—	
31	$\{4, 5, 6\} \cup \{R \ge 30\}$	—	

Proof. These six equalities are six of the rows of Theorems 5.2 and 5.3, each an iff for every positive R, so each is a genuine equality of subsets of $\mathbb{Z}^+$ and not merely an agreement on a window. □

The conjecture is thus not false by accident of its starting point. It is exactly right on the six wheels below the counterexamples that its hypotheses cover. Past its own computational range it is wrong at every single order we can reach, and the failure widens.

Theorem 6.4. *For every n with $27 \le n \le 31$ one has $5 \in \mathcal{E}(W_n)$ and $5 < n - 1$. Hence Conjecture 1.1 fails at every n in that range: its even half at $n = 28, 30$, where $5 \ne 3$, and its odd half at $n = 27, 29, 31$, where $5 \ne 4$. Moreover the deficit grows: the conjectured set omits $\{4, 5, 6\}$ at $n = 28, 30$ and $\{5, 6\}$ at $n = 27, 29, 31$.*

Proof. Read off Theorems 5.5 and 5.6; the omitted values are the elements of $\mathcal{E}(W_n) \cap [1, n-2]$ other than the single element the conjecture retains. □

Section 8 explains the pattern of the failure.

7. Seven rows of Table 3

The disagreements of Table 1 concern the range that [1] itself marks as experimental, so they are corrections to a hedged table and not to a theorem. Collected:

Corollary 7.1. *Of the sixteen rows $11 \le n \le 26$ of Table 3 of [1], nine are complete and seven are not. Precisely, the rows $n = 18, 20$ each omit the value $R = 4$; the rows $n = 22, 24, 26$ each omit both $R = 4$ and $R = 5$; and the rows $n = 23, 25$ each omit $R = 5$. No row contains a value that does not belong.*

Proof. Compare Theorems 5.2, 5.3 and 5.4 with the printed table, row by row. □

The source's own exhaustive range $4 \le n \le 10$ is, by Theorem 5.1, complete and correct.

Corollary 7.2. *The two extinction parameters printed by* [1] *are unaffected by Corollary 7.1. For* $12 \le n \le 26$ *with* n *even,* $3 \in \mathcal{E}(W_n)$ *while* $1, 2 \notin \mathcal{E}(W_n)$*, so* $\mathrm{hiv}(W_n) = 3$*. For* $17 \le n \le 26$ *with* n *odd,* $4 \in \mathcal{E}(W_n)$ *while* $1, 2, 3 \notin \mathcal{E}(W_n)$*, so* $\mathrm{hiv}(W_n) = 4$*. For the three remaining rows* $n = 11, 13, 15$*, Theorem 5.2 gives* $\mathcal{E}(W_n) = \{R : R \ge n-1\}$ *and hence* $\mathrm{hiv}(W_n) = n - 1$*. For* $11 \le n \le 26$*,* $n - 2 \notin \mathcal{E}(W_n)$ *and* $R \in \mathcal{E}(W_n)$ *for every* $R \ge n-1$*, so* $\mathrm{HIV}(W_n) = n - 1$*.*

Proof. Read off Theorems 5.2–5.4; equivalently, evaluate Theorem 8.1 below at $R = 1, 2, 3, 4$ and at $R = n - 2$. Every omitted value in Corollary 7.1 lies strictly between the least element of $\mathcal{E}(W_n)$ and $n - 2$, which is why neither parameter can see it. □

So the conclusions that [1] draws from its table are sound; what its table under-reports is the middle of the set, in a regime the paper had explicitly flagged. It is worth saying why the flag was there.

Remark 7.3 (why the exhaustive range stops at $n = 10$). The source links a public implementation [3]. Running it unmodified reproduces the values reported here for $4 \le n \le 11$ exactly, and does not finish $n = 12$ in an hour and a half on one core. The reason is structural rather than incidental: the program builds the transition *graph* of the system, materialising the whole reachable state space and storing a copy of a graph object for each distinct state encountered, so its memory grows with the number of reachable states rather than staying constant. Its committed notebook output stops part-way through $n = 12$. A search that keeps only the current state — which is all either certificate of Section 3 needs — reaches $n = 26$ in minutes. The disclaimer in [1] about the range $11 \le n \le 26$ is therefore exactly right, and Corollary 7.1 should be read as what that disclaimer anticipated.

We also fed the authors' own transition function every one of the 276 latent initial states underlying the negative certificates for $4 \le n \le 26$, and it confirms all 276 as latent. The negative half of every row of Table 1 in that range is thus corroborated cell by cell against an independent implementation of (1) written by the authors of [1].

8. The law that fits, and the question it raises

Theorem 8.1. *For every* n *with* $11 \le n \le 31$ *and every* $R \ge 1$*,*

$$R \in \mathcal{E}(W_n) \quad \Longleftrightarrow \quad \big(R = 3,\ n \text{ even},\ n \ge 12\big) \ \text{ or } \ \big(R \ge 4 \text{ and } 5R \le n + 3\big) \ \text{ or } \ R \ge n - 1.$$

Equivalently,

$$\mathcal{E}(W_n) \ = \ \{\, 3 : n \equiv 0 \pmod 2 \,\} \ \cup \ \{\, R : 4 \le R \le \lfloor (n+3)/5 \rfloor \,\} \ \cup \ \{\, R : R \ge n - 1 \,\},$$

and, for $4 \le R \le n - 2$ *— that is, for every* $R \ge 4$ *below the tail* (3) *—*

$$R \in \mathcal{E}(W_n) \iff n \ge 5R - 3. \tag{5}$$

Proof. A case distinction on the twenty-one values of n, each case being the corresponding row of Theorem 5.2, 5.3, 5.4, 5.5 or 5.6 together with linear arithmetic in R. No computation beyond those already performed for Section 5 enters. In the formal development the range splits into $11 \le n \le 26$ and $27 \le n \le 31$, one statement each. □

The restriction $R \le n - 2$ in (5) cannot be dropped: the tail (3) puts $R = n - 1$ and $R = n$ into $\mathcal{E}(W_n)$ whatever n is, so no condition of the form $n \ge 5R - 3$ can describe membership there. Below the tail, and only there, the answer is that single inequality — and it is so at every n we have computed, the seven wheels $4 \le n \le 10$ included, where the part of $\mathcal{E}(W_n)$ below the tail is empty and $\lfloor (n+3)/5 \rfloor < 4$.

The thresholds of (5) are the points where the middle block grows, and two of them were computed only after the law had been read off the other three.

Theorem 8.2. *$6 \notin \mathcal{E}(W_{26})$ and $6 \in \mathcal{E}(W_n)$ for every n with $27 \le n \le 31$; and $7 \notin \mathcal{E}(W_{31})$ while $7 \in \mathcal{E}(W_{32})$. So the least n at which 6 enters $\mathcal{E}(W_n)$ below the tail is exactly $n = 27 = 5 \cdot 6 - 3$, and the least n at which 7 does is exactly $n = 32 = 5 \cdot 7 - 3$.*

Proof sketch. The two non-membership statements are latency certificates (Proposition 3.2); the memberships at $27 \le n \le 31$ are cells of Theorems 5.5 and 5.6. The membership $7 \in \mathcal{E}(W_{32})$ is one further quotient sweep, over the dihedral-canonical masks of W_{32} — about $2^{32}/62$ of them — each reaching $\mathbf{0}$ within 44 steps; it is the largest single search in this paper and the only statement here about a wheel of order 32. That no smaller order does the same is Theorems 5.1–5.6: below the tail, $\mathcal{E}(W_n)$ stops at $\lfloor (n+3)/5 \rfloor$, which is at most 5 for $n \le 26$ and at most 6 for $n \le 31$. Both thresholds were predicted by (5) from the data on $11 \le n \le 26$ before the corresponding searches were run. $\square$

Three comments on the shape of (5).

First, it explains the pattern of agreement in Table 1. Write the middle block of Theorem 8.1 as $M_n = \{3 : n \text{ even}\} \cup \{R : 4 \le R \le \lfloor (n+3)/5 \rfloor\}$; its least element is 3 for even n and 4 for odd n. Conjecture 1.1 keeps that least element and drops everything else in M_n. So for even n the conjecture is correct exactly while the dropped part $\{R : 4 \le R \le \lfloor (n+3)/5 \rfloor\}$ is empty, that is while $\lfloor (n+3)/5 \rfloor < 4$, i.e. $n \le 16$; and for odd n exactly while M_n is the singleton $\{4\}$, that is while $\lfloor (n+3)/5 \rfloor = 4$, i.e. $17 \le n \le 21$. The first even n outside the first range is 18 and the first odd n outside the second is 23: precisely the two least counterexamples of Theorems 6.1 and 6.2, and precisely the seven rows of Corollary 7.1. Since $|M_n|$ grows like $n/5$, the conjecture is not merely wrong from $n = 18$ on but increasingly wrong, which is Theorem 6.4.

Second, the split of the description into a parity clause at $R = 3$ and a linear clause above it is exactly the split of Lemma 2.2: $R \ge 4$ is the regime in which the rim cannot feed itself, and it is precisely in that regime, below the tail, that the answer collapses to the single condition (5). The one value at which a dead rim vertex can be born infected is the one value whose membership depends on the parity of n.

Third, and this is what we cannot do:

Question 8.3. Why 5? Prove that $R \in \mathcal{E}(W_n) \iff n \ge 5R - 3$ for every n and every R with $4 \le R \le n-2$; or find the first n at which (5) fails. Equally open, and presumably easier, is a direct proof that $3 \in \mathcal{E}(W_n)$ for every even $n \ge 12$, which with $2 \notin \mathcal{E}(W_n)$ — itself unproved outside the computed range — would give $\mathrm{hiv}(W_n) = 3$ on the whole even range rather than on a finite one.

We have no structural account of the constant 5, and no reduction of (5) to a statement about a single front on the rim, which is what Lemma 2.2 suggests should exist. What we can report is that five thresholds now lie on the line, that the last two of them were predictions, and that the next one is out of reach: $R = 8$ would have its threshold at $n = 37$, and a single cell at that order is beyond the compute budget we allow one statement, even with Section 4.

Corollary 8.4. *On the new range the two extinction parameters continue to take the values Conjecture 1.1 predicts for them, even though the conjecture's description of the set is false there. For $27 \le n \le 31$: $1, 2 \notin \mathcal{E}(W_n)$; $3 \in \mathcal{E}(W_n)$ if and only if n is even; $4 \in \mathcal{E}(W_n)$ always; so $\mathrm{hiv}(W_n) = 3$ for even n and 4 for odd n. And $n - 2 \notin \mathcal{E}(W_n)$ while every $R \ge n - 1$ lies in $\mathcal{E}(W_n)$, so $\mathrm{HIV}(W_n) = n - 1$. In particular the gap $\mathrm{HIV}(W_n) - \mathrm{hiv}(W_n)$ equals $n - 4$ for even n and $n - 5$ for odd n.*

Proof. Evaluate Theorem 8.1 at $R = 1, 2, 3, 4$ and at $R = n - 2$, using $\lfloor (n+3)/5 \rfloor \ge 6$ for $n \ge 27$. In the formal development this is stated for $27 \le n \le 30$; the case $n = 31$ is the same arithmetic applied to the verified row for $n = 31$ in Theorem 5.6. $\square$

Together with Corollary 7.2 this gives the gap $\mathrm{HIV}(W_n) - \mathrm{hiv}(W_n)$ on all of $11 \le n \le 31$: it is $n-4$ for even $n \ge 12$ and $n - 5$ for odd $n \ge 17$, so on the range determined here it grows linearly in the order of the graph. That the gap is unbounded is not new — it is [1, Theorem 2.11], which for each

$k \geq 2$ exhibits a connected graph G, the flag graph F_k on $2k+2$ vertices, with $\mathrm{HIV}(F_k) - \mathrm{hiv}(F_k) = k$. What the wheels add is that the gap can be nearly as large as the order permits: for any graph of order n with at least one edge, [1, Theorem 2.8] gives $\mathrm{HIV} - \mathrm{hiv} \leq \Delta(G) - 1 \leq n - 2$, and on the wheels computed here the gap is $n - 4$ or $n - 5$, against $(n-2)/2$ for the flag graph of the same order. We prove nothing about wheels of order greater than 31, so this is a statement about a finite range and not a second construction of an unbounded gap.

Remark 8.5 (what is left as data). Outside the formal development, and recorded here as data rather than as a theorem, the same exhaustive search carried out by a separate implementation in C gives the full row

$$\mathcal{E}(W_{32}) = \{3, 4, 5, 6, 7\} \cup \{R \geq 31\},$$

of which only the entry $7 \in \mathcal{E}(W_{32})$ is a theorem (Theorem 8.2). It agrees with (5). That implementation runs over all 2^n admissible states, or over the rotation classes only, and it agrees with the verified answers at every one of the 28 orders $4 \leq n \leq 31$ where both have been run; it is also the source of every numerical constant used in the proofs, as described in Section 10. Nothing in this paper depends on the six unverified entries of the $n = 32$ row.

8.1. **How many intervals?** The conclusions of [1] pose a second question, about arbitrary graphs rather than wheels. Since $\mathcal{E}(G)$ is not upward closed, it is a finite union of maximal intervals of $\mathbb{Z}^+$, one of which is the infinite tail (3); how many can there be? In Section 5 of [1] the authors record the flag graphs F_k, for which $\mathcal{E}(F_k) = \{k+1\} \cup [2k+1, \infty)$, and one graph on nine vertices — their [1, Figure 7] — with $\mathcal{E}(G) = \{3\} \cup \{5\} \cup [7, \infty)$, three intervals. Of that graph they write that it "is the only one with at most 9 vertices where $\mathcal{E}$ is conformed by three intervals"; and they add "We suspect that the number of intervals of $\mathcal{E}(G)$ is not bounded" together with "we did not find examples for which $\mathcal{E}(G)$ was conformed by 4 or more intervals". Wheels do not answer it.

Corollary 8.6. *For $11 \leq n \leq 31$, $\mathcal{E}(W_n)$ is a union of at most two maximal intervals of $\mathbb{Z}^+$: the tail $\{R : R \geq n - 1\}$, and — when $n \geq 12$ is even or $n \geq 17$ is odd, and only then — the interval $\{R : \mathrm{hiv}(W_n) \leq R \leq \lfloor (n+3)/5 \rfloor\}$, which is disjoint from the tail. For $4 \leq n \leq 10$ there is a single interval.*

Proof. By Theorem 8.1 the part of $\mathcal{E}(W_n)$ below the tail is $\{3 : n \text{ even}\} \cup \{R : 4 \leq R \leq \lfloor (n+3)/5 \rfloor\}$. If n is even this is $\{3\} \cup [4, \lfloor (n+3)/5 \rfloor] = [3, \lfloor (n+3)/5 \rfloor]$, an interval containing 3; if n is odd it is $[4, \lfloor (n+3)/5 \rfloor]$, an interval, empty exactly when $\lfloor (n+3)/5 \rfloor < 4$, i.e. $n \leq 16$. In either case $\lfloor (n+3)/5 \rfloor \leq (n+3)/5 < n - 2$ for $n \geq 4$, so the block does not touch the tail. The case $4 \leq n \leq 10$ is Theorem 5.1. □

The middle block is contiguous whenever it contains more than one value, which is exactly what a linear law such as (5) forces, so the wheels cannot produce a third interval no matter how far they are computed. A graph with four or more intervals, if one exists, has to be looked for elsewhere; the search is well posed and cheap per graph — one sweep of the kind used here costs milliseconds at order 10 — but needs isomorphism-free generation of the graphs themselves, which is a different instrument from the one used in this paper.

9. Controls

Exhaustive-search papers fail silently when the object searched is not the object intended, so we record the checks that would have caught the standard failures. The first two say that the answer has content in both directions; the third says that the quantifier defining $\mathcal{E}$ is a real constraint.

Proposition 9.1. (i) *$\mathcal{E}(W_n)$ is not an up-set: $3 \in \mathcal{E}(W_{18})$ but $5 \notin \mathcal{E}(W_{18})$.*

(ii) *$\mathcal{E}(W_n)$ is strictly larger than its tail from $n = 12$ on: $\mathcal{E}(W_{11}) = \{R : R \geq 10\}$ exactly, while $3 \in \mathcal{E}(W_{12})$ and $3 < 11$.*

(iii) *The universal quantifier in* (2) *is not vacuous: some admissible initial state of* W_{18} *is latent at* $R = 5$*. Moreover the dynamics is neither everywhere trivial nor everywhere moving —* $\Phi_5(\mathbf{0}) = \mathbf{0}$ *on* W_{18}*, while the state with a single infected rim vertex is not fixed.*

Proof. (i) and (ii) are cells of Theorems 5.2 and 5.3. For (iii), the admissible initial state of W_{18} whose infected set is the 18-bit mask 135 301 is latent at $R = 5$ by Proposition 3.2 with $K = 8$; the two remaining assertions are single evaluations of (1). □

Item (i) matters most. Because $\mathcal{E}(W_n)$ is not upward closed, no theorem above can be secretly proving the weaker statement "every R at least $\mathrm{hiv}(W_n)$ lies in $\mathcal{E}(W_n)$", which is the shape a misplaced quantifier would most naturally take; and the non-monotonicity is also what makes Conjecture 1.1 a substantive guess rather than a normalisation. Item (iii) rules out the degenerate reading in which Ext_R holds trivially. On the new range the same two directions are controlled by Theorem 6.4, which says the published description is too small at every $n \in [27, 31]$, and by Theorem 8.2, which says the middle block does not run all the way to $n-2$: $7 \notin \mathcal{E}(W_{31})$ although $7 < 30$. Finally, Theorem 5.1 is itself a control: it reproduces seven rows that [1] computed exhaustively and published, so a misreading of (1) would have to agree with all seven and then disagree with seven of the sampled rows. The controls proper to the symmetry reduction are Remark 4.5 and Remark 5.8.

10. Verification

Every claim of the form $R \in \mathcal{E}(W_n)$ or $R \notin \mathcal{E}(W_n)$ made in this paper has been formally verified in Lean 4 (toolchain `v4.32.0`) against *Mathlib* [4, 5], and with them the three reductions of Section 3, the whole of Section 4, Theorems 5.1–5.6, 6.1, 6.2, 6.3, 6.4, 8.1 and 8.2, the parameter values of Corollary 7.2 and, on $27 \le n \le 30$, those of Corollary 8.4, and Proposition 9.1; the development contains no `sorry`. The following are not Lean theorems, and none of them carries a computation: the two structural lemmas of Section 2, which are proved on paper and used in no proof; the case $n = 31$ of Corollary 8.4, which is the arithmetic of that corollary applied to the verified row $\mathcal{E}(W_{31})$ of Theorem 5.6; Corollary 8.6, which is arithmetic on top of Theorem 8.1; the row-by-row comparisons with the printed table, which are a reading of [1]; and the six unverified entries of the $n = 32$ row in Remark 8.5. Proposition 4.4 is a timing measurement, not a mathematical statement, although the propositions it times are verified both ways. The development consists of an executable model — the rule (1), the wheel neighbour lists, the dihedral action, and the certificate checkers, on states encoded as a pair of machine words — and a reasoning layer proving Propositions 3.1, 3.2, 3.3, 4.1, 4.2 and 4.3 about it, followed by one theorem per wheel and the summary statements above; it is 359 theorems in 22 modules and rebuilds in a little over two hours of processor time.

The axiom profile is the point worth stating. The soundness of both certificates, the window-splitting lemma, the R-independence of Proposition 3.3, the bit laws (4), the equivariance of Proposition 4.1, the covering argument of Proposition 4.2, the quotient certificate of Proposition 4.3, the automorphism checks and every non-membership assertion in the paper depend on no axioms beyond `propext`, `Classical.choice` and `Quot.sound`: all 394 latency certificates, and both automorphism checks of every quotient-swept wheel, are replayed by the proof kernel itself, by evaluation, and carry no compiled-evaluation axiom at all. In particular *nothing in the argument that a restricted sweep suffices* is delegated. Compiled evaluation, through Lean's `native_decide`, is used for exactly one thing — the sweeps that witness membership — and the dependency is fine-grained: the even half of the refutation, Theorem 6.1, rests on the single sweep $\mathrm{sweep}(18, 4, 25)$ and on nothing else beyond the three standard axioms, the odd half, Theorem 6.2, on the single sweep $\mathrm{sweep}(23, 5, 34)$, and the $R = 7$ threshold of Theorem 8.2 on the single quotient sweep at $n = 32$. Theorem 8.1 depends on all 55 full sweeps for $11 \le n \le 26$ and all 27 quotient sweeps for $27 \le n \le 31$; Corollary 8.4 on the 22 of those with $n \le 30$. The single largest verified item whose cost was measured on its own is the $n = 32$, $R = 7$ quotient sweep, at 804 seconds.

Every numerical constant that appears in a proof — each fuel bound K, each latency witness w — was produced by a separate C implementation of (1), is transcribed into the formal development by a generator rather than by hand, and is then checked rather than merely used: a fuel one smaller makes the corresponding sweep return false and the proof fail, and a wrong witness makes a latency certificate fail. The C tool and the verified model therefore agree on maximum transients and on witnesses, and not only on the verdicts. A third implementation, a deliberately naive Python one — explicit adjacency lists, states as tuples in $\{0, 1, 2\}^V$, a hash set for cycle detection, no bitmasks and no shared code with either the C tool or the formal development — reports verdicts only; it was run on the source's own exhaustive range $4 \le n \le 10$ and on six further cells, finding $(n, R) = (18, 3)$ and $(18, 4)$ in $\mathcal{E}$ and $(12, 4)$, $(13, 3)$, $(17, 3)$ and $(18, 5)$ outside it, in agreement throughout. The authors' own implementation [3] supplies a fourth opinion, described in Remark 7.3. Two caveats are recorded explicitly: the mask encoding of states is faithful only up to the word length — $n \le 32$ for the searches behind Theorems 5.1–5.4 and $n \le 64$ for the quotient searches — and the quotient certificate of Proposition 4.3 is stated for $4 \le n \le 63$; and nothing in the formal development is asserted for $n > 32$.

The repository holding the Lean development, the C sweep and the Python cross-check is private at the time of writing, so no archival identifier for it is quoted here; a repository reference is to be added at submission.

References

[1] Manuel A. Espinosa-García, Ana Paulina Figueroa, Julián A. Fresán-Figueroa, Gerardo L. Maldonado and L. Ariadna Sánchez-Solís, *Extinction thresholds in a graph-based model of HIV infection dynamics*, arXiv:2608.00340v1 [math.DS] (submitted 31 July 2026).

[2] Simon Mukwembi, *A note on the effects of replenishment of depleted cells on HIV infection dynamics: A graph-theoretic approach*, Physica A: Statistical Mechanics and its Applications **387** (2008), no. 5–6, 1200–1204; `doi:10.1016/j.physa.2007.10.038`.

[3] Gerardo L. Maldonado (GitHub user `XGEu2X`), *GraphDynamicSpreading*, software repository, `https://github.com/XGEu2X/GraphDynamicSpreading`, linked from the conclusions of [1]; commit `5a7ab33` of 31 July 2026, the repository head, consulted 2026-08-28.

[4] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: Automated Deduction — CADE 28, Lecture Notes in Computer Science 12699, Springer, 2021, pp. 625–635.

[5] The mathlib Community, *The Lean mathematical library*, in: Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020), ACM, 2020, pp. 367–381; arXiv:1910.09336.