

How to cool a graph^{*}

Anthony Bonato¹, Holden Milne¹, Trent G. Marbach¹, Teddy Mishura¹

Toronto Metropolitan University

Abstract. We introduce a new graph parameter called the cooling number, inspired by the spread of influence in networks and its predecessor, the burning number. The cooling number measures the speed of a slow-moving contagion in a graph; the lower the cooling number, the faster the contagion spreads. We provide tight bounds on the cooling number via a graph's order and diameter. Using isoperimetric results, we derive the cooling number of Cartesian grids. The cooling number is studied in graphs generated by the Iterated Local Transitivity model for social networks. We conclude with open problems.

1 Introduction

The spread of influence has been studied since the early days of modern network science; see [14–16, 18]. From the spread of memes and disinformation in social networks like X, TikTok, and Instagram to the spread of viruses such as COVID-19 and influenza in human contact networks, the spread of influence is a central topic. Common features of most influence spreading models include nodes infecting their neighbors, with the spread governed by various deterministic or stochastic rules. The simplest form of influence spreading is for a node to infect all of its neighbors.

Inspired by a desire for a simplified, deterministic model for influence spreading and by pursuit-evasion games and processes such as Firefighter played on graphs, in [11, 12], burning was introduced. Unknown to the authors then, a similar problem was studied in the context of hypercubes much earlier by Noga Alon [1]. In burning, nodes are either burning or not burning, with all nodes initially labeled as not burning. Burning plays out in discrete rounds or time-steps; we choose one node to burn in the first round. In subsequent rounds, neighbors of burning nodes themselves become burning, and in each round, we choose an additional source of burning from the nodes that are not burned (if such a node exists). Those sources taken in order of selection form a *burning sequence*. Burning ends on a graph when all nodes are burning, and the minimum length of a burning sequence is the *burning number* of G , denoted $b(G)$.

Much is now known about the burning number of a graph. For example, burning a graph reduces to burning a spanning tree. The burning problem is **NP**-complete on many graph families, such as the disjoint union of paths, caterpillars with maximum degree three, and spiders, which are trees where there is a unique node of degree at least three. For a connected graph G of

^{*} Research supported by a grant of the first author from NSERC.

order n , it is known that $b(G) \leq \sqrt{4n/3} + 1$, and it is conjectured that the bound can be improved to $\lceil \sqrt{n} \rceil$ (which is the burning number of a path of order n). For more on burning, see the survey [4] and book [5].

Burning models explosive spread in networks, as in the case of a rapidly spreading meme on social media. However, as is the case for viral outbreaks such as COVID-19, while the spread may happen rapidly based on close contact, it can be mitigated by social distancing and other measures such as ventilation and vaccination. In a certain sense, a dual problem to burning is how to slow an infection as much as possible, as one would attempt to do in a pandemic. For this, we introduce a new contact process called cooling. Cooling spreads analogous to burning, with cooled neighbors spreading their infection to neighboring nodes, and a new cooling source is chosen in each round. However, where burning seeks to minimize the number of rounds to burn all nodes, cooling seeks to *maximize* the number of rounds to cool all nodes.

More formally, given a finite, simple, undirected graph G , the cooling process on G is a discrete-time process. Nodes may be either *uncooled* or *cooled* throughout the process. Initially, in the first round, all nodes are uncooled. At each round $t \geq 1$, one new uncooled node is chosen to cool if such a node is available. We call such a chosen node a *source*. If a node is *cooled*, then it remains in that state until the end of the process. Once a node is cooled in round t , in round $t + 1$, its uncooled neighbors become cooled. The process ends in a given round when all nodes of G are cooled. Sources are chosen in each round for which they are available.

We define the *cooling number* of G , written $CL(G)$, to be the maximum number of rounds for the cooling process to end. A *cooling sequence* is the set of sources taken in order during cooling. We have that $b(G) \leq CL(G)$, and in some cases, the cooling number is much larger than the burning number. Note that while a choice of sources that burns the graph gives an upper bound to the burning number, a choice of sources that cools the graph gives a lower bound to the cooling number.

For all graphs with diameter at most two, we have that $b(G) = CL(G)$. Consider a cycle C_8 with eight nodes for an elementary example of cooling. By symmetry, we may choose any node as the initial source. See Figure 1 for a cooling sequence of length four. There is no smaller cooling sequence, and so $CL(C_8) = 4$.

The present paper aims to introduce cooling, provide bounds, and consider its value on various graph families. In Section 2, we will discuss cooling's relation to well-known graph parameters that provide various bounds on the cooling number. Using isoperimetric properties, we determine the cooling number of Cartesian grid graphs in the next section. The Iterated Local Transitivity (ILT) model, introduced in [10] and further studied in [2, 6–8], simulates structural properties in complex networks emerging from transitivity. The ILT model simulates many properties of social networks. For example, as shown in [10], graphs generated by the model densify over time, have small diameter,




Fig. 1: An example of cooling on the cycle of length 8. Black labels indicate the nodes of the cooling sequence in increasing order. Blue labels indicate the round that the corresponding node was cooled.

high local clustering, and exhibit bad spectral expansion. We derive results for cooling ILT graphs in Section 4 and prove that the cooling number of ILT graphs is dependent on the cooling number after two time-steps of the model. We finish with open problems on the cooling number.

All graphs we consider are finite, simple, and undirected. We only consider connected graphs unless otherwise stated. For further background on graph theory, see [19].

2 Bounds on the cooling number

As a warm-up, we consider bounds on the cooling number in terms of various graph parameters. We apply these to derive the cooling number of various graph families, such as paths, cycles, and certain caterpillars. The following elementary theorem bounds the cooling number via a graph's order.

Theorem 1. *For a graph G on n nodes, we have that*

$$\text{CL}(G) \leq \left\lceil \frac{n+1}{2} \right\rceil.$$

Proof. One uncooled node is cooled from the cooling sequence during each round, except during the last round if all nodes are cooled. The cooling will spread to at least one additional uncooled node per round, except for the first round. This implies that there are

$$\left\lceil \frac{n+1}{2} \right\rceil + \left\lceil \frac{n+1}{2} \right\rceil - 1 \geq n$$

cooled nodes at the end of round $\lceil \frac{n+1}{2} \rceil$, and the result follows. \square

We next bound the cooling number by the diameter.

Theorem 2. *For a graph G , we have that*

$$\left\lceil \frac{\text{diam}(G) + 2}{2} \right\rceil \leq \text{CL}(G) \leq \text{diam}(G) + 1.$$

Proof. For the upper bound, if some node v is cooled during the first round, then all nodes of distance at most i from v will be cooled by the end of round $i + 1$. All nodes have distance at most $\text{diam}(G)$ from v , so all nodes will be cooled by the end of round $\text{diam}(G) + 1$.

For the lower bound, we provide a cooling sequence that cools G in at least $\left\lceil \frac{\text{diam}(G)+2}{2} \right\rceil$ rounds. Let (v_0, v_1, \dots, v_d) be a path of diameter length in G . The cooling sequence will be $(v_{2i-2} : 1 \leq i \leq \left\lceil \frac{\text{diam}(G)+2}{2} \right\rceil)$.

Assume that at the start of round $i \geq 2$, each cooled node has distance at most $2i - 4$ from v_0 . After the cooling spreads in this round, every cooled node has distance at most $2i - 3$ from v_0 . As such, the node v_{2i-2} is uncooled and a possible choice for the next node in the cooling sequence. This node is now cooled, so starting round $i+1$, each cooled node has distance at most $2(i+1) - 4$ from v_0 . This sets up the recursion, where we note that the condition holds at the start of the second round. Since $2 \left\lceil \frac{\text{diam}(G)+2}{2} \right\rceil - 2 \leq \text{diam}(G) + 1$, the sequence provided is a cooling sequence. \square

We now apply these bounds to give that for the path P_n of order n ,

$$\text{CL}(P_n) = \left\lceil \frac{n+1}{2} \right\rceil = \left\lceil \frac{\text{diam}(P_n) + 2}{2} \right\rceil.$$

The upper bound follows from Theorem 1 and the lower bound from Theorem 2. In particular, the upper bound of Theorem 1 is tight. In passing, we note (with proof omitted) that for a cycle C_n , $\text{CL}(C_n) = \left\lceil \frac{n+2}{3} \right\rceil$.

Using a more complicated example, we can show that the upper bound of Theorem 2 is also tight. Define the *complete caterpillar* of length d , CC_d , as the graph formed by appending one node to each non-leaf node of P_d . We call the nodes of P_d in CC_d the *spine* of the caterpillar. Note that CC_d has $n = 2d - 2$ nodes.

Theorem 3. *We have that*

$$\text{CL}(\text{CC}_d) = \left\lceil \frac{n+1}{2} \right\rceil = \text{diam}(P_n) + 1.$$

Proof. The upper bound follows from Theorem 1, and so to find the lower bound, we provide a strategy that takes $\left\lceil \frac{n+1}{2} \right\rceil$ rounds. Let (v_1, v_2, \dots, v_d) be the spine of the caterpillar, and let v'_i be the additional node that is appended to v_i , for $2 \leq i \leq d-1$. The cooling sequence will be v_1 followed by $(v'_i : 2 \leq i \leq d-1)$.




Fig. 2: An example of cooling on the complete caterpillar of length 6. Black labels indicate the nodes of the cooling sequence in increasing order. Blue labels indicate the round that the corresponding node was cooled.

On round 1, the first node in the cooling sequence v_1 is cooled. At the start of round $i \geq 2$, the nodes $\{v_1, \dots, v_{i-1}\} \cup \{v'_2, \dots, v'_{i-1}\}$ have been cooled. The cooling spreads, which cools only the node v_i . The next node in the cooling sequence is v'_i . At the start of round $i + 1$, the nodes $\{v_1, \dots, v_i\} \cup \{v'_2, \dots, v'_i\}$ have been cooled. This argument recursively repeats and ends on round d when v_d is cooled as cooling spreads. We then have that $\text{CL}(\text{CC}_d) = d = \lceil \frac{n+1}{2} \rceil$, as required. \square

We finish the section by noting that determining the cooling number of certain graph families appears challenging. Even for spiders in general, determining the exact cooling number is not obvious. The following result provides bounds and exact values of the cooling numbers of certain spiders. We refer to the paths attached to the root node of a spider as its *legs*.

Theorem 4. *Let T be a spider with $2m$ legs, each of length r . If we have that $m < \lceil \log_2 r + 1 \rceil$, then*

$$\text{CL}(T) \geq 2 \sum_{1 \leq i \leq m} \left\lfloor \frac{r+1}{2^i} \right\rfloor \sim (1 - 1/2^m)2r.$$

Otherwise,

$$\text{CL}(T) = \text{diam}(S) + 1.$$

Proof. We partition the cooling process into two major phases: before the head of the spider is cooled and after it is cooled. Let

$$m' = \min\{m, \lceil \log_2 r + 1 \rceil\}.$$

We further split each phase into m' subphases.

We start in the first phase. Suppose we enter subphase i , for i from 1 up to m' . This subphase runs for $\left\lfloor \frac{r+1}{2^{(m'+1-i)}} \right\rfloor$ rounds. In each round, we add the uncooled node in leg i furthest from the head to the cooling sequence. Note that after subphase i , each leg $i' < i$ will have at most $\frac{r+1}{2^{(m'-i)}} \leq r$ cooled nodes. At the end of this phase, there are m' legs with at most r cooled nodes, and

the remaining at least m' legs do not have any cooled nodes. Note that the head of the spider has not yet been cooled.

The next phase begins. Suppose we enter subphase i of the second phase, for i from 1 up to m' . This subphase runs for $\lfloor \frac{r+1}{2^i} \rfloor$ rounds. In each round, we add the uncooled node in leg i to the cooling sequence that is closest to the head. Note that after subphase i , each leg $i' \leq i$ can be assumed to be completely cooled, while legs $i' > i$ will have $\sum_{1 \leq j \leq i'} (\lfloor \frac{r+1}{2^j} \rfloor) - 1 < r$ cooled nodes (recalling that the cooling spreads at the start of a round).

A total of $2 \sum_{1 \leq i \leq m'} (\lfloor \frac{r+1}{2^i} \rfloor)$ rounds have been played, completing the result when $m < \lceil \log_2 r + 1 \rceil$. If $m \geq \lceil \log_2 r + 1 \rceil$, then note that the head was cooled in round $r + 1$ and that leg m did not contain a node in the cooling sequence. Therefore, the leaf of leg m was cooled on round $2r + 1 = \text{diam}(S) + 1$. \square

3 Isoperimetric results and grids

This section studies cooling on Cartesian grids using isoperimetric results. Burning on Cartesian grids remains a difficult problem, with only bounds available in many cases. See [9, 17] for results on burning Cartesian grids.

Suppose G is a graph. For a $S \subseteq V(G)$, define its *node border*, $N(S)$, to be the set of nodes in $V(G) \setminus S$ that neighbor nodes in S . We then have the *node-isoperimetric parameter* of G at s as $\Phi_V(G, s) = \min_{S: |S|=s} |N(S)|$, and the *isoperimetric peak* of G as $\Phi_V(G) = \max_s \{\Phi_V(G, s)\}$. Note that in other work, it is common to use $\delta(S)$ in place of $N(S)$, but we use the chosen notation to prevent confusion with the minimum degree of a graph, $\delta(G)$.

Literature around node borders often either focuses on Cheeger's inequality, which is based on the ratio of $|N(S)|$ to $|S|$, or focuses on isoperimetric inequalities, which are bounds on $\Phi_V(G, s)$ from below. In recent work, for positive x and y , an inequality that bounds the maximum difference between the node-isoperimetric parameter at x and at $x + y$ was given while proving a distinct result (see Theorem 5 of [13]). We provide an analogous, reworded version of this bound with full proof for completeness.

Lemma 1 ([13]). *For a graph G and integers $x, y \geq 0$,*

$$\Phi_V(G, x) - y \leq \Phi_V(G, x + y).$$

Proof. Let S_{x+y} be a set of nodes of cardinality $x + y$ with $|N(S_{x+y})| = \Phi_V(G, x + y)$. Note that if a node u is removed from S_{x+y} , then the only node that may be in the border of the new set $S_{x+y} \setminus \{u\}$ that was not in the border of S_{x+y} is the node u itself, as the nodes in the border are neighbors of nodes in the set. It follows that $N(S_{x+y} \setminus \{u\}) \subseteq N(S_{x+y}) \cup \{u\}$.

By a similar argument, if we remove any set $S_y \subseteq S_{x+y}$ of cardinality y from S_{x+y} , we have $N(S_{x+y} \setminus S_y) \subseteq N(S_{x+y}) \cup S_y$. But this gives $|N(S_{x+y} \setminus S_y)| \leq |N(S_{x+y})| + y$, which re-arranges to yield

$$|N(S_{x+y} \setminus S_y)| - y \leq |N(S_{x+y})| = \Phi_V(G, x + y).$$

This completes the proof since $\Phi_V(G, x) \leq |N(S_{x+y} \setminus S_y)|$. \square

Lemma 1 can be considered a relative isoperimetric inequality. When we define the set of cooled nodes at time i as S_i during cooling, this relative isoperimetric inequality yields a sequence of values, x_i , such that $|S_i| \geq x_i$ for all i , independent of the strategy that was used to cool the graph. This sequence of values derives a natural upper bound on $\text{CL}(G)$ as follows. We define $x_1 = 1$, and let

$$x_{i+1} = x_i + \Phi_V(G, x_i) + 1.$$

Suppose I is the smallest value with $x_I \geq |V(G)|$.

Theorem 5. *If G is a graph, then we have that*

$$\text{CL}(G) \leq I.$$

Proof. Suppose for the sake of contradiction that $\text{CL}(G) \geq I + 1$. For some optimal cooling strategy, let S_j be the set of cooled nodes at the end of round j . We then have that $|S_1| = 1$. Assume that there is some round J where $|S_{J-1}| \geq x_{J-1}$ but $|S_J| < x_J$. Define $y = |S_{J-1}| - x_{J-1} \geq 0$. It follows by the definition of the cooling process, the definition of Φ_V , and by Lemma 1 that

$$\begin{aligned} |S_J| &= |S_{J-1}| + |N(S_{J-1})| + 1 \\ &\geq x_{J-1} + y + \Phi_V(G, x_{J-1} + y) + 1 \\ &\geq x_{J-1} + \Phi_V(G, x_{J-1}) - y + y + 1 \\ &= x_{J-1} + \Phi_V(G, x_{J-1}) + 1 = x_J. \end{aligned}$$

This contradicts the fact that $|S_J| < x_J$, and so we are done. \square

The (Cartesian) grid graph of length n , written G_n , is the graph with nodes of the form (u_1, u_2) , where $1 \leq u_1, u_2 \leq n$, and an edge between (u_1, u_2) and (v_1, v_2) if $u_1 = v_1$ and $|u_2 - v_2| = 1$, or if $u_2 = v_2$ and $|u_1 - v_1| = 1$. The following total ordering of the nodes in the grid, called the *simplicial ordering*, can be found in [3]. For two nodes $u = (u_1, u_2)$ and $v = (v_1, v_2)$, define $u < v$ when either $u_1 + u_2 < v_1 + v_2$, or $u_1 + u_2 = v_1 + v_2$ and $u_1 < v_1$. Let S_i be the i smallest nodes under this total ordering.

Lemma 2 ([3]). *No i -subset of $V(G)$ has a smaller node border than S_i , and so $|N(S_i)| = \Phi_V(G, i)$.*

Under the simplicial ordering, adding the smallest node in $V(G) \setminus S_i$ to S_i yields the set S_{i+1} . The set of nodes $N(S_i)$ is the set containing the $|\Phi_V(G, i)|$ smallest nodes that are larger than the nodes in S_i . As such, $S_i \cup N(S_i) = S_{i+\Phi_V(G, i)}$ and it also follows that $|S_i \cup N(S_i)| = i + \Phi_V(G, i)$.

Theorem 6. *An optimal cooling strategy for a grid graph G_n is formed by choosing the next node in a cooling sequence to be the smallest node in the simplicial ordering that has not yet been cooled.*

Proof. Let U_i be the set of nodes cooled by the end of round i when this strategy was performed. Note that for each i , there exists an x_i with $S_{x_i} = U_i$. We then have that that $|U_i| = x_i$. During the next round, the cooling will spread to $N(U_i)$, meaning the nodes in $U_i \cup N(U_i) = S_{x_i + \Phi_V(G, x_i)}$ are now cooled.

We then cool the smallest uncooled node, implying that exactly the nodes $S_{x_i + \Phi_V(G, x_i) + 1}$ are cooled, and so $U_{i+1} = S_{x_i + \Phi_V(G, x_i) + 1}$. It follows that

$$x_{i+1} = |U_{i+1}| = x_i + \Phi_V(G, x_i) + 1.$$

If I is the smallest value with $x_I \geq |V(G)|$, then on all rounds $i < I$, $|U_i| < |V(G)|$, and so there is still some uncooled node at the end of round i , and so play continues into round I . The process will, therefore, terminate under this approach exactly on round I . This proves that $\text{CL}(G) \geq I$. Theorem 5 finishes the proof. \square

We explicitly determine the cooling number for Cartesian grids up to a small additive constant in the following result.




Fig. 3: An example of the cooling sequence for the 15×15 grid using the strategy of Theorem 7. The nodes on the first row with distance d_2, d_3, d_4 from $(1, 1)$ are also indicated.

Theorem 7. For each $n \geq 1$, there is a $\varepsilon \in \{0, 1, 2\}$ so that

$$\text{CL}(G_n) = 2n - 2\lceil \log_2(n + 3) \rceil + \varepsilon.$$

Proof. We analyze the strategy of Theorem 6 in phases, which we know is an optimal cooling strategy. Phase 1 is just the first round, where only the node

$(1, 1)$ is cooled. Phase p , with $2 \leq p \leq \lfloor \log_2(n+3) \rfloor$ starts with exactly those nodes of distance at most $d_p = 2^p - 4$ from $(1, 1)$ being cooled, and this phase will last for exactly $d_p + 3$ rounds.

In round t of phase p , with $1 \leq t \leq d_p + 3$, the cooling spreads to all uncooled nodes of distance $d_p + t$ from $(1, 1)$ and to the nodes $\{(t', d_p + t + 3 - t') : 1 \leq t' \leq 2t - 2\}$, which each have distance $d_p + t + 1$ from $(1, 1)$. The node $(2t - 1, d_p + 4 - t)$ is then chosen to be the next node of the cooling sequence, so the set of cooled nodes is exactly those nodes with distance at most $d_p + t$ from $(1, 1)$ and the nodes $\{(t', d_p + t + 3 - t') : 1 \leq t' \leq 2t - 1\}$, which have distance $d_p + t + 1$ from $(1, 1)$. At the end of round $d_p + 3$, this is exactly the set of nodes of distance $2d_p + 4 = 2^{p+1} - 4 = d_{p+1}$, and so we may iterate this procedure until phase $\lfloor \log_2(n+3) \rfloor$.

Note that between the start and end of phase p , the ball of cooled nodes about $(1, 1)$ grows in radius by 2^p even though only $2^p - 1$ rounds have occurred. Thus, if T rounds have occurred from the start of play until some round during phase p , then the cooled nodes all sit within a ball of radius $T + p - 2$ around $(1, 1)$, and all the nodes in this ball will be cooled if the round is the last of phase p .

At the end of phase $p = \lfloor \log_2(n+3) \rfloor - 1$, note that the cooled nodes form a ball of radius r around $(1, 1)$, where $\frac{n}{2} - 2 \leq r \leq n - 1$, and after phase p , the cooled nodes form a ball of radius r around $(1, 1)$, where $r \geq n$.

We will briefly discuss a different strategy that will help us bound the time the optimal cooling process will take. Suppose we have played the above strategy for the first $\lfloor \log_2(n+3) \rfloor - 1$ phases, and for phase $\lfloor \log_2(n+3) \rfloor$, we terminate the phase after the first round where a node of distance $n - 2$ from $(1, 1)$ was cooled.

Let T denote the number of rounds that have occurred by the end of this last modified phase. Note that $T + \lfloor \log_2(n+3) \rfloor - 2 = n - 2$. We construct an additional T rounds for our modified strategy, which naturally split into $\lfloor \log_2(n+3) \rfloor$ phases. If node (r_i, c_i) was played on round i with $1 \leq i \leq T$, then we play node $(n + 1 - r_i, n + 1 - c_i)$ on round $2T + 1 - i$, for $1 \leq i \leq T$.

If phase p consisted of rounds t_1 through to t_2 , then phase $2(\lfloor \log_2(n+3) \rfloor - 1) + 1 - p$ consists of rounds $2T + 1 - t_2$ through to $2T + 1 - t_1$. A similar analysis yields that each such node in the cooling sequence is uncooled before we cool it; hence, this is a cooling sequence, and cooling on G_n using this modified approach lasts for at least $2T$ rounds. If we proceed optimally, then the graph would take at least $2T$ rounds to cool. Noting that $T + \lfloor \log_2(n+3) \rfloor - 2 = n - 2$, we have that $2T = 2n - 2\lfloor \log_2(n+3) \rfloor$.

Suppose, for the sake of contradiction, that the optimal approach lasts for at least $2T + 3$ rounds. Note that after $T + 3$ rounds, some nodes of distance $n + 1$ from $(1, 1)$ must have been cooled, and all nodes of distance n from $(1, 1)$ must have been cooled.

We now consider the reverse strategy, which first plays the last choice made in the optimal strategy, and so on. Note that by symmetry, after the first T

rounds of this reverse strategy, the node that was chosen to be cooled has distance at least $n - 2$ from (n, n) . However, this means that the node chosen to be cooled on round T of this reverse strategy has distance at most n from $(1, 1)$. But going back to the original optimal strategy, this means that round $T + 4$ must have been played at a distance of at most n from $(1, 1)$, but all of these nodes were cooled by round $T + 3$. This gives us the desired contradiction, so the optimal strategy must last for at most $2T + 2$ rounds. Therefore, the optimal strategy lasts for $2T$, $2T + 1$, or $2T + 2$ rounds, and the proof is complete. \square

4 Cooling the ILT model

Motivated by structural balance theory, the *Iterated Local Transitivity* (or *ILT* model) iteratively adds transitive triangles over time; see [10]. Graphs generated by these models exhibit several properties observed in complex networks, such as densification, small-world properties, and bad spectral expansion. The ILT model takes a graph $G = G_0$ as input and defines $\text{ILT}(G)$ by adding a cloned node x' for each node x in G , and making x' adjacent to the original node x and the neighbors of x . The set of cloned nodes is referred to as $\text{CLONE}(G)$, and we write $\text{ILT}_t(G)$ to denote the graph obtained by applying the ILT process t times to G . When $t = 1$, we typically write $\text{ILT}(G)$. We call any graph produced using at least one iteration of the ILT process an *ILT graph*. The diameter of $\text{ILT}(G)$ is the same as that of G unless the diameter of G is 1, in which case the diameter of $\text{ILT}(G)$ is 2.

The burning numbers of an ILT graph either equals $b(G)$ or $b(G) + 1$; see [12]. Even though the order of the graphs generated by the ILT model grows exponentially with time, the burning number of ILT graphs remains constant. We now consider the cooling number of ILT graphs.

Theorem 8. *If G is a graph, then $\text{CL}(\text{ILT}(G)) \geq \text{CL}(G)$.*

Proof. Let (v_1, v_2, \dots, v_k) be a maximum-length cooling sequence for G . Since the distance between v_i and v_j is the same in both G and in $\text{ILT}(G)$, the sequence (v_1, v_2, \dots, v_k) is also a cooling sequence for $\text{ILT}(G)$. If $\text{CL}(G) = k$, then as the cooling process on $\text{ILT}(G)$ with the given cooling sequence lasted k rounds, we have the result.

Otherwise, we may assume that $\text{CL}(G) = k + 1$ and that some node, say v , was one of the nodes cooled during the last round when cooling G . Since v must have a distance at least $k + 1 - i$ from v_i in G for this to occur, and since the distance from v to v_i is the same as this in $\text{ILT}(G)$, it follows that v must have a distance at least $k + 1 - i$ from v_i in $\text{ILT}(G)$. However, this implies that at the end of round k of the cooling process on $\text{ILT}(G)$, node v has not yet been cooled. Thus, the cooling process on $\text{ILT}(G)$ lasts at least $k + 1 = \text{CL}(G)$ rounds, and so the proof is complete. \square




Fig. 4: An example of cooling on $\text{ILT}(P_6)$. Black labels indicate the nodes of the cooling sequence in increasing order. Blue labels indicate the round that the corresponding node was cooled.

We next investigate the cooling number of the graphs $\text{ILT}_t(P_n)$, where $t \geq 1$ and $P_n = (p_1, \dots, p_n)$ is the path graph on n nodes. See Figure 4 for an illustration of $\text{ILT}_1(P_6)$. For $1 \leq k \leq t$, the k -th layer set of p_i is the set $L_k(p_i) = L_{k-1}(p_i) \cup \text{CLONE}(L_{k-1}(p_i))$, where $L_0(p_i) = \{p_i\}$. When $k = t$, we write $L(p_i)$ instead.

Theorem 9. *If $n \geq 3$ and $t \geq 1$, then*

$$\text{CL}(\text{ILT}_t(P_n)) = \begin{cases} \lceil 2n/3 \rceil & \text{if } t = 1 \text{ and } n \equiv 2 \pmod{3}; \\ \lceil 2n/3 \rceil + 1 & \text{otherwise.} \end{cases}$$

Proof. Let $G = \text{ILT}_t(P_n)$ and let (v_1, \dots, v_k) be a cooling sequence for G . We begin with an upper bound on k . As the layer sets $\{L(p_i)\}$ form a partition of $V(G)$, each element v_i of the cooling sequence is in exactly one $L(p_j)$, for some $1 \leq j \leq n$. We claim that any three consecutive layer sets, those of the form $L(p_j), L(p_{j+1}), L(p_{j+2})$, can contain at most two sequence elements total.

Indeed, suppose that v, w are nodes belonging to any of these three layer sets. Since p_j and p_{j+2} are adjacent to p_{j+1} in P_n and by definition of the ILT process, v and w are adjacent to p_{j+1} in G . As such, $d(v, w) \leq 2$. If a node u in any of those layer sets were cooled, all three layer sets would be fully cooled in two rounds. In each round, the selection of a node in the cooling sequence happens after the cooling has spread from the nodes that were cool by the end of the previous round, so only one node in the three layer sets could be selected after u was cooled, proving the claim. Thus, as there are n layer sets, we can separate them into $\lfloor n/3 \rfloor$ consecutive triples, each of which have at most two cooled members. If n is not divisible by 3, then the remaining layer sets may also each have a cooled element. Now we have that $k \leq \lceil 2n/3 \rceil$.

This upper bound is also achieved, which we show by constructing a cooling sequence (w_1, \dots, w_k) of the desired length. Let the clone of the node $p_i \in V(\text{ILT}_{t-1}(P_n))$ that was created in the t -th application of the ILT process be denoted p'_i . For $i \geq 1$, define $w_i = p'_{i + \lfloor (i-1)/2 \rfloor}$. This sequence has length $\lceil 2n/3 \rceil$, and we claim it is a cooling sequence. Indeed, for any $1 \leq i < j \leq n$ we have that $d(p'_i, p'_j) = \max(2, |i - j|)$, as the path $(p'_i, p_{i+1}, p_{i+2}, \dots, p_{j-1}, p'_j)$ is of this

length and any path between nodes in layer sets $L(p_i)$ and $L(p_j)$ must cross through every layer set $L(p_k)$ with $i < k < j$, implying no shorter path can exist. Thus, the sequence is a cooling sequence of length $\lceil 2n/3 \rceil$, and as this matches the upper bound, this cooling sequence is optimal.

We have shown that the length of an optimal cooling sequence of G is $\lceil 2n/3 \rceil$. However, it is possible that the cooling process could last one round longer: this breaks into two cases, with whose study we conclude the proof.

Case 1: $t = 1$ and $n \equiv 2 \pmod{3}$. Consider the cooling sequence (w_1, w_2, \dots, w_k) as defined above. As $n \equiv 2 \pmod{3}$, the final cooled node w_k is the node p'_n , and the node p_n is cooled at the beginning of round k . Thus, there is no additional round of cooling, showing that $\text{CL}(G) = \lceil 2n/3 \rceil$.

Case 2: Otherwise. Consider the cooling sequence (w_1, w_2, \dots, w_k) as defined above. Here, we have that either $w_k \neq p'_n$ or that $\text{CLONE}(L_{t-1}(p_n)) \setminus \{p'_n\}$ is nonempty. This implies that after the cooling of w_k , there is at least one uncooled element of $L(p_n)$, and an extra round of cooling is needed. \square

The next result shows that the second time-step determines the cooling number of ILT graphs.

Theorem 10. *For any graph G , the maximum length of a cooling sequence in $\text{ILT}_2(G)$ and $\text{ILT}_t(G)$ are the same.*

Proof. Suppose that (u_1, u_2, \dots, u_d) is an optimal cooling sequence in $\text{ILT}_t(G)$. Given a node u in $\text{ILT}_t(G)$, let $f(u)$ be the node in G that was cloned to form u . Since $t \geq 2$, for each $x \in V(G)$, there are at least two distinct clone nodes in $\text{ILT}_t(G)$ that were created in the latest iteration of the ILT process, say x' and x'' , such that $f(x') = f(x'') = x$.

We define a sequence of nodes in $\text{ILT}_2(G)$ of length d , (v_1, \dots, v_d) , as follows. If $f(u_i) \neq f(u_{i-1})$, then define $v_i = f(u_i)'$; otherwise, define $v_i = f(u_i)''$. We will show that this is an optimal cooling sequence for $\text{ILT}_2(G)$.

To begin, we show that this is a valid cooling sequence. Assume for a contradiction that this sequence is not a valid cooling sequence. As such, there must be some v_j that will already be cool when we try to cool it on round j . For this to have occurred, there must be some $i < j$ and some path from v_i to v_j of length at most $j - i$, say $(p_i, p_{i+1}, \dots, p_j)$, such that node p_r was cooled on round r for each $r \in \{i, i+1, \dots, j\}$. The walk $(f(p_i), f(p_{i+1}), \dots, f(p_j))$ in $\text{ILT}_2(G)$ has length $j - i$, and so contains a path of length $r' \leq j - i$. For simplicity, we may assume that $(f(p_i), f(p_{i+1}), \dots, f(p_j))$ is a path in $\text{ILT}_2(G)$. Since $f(p_i) = f(v_i) = f(u_i)$ and $f(p_j) = f(v_j) = f(u_j)$, it follows that $(u_i, f(p_{i+1}), \dots, f(p_{j-1}), u_j)$ is a path of length $j - i$ in $\text{ILT}_t(G)$. Since u_i was cooled in round i while cooling the graph $\text{ILT}_t(G)$, and the cooling spreads to each uncooled neighbor over each round, it follows that u_j must have been cooled within $j - i$ rounds. Then u_j was cooled by round $i + (j - i) = j$, so u_j was already cooled by round j . However, (u_1, u_2, \dots, u_d) is a cooling sequence in $\text{ILT}_t(G)$, and so we have the required contradiction.

To show that the cooling sequence (v_1, \dots, v_d) in $\text{ILT}_2(G)$ is optimal, assume that there is some cooling sequence (v_1, \dots, v_{d+1}) in $\text{ILT}_2(G)$. This sequence of nodes is also a sequence of nodes in $\text{ILT}_t(G)$, and further, the distance between any of these nodes is the same in both $\text{ILT}_2(G)$ and $\text{ILT}_t(G)$. In $\text{ILT}_t(G)$, since no cooling sequence can have length $d + 1$, there must be a pair of nodes v_i and v_j , $i < j$, with $d(v_i, v_j) \leq j - i$. (This follows by a similar argument to the first part of the proof.) We then have that v_j will already be cooled by the time we try to cool it on turn j , contradicting that this is a cooling sequence. Thus, the maximum length of a cooling sequence in $\text{ILT}_2(G)$ is d , and the proof follows. \square

If a graph G has a maximum cooling sequence of length s , then $\text{CL}(G) \in \{s, s + 1\}$. Theorem 8 assures us that applying the ILT process only increases the cooling value. We thus have the following result proving that, as in the case of burning, the cooling number remains bounded by a constant throughout the ILT process.

Corollary 1. *For any graph G , we have that*

$$\text{CL}(\text{ILT}_2(G)) \leq \text{CL}(\text{ILT}_t(G)) \leq \text{CL}(\text{ILT}_2(G)) + 1.$$

In particular, Theorem 8 and Corollary 1 together imply that for every graph G , either $\text{CL}(\text{ILT}_t(G)) = \text{CL}(\text{ILT}_2(G))$ for all $t \geq 2$, or there exists a threshold value t_0 such that the cooling number of $\text{ILT}_t(G)$ is equal to $\text{CL}(\text{ILT}_2(G)) + 1$ if and only if $t \geq t_0$.

5 Conclusion and further directions

We introduced the cooling number of a graph, which quantifies the spread of a slow-moving contagion in a network. We gave tight bounds on the cooling number as functions of the order and diameter of the graph. Using isoperimetric techniques, we determined the cooling number of Cartesian grids. The cooling number of ILT graphs was considered in the previous section.

Several questions remain on the cooling number. Determining the exact value of the cooling number in various graph families, such as spiders and, more generally, trees, remains open. In the full version of the paper, we will consider the cooling number of other grids, such as strong or hexagonal grids. We want to classify the cooling number of ILT graphs where the initial graphs are not paths. Another direction to consider is the complexity of cooling, which is likely **NP**-hard.

References

1. N. Alon, Transmitting in the n -dimensional cube, *Discrete Applied Mathematics* **37** (1992) 9–11.

2. N. Behague, A. Bonato, M.A. Huggan, R. Malik, T.G. Marbach, The iterated local transitivity model for hypergraphs, *Discrete Applied Mathematics* **337** (2023) 106–119.
3. B. Bollobás, I. Leader, Compressions and isoperimetric inequalities, *Journal of Combinatorial Theory Series A* **56** (1991) 47–62 (1991).
4. A. Bonato, A survey of graph burning, *Contributions to Discrete Mathematics* **16** (2021) 185–197.
5. A. Bonato, *An Invitation to Pursuit-Evasion Games and Graph Theory*, American Mathematical Society, Providence, Rhode Island, 2022.
6. A. Bonato, K. Chaudhary, The Iterated Local Transitivity model for tournaments, In: *Proceedings of WAW'23*, 2023.
7. A. Bonato, H. Chuangpishit, S. English, B. Kay, E. Meger, The iterated local model for social networks, *Discrete Applied Mathematics* **284** (2020) 555–571.
8. A. Bonato, D.W. Cranston, M.A. Huggan, T.G. Marbach, R. Mutharasan, The Iterated Local Directed Transitivity model for social networks, In: *Proceedings of WAW'20*, 2020.
9. A. Bonato, S. English, B. Kay, D. Moghbel, Improved bounds for burning fence graphs, *Graphs and Combinatorics* **37** (2021) 2761–2773.
10. A. Bonato, N. Hadi, P. Horn, P. Prałat, C. Wang, Models of on-line social networks, *Internet Mathematics* **6** (2011) 285–313.
11. A. Bonato, J. Janssen, E. Roshanbin, Burning a graph as a model of social contagion, In: *Proceedings of WAW'14*, 2014.
12. A. Bonato, J. Janssen, E. Roshanbin, How to burn a graph, *Internet Mathematics* **1-2** (2016) 85–100.
13. A. Bonato, T. Marbach, J. Marcoux, JD Nir, The k -visibility Localization game, Preprint 2023.
14. P. Domingos, M. Richardson, Mining the network value of customers, In: *Proceedings of the 7th International Conference on Knowledge Discovery and Data Mining (KDD)*, 2001.
15. D. Kempe, J. Kleinberg, E. Tardos, Maximizing the spread of influence through a social network, In: *Proceedings of the 9th International Conference on Knowledge Discovery and Data Mining (KDD)*, 2003.
16. D. Kempe, J. Kleinberg, E. Tardos, Influential nodes in a diffusion model for social networks, In: *Proceedings 32nd International Colloquium on Automata, Languages and Programming (ICALP)*, 2005.
17. D. Mitsche, P. Prałat, E. Roshanbin, Burning graphs—a probabilistic perspective, *Graphs and Combinatorics* **33** (2017) 449–471.
18. M. Richardson, P. Domingos, Mining knowledge-sharing sites for viral marketing, In: *Proceedings of the 8th International Conference on Knowledge Discovery and Data Mining (KDD)*, 2002.
19. D.B. West, *Introduction to Graph Theory*, 2nd edition, Prentice Hall, 2001.