Proc. 9th Int. Coll. on Structural Information & Communication Complexity (2002), eds. C. Kaklamanis and L. Kirousis.
Proceedings in Informatics 13 pp. 73–85 (2002). ISBN 1-894145-12-7

Broadcasting in Small-World Communication Networks*

F. COMELLAS, M. MITJANA
Universitat Politècnica de Catalunya, Spain

J.G. PETERS
Simon Fraser University, Canada

Abstract

The two main characteristics of *small-world networks* are strong local clustering, and small diameter. These two characteristics are desirable properties in communication networks since typical communication patterns show large amounts of local communication and a small amount of non-local communication that must be completed quickly. In this paper, we study variants of *broadcasting* that resemble the spread of computer viruses in networks. Our deterministic results exhibit rates of "infection" that are similar to previously obtained probabilistic results for the spread of contagious diseases in populations.

Keywords

Small-world networks. Communication networks. Broadcasting.

1 Introduction

Many real world situations, including the world wide web [1], electric power grids [18], and the network of mathematicians with finite Erdős numbers [6], can be modeled as *small-world networks*. The two distinguishing characteristics of small-world networks are strong local *clustering* (nodes have many mutual neighbours), and small average or maximum distances between pairs of nodes. Informally, clustering is the fraction of possible edges among neighbours of a node that are actually present, averaged over all nodes. Clustering, or locality, is a common property of many computations that require the exchange of information

*Research supported by a NATO Collaborative Research Grant (CRG 972241).

Additional support for the first two authors provided by the Ministry of Science and Technology, Spain, and the European Regional Development Fund (ERDF) under projects TIC-2001-2171 and TIC-1997-0963).

Additional support for the third author provided by NSERC of Canada.

among processors, and clustering of communication is typical of many applications of communication networks. The need to communicate quickly between an arbitrary pair of nodes is also a desirable property of communication networks, so small-world networks are interesting candidates for communication networks.

Many classes of *structured* networks, including the restricted class of *circulant graphs* studied in [18] and subsequent papers, have strong local clustering, but large distances between pairs of nodes. The opposite extreme is *random* networks which have small distances between nodes but exhibit very little clustering. Networks between these two extremes were constructed in [7] by starting with a structured network with n nodes, choosing a small number h of evenly-spaced nodes to be *hubs*, and interconnecting the hubs. Small-world networks occur when approximately 1% of the nodes are chosen to be hubs ($h \approx .01n$). The maximum distance between a pair of nodes in these small-world networks is approximately 20% of the value in the original network, but 95% of the clustering remains [7].

In this paper, we investigate *broadcasting* in the small-world networks that were introduced in [7]. The broadcasting communication pattern, in which the information of an originating node is distributed to the other nodes of a network, is similar to the pattern by which contagious diseases spread. Computer viruses spread by a broadcast-like pattern, and cascaded network failures also follow a broadcast-like pattern.

There has been considerable recent study of the spread of real diseases and computer viruses in small-world networks [10, 13, 14, 15, 16, 17, 19, 20, 21]. Most of the results use probabilistic methods based on the model in [18]. In [17], it is shown that random immunizations do not prevent the spread of contagious diseases. In many networks, there is a threshold below which epidemics die out before spreading to the entire population. In [15], it is shown that there is no threshold for epidemics in scale-free networks.

In this paper, we introduce a parameterized model that can be adjusted to model different network situations. At time $t = 0$, a single originating node begins to infect some of its neighbours, these neighbours infect some of their neighbours, and so on. An infected node remains *active* (i.e., contagious) for $\mathcal{A}$ times units and can infect k of its neighbours during each time unit that it is active. Broadcasting in the case $\mathcal{A} = \infty$ has been studied in [11, 9]. The probabilistic results about the spread of diseases in [18] and most of the other papers mentioned above use $\mathcal{A} = 1$. In this paper, we investigate the case $\mathcal{A} = 1$ using deterministic methods based on [7]. We show that our results about the spread of computer viruses are similar to the probabilistic results about the spread of diseases in [18] and the more recent papers cited above. We also show that there is no threshold for epidemics in small-world networks. Unlike the results in [15], we do not require the assumption that the network is scale-free.

2 Notation

In this paper, a network is represented by a graph $G=(V,E)$ of *order* $n=|V|$. We use standard graph theory terminology. The *degree* of a node x, denoted $\delta(x)$, is the number of edges incident on x and the degree of a graph G is $\Delta=\max_{x\in V}\delta(x)$. A graph is Δ-*regular* if the degree of every node is Δ. The *distance* between two nodes x and y, $d(x,y)$, is the number of edges of a shortest path between x and y. The maximum distance over all pairs of nodes, $D=\max\limits_{x,y\in V} d(x,y)$, is the *diameter* of the graph. *Clustering* is a measure of the connectedness of a graph and is one of the parameters used to characterize small-world networks. For each node x of a graph G, let $\mathcal{C}_x$ be the fraction of the $\frac{\delta(x)(\delta(x)-1)}{2}$ possible edges among the neighbours of x that are present in G. The *clustering parameter* of G, denoted $\mathcal{C}_G$, is the average over all nodes x of $\mathcal{C}_x$.

The basic family of graphs considered here is circulant graphs. We will use $C_{n,\Delta}$, Δ even, to denote the circulant graph $C(n;1,2,\cdots,\frac{\Delta}{2})$ with n nodes labelled with integers modulo n, and Δ links per node such that each node i is adjacent to the nodes $i\pm 1, i\pm 2,\cdots,i\pm\frac{\Delta}{2} \pmod n$. This graph has diameter $\lceil\frac{n-1}{\Delta}\rceil$. We will refer to edges between node i and nodes $i\pm\ell$, $\ell>1$, as *chords* of length ℓ. We will also use *double-step graphs*, $C(n;a,b)$, which are circulant graphs such that each node i is adjacent to the four nodes $i\pm a, i\pm b \pmod n$. The minimum possible diameter of $C(n;a,b)$ is $D=\left\lceil\frac{-1+\sqrt{2n-1}}{2}\right\rceil$ and occurs when $a=D$ and $b=D+1$ [5, 3]. In [7], small-world networks were constructed by choosing h nodes of $C_{n,\Delta}$ to be *hubs* and then using an optimal diameter double-step graph of order h to interconnect the hubs. In this way, the clustering parameter of the final graph is high and very near to that of the original graph while the diameter is reduced considerably. We use the notation $C_{n,\Delta,h}$ for these small-world graphs. The *broadcast time* $\mathcal{T}_{k,\mathcal{A}}(u)$ of an originating node u is the minimum number of time units (*rounds*) needed to inform all other nodes under the conditions that a node remains *active* for $\mathcal{A}$ rounds after it has been informed, and a node can inform $k\le\Delta$ of its neighbours during each round that it is active. The broadcast time of a graph G is $\mathcal{T}_{k,\mathcal{A}}(G)=\max\{\mathcal{T}_{k,\mathcal{A}}(u)|u\in V(G)\}$. Since we only consider the case $\mathcal{A}=1$ in this paper, we simplify the notation to $\mathcal{T}_k(G)$.

3 Broadcasting in the circulant graph $C_{n,\Delta}$

In this section we determine $\mathcal{T}_k(C_{n,\Delta})$. We are particularly interested in the values of k, under this model, which allow a broadcasting time equal to the diameter $\lceil\frac{n-1}{\Delta}\rceil$. We would also like to know $\mathcal{T}$ as a function of k, Δ, and n (the order of $C_{n,\Delta}$).

Theorem 1 *The broadcast time in $C_{n,\Delta}$, $n>2\Delta$ is:*

$$\mathcal{T}_k(C_{n,\Delta})=n-1 \qquad \text{if } k=1$$

$$\begin{array}{ll} \mathcal{T}_k(C_{n,\Delta}) = \lceil \frac{n-1}{\Delta} \rceil & \text{if } k = \Delta \\ \mathcal{T}_k(C_{n,\Delta}) = \lceil \frac{n-1}{\Delta} \rceil & \text{if} \frac{\Delta}{2} < k < \Delta \\ \mathcal{T}_k(C_{n,\Delta}) \leq \lceil \log_k \Delta \rceil + \lceil \frac{n-1}{\Delta} \rceil - 1 & \text{if } 2 \leq k \leq \frac{\Delta}{2} \end{array}$$

Proof.
Case $k = 1$: During each round, exactly one informed node is active, so $n-1$ rounds are necessary. Since $C_{n,\Delta}$ is Hamiltonian, $n-1$ rounds suffice.
Case $k = \Delta$: The broadcasting scheme is *flooding*. Each node informs all of its uninformed neighbours during the round after it is informed.
Case $\frac{\Delta}{2} < k < \Delta$: The diameter of $C_{n,\Delta}$ is $\lceil \frac{n-1}{\Delta} \rceil$, so $\lceil \frac{n-1}{\Delta} \rceil$ rounds are necessary. To achieve this bound, the originator i informs the k nodes $i \pm (\frac{\Delta}{2} - \lfloor \frac{k}{2} \rfloor + 1), \cdots, i \pm \frac{\Delta}{2} (\bmod n)$ during the first round. In the second round, the nodes $i \pm 1, \cdots, i \pm (\frac{\Delta}{2} - \lfloor \frac{k}{2} \rfloor)$ and $i \pm (\frac{\Delta}{2} + 1), \cdots, i \pm \Delta$ are informed by the nodes that were informed during the first round. The total number of informed nodes after two rounds including the originator is $2\Delta + 1$. Each subsequent round (except possibly the last round) will reach Δ more nodes, $\frac{\Delta}{2}$ in each direction. An example is shown in the left half of Figure 1.
Case $2 \leq k \leq \frac{\Delta}{2}$: The first round will inform $\lfloor \frac{k}{2} \rfloor$ nodes on one side of the originator and $\lceil \frac{k}{2} \rceil$ nodes on the other side. In the second round, each active node will inform k new nodes, so the total number of informed nodes including the originator is $1 + k + k^2$. Let α be the smallest number of rounds such that $k^\alpha \geq \Delta$. That is, $1 + \cdots + k^\alpha \geq 1 + \cdots + k^{\alpha-1} + \Delta$, so $\alpha = \lceil \log_k \Delta \rceil$. After round $\alpha - 1$, at most $\frac{\Delta}{2}$ new nodes on each side of the originator can be informed and we need $\left\lceil \frac{n-(1+k+\cdots+k^{\alpha-1})}{\Delta} \right\rceil \leq \lceil \frac{n-1}{\Delta} \rceil$ rounds to finish the broadcast. We obtain $\mathcal{T}_k(C_{n,\Delta}) \leq \lceil \log_k \Delta \rceil + \lceil \frac{n-1}{\Delta} \rceil - 1$. An example is shown in the right half of Figure 1. □

4 Broadcasting in the small-world graph $C_{n,\Delta,h}$

It is possible to reduce the diameter of $C_{n,\Delta}$ by selecting h nodes of $C_{n,\Delta}$ to be *hubs* and interconnecting them using a graph H. This will result in a small-world graph, $C_{n,\Delta,h}$, with strong local clustering and small diameter, as was shown in [7]. We are interested in the broadcasting time, $\mathcal{T}_k(C_{n,\Delta,h})$, for these graphs for different values of h and k.

Definition 1 *A* segment $\mathcal{S}$ *of* $C_{n,\Delta}$ *is the subgraph of* $C_{n,\Delta}$ *induced by two consecutive hubs* i *and* j *and all nodes between* i *and* j*. The* length *of* $\mathcal{S}$ *is* $\ell_{\mathcal{S}} = \min(|i-j|, n-|i-j|)$.

The following result gives the broadcast time for a segment and will be used in the broadcast strategy for $C_{n,\Delta,h}$.

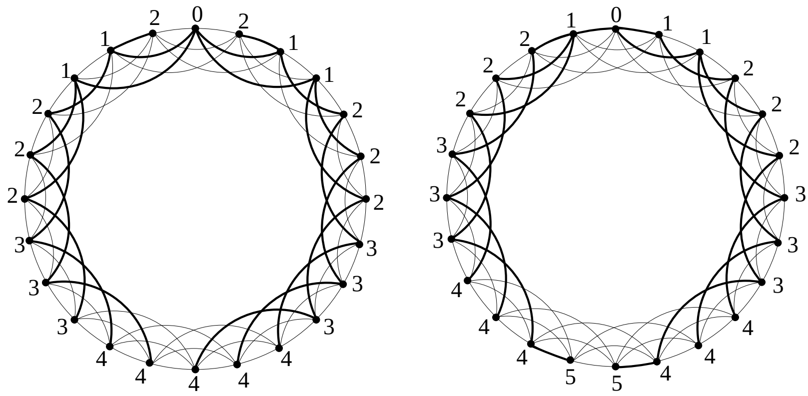

Figure 1: Optimal broadcasts in $C_{n,\Delta}$ for $\frac{\Delta}{2} < k = 4 < \Delta$ and $2 \leq k = 3 \leq \frac{\Delta}{2}$ ($n = 24$, $\Delta = 6$). Edges used in the broadcasts are marked heavier.

Theorem 2 *Let $\mathcal{S}$ be a segment of $C_{n,\Delta}$ with length $(m-1)\Delta + 1 \leq \ell_{\mathcal{S}} \leq m\Delta$, $m \geq 1$. The number of rounds to broadcast in $\mathcal{S}$ is:*

$$\mathcal{T}_k(\mathcal{S}) \leq 2m \qquad \text{if } k \geq \frac{\Delta}{2}$$

$$\mathcal{T}_k(\mathcal{S}) \leq \left\lceil \log_k \frac{\Delta}{2} \right\rceil + \left\lceil \frac{\ell_{\mathcal{S}}}{\frac{\Delta}{2}} \right\rceil - 1 \qquad \text{if } 2 \leq k < \frac{\Delta}{2}$$

Proof.
Case $k \geq \frac{\Delta}{2}$: During the first round, the originator informs k nodes. In each subsequent round, at most $\frac{\Delta}{2}$ more nodes can be informed in each direction. If the originator is close to one end of the segment, then $2m$ rounds are needed.
Case $k < \frac{\Delta}{2}$: Let us suppose that the originator is one of the end nodes of the segment. In the first round, the originator informs k nodes, in the second round k^2 new nodes are informed, and so. Let $\alpha = \left\lceil \log_k \frac{\Delta}{2} \right\rceil$. After round α, $1 + k + k^2 + \cdots + k^{\alpha} \geq 1 + k + k^2 + \cdots + k^{\alpha-1} + \frac{\Delta}{2}$ and there are $1 + k + k^2 + \cdots + k^{\alpha-1} + \frac{\Delta}{2}$ informed nodes. Notice that in round α and in all subsequent rounds, at most $\frac{\Delta}{2}$ new nodes are informed, so we need $\left\lceil \frac{\ell_{\mathcal{S}} - (1+k+\cdots+k^{\alpha-1})}{\frac{\Delta}{2}} \right\rceil \leq \left\lceil \frac{\ell_{\mathcal{S}}}{\frac{\Delta}{2}} \right\rceil$ more rounds to finish the broadcast.

A second strategy is possible, but the broadcast time is the same as the first strategy. Suppose that the originator informs its $k-1$ closest neighbors and one neighbour at distance $\frac{\Delta}{2}$. We will need $\left\lceil \frac{\ell_{\mathcal{S}}}{\frac{\Delta}{2}} \right\rceil$ rounds to inform the end node of the segment as quickly as possible and α more rounds to finish the broadcast in the last piece of length $\frac{\Delta}{2}$ of the segment. □

Remark: If the originator is in the middle of a segment $\mathcal{S}$ of length $\ell_{\mathcal{S}}$, then the broadcast time is $\mathcal{T}_k(\mathcal{S}) \leq \lceil \log_k \Delta \rceil + \left\lceil \frac{\ell_{\mathcal{S}}}{\Delta} \right\rceil - 1$. The proof is similar to the $k < \frac{\Delta}{2}$ case of the previous proof. In the first round, the originator informs $\frac{k}{2}$ nodes on each side. During the second round, each of these nodes informs k new nodes, and so on until round α' when $1 + \frac{k}{2} + \frac{k^2}{2} + \cdots + \frac{k^{\alpha'}}{2} \geq 1 + \frac{k}{2} + \cdots + \frac{k^{\alpha'-1}}{2} + \frac{\Delta}{2}$. From now on, at most $\frac{\Delta}{2}$ new nodes can be informed in each round on each side of the originator, so m'' more rounds are needed to inform the two end nodes of the segment, where $m'' = \left\lceil \frac{\frac{\ell_{\mathcal{S}}}{2} - (1 + \cdots + \frac{k^{\alpha-1}}{2})}{\frac{\Delta}{2}} \right\rceil = \left\lceil \frac{\ell_{\mathcal{S}} - (2 + \cdots + k^{\alpha-1})}{\Delta} \right\rceil \leq \left\lceil \frac{\ell_{\mathcal{S}}}{\Delta} \right\rceil$.

Theorem 3 *The broadcasting time in a double-step graph of order n and optimal diameter D is at most $D+2$ for $k \geq 2$.*

Proof. Any double-step graph of diameter D can be represented as a tile which tessellates the plane. See [3, 4]. This tile will fit inside the tile that corresponds to a double-step graph with the maximum possible order for the same given diameter D. Broadcasting in double-step graphs with optimal diameter was studied in [12] for the case $k = 1$. Using similar arguments for the case $k = 2$, broadcasting in optimal time can be completed in $D+2$ rounds. The broadcast tree for an optimal tile and $k = 2$ is shown in Figure 2. Broadcasting can be completed in $D+2$ rounds for any graph with the same diameter D, order less than the optimal, and chords of length D and $D+1$ using the same scheme. The result remains true for $k > 2$ because a broadcast tree for $k = 2$ is a valid broadcast tree for $k > 2$. □

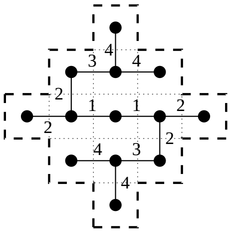

Figure 2: Broadcast in a tile with $k = 2$.

Broadcasting in $C_{n,\Delta,h}$ will be done in three phases.

1. The originator informs the nearest hub as quickly as possible.

2. Broadcast in the graph H that connects the hubs.

3. Broadcast from the hubs to the segments.

Theorem 4 *Given* $C_{n,\Delta,h}$ *with* $\Delta \geq 4$, $3 \leq k \leq \frac{\Delta}{2}$. *Then* $\mathcal{T}_k(C_{n,\Delta,h}) \leq \lceil \log_k \frac{\Delta}{2} \rceil + 2 \left\lceil \frac{\lceil \frac{n}{h} \rceil}{\Delta} \right\rceil + 1 + D_H$.

Proof. In the first phase of the broadcast, the originator sends the message in the direction of the nearest hub using chords of length $\frac{\Delta}{2}$. The worst case occurs if the originator is half-way between two hubs. If the length of the segment is $\ell_{\mathcal{S}} = \lceil \frac{n}{h} \rceil$ then the number of rounds for the message to reach the nearest hub is at most $\left\lceil \frac{\frac{1}{2} \lceil \frac{n}{h} \rceil}{\frac{\Delta}{2}} \right\rceil = \left\lceil \frac{\lceil \frac{n}{h} \rceil}{\Delta} \right\rceil$.

In the second phase, the informed hub originates a broadcast in the graph H that interconnects the hubs. Suppose that H is a double-step graph with optimal diameter $D_H = \left\lceil \frac{-1+\sqrt{2h-1}}{2} \right\rceil$. If each node of H informs at least two of its uninformed neighbouring hubs in one round, then the broadcast in H will take $D_H + 2$ rounds by Theorem 3.

In the third phase, each hub broadcasts to the two segments on either side of it. This third phase overlaps with the second phase, so that each hub informs at least one neighbour that is not a hub and at least two of its neighbouring hubs in the round after it is informed. The non-hub neighbours begin broadcasts in the segments on both sides of the hub. The worst case will happen when the last informed hubs are consecutive. In this situation, at the end of the second phase there will be one whole uninformed segment and it will take $\mathcal{T}_k(\mathcal{S})$ more rounds to complete the broadcast, where $\mathcal{S}$ has length $\lceil \frac{n}{2h} \rceil$. The total time is

$$T_k(C_{n,\Delta,h}) \leq \left\lceil \frac{\lceil \frac{n}{h} \rceil}{\Delta} \right\rceil + D_H + 2 + \lceil \log_k \frac{\Delta}{2} \rceil + \left\lceil \frac{\lceil \frac{n}{2h} \rceil}{\frac{\Delta}{2}} \right\rceil - 1 = \lceil \log_k \frac{\Delta}{2} \rceil + 2 \left\lceil \frac{\lceil \frac{n}{h} \rceil}{\Delta} \right\rceil + 1 + D_H.$$

□

Remark: *Strategy for the case* $k = 2$*:* Suppose that the first hub is informed in round t_0. The non-hub node **x** that is informed in round t_0 can belong to either the segment that contains the originator (and other informed nodes) or to the uninformed segment on the other side of the hub. In either case, in round $t_0 + 1$ the hub informs two more hubs, while **x** informs two nodes in the segment on the other side of the hub from the originator.

Remark: The graph $C_{n,\Delta,h}$ is not Δ–regular as the degree of the hubs is $\Delta + 4$. In [7] it was shown that a regular graph can be obtained from $C_{n,\Delta,h}$ by reconnecting some nodes.

5 Comparison of analytical and numerical approaches

In Figure 3, we compare our analytical results for broadcasting times with numerical values obtained from computer simulations on small-world graphs constructed using the method of [18]. In all cases, we started with an initial circulant graph $C_{1000,10}$ and performed random broadcasts in the simulations for several values of k. All the numerical results are averages of 20 runs.

The analytical and numerical broadcast curves of Figure 3 (square and circle symbols, respectively) are for the graphs $C_{1000,10,h}$ obtained using the methods of this paper and double-step graphs to interconnect the hubs. For reference we have included the numerical clustering (upper curve) and diameter (lower curve) obtained using the method of [18]. We present the case $k = 3$. The results are similar for other values of k (see Theorem 4).

The parameter p (probability that an edge is randomly rerouted) for the numerical results corresponds to the ratio of the number of edges added (i.e., the number of edges in the double-step graph) to the total number of edges in our analytical models. For $C_{n,\Delta,h}$ the ratio is $p = \frac{4h}{|E|}$. In this example, $|E| = n\Delta/2 = 5000$ and $p = \frac{4h}{5000}$. All of the curves in Figure 3 are normalized with respect to the graph $C_{1000,10}$; they show the diameter, clustering, and broadcast time as fractions of the values at $p = 0$.

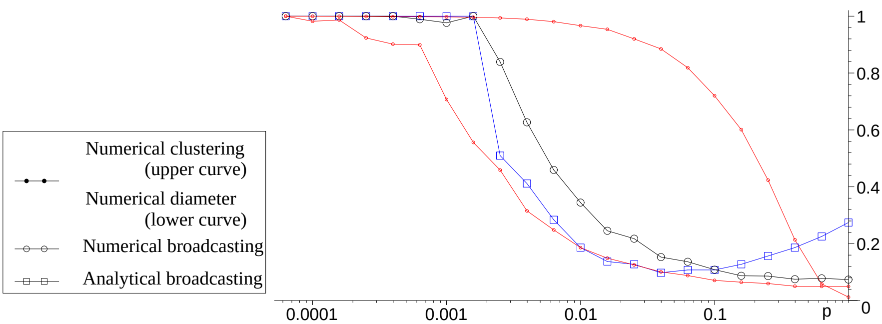

Figure 3: Comparison of analytical results and numerical simulations for $k = 3$.

The small-world region in Figure 3 occurs around $p = 0.01$ where clustering remains above 95% of the value for $p = 0$ and the diameter is less than 20% of the value for $p = 0$. In our graph-theoretic model, $p = 0.01$ corresponds to $h = 12$ hubs for this example.

The figure show that the broadcasting time is smaller in our graph-theoretic model than in the probabilistic model used for the simulations. This is because our analytical approach chooses the hubs in the best way to obtain a small-world graph whereas the probabilistic approach makes random changes to the edges.

For the graphs in our example, the brodcasting time begins to increase around

$p = 0.08$ where the segments have length at most $\Delta + 1$ and the distance from any node to a hub is 1. Adding more hubs increases the diameter of the graph and therefore the broadcasting time. However, this increase of the broadcasting time occurs far from the small-world region.

6 Minimizing the Infection

Infections can spread in a population in a broadcast-like pattern. In this section, we investigate strategies to stop infections from spreading to an entire population. We first show that if all messages of a broadcast in $C_{n,\Delta}$ are successfully delivered (i.e., each node successfully infects its neighbours), then an infection is guaranteed to spread to the entire population for any k and any Δ. We then determine the minimum number of messages that must be lost (or destroyed) to stop an infection. Finally, we determine the effects of hubs in $C_{n,\Delta,h}$ on the spread of infections.

Theorem 5 *If no messages are lost, then all nodes of $C_{n,\Delta}$ will be infected for any $2 \leq k \leq \Delta$.*

Proof. If $\frac{\Delta}{2} < k \leq \Delta$, then the infection spreads at the maximum possible rate and it is impossible to leave gaps by arguments similar to the proofs of the cases $\frac{\Delta}{2} < k < \Delta$ and $k = \Delta$ in Theorem 1. We will now concentrate on the case $2 \leq k \leq \frac{\Delta}{2}$. The only way that a node can be left uninformed permanently is for it to become unreachable from all active nodes. In other words, there can be no path of uninformed nodes from an active node to a node that will be uninformed permanently. For this to happen, there must by two blocks of inactive informed nodes, one block on each side of the uninformed nodes that we wish to isolate. These blocks must contain at least Δ nodes to prevent all paths from active nodes on one side of the blocks to isolated nodes on the other side of the blocks. There are two strategies that we can use to try to construct these blocks. As we shall see, both strategies fail. We will describe the strategies for the case $k = 2$ and then generalize to larger k. The first strategy is to isolate nodes as far as possible away from the originator. During the first phase, the originator i informs nodes $i \pm \frac{\Delta}{2}$. These nodes inform $i \pm \Delta - 1$ and $i \pm \Delta$, and so on. After one or more of these rounds, we start the second phase in which we build the blocks. Suppose that we start the second phase after one round. The node $i + \frac{\Delta}{2}$ informs $i + \frac{\Delta}{2} + 1$ and $i + \frac{\Delta}{2} + 2$, then $i + \frac{\Delta}{2} + 3$ through $i + \frac{\Delta}{2} + 6$ are informed, and so on, with the number of active nodes doubling at each round. The node $i - \frac{\Delta}{2}$ behaves similarly on the other side of the originator to build the other block. In round $\beta = \lfloor \log_2 \frac{\Delta}{2} \rfloor$, the active nodes will try to inform $2^\beta > \frac{\Delta}{2}$ nodes. Since only $\frac{\Delta}{2}$ nodes are reachable in the direction away from the originator, at least one active node is forced to inform a node in the direction towards the originator. Since the block contains

$2^\beta - 1$ nodes, it is not big enough to prevent this from happening and we have failed to isolate the active nodes between two blocks of informed inactive nodes. If the second phase is started after $j > 1$ rounds, then the blocks contain less than $2^\beta - 1$ nodes. If $k > 2$, then $\beta = \lfloor \log_k \frac{\Delta}{2} \rfloor$, and one of the blocks contains at most $\frac{k^\beta - 1}{k-1} < \frac{\Delta}{2}$ nodes. An example is shown in Figure 4 with $k = 2$ and $\Delta = 8$.

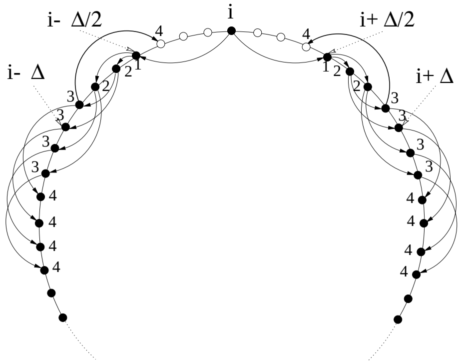

Figure 4: Attempt to isolate nodes close to the originator.

The second strategy attempts to isolate the active nodes in the region around the originator. The first phase is the same as in the first strategy. In the second phase, the active nodes attempt to inform nodes in the direction of the originator instead of in the direction away from the originator. The argument is very similar to the argument for the first strategy. The only differences are that the first phase must have at least two rounds, and there might be one or more inactive informed nodes in the region around the originator that we are trying to isolate. If $k = 2$, the blocks contain at most $2^\beta - 1$ nodes, and the region between the two blocks contains at most $2^{\beta+1} - 2$ uninformed nodes (because the originator is between the two blocks). In round β, the active nodes will try to inform $2^{\beta+1}$ nodes and the attempt to isolate the active nodes fails again. For $k > 2$, the argument can be generalized as in the first strategy. □

Theorem 6 *If four or more messages are lost, then the broadcast process in $C_{n,\Delta}$ can stop after as few as $2 \cdot \Delta - 1$ nodes are informed, $k = 2$.*

Proof. The worst case occurs when $\Delta = 2^j$, $j > 2$. The broadcast starts the same way as strategy 2 in the proof of Theorem 5. The first phase has two rounds during which active nodes inform nodes in the direction away from the originator. To be precise, the originator informs nodes $i \pm \frac{\Delta}{2}$, and these nodes inform nodes $i \pm 2 \cdot (\frac{\Delta}{2} - 1)$ and $i \pm 2 \cdot (\frac{\Delta}{2} - 2)$. Phase 2 now begins and active nodes inform the nodes closest to them in the direction of the originator. During round $j - 1$ of phase 2, the nodes $i \pm 1, i \pm 2, \dots i \pm (\frac{\Delta}{2} - 1)$ and nodes $i \pm (\frac{\Delta}{2} + 1)$ are informed, so all nodes between $i + 2 \cdot (\frac{\Delta}{2} - 1)$ and $i - 2 \cdot (\frac{\Delta}{2} - 1)$ are informed. The only active nodes that can reach uninformed nodes are $i \pm (\frac{\Delta}{2} + 1)$. If all four of the messages of these two nodes fail to reach their destinations, then the infection cannot spread further. □

Notice that with the loss of only a small constant number of carefully chosen messages, the infection stops spreading and the number of infected nodes is independent of *n*. Theorem 6 provides an extreme case. The proof depends on specific nodes failing to deliver their messages and also depends on the pattern by which nodes are informed in the rounds preceding the failures. If different messages are lost or the pattern is different, then it is still possible to stop the spread of the disease but more messages must fail.

We now consider the effects of the addition of hubs on the spread of the disease. Hubs allow the disease to spread to distant parts of the graph and also allow the disease to escape from an isolated region even if messages fail. In a sense, hubs have the opposite effect of message failures, so a less contagious disease (which loses more messages) can still spread to infect the entire graph.

Observation 1 *If the hubs are $\ell < \Delta$ nodes apart, then at least $2 \cdot \lfloor \frac{\Delta}{\ell} \rfloor$ additional message failures may be required to stop the spread of the disease, $k = 2$.*

Proof. In the proof of Theorem 6, the nodes $i \pm 1, i \pm 2, \dots i \pm (\frac{\Delta}{2} - 1)$ are active during the last round that the disease spreads but cannot reach any uninformed nodes. If the hubs are $\ell < \Delta$ nodes apart, then there are $\lfloor \frac{\Delta}{\ell} \rfloor$ or $\lfloor \frac{\Delta}{\ell} \rfloor + 1$ hubs in the region between nodes $i + (\frac{\Delta}{2} - 1)$ and $i - (\frac{\Delta}{2} - 1)$. If there is only one hub in this region and it is the originator, then the hub has no effect on the spread of the disease. In any other case, the disease can escape to other parts of the graph. □

References

[1] R. Albert, H. Jeong, and A.-L. Barabási, Diameter of the world wide web, *Nature*, **vol. 401**(1999) 130—131.

[2] M. Barthelemy, and L.A.N. Amaral, Small-world networks: Evidence for a crossover picture, *Phys. Rev. Lett.* **82(15)** (1999) 3180—3183.

[3] J.-C. Bermond, F. Comellas, and D.F. Hsu, Distributed loop computer networks: A survey, *J. Parallel Distrib. Comput.* **24** (1995) 2–10.

[4] J.-C. Bermond, G. Illiades, and C. Peyrat, An optimization problem in distributed loop computer networks, *Third International Conference on Combinatorial Mathematics*, June 1985. *Ann. New York Acad. Sci.*, **555** (1989) 45–55.

[5] F.T. Boesch, and J.K. Wang, Reliable circulant networks with minimum transmission delay. *IEEE Trans. Circuits & Systems* **CAS-32** (1985) 1286—1291.

[6] R. de Castro and J.W. Grossman, Famous trails to Paul Erdős, *Math. Intelligencer* **21 (3)** (1999) 51—63.

[7] F. Comellas, J. Ozón and J.G. Peters, Deterministic small-world communication networks, *Inform. Process. Lett.* **76** (2000) 83–90.

[8] F. Comellas and M. Sampels, Deterministic small-world networks, *Physica A*, **309** (2002) 231–235.

[9] J.-C. Konig and E. Lazard, Minimum *k*-broadcast graphs, *Discrete Appl. Math.* **53** (1994) 199–209.

[10] M. Kuperman and G. Abramson, Small world effect in an epidemiological model, *Phys. Rev. Lett.* **86** (2001) 2909–2912.

[11] E. Lazard, Broadcasting in DMA-bound bounded-degree graphs, *Discrete Appl. Math.* **37/38** (1992) 387–400.

[12] A.L. Liestman, J. Opatrny, and M. Zaragozá, Network properties of double and triple loop graphs, *Internat. J. Found. Comput. Sci.* **9 (1)** (1998) 57–76.

[13] M. E. J. Newman, Exact solutions of epidemic models on networks, arXiv:cond-mat/0201433

[14] M. E. J. Newman, I. Jensen, and R. M. Ziff, Percolation and epidemics in a two-dimensional small world, *Phys. Rev. E* **65** (2002) 021904.

[15] R. Pastor-Satorras and A. Vespignani, Epidemic spreading in scale-free networks, *Phys. Rev. Lett.* **86** (2001) 3200–3203.

[16] R. Pastor-Satorras and A. Vespignani, Epidemic dynamics and endemic states in complex networks, *Phys. Rev. E* **63** (2001) 066117.

[17] R. Pastor-Satorras and A. Vespignani, Immunization of complex networks. *Phys. Rev. E* **65** (2001) 036104.

[18] D.J. Watts, and H. Strogatz, Collective dynamics of ‘small-world’ networks, *Nature* **393** (1998) 440—442.

[19] D. H. Zanette, Critical behavior of propagation on small-world networks, *Phys. Rev. E* **64** (2001) 050901.

[20] D. H. Zanette and M. N. Kuperman, Effects of immunization on small-world epidemics, *Physica A* (2002), in press.

[21] D. H. Zanette, Dynamics of rumor propagation on small-world networks, *Phys. Rev. E* **65** (2002) 041908.

Francesc Comellas is with the Departament de Matemàtica Aplicada IV, Universitat Politècnica de Catalunya, Escola Politècnica Superior de Castelldefels, Avda. del Canal Olímpic s/n, E-08860 Castelldefels, Catalonia, Spain. E-mail: comellas@mat.upc.es

Margarida Mitjana is with the Departament de Matemàtica Aplicada I, Universitat Politècnica de Catalunya, Gregorio Marañon 44, E-08028 Barcelona, Catalonia, Spain. E-mail: margarida.mitjana@upc.es

Joseph G. Peters is with the School of Computing Science, Simon Fraser University, Burnaby, British Columbia, V5A 1S6, Canada. E-mail: peters@cs.sfu.ca