

Broadcasting in Conflict-Aware Multi-Channel Networks

Francisco Claude¹, Reza Dorrigiv², Shahin Kamali¹, Alejandro López-Ortiz¹,
Paweł Prałat³, Jazmín Romero¹, Alejandro Salinger¹, and Diego Seco⁴

¹ *David R. Cheriton School of Computer Science, University of Waterloo, Canada.*

² *Faculty of Computer Science, Dalhousie University, Canada.*

³ *Department of Mathematics, Ryerson University, Canada.*

⁴ *Department of Computer Science, University of A Coruña, Spain.*

Abstract. The broadcasting problem asks for the fastest way of transmitting a message to all nodes of a communication network. We consider the broadcasting problem in conflict-aware multi-channel networks. These networks can be modeled as undirected graphs in which each edge is labeled with a set of available channels to transmit data between its endpoints. Each node can send and receive data through any channel on its incident edges, with the restriction that it cannot successfully receive through a channel when multiple neighbors send data via that channel simultaneously.

We present efficient algorithms as well as hardness results for the broadcasting problem on various network topologies. We propose polynomial time algorithms for optimal broadcasting in grids, and also for trees when there is only one channel on each edge. Nevertheless, we show that the problem is NP-hard for trees in general, as well as for complete graphs. In addition, we consider balanced complete graphs as an important family of multi-channel networks; we propose a policy for assigning channels to these graphs that, together with the presented broadcasting schemes, provide fault-tolerant networks with optimal broadcasting time.

1 Introduction

Multi-channel networks constitute a class of networks in which communication is achieved via a set of orthogonal *channels*. Two nodes of a multi-channel network can directly communicate if they share at least one common channel. Channels may represent different frequencies in Multi-radio Wireless Networks [12, 15], different wavelengths in Free Space Optical Networks (FSON) [2], or different communication buffers in parallel computers [17].

A multi-channel network can be modeled as an undirected graph with multiple labels on edges, where vertices represent nodes in the network and labels represent available channels between connected nodes. Communication is assumed to occur in discrete *rounds* in which a node can transmit data through one of its channels. For a node u and channel c , we say that a *conflict* occurs when two or more neighbors of u send data to u through channel c at the same round, in which case u does not receive data through this channel. This definition of conflict arises in many practical scenarios; for example, in wireless networks, conflicts can represent the interference of radio waves with the same frequency.

Multi-channel networks have been studied in the context of wireless networks, in which the underlying network is usually modeled as a geometric graph in Euclidean metric space. Most research works on this trend do not focus on theoretical aspects of the problem. Besides, geometric graphs are not good representatives of all types of wireless networks. For example, consider indoor networks in which walls can block transmissions between pairs of nodes that fall

in the transmission range of each other; the underlying network in these cases can form any graph topology [18]. There are several works that provide heuristics for information dissemination in wireless multi-channel networks assuming that conflicts do not occur [8, 10, 11].

A more realistic model that considers conflicts is known as the *conflict-aware* model [1, 19]. In this paper, we present the *conflict-aware multi-channel model*, a comprehensive model that captures several aspects of multi-channel networks that are tied to existing network technologies, in particular conflict awareness and the advantage of simultaneous communication through one channel. Theoretical analysis of this model can provide insights into the capabilities of multi-channel networks for future technology advances; particularly because the model represents a broad spectrum of network technologies, such as wireless mesh networks, FSONs, and parallel computers.

The focus of this work is on the *Broadcasting Problem* in multi-channel networks, in which the goal is to transmit one message from a given source node to all other nodes in the minimum number of rounds. In the classical model of broadcasting, each node can send data to at most one of its neighbors via a *telephone call* (hence the model is called telephone model). In contrast, in multi-channel networks, when a node u transmits through one channel c , all the nodes connected to u via channel c will receive the message (if no conflicts occur). In wireless networks, this is termed the *Wireless Broadcast Advantage* [16], and makes the broadcasting problem more complicated compared to broadcasting under the telephone model [3, 6]. Note that the telephone model can be considered as a restricted version of the multi-channel model in which there is a single and unique channel associated with each edge, i.e., each node is connected to its neighbors via distinct channels. Similarly to the related works in multi-channel wireless networks [9, 12, 13, 15], we assume the broadcasting algorithms are centralized, i.e., the algorithms know the topology of the network and the channels available for each edge.

Channel Assignment is another problem that has been studied for multi-channel wireless networks [12, 13]. This problem asks to assign channels to the edges in a given network in order to optimize the network performance – where the performance can refer to network *goodput* or traffic [13] – or the number of conflicts or interference within the network [12]. We consider this problem for multi-channel networks when the goal is to assign channels in a way to perform broadcasting in minimum time. It is desirable that such channel assignment enables broadcasting of more than a single message in parallel. We consider this problem for homogeneous multi-channel networks that can be modeled by complete graphs with a balanced distribution of channels for each node.

Summary of Results. In Section 2, we describe the conflict-aware multi-channel model. In general, it is assumed that there can be any number of channels between a pair of nodes; however in some occasions, we consider the case when there is only one channel on each edge of the graph. The broadcasting problem seems to be much easier for this restricted case. In Section 3, we show that the broadcasting problem is NP-hard for trees in the general case, while we describe a polynomial time algorithm when there exists only one channel on each edge of the tree. We also provide a polynomial time algorithm for optimal broadcasting in grids (in the general case). In Section 4, we show that the broadcasting problem is NP-hard for complete graphs, even if restricted to graphs with only one channel on each edge. In Section 5, we focus on the special case of complete graphs when there is only one channel on each edge and the channel assignment is balanced, i.e., each node is connected to approximately same number of nodes with each channel. We refer to these graphs as *balanced complete graphs*, and show that broadcasting in these networks requires at least three rounds when the number of different channels does not grow too fast with the size of the network (which is the case in practical settings). On the positive side, we provide a channel assignment that yields a balanced complete network

for which broadcasting can always be completed in two rounds. This channel assignment also enables broadcasting of k messages simultaneously in three rounds, where k is the number of channels in the network.

2 Conflict-Aware Multi-Channel Model


A multi-channel network is modeled as an undirected graph $G = (V, E)$ where V is the set of nodes and E the set of edges. Each edge $e \in E$ has a set of labels $C(e) \subseteq \{c_1, c_2, \dots, c_k\}$ that denotes its set of available channels.

The communication of messages through the network occurs in discrete rounds and is governed by the following assumptions and restrictions. In any given round, a node may be involved in receiving and/or transmitting (sending) messages through the channels on its incident edges. If a node u transmits through a channel c , it cannot transmit through any other channel in the same round, and also cannot receive through channel c . When u sends a message through channel c , the message is simultaneously transmitted through all incident edges of u that have channel c in their set of labels. A key restriction is that a node cannot successfully receive any data through a channel when more than one of its neighbors send data through that channel. More precisely, a node v can only receive a message through channel c in round r if exactly one of the nodes that are adjacent to it with edges labeled with channel c is transmitting through channel c in round r . Otherwise we say there is a conflict at node v on channel c . A node will successfully receive the message if it is transmitted by any of its neighbors through a channel without conflict.

The transmission of a message on any edge completes in one round: if in round r node u transmits a message through channel c , then every node v such that $e = (u, v) \in E$ and $c \in C(e)$ will receive the message during this round, provided that there is no conflict at v on channel c . In this case we say that u *informs* v during round r , and node v is ready to transmit in round $r + 1$ if desired. For any round r during the execution of the broadcast, we say that a node is *active* if it is transmitting the message in round r and *inactive* otherwise.

Given a network represented by a graph G , the broadcasting problem is defined as follows. At the beginning, a single node, called the *source*, has a message. In each round, those vertices that have the message can transmit through one channel to inform some uninformed vertices. The broadcasting completes when all vertices successfully receive the message. The broadcasting problem asks for a scheme that completes this procedure in minimum time. We are interested in centralized broadcasting schemes, i.e., we assume the broadcasting algorithm can be determined in advance and with full knowledge of the network topology. Finally, we assume that the network is static, thus nodes, edges, and their channel assignment remain fixed during the broadcast. Figure 1 illustrates broadcasting in this model.

Figure 1: An illustration of broadcasting in the conflict-aware multi-channel model. Assume node A is the source and sends the message through channel 1 in the first round. Hence, at the beginning of the second round, A , B , and C have the message. Assume B sends the message through channel 1, and A and C through channel 2. At the end of the round, node D receives the message through both channels, node E receives the message through channel 1 (via B); there is a conflict on channel 2 at node F , which does not receive the message in this round.



3 Basic Topologies

3.1 Trees

In this section, we show that the broadcasting problem in the general case (i.e., when there is more than one channel between two nodes) is NP-hard even if the network topology is a tree. On the positive side, we show that when there is a single channel on each edge of the tree, there is an algorithm that finds the optimal broadcasting scheme in polynomial time.

Theorem 3.1. *The broadcasting problem in the conflict-aware multi-channel model is NP-hard for trees.*

Proof. We use a reduction from the set cover problem, which is NP-hard [5]. Recall that an instance of set cover includes a collection of subsets of a universe U , and the goal is to find the minimum number of subsets that cover the universe. Given an instance I of set cover, we create an instance of the broadcasting problem in a tree as follows. We create a tree T with a root node and u children, where $u = |U|$ is the size of the universe. Each child of the root is a leaf of the tree and represents a member of the universe (hence T is a star). Each subset S in I is assigned a label that represents a channel in the broadcast instance. For each member of S , the label of S is added to the edge that connects the root with that member. For example, if $S = \{x, y\}$, the label of S is added to the edges that connect the root to the leaves x and y (See Figure 2). It is not hard to see that there is a set cover of size k if and only if the broadcast finishes in k rounds: assume there is a set cover of size k , then if the root sends the message through the k channels associated with the k subsets (in any order), after k rounds all the nodes of T are informed. This is because there are no conflicts (one channel is used at each round), and all the nodes are covered by k channels. Similarly, if there is a broadcasting scheme that completes in k rounds, the subsets associated with the k channels used by the root cover the universe. \square


The problem becomes easy when there is a single channel on each edge. Consider a tree of n nodes with only one channel on each edge. We show that the optimal broadcasting scheme can be obtained in $O(n \log n)$ time with a simple recursive algorithm.

Given a root node v , we compute the cost (number of rounds) of broadcasting from each of v 's children recursively, and associate with each outgoing channel of v the cost of the most expensive child connected to v with that channel. We then sort these channels in decreasing order of associated cost and transmit through each one following this order. It is not hard to see that this strategy is optimal. Note as well that there are no conflicts in this topology. Algorithm 1 shows the procedure for computing the optimal broadcast scheme for a tree \mathcal{T} and channel assignment \mathcal{C} . A simple implementation of the algorithm runs in $O(n \log n)$ time.

3.2 Grids

Unlike trees, the broadcasting problem can be solved in polynomial time for grids, even if there are multiple channels on edges. In what follows, we describe a scheme for optimal broadcasting

Figure 2: The instance of the broadcasting problem when the set cover instance I contains the subsets $W = \{1, 2, 3\}$, $X = \{2, 4\}$, $Y = \{3, 4\}$, and $Z = \{4, 5\}$.



Algorithm 1 $\text{treeCost}(\text{root}, \mathcal{T} = (\mathbf{V}, \mathbf{E}), \mathcal{C})$

```

time  $\leftarrow$  0
for  $v$  child of  $\text{root}$  do
    Cost[ $v$ ]  $\leftarrow$   $\text{treeCost}(v, \mathcal{T}_v, \mathcal{C})$  { $\mathcal{T}_v$  denotes the tree rooted at  $v$ }
     $p \leftarrow 1$  { $p$  counts the number of different channels seen so far}
    for  $v$  in children of  $\text{root}$  sorted by Cost[ $v$ ] do
        if we have not sent the message through channel  $\mathcal{C}((\text{root}, v))$  then
            time = max(time,  $p + \text{Cost}[v]$ )
             $p \leftarrow p + 1$ 
            inform through channel  $\mathcal{C}((\text{root}, v))$ 
    return time

```

in a grid of size $n \times m$.

Consider first the simple case when the source is one of the corner nodes. W.l.o.g., assume the source is on the upper-leftmost node. A simple scheme is to send the message to the nodes in the first row: after receiving the message, each node transmits to its right neighbor through any one of the available channels. This takes $n - 1$ rounds. Then, in parallel, the message is transmitted in each column downwards, again through any available channel. The broadcast finishes in $m + n - 2$ rounds, which matches a trivial lower bound determined by the diameter of the grid. Note that conflicts do not arise in this strategy.

Combinations of small variations of the strategy described above will serve for the general case in which the source is any node (i, j) in the grid. Consider the set of nodes $N = \{(k, \ell) \mid k = i, \ell \neq j \text{ or } \ell = j, k \neq i\}$, i.e., nodes in the same row or column as the source node, not including the source. Let Q_i denote the i -th quadrant defined by N in G , in clockwise order starting from the upper-left quadrant (See Figure 3 (a)). We say that a node $u \in N$ is a *splitter* if it is connected to neighbors in two different quadrants with at least one channel in common. Similarly, we say that the source is a vertical (resp. horizontal) splitter if it is connected to neighbors above and below (resp. to the left and right) with at least one common channel.

Broadcasting schemes may differ depending on the availability of splitters and the relative sizes of the quadrants. If there are no splitters or the sizes of all the quadrants are different, then optimal strategies for broadcasting in grids in the telephone model [4] apply to our model

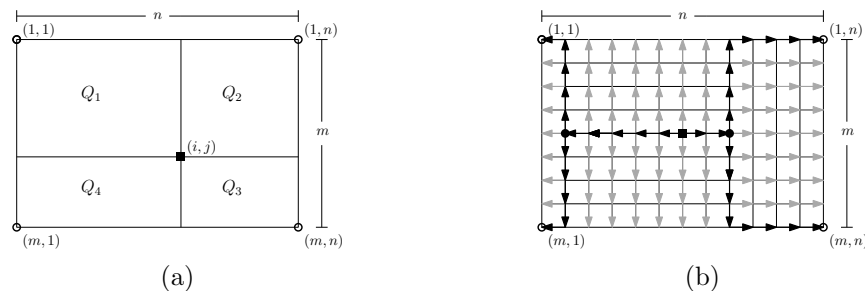


Figure 3: (a) Quadrants defined by the source (i, j) in a grid of size $m \times n$. (b) Example of a broadcast from a source in the center of the grid. The source is a horizontal splitter and there are two splitters in row i , depicted by black discs. Arrows indicate the route of the message to any node, in particular black arrows show the first direction of transmission from each node on the critical path of the scheme. Note that splitters have two black arrows. The broadcasting completes in optimal $(n + m - 2)/2$ rounds.

as well. For other cases, we derive optimal strategies by taking advantage of the splitters (See Figure 3 (b) for an example), thus proving the following theorem. The proof requires a tedious case analysis, hence we provide a sketch of it in Appendix A.

Theorem 3.2. *Given an $m \times n$ grid G with k channels and a source node (i, j) , where $1 \leq i \leq m$, $1 \leq j \leq n$, an optimal broadcast scheme can be computed in $O((n + m)k)$ time.*

4 Complete Graphs

In this section we show that the broadcasting problem in multi-channel networks is NP-hard for complete bipartite graphs and complete graphs. Through this section, we assume there is a single channel on each edge of concerned graphs. Using a reduction from the exact cover problem, we show that the broadcasting problem is NP-hard for complete bipartite graphs; then we provide a reduction from the broadcasting problem in complete bipartite graphs to the same problem in complete graphs. The following lemma illustrates the above claim for complete bipartite graphs. The proof appears in Appendix B.

Lemma 4.1. *The broadcasting problem is NP-hard for complete bipartite graphs in the conflict-aware multi-channel model, even in the special case when there are 2 channels and the source is connected to all its neighbors with the same channel.*

We provide a reduction from the instance described in the above lemma to the broadcasting problem in complete graphs. The reduction uses a construction that we call a *ladder bipartite graph* (See Figure 4).

Definition 1. *A ladder bipartite graph with channels i, j is a balanced complete bipartite graph with n vertices on each side. There is a one-to-one mapping between the vertices of two sides such that the edge connecting a vertex u to its mapped vertex u' has channel j and all the other edges incident to u have channel i .*

Lemma 4.2. *Assume all vertices on one side of a ladder bipartite graph with channels i, j have received the message. If these vertices need to inform the vertices on the other side in one round, all the vertices should be active in that round, i.e., they need to transmit the message either through channel i or j .*

Proof. W.l.o.g., we may assume that i and j are 1 and 2 respectively. By contradiction, suppose a vertex u is inactive, so its opposite vertex u' should receive the message through channel 1 from another vertex v . Thus v uses channel 1 and its opposite vertex v' should receive from another vertex x via channel 1. Then there will be a conflict in u' since x and v both use channel 1, and thus not all vertices can be informed in one round, which is a contradiction. \square





Figure 4: (left) A ladder bipartite graph with channels 1, 2; grey edges are labeled with channel 1. (right) A schematic representation of a ladder graph.

Theorem 4.1. *The broadcasting problem in the conflict-aware multi-channel model is NP-hard for complete graphs, when there are at least 8 channels in the network.*

Proof. Given an instance (G, r_0) of the broadcasting problem in a complete bipartite graph in which there are two channels and the edges adjacent to the source r_0 are labeled with the same channel, we create an instance of the broadcasting problem in a complete graph in which there are 8 channels. Let L and R denote the two partitions of the vertices of G so that $r_0 \in R$. We create a complete graph H as follows (See Figure 5). We take two copies of L and three copies of $R - \{r_0\}$ (r_0 is the source in the original instance). Call these *components* L_1, L_2 , and R_1, R_2, R_3 , respectively, and also add a new vertex r as the new source. The channels of edges connecting vertices in L_1 and L_2 to any of R_1, R_2, R_3 are copied from the original bipartite graph G . Let vertex r be connected to the 5 components via 5 different channels so that the edges connecting r to the vertices in the same component have the same channel. Moreover, we assign the channels to the edges connecting vertices in R_1 to vertices in R_2 in a way that these edges form a ladder bipartite graph with channels 3, 4. Similarly, we set the edges between R_2 and R_3 to form a ladder graph with channels 5, 6, and between R_1 and R_3 to form a ladder with channels 7, 8. The edges connecting vertices in L_1 and L_2 get channel 3 and all other edges (the edges inside components) get arbitrary channels. We claim that there is a broadcasting scheme for the instance (G, r_0) that completes in 2 rounds if and only if there is a broadcasting scheme for (H, r) that also takes 2 rounds.

Assume there is a broadcasting scheme for (G, r_0) that completes in 2 rounds. In the first round r_0 informs the vertices of L via its single channel, so in the new instance r can inform the vertices of L_1 via the single channel that connects them (channel 1 in Figure 5). In the second round of the broadcast in (G, r_0) , a subset of L informs all vertices of R . In the new instance the same subset can inform all vertices of R_1, R_2, R_3 (via the same edges used in the first instance), while r informs L_2 via the unique connecting channel (channel 2 in Figure 5). Hence, the broadcast completes in 2 rounds.



(a) The original instance in complete bipartite graph G .

(b) The instance of the broadcasting problem in complete graph H .

Figure 5: The broadcasting problem in complete bipartite graphs (with one channel in source incident edges) reduces to the broadcasting problem in complete graphs. Here, a number i on the solid edge connecting two components indicate that all edges between the vertices of the two components are labeled with channel i . The channels of the edges between two components connected by curved blue edges are copied from the reduced bipartite graph. Solid and dashed paired lines indicate that the components form a ladder bipartite graph.

Now assume that there is a broadcasting scheme for (H, r) that completes in 2 rounds. First, we show that r cannot inform any of R_1, R_2, R_3 in the first round. For the sake of contradiction, suppose r informs R_1 in the first round (the same reasoning holds for R_2 and R_3); in the second round r cannot inform both L_1 and L_2 . Thus, at least one vertex in R_1 should use channels 1 or 2 to inform some vertices of L_1 and L_2 . Since the edges between R_1 and R_2 form a ladder bipartite graph and at least one vertex of R_1 is busy informing vertices of L_1 and L_2 , by Lemma 4.2, R_1 cannot inform all vertices of R_2 . Thus, some vertices of R_2 are to be informed by the source. Similarly, some vertices of R_3 are also left for the source to inform them. However, the source is connected to R_2 and R_3 with two different channels, thus it cannot inform both in a single round. Hence, the broadcast cannot be completed in 2 rounds and we get a contradiction. As a result, we may assume that in the first round r informs either L_1 or L_2 .

Assume r informs L_1 in the first round (the same reasoning holds for L_2). Since in the second round r can inform at most one of the R_i 's, the other two should be informed via L_1 , which implies a subset of vertices in L_1 can inform all vertices in two R_i 's. The same subset can be used for the instance (G, r_0) to inform all the vertices on the right in the second round. Therefore, there is a broadcasting scheme for (G, r_0) that completes in 2 rounds. \square

5 Balanced Complete Graphs

As the problem is NP-hard for complete graphs, we consider a particular case where there is a single channel on each edge and every node is connected to $\frac{n-1}{k}$ nodes through edges with the same channel. Thus, all the nodes use k different channels. We refer to this subset of complete graphs as *balanced complete graphs*. Since this would restrict us from considering networks where n is not congruent to one modulo k , we relax the condition slightly in order to include almost balanced assignments. For a given $\epsilon \geq 0$, we require that for every node v and every channel i , the number of nodes connected to v using channel i is at least $(1 - \epsilon)(n - 1)/k$ and at most $(1 + \epsilon)(n - 1)/k$. We call this family of graphs ϵ -*balanced complete graphs*.

In this setting, k corresponds to a trivial upper bound on the broadcast time. It suffices that the source transmits once through each channel and, since the graph is complete, the broadcasting is done. If we ignore all possible conflicts, it is easy to obtain a simple lower bound on the transmission time. Consider a graph where at any round a node can transmit to at most $(1 + \epsilon)(n - 1)/k$ nodes without conflicts. It is clear then that after the first round, we have at most $(1 + \epsilon)(n - 1)/k + 1$ nodes informed. The general formula for an upper bound for the number of nodes that have been informed after i rounds is $((1 + \epsilon)(n - 1)/k + 1)^i$, and thus we get a lower bound for the total number of rounds to inform all nodes.

Lemma 5.1. *Let $\epsilon \geq 0$. For ϵ -balanced complete graphs, at least $\lceil \log n / \log((1 + \epsilon)(n - 1)/k + 1) \rceil$ rounds are required to complete a broadcast.*

When $k = n - 1$ and $\epsilon < 1$ (i.e., each node is connected to exactly one node using each channel) a simple greedy algorithm finds the optimal broadcast scheme and it takes $\lceil \log_2 n \rceil$ rounds. This is because there are no conflicts when receiving the message, since all channels are different. The solution matches the lower bound in Lemma 5.1. This example shows that there are some cases where the broadcast time is not as bad as the trivial upper bound of k .

When aiming at practical applications, a more interesting question is when the number of channels is relatively small compared to the number of nodes. Note that for $k \leq (1 + \epsilon)(n - 1)/(\sqrt{n} - 1) = O(\sqrt{n})$ the lower bound in Lemma 5.1 asserts that the broadcast requires at least 2 rounds. Therefore, it would be desirable to have the property that there exists a constant $C > 0$

such that for every ϵ -balanced complete graph G with at most $C\sqrt{n}$ channels, a broadcast can always be completed in 2 rounds. Unfortunately, we can show that this is not true by constructing a counterexample using a random assignment of channels.

For given natural numbers n and k , let $G(n, k)$ be a complete graph with node set $[n] = \{1, 2, \dots, n\}$ in which two nodes are connected via channel $c \in [k]$ with probability $1/k$, independently for each such a pair. As is typical in random graph theory, we shall consider only asymptotic properties of $G(n, k)$ as $n \rightarrow \infty$, where k may and usually does depend on n . We say that an event in a probability space holds *asymptotically almost surely (a.a.s.)* if its probability tends to one as n goes to infinity. The following theorem shows that we cannot hope for the desired property when k is small enough (in particular, if k is a constant).

Theorem 5.1. *Let $\epsilon > 0$, $c_0 = \min_{c>0}(1 - ce^{-c}) = 1 - 1/e \in (0, 1)$, $f = f(n)$ is any function tending to infinity together with n , $k' = k'(M) = \log_{1/c_0} n - 3 \log_{1/c_0} \log n - M$, $k'' = \log_{1/c_0} n + f$, and $k''' = \sqrt{n/(2 \log n)}$. Then, there exists a sufficiently large constant M such that the following holds a.a.s.:*

- $G(n, k)$ is an ϵ -balanced complete graph for any k such that $2 \leq k \leq k'''$,
- Broadcasting in $G(n, k)$ requires at least 3 rounds for any k such that $3 \leq k \leq k'(M)$,
- Broadcasting in $G(n, k)$ requires 2 rounds for any k such that $k'' \leq k \leq k'''$.

The proof can be found in Appendix C. Although this is an asymptotic result, one can take n large enough to get that with probability at least $1/2$ the first and one of the other two properties hold, and thus we get a counterexample of this order. We conjecture that there exist constants $c_1, c_2 > 0$ such that for any ϵ -balanced complete graph G with $c_1 \log n \leq k \leq c_2 \sqrt{n/\log n}$ channels, a broadcast can always be completed in 2 rounds. Since a random assignment of channels seems to be a natural candidate for a counterexample, the fact that the conjecture holds a.a.s. for $G(n, k)$ strongly supports it.

Moreover, since there are ϵ -balanced complete graphs with bounded number of channels ($k = O(1)$) in which broadcasting requires at least 3 rounds, it is important to be able to design ϵ -balanced complete graphs (or even, balanced complete graphs) that can be broadcasted in 2 rounds. Next we present a construction algorithm together with a broadcasting scheme that completes the broadcast in 2 rounds. Note that since the topology is fixed (a complete graph), finding a construction algorithm is equivalent to finding a channel assignment. Our channel assignment algorithm relies on the following known result in graph theory for edge coloring.

Lemma 5.2. *[14, problem 16.5, p. 133] The minimum number of colors required for an edge coloring of a complete graph K_n is $n - 1$ if n is even, and n otherwise.*

A constructive proof of this lemma leads to the following edge-coloring algorithm. When n is odd, assign the color $((i + j) \bmod n) + 1$ to each edge $e = (v_i, v_j)$ with $v_i, v_j \in V = \{v_1, v_2, \dots, v_n\}$ (an edge-coloring for K_3 is shown in Figure 6(b)). We say that a node v_i *uses a color c* if there is an edge (v_i, v_j) , $i \neq j$ colored with c . For n even, the graph K_{n-1} is colored using the above method. Then, the remaining vertex v_n is added and each edge $e = (v_i, v_n)$ is colored with the color not used by v_i (See an edge-coloring for K_4 in Figure 6(a)). Note that *coloring an edge e* and *assigning a channel to e* are equivalent terms that define the same action in different problems. Thus, from now on we shall refer Lemma 5.2 using the terminology defined for our problem. Relying on this result, we obtain the following theorem (our construction algorithm follows immediately from its constructive proof).

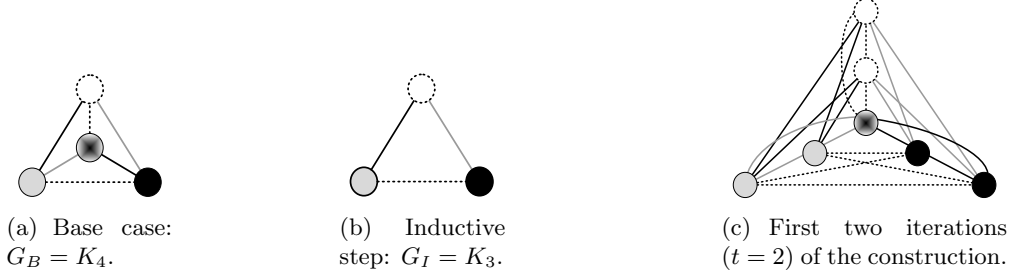


Figure 6: Construction example using K_4 as base case and K_3 in the inductive step.

Theorem 5.2. *Given an odd number of channels k and a positive integer t , it is possible to construct a balanced complete graph with $kt + 1$ nodes (i.e., K_{kt+1}).*

Proof. We prove the theorem by induction on t . For our construction, we use a more restricted hypothesis as follows. Given k and t as in the statement of the theorem, we can construct a complete graph with $kt + 1$ nodes that satisfies the following properties: (i) the vertices of the graph can be classified in k classes with t vertices in each class (and one root vertex in no class), (ii) vertices in the same class are all connected with one channel, and are connected to the root with that same channel, and (iii) for each pair of classes, all edges connecting vertices in the two classes are connected with the same channel. It is not hard to see that proving this hypothesis proves the theorem.

Let $G_B = K_{k+1}$ be the base case. As we define k to be odd, G_B is a complete graph with an even number of nodes. Then, we can assign k different channels to G_B in such a way that no two edges adjacent to the same node use the same channel (by Lemma 5.2). Note that each node in G_B uses a different channel to connect with the other k nodes. Let us call the last node added by the coloring given by Lemma 5.2 the *root*. We assign each non-root node to a class defined by the channel that connects it to the root. In Figure 6(a), for example, the root is the center node, and we name each non-root node with one of the 3 channels (black, gray, and dashed).

For the inductive step, assume G_t is a complete graph with $kt + 1$ nodes satisfying the desired properties. We add k new nodes to G_t to form G_{t+1} . For this sake, we connect all vertices of G_t to the vertices of a complete graph $G_I = K_k$. Thus G_{t+1} is a complete graph with $kt + 1 + k = k(t + 1) + 1$ vertices. Since k is odd, we can assign k different channels to G_I in such a way that no two edges adjacent to the same node use the same channel (by Lemma 5.2). By construction, each node in G_I uses $k - 1$ different channels. We assign each node to the class corresponding to the channel it does not use. Thus G_{t+1} satisfies (i).

Let $class(c)$ be the set of nodes in G_{t+1} that belong to the class corresponding to channel c . We assign channel c to each edge (u, v) such that $u, v \in class(c)$, and also to each edge $(u, root), \forall u \in class(c)$. Thus G_{t+1} satisfies (ii), and all the nodes in the same class are interconnected and connected with the root using the channel that defines the class. The remaining step is to assign channels to edges with end-points in different classes. Consider two classes c_1 and c_2 . By property (iii) all edges in G_t connecting nodes in these classes are labeled with the same channel. We assign this channel to all edges (u, v) such that $u \in class(c_1)$ and $v \in class(c_2)$, with $u \in G_I$ and $v \in G_t$. This step is repeated for all pairs of classes. Finally, observe that since the color assignment for G_B given by Lemma 5.2 builds on the assignment for G_I , for any pair of classes, edges connecting vertices in these classes have the same color in both G_B and G_I . Thus for all pairs of classes c_1 and c_2 , the edge (u, v) with $u, v \in G_I$ and $u \in class(c_1)$ and

$v \in \text{class}(c_2)$ has the same color of the edges in G_t connecting vertices in $\text{class}(c_1)$ to vertices in $\text{class}(c_2)$. Hence G_{t+1} satisfies (iii), completing the proof. \square

Figure 6 shows an example of the construction algorithm with $k = 3$ channels (thus, a balanced complete graph with $3t + 1$ nodes). K_4 with 3 different channels is used as the base case in the inductive construction. The graph used in the inductive steps is a K_3 designed using a channel assignment with 3 different channels. The algorithm iteratively adds K_3 at each step. Figure 6(c) shows how the construction algorithm connects G_I and G_B to obtain the final graph.

We now analyze the number of rounds required to broadcast in a graph constructed with our algorithm.

Theorem 5.3. *Let G be a complete graph with k channels and at least $k^2 - 2k + 1$ nodes constructed according to the inductive algorithm described in Theorem 5.2. Then, a broadcast in G from any node can be completed in 2 rounds.*

Proof. First, consider the case when the source of the broadcast is the root. In the first round, it informs all the nodes in one class (using for example channel 1). The number of nodes in the graph guarantees that there are at least $k - 2$ nodes in that class. Thus, after the first round, at least $k - 1$ nodes are informed. In the second round, it suffices that the root transmits through a channel different from the one used in round 1 (e.g., channel 2) to inform a new class, and other $k - 2$ nodes among the informed ones take care of the remaining $k - 2$ classes. Note that conflicts do not occur because all the nodes informed in the first round belong to the same class.

When the source is not the root, in the first round the source transmits through the channel that defines its class. This results in at least $k - 1$ informed nodes (including the source and the root). In the second round, the source transmits through a different channel to inform one class, while other $k - 2$ nodes among the informed ones take care of the remaining $k - 2$ classes (one of those nodes may be the root). Note that conflicts do not occur even when the root is transmitting (no other node is transmitting to the same class). \square

The broadcasting scheme follows from this constructive proof. Notice also that the channel assignment together with the broadcasting scheme constitute a fault-tolerant system. The network may be much larger than $k^2 - 2k + 1$ nodes, and this broadcast scheme will still work when some of the nodes fail. More precisely, if the root and $k - 2$ nodes in each class do not fail, a message can be broadcasted to all the nodes that are working properly in 2 rounds.

Finally, we prove that our channel assignment is also efficient when several messages need to be broadcasted from different sources at the same time. Specifically, up to k messages can be broadcasted simultaneously, and all the broadcasts complete in 3 rounds. We formalize this in Theorem 5.4. The fault-tolerance property that holds for the broadcast of one message holds as well for this scheme. The proof appears in Appendix D.

Theorem 5.4. *Let G be a complete graph with k channels and at least $k^2 - 2k + 1$ nodes constructed according to the inductive algorithm described in Theorem 5.2. Then, a broadcast in G of k messages from any k different nodes can be completed in 3 rounds.*

6 Conclusions

We studied the broadcasting problem in conflict-aware multi-channel networks, and presented positive and negative results for various network topologies. These include polynomial time algorithms that give optimal broadcast schemes for grids, and also for trees when there is a

single channel on each edge. We proved that the problem is NP-hard for trees in general case, and also for complete graphs even in the restricted case with only one channel on each edge. Moreover, we studied balanced complete graphs as a subclass of complete graphs in which each node is connected to approximately same number of nodes with each channel. In this setting, we proposed a channel assignment that results in broadcast schemes that complete in two rounds, which is optimal for non-trivial networks. Besides, we proved that broadcasting in some balanced complete graphs requires at least three rounds, thus justifying the significance of our construction. The construction results in fault-tolerant networks that enable efficient broadcasting of multiple messages simultaneously.

References

- [1] Arachchige, C.L., Venkatesan, S., Chandrasekaran, R., Mittal, N.: Minimal time broadcasting in cognitive radio networks. In: ICDCN'11. pp. 364–375 (2011)
- [2] Bloom, S., Korevaar, E., Schuster, J., Willebrand, H.: Understanding the performance of free-space optics. *Journal of Optical Networking* 2(6), 178–200 (2003)
- [3] Elkin, M., Kortsarz, G.: Sublogarithmic approximation for telephone multicast: Path out of jungle. In: SODA'03. pp. 76–85 (2003)
- [4] Farley, A.M., Hedetniemi, S.T.: Broadcasting in grid graphs. In: Proc. 9th S-E Conf. Combinatorics, Graph Theory, and Computing. pp. 275–288 (1978)
- [5] Garey, M.R., Johnson, D.S.: *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA (1990)
- [6] Hedetniemi, S., Hedetniemi, S., Liestman, A.: A survey of gossiping and broadcasting in communication networks. *Networks* 18(4), 319–349 (1988)
- [7] Janson, S., Łuczak, T., Ruciński, A.: *Random Graphs*. Wiley, New York (2000)
- [8] Kondareddy, Y., Agrawal, P.: Selective broadcasting in multi-hop cognitive radio networks. In: IEEE Sarnoff Symposium. pp. 1–5 (2008)
- [9] Kyasanur, P., Vaidya, N.: Routing and interface assignment in multi-channel wireless networks. In: WCNC'05. vol. 4, pp. 2051–2056 (2005)
- [10] Li, L., Qin, B., Zhang, C., Li, H.: Efficient broadcasting in multi-radio multi-channel and multi-hop wireless networks based on self-pruning. In: HPCC'07. pp. 484–495 (2007)
- [11] Qadir, J., Chou, C.T., Misra, A., Lim, J.G.: Localized minimum-latency broadcasting in multi-radio multi-rate wireless mesh networks. In: WOWMOM'08. pp. 1–12 (2008)
- [12] Ramachandran, K.N., Belding, E.M., Almeroth, K.C., Buddhikot, M.M.: Interference-aware channel assignment in multi-radio wireless mesh networks. In: INFOCOM'06. pp. 1–12 (2006)
- [13] Raniwala, A., Gopalan, K., cker Chiueh, T.: Centralized channel assignment and routing algorithms for multi-channel wireless mesh networks. *Mobile Computing and Communications Review* 8(2), 50–65 (2004)
- [14] Soifer, A.: *The Mathematical Coloring Book*. Springer (2009)
- [15] Subramanian, A., Buddhikot, M., Miller, S.: Interference aware routing in multi-radio wireless mesh networks. In: WiMesh'06. pp. 55–63 (2006)
- [16] Wieselthier, J.E., Nguyen, G.D., Ephremides, A.: Energy-efficient broadcast and multicast trees in wireless networks. *MONET* 7(6), 481–492 (2002)
- [17] Wilkinson, B., Allen, M.: *Parallel Programming: Techniques and Applications Using Networked Workstations and Parallel Computers*. Prentice Hall (2005)
- [18] Zagalj, M., Hubaux, J.P., Enz, C.C.: Minimum-energy broadcast in all-wireless networks: NP-completeness and distribution issues. In: MOBICOM'02. pp. 172–182 (2002)
- [19] Zeng, G., Wang, B., Ding, Y., Xiao, L., Mutka, M.W.: Efficient multicast algorithms for multichannel wireless mesh networks. *IEEE Transactions on Parallel and Distributed Systems* 21(1), 86–99 (2010)

A Proof (sketch) of Theorem 3.2

Assume first that no node is a splitter in G for the given source (i, j) . In this case, lower and upper bounds for broadcasting coincide with the ones in the telephone model. Broadcasting in this model for grid graphs was studied in [4], in which broadcast times are shown for any location of the source. In the case of no splitters, the following simple strategy achieves the lower bounds in [4]. Let $|Q_i|$ denote the diameter of quadrant i , which equals the minimum distance from (i, j) to the corner of Q_i . Let i_1, i_2, i_3, i_4 be indices such that $|Q_{i_1}| \geq |Q_{i_2}| \geq |Q_{i_3}| \geq |Q_{i_4}|$. The message is transmitted from (i, j) along row i in the direction corresponding to the corner of Q_{i_1} , and simultaneously (starting in the second round) in the same row but in the opposite direction. When an end node of the row is reached, the nodes in row i between the source and the end node start informing the nodes in their corresponding columns, first in the direction of the quadrant with the largest diameter. Let $t(Q_i)$ denote the time that it takes to inform all nodes in quadrant Q_i . This strategy achieves $t(Q_{i_1}) = |Q_{i_1}|$, $t(Q_{i_2}) \leq |Q_{i_1}| + 1$, $t(Q_{i_3}) \leq |Q_{i_1}| + 1$, and $t(Q_{i_4}) \leq |Q_{i_1}| + 2$.

Broadcasting times may differ with respect to the telephone model when there are splitters. Note, however, that the strategy above is optimal even when there are splitters if the sizes of all quadrants are different (i.e., the source is not in the middle row or column): in this case $t(Q_i) \leq |Q_{i_1}|$ for all $i = 1 \dots 4$, and $|Q_{i_1}|$ is a lower bound.

Splitters make a difference in the case when the source is in the middle row or middle column, or both. For the sake of brevity, we describe here only the case in which the source is in the middle of the grid. Assume that the source is at $(i, j) = ((m + 1)/2, (n + 1)/2)$, and thus all quadrants have the same diameter $d = |Q_1|$.

The following strategy takes advantage of splitters to reduce the broadcasting time by transmitting the message simultaneously through the critical path of two quadrants. Suppose the source is a splitter. If it is a vertical splitter but not a horizontal one, we start informing nodes in column j above and below simultaneously. If it is only a horizontal splitter, we inform nodes in row i , to the left and right. If it is both, we pick an orientation based on other splitters: inform vertically if there is at least one splitter between Q_1 and Q_2 , and one between Q_3 and Q_4 , and inform horizontally otherwise. If there are no splitters, just like before, when messages reach the end of the row, each node in a row informs nodes above and below in their corresponding column. Suppose, for example, that the message reaches a splitter (i, j') dividing Q_2 and Q_3 . Nodes on row i that have the message act as if the message had reached the end of the row: they inform above and below in their column in two rounds. The splitter informs $(i - 1, j')$ and $(i + 1, j')$ in one round. When the message reaches the beginning and end of column j' , each of the nodes in the column informs to the right, completing the broadcast in quadrants Q_2 and Q_3 at the same time (See Figure 3 (b)). The broadcast is analogous for other splitters.

It is not hard to see that if there is at least one splitter in each of the directions in which the source informs, then the broadcasting time is d , and it is $d + 1$ otherwise. These are optimal. To see this, note that nodes in N (i.e., nodes in the same row or column as the source) are in the critical path from the source to two corners. Since at least one of these nodes must send messages to neighbors in different quadrants, if this cannot be done simultaneously one of the quadrants will suffer a delay of one round.

Suppose now that the source is not a splitter. The source informs vertically in column j if there is at least one splitter in this column, and horizontally otherwise, sending the message first in the direction opposite of the splitter (or any if both directions have splitters). For example, if the only splitter is in row i between quadrants Q_2 and Q_3 , the source informs first to its left neighbor and then to its right neighbor. The broadcast is completed in $d + 1$ rounds, which is




Figure 7: An instance of exact cover with neighborhood. ($X = \{a_1, a_3\}, u$) is a solution.




Figure 8: Reduction from exact cover to exact cover with neighborhood.

again optimal: since the source is not a splitter, none of the nodes of two quadrants will have received a message in the first round. This inevitably adds one round to the lower bound of d .

The same arguments can be used to show that the broadcast can be completed optimally in cases when the source is in the middle row or column but not both. Note as well that the strategies above also apply in the cases when there are two or three empty quadrants. Computing the optimal scheme for each case requires calculating the quadrants' diameters (which takes constant time), and possibly searching for splitters, which can be done in $O((n+m)k)$ time.

B Proof of Lemma 4.1

We use a reduction from the exact cover problem. To simplify the proof, we define *exact cover with neighborhood* as a variant of exact cover. We show that exact cover reduces to exact cover with neighborhood and then show a reduction from exact cover with neighborhood to the broadcasting problem in complete bipartite graphs. Recall that the exact cover problem is defined over bipartite graphs; given a bipartite graph in which the vertex set is partitioned into *left* and *right* subsets, the exact cover problem asks if there is a subset S of vertices on the left such that all vertices on the right are connected to exactly one of the vertices in S . This problem is a classical NP-hard problem [5].

Definition 2. *The exact cover with neighborhood problem is a decision problem, which given a bipartite graph $G = (V, E)$ [$V = L \cup R$, where L and R are vertices on the left and right sides respectively] asks if there exists a vertex $u \in L$ and also a subset $X \subseteq L, u \notin X$, such that all neighbors of u are exactly covered by X , i.e., any neighbor of u is connected to exactly one vertex in X (See Figure 7).*

Lemma B.1. *Exact cover with neighborhood is NP-hard.*

Proof. We show a reduction from the exact cover problem. Given a bipartite graph G as an instance of exact cover, let $L = \{a_1, a_2, \dots, a_n\}$ and $R = \{b_1, b_2, \dots, b_m\}$ be the vertices of G on the left and right sides, respectively. We create a bipartite graph H as follows: start with a copy of G and add n vertices x_1, x_2, \dots, x_n on the right. Connect each x_i to all vertices on the left except a_i . Moreover, add a single vertex y on the left, and connect it to all b_i 's and none of x_i 's (See Figure 8).

We claim that there is an exact cover for G if and only if there is an exact cover with neighborhood for H . Assume there is an exact cover for G , i.e., there is a subset of L that exactly covers all members of R . Note that the members of R are exactly the neighbors of y ; so all neighbors of y are covered and we are done.

Now assume there is an exact cover with neighborhood for H , so there is a vertex u such that all its neighbors on the right are exactly covered by some other vertices on the left. We claim that y is the only vertex with this property. For the sake of contradiction, suppose that a_t ($t \leq n$) is such a vertex. All x_i 's except x_t are connected to a_t , hence they should be covered by some of the vertices in $L - \{a_t, y\}$ (note that y is not connected to x_i 's). Moreover, at most one of the vertices in $L - \{a_t, y\}$ can be selected since any two vertices of this set will conflict in $n - 2$ of x_i 's. Let a_α ($\alpha \neq t$) be the selected vertex, i.e., a_α should cover all x_i 's but x_t . In particular a_α should cover x_α , which is not possible as they are not connected by definition. So y is the vertex with all its neighbors covered by other vertices on the left. Since y is connected to all vertices of G on the right, there is an exact cover for G , which completes the proof. \square

Next, we show a reduction from the exact cover with neighborhood problem to broadcasting in complete bipartite graphs. Given an instance of exact cover with neighborhood, which is a bipartite graph $H = (L, R)$, we create a complete bipartite graph with a single channel on each edge as follows. We start with a copy of H with channel 1 on all edges and add two vertices r and v on the right side. In order to form a complete bipartite graph G , we add all missing edges and assign channel 2 to them (See Figure 9). Also, let r be the source in an instance of the broadcasting problem defined over G . We claim that there is an exact cover with neighborhood for H if and only if there is a broadcasting scheme that completes in 2 rounds for G .

Assume there is a subset $X \subseteq L$ that exactly covers all neighbors of a vertex $u \in L$. In the broadcast instance, in the first round r informs all vertices on the left (using channel 2). In the second round, the vertices in X use channel 1 to inform the neighbors of u (there will be no conflict by definition of exact cover), and u uses channel 2 to inform other vertices including v . Hence, the broadcast completes in 2 rounds.

Now assume there is a broadcasting scheme that completes in 2 rounds. So v receives the message in round 2 as it is not connected to r . Also, since all incident edges of v are labeled with channel 2, there is exactly one vertex on the left that can be used to transmit on channel 2 in round 2, otherwise there will be a conflict in v . Let u be the vertex on the left that uses channel 2 to inform v and its own non-neighbors in H . The other vertices on the right are neighbors of u in H , which receive the message via some other vertices on the left. This set of vertices forms an exact cover for neighbors of u in H .

C Proof of Theorem 5.1

In this section, we will use the following version of the Chernoff bound (See for example Theorem 2.8 [7]). Let Z be a random variable that can be expressed as a sum $Z = \sum_{i=1}^n Z_i$ of independent random indicator variables where $Z_i \in \text{Be}(p_i)$ with (possibly) different $p_i = \mathbb{P}(Z_i = 1) = \mathbb{E}Z_i$. Then the following holds for $t \geq 0$:

$$\begin{aligned} \mathbb{P}(Z \geq \mathbb{E}Z + t) &\leq \exp\left(-\frac{t^2}{2(\mathbb{E}Z + t/3)}\right), \\ \mathbb{P}(Z \leq \mathbb{E}Z - t) &\leq \exp\left(-\frac{t^2}{2\mathbb{E}Z}\right). \end{aligned}$$




Figure 9: The exact cover with neighborhood instance and the resulting broadcasting scheme (the bold edges have channel 1 and the rest have channel 2).

The proof that $G(n, k)$ is an ϵ -balanced complete graph a.a.s. follows immediately from the Chernoff bound. Note that for a given node v and channel c , we expect $(n-1)/k = \Omega(\sqrt{n})$ neighbors connected through channel c . Hence, with probability $1 - o(n^{-2})$ the number of edges of this type is $n/k + O(\sqrt{n/k} \log n) = (1 + o(1))n/k$. We get that a.a.s. for every node v and every channel c the number of edges adjacent to v with channel c assigned is $(1 + o(1))\frac{n}{k}$. Hence, the property holds for any $\epsilon > 0$.

Let $3 \leq k \leq k'$. Now, we will show that broadcasting in $G(n, k)$ requires at least 3 rounds a.a.s. Fix a node v that transmits at the first round, nodes from the set X receive the message ($|X| = (1 + o(1))n/k$). Now, fix nodes that are going to transmit the message during the second round, and assign channels to them. In particular, y_i nodes of X use channel $i \in [k]$.

We have two possibilities to consider: a) v transmits during the second round, and b) v does not transmit during the second round. For a given sequence (y_1, y_2, \dots, y_k) , the total number of configurations to consider (that is, the number of broadcasting schemes) is at most

$$nk(k+1) \prod_i \binom{n}{y_i} \leq n^{3 + \sum_i y_i}.$$

(There are n choices for v , k choices for a channel used by v at the first round, $k+1$ choices for a behavior of n in the second round. Finally, $\binom{n}{y_i}$ nodes transmit on channel i at the second round.)

Fix any configuration of nodes of X sending the message during the second round. Let u be a node in $V \setminus (X \cup \{v\})$. The probability that node u does not receive the message from X is

$$\prod_{i=1}^k \left(1 - \frac{y_i}{k} \left(1 - \frac{1}{k} \right)^{y_i - 1} \right).$$

(We select one node that sends the message to u via channel i (term y_i). The probability that the edge from this node to u has label i is $1/k$. Nobody else sending on channel i at this round can reach u (term $(1 - \frac{1}{k})^{y_i - 1}$.) So in case a) we expect to see

$$(1 + o(1)) \left(1 - \frac{1}{k} \right) n \prod_{i=1}^k \left(1 - \frac{y_i}{k} \left(1 - \frac{1}{k} \right)^{y_i - 1} \right)$$

nodes not informed after two rounds. Case b) is slightly more complicated but it is easy to provide a lower bound on the number of nodes not receiving the message. Some of the nodes

receiving message from v (via channel i) at the second round can have conflicts from nodes of X sending on the same channel (unless $y_i = 0$). Therefore, perhaps more nodes do not get the message but, in any case, we expect to see at least

$$(1 + o(1)) \left(1 - \frac{2}{k}\right) n \prod_{i=1}^k \left(1 - \frac{y_i}{k} \left(1 - \frac{1}{k}\right)^{y_i-1}\right)$$

nodes not informed after two rounds.

If $k \geq 3$ is a constant, then clearly a positive fraction of nodes remain uniformed after two rounds a.a.s., and we are done. Suppose then that $k = k(n)$ grows together with n . Due to the symmetry, the probability is minimized when $y_i = y/k + O(1)$ ($y = \sum_i y_i$) for all values i , and so the probability of not being informed is at least

$$(1 + o(1)) \left(1 - (1 + O(k^{-1})) \frac{y}{k^2} \exp\left(-\frac{y}{k^2}\right)\right)^k.$$

If $y/k \rightarrow c$, then this is asymptotic to $\Theta((1 - ce^{-c})^k)$ and so we get that the probability of being not informed is of order at least c_0^k . (Recall that $c_0 = \min_{c>0}(1 - ce^{-c}) = 1 - 1/e$.) Therefore, the expected number of nodes not informed is at least

$$(1 + o(1))(1 - 2/k)n\Omega(c_0^k) \geq C(M) \log^3 n,$$

where $C(M)$ is a function of M that grows to infinity together with M . (Recall that $k \leq \log_{1/c_0} n - 3 \log_{1/c_0} \log n - M$.) It follows from the Chernoff bound that the probability that every node is informed is at most $\exp(-C(M)(\log^3 n)/2)$.

Since the total number of configurations with $y \leq 2k^2$ is at most

$$y^k n^{3+y} \leq \exp((3 + k + y) \log n) \leq \exp(O(\log^3 n)),$$

we can use the union bound (with $C(M)$ large enough) to get that a.a.s. there is no broadcasting scheme that informs every node after two rounds, provided that $y \leq 2k^2$ nodes are active at the second round. For $y > 2k^2$ we use the fact that the probability of u not receiving a message via a random channel i is not $(1 + o(1))c_0$ anymore but slightly larger (for $k^2/y = o(1)$ it is, in fact, tending to one). It is straightforward to check that there is no chance to have more active nodes ($y > 2k^2$) in the two-round broadcasting scheme.

Suppose now that $k'' \leq k \leq k'''$. We will show that this time there is a broadcasting scheme that informs all nodes in two rounds a.a.s. Choose any node v and any channel c . At the first round, v transmits on channel c , X nodes are informed ($|X| = (1 + o(1))n/k$). The probability of receiving a message at the second round is maximized when exactly k nodes transmit on each of k channels. However, this is not possible for large values of k . Therefore, we need to consider three sub-cases depending on the range for k .

Suppose first that $k \leq \frac{1}{2}n^{1/3}$. Since $(2 + o(1))n^{2/3}$ nodes are informed after the first round, there is no problem with desired assignment. Choose any k^2 nodes of X , partition them such that exactly k nodes transmit on each channel. The expected number of nodes not informed after the second round is

$$O(n) \left(1 - (1 + O(k^{-1}))/e\right)^k = O(nc_0^k) = o(1),$$

and so a.a.s. every node receives the message by Markov's inequality.

Suppose now that $(1/2)n^{1/3} < k \leq n^{2/5}$. This time all nodes of X transmit the message: $(1 + o(1))n/k^2$ nodes transmit on each channel. The expectation is now

$$O(n) \left(1 - (1 + o(1)) \frac{n/k^2}{k} \left(1 - \frac{1}{k} \right)^{n/k^2} \right)^k \leq \exp(\log n - O(n/k^2)) = o(1).$$

Finally, for $n^{2/5} < k \leq k'''$, the calculations can be done slightly more carefully to get that the expectation is

$$O(n) \left(1 - (1 + o(1)) \frac{n/k^2}{k} \right)^k \leq \exp\left(\log n - (1 + o(1)) \frac{n}{k^2}\right) = o(1).$$

The result holds by Markov's inequality, as before.

D Proof of Theorem 5.4

First, consider the case when the k sources of the broadcast belong to different classes. Let m_i be the message that must be broadcasted from a source node in $class(c_i)$. Initially, each source transmits through the channel that defines its class. Therefore, after the first round all the nodes in $class(c_i)$ have been informed of message m_i . Note that the root receives the k messages in the first round. W.l.o.g., let us define a total order in the classes using the name of the channel that defines the class (i.e., $class(c_1) \prec \dots \prec class(c_k)$). We say that $class(c_i)$ transmits a message to $class(c_j)$ if a node in $class(c_i)$ transmits through the channel that connects it with all the nodes in $class(c_j)$. In the second round, each $class(c_i)$ transmits the message m_i to the $k - i$ higher classes. Note that $k - i$ nodes in $class(c_i)$ transmit in this second round. As the graph has at least $k^2 - 2k + 1$ nodes, every class has at least $k - 2$ nodes, and $class(c_1)$ can use the root to complete the $k - 1$ necessary transmissions. After the second round, all the nodes in $class(c_i)$ have been informed of messages m_1, \dots, m_i . In the third round, each $class(c_i)$ transmits m_i to the $i - 1$ lower classes. Analogously to the previous round, $i - 1$ nodes in $class(c_i)$ transmit in this third round, and the number of nodes in the graph guarantees that the class have enough nodes for all classes but c_k , which uses the root to complete the $k - 1$ necessary transmissions. Observe that the order in the transmissions avoids conflicts in the second and third round.

In general, the k sources may not be distributed one in each class. For each class with more than one source we arbitrarily choose one of the sources and call it a proper source. All other sources in the class are non-proper. A class without any source is called *orphan*. Note that there are as many orphan classes as non-proper sources. In the first round, each proper source transmits through the channel that defines its class, thus informing all the nodes in its class, and each non-proper source *adopts* one of the orphan classes (i.e., it transmits through the channel that connects it with all the nodes in the orphan class). After this round the situation is similar to the case where each source belongs to one class. The main difference is that the root has not been informed of any of the messages broadcasted from non-proper sources. We now define a total order among classes by choosing any of the classes with a proper source as the first class, and we assign any arbitrary order to the rest of the classes. Note that there is at least one proper source. Then, the second and third rounds are analogous to the case in which each class has one source, with the difference that in the second round one of the nodes of each adopted class informs the root. Note that the first class is the only one that needs to transmit through all the channels in the second round, which is why we choose a non-orphan class to be the first class. The special case where the root is one of the sources does not imply any difference in the algorithm because the root can be considered as a proper source of one of the orphan classes.