

Probabilistic Forwarding of Coded Packets on Networks

B. R. Vinay Kumar and Navin Kashyap

Department of Electrical Communication Engineering
Indian Institute of Science, Bengaluru, India
Email: {vinaykb, nkashyap}@iisc.ac.in

Abstract—We consider a scenario of broadcasting information over a network of nodes connected by noiseless communication links. A source node in the network has k data packets to broadcast, and it suffices that a large fraction of the network nodes receives the broadcast. The source encodes the k data packets into $n \geq k$ coded packets using a maximum distance separable (MDS) code, and transmits them to its one-hop neighbours. Every other node in the network follows a probabilistic forwarding protocol, in which it forwards a previously unreceived packet to all its neighbours with a certain probability p . A “near-broadcast” is when the expected fraction of nodes that receive at least k of the n coded packets is close to 1. The forwarding probability p is chosen so as to minimize the expected total number of transmissions needed for a near-broadcast. In this paper, we analyze the probabilistic forwarding of coded packets on two specific network topologies: binary trees and square grids. For trees, our analysis shows that for fixed k , the expected total number of transmissions increases with n . On the other hand, on grids, we use ideas from percolation theory to show that a judicious choice of n will significantly reduce the expected total number of transmissions needed for a near-broadcast.

I. INTRODUCTION

The Internet of Things (IoT) involves different types of physical devices — sensors, actuators, routers, mobiles etc. — communicating with each other over a network. Each node in the network has minimal computational ability and limited knowledge of the network topology. Broadcast mechanisms are often required in such ad-hoc networks to disburse important network-related information, for example, to carry out over-the-air programming of the IoT nodes. Further, these mechanisms need to be completely distributed and must impose minimal computational burden on the nodes. Broadcast mechanisms such as flooding, although being distributed and reliable, are not efficient, since there are excessive transmissions and consequently a high energy expenditure [1]. To overcome this, probabilistic forwarding of received packets may be employed (see [2], [3]), wherein each node either forwards a previously unreceived packet to all its one-hop neighbors with probability p or takes no action with probability $1 - p$.

In a previous paper [4], we studied the effect of introducing redundancy in the form of coded packets into this probabilistic forwarding protocol. We describe the setup here. Consider a large network with a particular node designated as the source. The source has k message packets to send to a large fraction of nodes in the network. The k message packets are first encoded into $n \geq k$ coded packets using a maximum distance separable

(MDS) code (see e.g., [5, Ch. 11]). The source transmits each of the n coded packets to all its one-hop neighbours. All the other nodes in the network use the probabilistic forwarding mechanism as described before. Subsequent receptions of the same packet are ignored.

Our goal is to analyze the performance of the above algorithm. In particular, we wish to find the minimum retransmission probability p for which the expected fraction of nodes receiving at least k out of the n coded packets is close to 1, which we deem a “near-broadcast”. (The condition for a near-broadcast stems from our use of an MDS code.) This probability results in the least expected total number of transmissions, across all the network nodes, needed for a near-broadcast. The expected total number of transmissions is a measure of the energy expenditure in the network.

Simulation results presented in [4] indicate that over a wide range of network topologies (including the important case of random geometric graphs, but not including tree-like topologies), the expected total number of transmissions initially decreases to a minimum and then increases with n . Our aim is to understand this behaviour and predict, via analysis, the value of n that minimizes the expected number of transmissions. While we would ultimately like to explain this behaviour on random geometric graphs, which constitute an important model for wireless ad-hoc networks [6], we have not yet developed the tools required for the analysis there. In this paper, we present an analysis for trees and grids.

The rest of the paper is organized as follows. Section II provides a theoretical formulation of the problem. In Sections III and IV, we consider the problem on rooted binary trees and grids, respectively, and provide bounds and estimates for the expected number of transmissions. Section V concludes the paper with a discussion on other graphs where our ideas can be applied. Proofs for most of our results are provided in the full version of this paper [7].

II. PROBLEM FORMULATION

The problem formulation here is essentially reproduced from [4]. Consider a graph $G = (V, E)$, where V is a set of N vertices (nodes), and E is a set of edges (noiseless communication links). A source node $s \in V$ has k message packets which need to be broadcast over the network. The source s encodes the k messages into n coded packets using an MDS code. The MDS code ensures that any node that receives

at least k of the n coded packets can retrieve the original k message packets by treating the unreceived packets as erasures. It is assumed that all the required encoding/decoding operations are carried out over a sufficiently large field, so that an MDS code with the necessary parameters exists. It is also assumed that each packet has a header which identifies the packet index $j \in [n] := \{1, 2, \dots, n\}$.

The source node broadcasts all n coded packets to its one-hop neighbours, after which the probabilistic forwarding protocol takes over. When a packet (say, packet # j) is received by a node for the first time, it either transmits it to all its one-hop neighbours with probability p or does nothing with probability $1 - p$. The node ignores all subsequent receptions of packet # j , irrespective of the decision it took at the time of first reception. Each packet is forwarded independently of other packets and other nodes. This probabilistic forwarding continues until there are no further transmissions in the system. The protocol indeed must terminate after finitely many transmissions since each node in the network chooses to forward a particular coded packet only the first time it is received.

We are interested in the following scenario. Let $R_{k,n}$ be the number of nodes, including the source node, that receive at least k out of the n coded packets. Given a $\delta \in (0, 1)$, let $p_{k,n,\delta}$ be the minimum forwarding probability p for a near-broadcast, i.e.,

$$p_{k,n,\delta} := \inf \left\{ p \mid \mathbb{E} \left[\frac{R_{k,n}}{N} \right] \geq 1 - \delta \right\}. \quad (1)$$

The performance measure of interest, denoted by $\tau_{k,n,\delta}$, is the expected total number of transmissions across all nodes when the forwarding probability is set to $p_{k,n,\delta}$. Here, it should be clarified that whenever a node forwards (broadcasts) a packet to all its one-hop neighbours, it is counted as a single (simulcast) transmission. Our aim is to determine, for a given k and δ , how $\tau_{k,n,\delta}$ varies with n , and the value of n at which it is minimized. To this end, it is necessary to first understand the behaviour of $p_{k,n,\delta}$ as a function of n . The following simple lemma is valid for any connected graph $G = (V, E)$.

Lemma II.1. *For fixed values of k and δ ,*

- (a) $p_{k,n,\delta}$ is a non-increasing function of n .
- (b) $p_{k,n,\delta} \rightarrow 0$ as $n \rightarrow \infty$.

Proof: (a) For any $n > 0$, the random variables $R_{k,n}$ and $R_{k,n-1}$ can be coupled as follows: If the k data packets are encoded into n coded packets, then $R_{k,n-1}$ (resp. $R_{k,n}$) is realized as the number of nodes, including the source node, that receive at least k of the first $n-1$ (resp. at least k of the n) coded packets. It is then clear that $\mathbb{E}[\frac{1}{N}R_{k,n}] \geq \mathbb{E}[\frac{1}{N}R_{k,n-1}]$, and hence, by (1), we have $p_{k,n,\delta} \leq p_{k,n-1,\delta}$.

(b) From the n coded packets, create $\lfloor \frac{n}{k} \rfloor$ non-overlapping (i.e., disjoint) groups of k packets each. For $i = 1, 2, \dots, \lfloor \frac{n}{k} \rfloor$, let A_i be the event that the i th group of k coded packets is received by at least $(1 - \delta/2)N$ nodes. The events A_i are mutually independent and have the same probability of occurrence. For any $p > 0$, we have $\mathbb{P}(A_i)$ being strictly positive (but perhaps small). Hence,

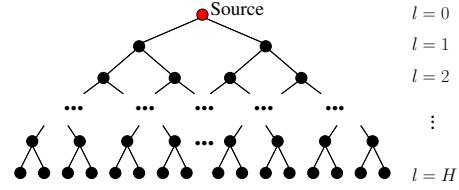


Fig. 1. A rooted binary tree of height H .

$$\mathbb{P}(\text{at least one } A_i \text{ occurs}) = 1 - (1 - \mathbb{P}(A_1))^{\lfloor \frac{n}{k} \rfloor} \geq 1 - \frac{\delta}{2}$$

for all sufficiently large n , so that $\mathbb{P}\left(\frac{R_{k,n}}{N} \geq 1 - \frac{\delta}{2}\right) \geq 1 - \frac{\delta}{2}$. Consequently, $\frac{\mathbb{E}[R_{k,n}]}{N} \geq (1 - \frac{\delta}{2})(1 - \frac{\delta}{2}) \geq 1 - \delta$. Thus, for any $p > 0$, we have $p_{k,n,\delta} \leq p$ for all sufficiently large n . ■

On the other hand, $\tau_{k,n,\delta}$ typically exhibits more complex behaviour. As demonstrated via simulations in [4], over a wide range of graph topologies (both deterministic and random), except notably for trees (see Section III), $\tau_{k,n,\delta}$ initially decreases and then grows gradually as n increases. This trend can be seen most clearly in a grid topology — see Section IV. Thus, there typically is an optimal value of n that minimizes $\tau_{k,n,\delta}$. This happens due to an interplay between two opposing factors: as n increases, $p_{k,n,\delta}$ decreases (Lemma II.1), which contributes towards a decrease in $\tau_{k,n,\delta}$. But this is opposed by the fact that the overall number of transmissions tends to increase when there are a larger number of packets traversing the network. To determine the value of n that minimizes $\tau_{k,n,\delta}$, we need more precise estimates of $p_{k,n,\delta}$, and subsequently, of $\tau_{k,n,\delta}$. For the specific graph topologies of binary trees and grids, we are able to obtain such estimates using methods tailored to those topologies, as we demonstrate in the sequel.

III. ROOTED BINARY TREES

Consider a rooted binary tree of height $H \geq 2$ as depicted in Fig. 1. The root of the tree is the source node and is at level $l = 0$. The root node encodes the k data packets into n coded packets and transmits them to its children. Every other node on the tree follows the probabilistic forwarding strategy with some fixed forwarding probability $p > 0$. We will assume that the nodes at level H (i.e., the leaf nodes) do not transmit, as there is nothing to be gained in allowing them to do so.

On a rooted binary tree of height $H \geq 2$, when k is fixed and $\delta > 0$ is sufficiently small, simulation results presented in [4] show that $\tau_{k,n,\delta}$ increases with n . A large-deviations analysis aiming to explain these results was also attempted in [4]. However, this analysis is only valid in the regime where k and n are both large but the ratio k/n is fixed. Thus, the analysis in [4] does not in fact explain the simulation results.

In this section, we present results for the regime when k is fixed and n is allowed to vary. We also fix a $H \geq 2$ and a $\delta \in (0, \frac{1}{8})$. In the interest of brevity, we only provide the statements of the results here. Detailed derivations can be found in [7].

It was shown in [4] that $p_{k,n,\delta}$ is the least value of $p \in [0, 1]$ for which¹

¹This is a re-arrangement of Eq. (4) in [4].

$$\frac{\sum_{l=0}^{H-1} 2^{l+1} \mathbb{P}(Z_l \leq k-1)}{2^{H+1} - 1} \leq \delta, \quad (2)$$

where $Z_l \sim \text{Bin}(n, p^l)$ for $l = 0, 1, \dots, H-1$, and

$$\tau_{k,n,\delta} = n \left[\frac{(2p_{k,n,\delta})^H - 1}{2p_{k,n,\delta} - 1} \right]. \quad (3)$$

An analysis starting from (2) yields the two propositions below, which provide good lower and upper bounds on $p_{k,n,\delta}$. These bounds are plotted, for $k = 100$, $\delta = 0.1$ and $H = 50$, in Fig. 2(a) along with the exact values of $p_{k,n,\delta}$ obtained numerically from (2). The corresponding plots for $\tau_{k,n,\delta}$, obtained via (3), are shown in Fig. 2(b).

Proposition III.1. Let $k \geq 2$, $H \geq 2$, and $0 \leq \delta < \frac{1}{8}$ be fixed. For all $n \geq k$, we have $p_{k,n,\delta} > \left(\frac{k-1}{n}\right)^{\frac{1}{H-1}}$.

In the case of $k = 1$ and $n > 1$, the lower bound can be improved to $p_{k,n,\delta} > \left(\frac{1}{n}\right)^{\frac{1}{H-1}}$.

Proposition III.2. Let $k \geq 2$, $H \geq 2$, and $0 < \delta \leq 1$ be fixed, and let $\delta' := \min \left\{ \delta \left(\frac{2^{H+1}-1}{2^{H+1}-2} \right), 1 \right\}$. Then, for all $n \geq 1$, we have $p_{k,n,\delta} \leq \min \left\{ \left(\frac{k-1+t}{n} \right)^{\frac{1}{H-1}}, 1 \right\}$, where $t = \sqrt{2(k-1)(-\ln \delta') + (\ln \delta')^2} - \ln \delta'$.

In the case of $k = 1$, the bound $p_{k,n,\delta} \leq \min \left\{ \left(\frac{-\ln \delta'}{n} \right)^{\frac{1}{H-1}}, 1 \right\}$ holds for all $n \geq 1$.

The following theorem summarizing the behaviour of $p_{k,n,\delta}$ on binary trees is a direct consequence of these propositions.

Theorem III.3. Let $k \geq 2$, $H \geq 2$ and $0 < \delta < \frac{1}{8}$ be fixed. We then have $p_{k,n,\delta} = \Theta \left(\left(\frac{k}{n} \right)^{\frac{1}{H-1}} \right)$, where the constants within the Θ -notation² may be chosen to depend only on H and δ .

The plots in Fig. 2 corroborate the simulation results reported in [4], thus providing a theoretical explanation for why $\tau_{k,n,\delta}$ increases with n . Another confirmation of this behaviour can be obtained by substituting $p_{k,n,\delta} \approx c \left(\frac{k}{n} \right)^{\frac{1}{H-1}}$, for any positive constant $c \equiv c(H, \delta)$ into the expression for $\tau_{k,n,\delta}$ in (3). This yields the approximation

$$\tau_{k,n,\delta} \approx n \left[\frac{(2c)^H \left(\frac{k}{n} \right)^{\frac{H}{H-1}} - 1}{2c \left(\frac{k}{n} \right)^{\frac{1}{H-1}} - 1} \right], \quad (4)$$

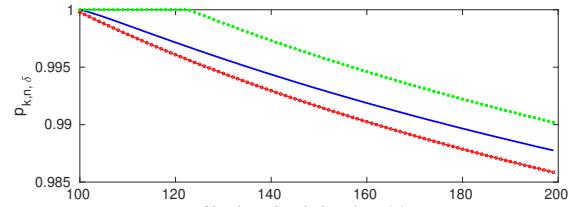
which can be shown to be increasing in n .

The analysis in this section extends easily to the case of rooted d -ary trees, for any $d \geq 2$. In summary, introducing redundancy in the form of coding to the probabilistic retransmission protocol on a rooted d -ary tree is not beneficial in terms of the overall energy expenditure in the network.

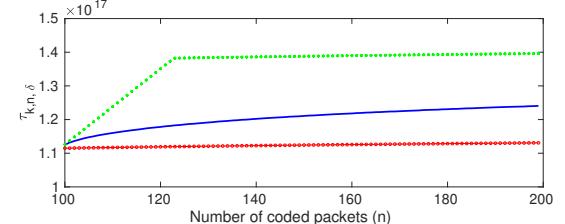
IV. GRIDS

For odd m , we consider the $m \times m$ grid $\Gamma_m := \left[-\frac{m-1}{2}, \frac{m-1}{2} \right]^2 \cap \mathbb{Z}^2$ centred at the origin. The source node

²The notation $a(n) = \Theta(b(n))$ means that there are positive constants c_1 and c_2 such that $c_1 b(n) \leq a(n) \leq c_2 b(n)$ for all sufficiently large n .



(a) Minimum retransmission probability



(b) Expected total number of transmissions

Fig. 2. The middle curves are plots of the true values of $p_{k,n,\delta}$ and $\tau_{k,n,\delta}$ obtained from (2) and (3), for $k = 100$, $\delta = 0.1$ and $H = 50$. The other curves are bounds obtained via Propositions III.1 and III.2, and (3).

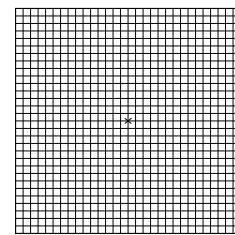


Fig. 3. The source node (\times) is at the centre of the 31×31 grid.

is assumed to be at the centre of the grid. Fig. 3 depicts this for $m = 31$. Simulation results for the probabilistic forwarding algorithm on grids of various sizes were presented in [4]. Some results from simulations on the Γ_{31} and Γ_{501} grids are shown in Figs. 5 and 6 in Section IV-C. In this section, we try to explain these observations by developing an analysis that is at least valid for large m . Specifically, we turn to the theory of site percolation on the integer lattice \mathbb{Z}^2 to explain the $p_{k,n,\delta}$ and $\tau_{k,n,\delta}$ curves obtained via simulations on large grids Γ_m .

A. Site percolation on \mathbb{Z}^2

We start with a brief description of the site percolation process (see e.g. [8]) on \mathbb{Z}^2 . This is an i.i.d. process $(X_u)_{u \in \mathbb{Z}^2}$, with $X_u \sim \text{Ber}(p)$ for each $u \in \mathbb{Z}^2$, where the probability $p \in [0, 1]$ is a parameter of the process. A node or *site* $u \in \mathbb{Z}^2$ is *open* if $X_u = 1$, and is *closed* otherwise. For $u = (u_x, u_y) \in \mathbb{Z}^2$, define $|u| := |u_x| + |u_y|$. Two sites u and v are joined by an edge, denoted by $u-v$, iff $|u-v| = 1$. The next few definitions are made with respect to a given realization of the process $(X_u)_{u \in \mathbb{Z}^2}$. Two sites u and v are connected by an *open path*, denoted by $u \longleftrightarrow v$, if there is a sequence of sites $u_0 = u, u_1, u_2, \dots, u_n = v$ such that u_k is open for all $k \in \{0, 1, \dots, n\}$ and $u_{k-1}-u_k$ for all $k \in [n]$. The *open cluster*, C_u , containing the site u is defined as $C_u = \{v \in \mathbb{Z}^2 | u \longleftrightarrow v\}$. Thus, C_u consists of all sites connected to u by open paths. In particular, $C_u = \emptyset$ if u is

itself closed. The *boundary*, ∂C_u , of a non-empty open cluster C_u is the set of all closed sites $v \in \mathbb{Z}^2$ such that $v - w$ for some $w \in C_u$. The set $C_u^+ := C_u \cup \partial C_u$ is called an *extended cluster*. The cluster C_u (resp. C_u^+) is termed an *infinite open cluster (IOC)* (resp. *infinite extended cluster (IEC)*) if it has infinite cardinality. Note that C_u^+ is infinite iff C_u is infinite.

It is well-known that there exists a *critical probability* $p_c \in (0, 1)$ such that for all $p < p_c$, there is almost surely³ no IOC, while for all $p > p_c$, there is almost surely a unique IOC. We do not know what happens at $p = p_c$, as the exact value of p_c is itself not known (for site percolation on \mathbb{Z}^2). It is believed that $p_c \approx 0.59$ [8, Ch. 1]. Another quantity of interest, which will play a crucial role in our analysis, is the *percolation probability* $\theta(p)$, defined to be the probability that the origin $\mathbf{0}$ is in an IOC. In our analysis, we also consider the probability, $\theta^+(p)$, of the origin $\mathbf{0}$ being in an IEC. Clearly, for $p < p_c$, we have $\theta^+(p) = \theta(p) = 0$; for $p > p_c$, it is not difficult to see that $\theta^+(p) \geq \theta(p) > 0$. The following lemma expresses $\theta^+(p)$ in terms of $\theta(p)$.

Lemma IV.1. *For any $p > p_c$, we have $\theta^+(p) = \frac{\theta(p)}{p}$.*

Proof. Let C and C^+ be the (unique) IOC and IEC, respectively. We then have $\theta(p) = \mathbb{P}(\mathbf{0} \in C) = \mathbb{P}(\mathbf{0} \in C^+ \text{ and } \mathbf{0} \text{ is open})$. Now, observe that the event $\{\mathbf{0} \in C^+\}$ is determined purely by the states of the nodes other than $\mathbf{0}$. Hence, this event is independent of the event that $\mathbf{0}$ is open. Thus, $\theta(p) = \mathbb{P}(\mathbf{0} \in C^+ \text{ and } \mathbf{0} \text{ is open}) = \theta^+(p) \cdot p$. \square

There is, unfortunately, no analytical expression known for $\theta(p)$. Fig. 4 plots $\theta(p)$ and $\theta^+(p)$ as functions of p , the former being obtained via simulations based on the theorem below.

Theorem IV.2. *Let $p > p_c$, and let C and C^+ , respectively, be the (almost surely) unique IOC and IEC of a site percolation process on \mathbb{Z}^2 with parameter p . Then, almost surely, we have*

$$\lim_{m \rightarrow \infty} \frac{|C \cap \Gamma_m|}{m^2} = \theta(p) \quad \text{and} \quad \lim_{m \rightarrow \infty} \frac{|C^+ \cap \Gamma_m|}{m^2} = \theta^+(p).$$

The theorem is obtained as a straightforward application of an ergodic theorem for a multi-dimensional i.i.d. random field [9, Proposition 8]. Using the dominated convergence theorem (DCT), we also have $\lim_{m \rightarrow \infty} \mathbb{E} \left[\frac{|C \cap \Gamma_m|}{m^2} \right] = \theta(p)$ and $\lim_{m \rightarrow \infty} \mathbb{E} \left[\frac{|C^+ \cap \Gamma_m|}{m^2} \right] = \theta^+(p)$. Based on this, to obtain an estimate of $\theta(p)$, the site percolation process with parameter p was simulated on a 501×501 grid and the average fraction of nodes (averaged over 100 realizations of the process) in the largest open cluster was taken to be the value of $\theta(p)$. These are the values of $\theta(p)$ plotted in Fig. 4. We would like to emphasize that the plots in the figure should only be trusted for $p > p_c$, as Theorem IV.2 is only valid in that range. However, as the exact value of p_c is unknown, simulation results are reported for the range of p values shown in the plot.

³with respect to the product measure $\otimes_u \nu_u$, with $\nu_u \sim \text{Ber}(p) \ \forall u \in \mathbb{Z}^2$.

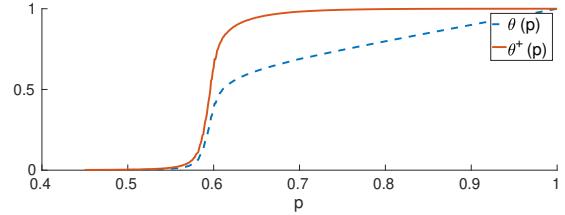


Fig. 4. $\theta(p)$ and $\theta^+(p)$ vs. p

B. Relating site percolation to probabilistic forwarding

Site percolation on \mathbb{Z}^2 is a faithful model for probabilistic forwarding of a single packet on the infinite lattice \mathbb{Z}^2 . The origin $\mathbf{0}$ is the source of the packet. The open cluster, C_0 , containing $\mathbf{0}$ corresponds to the set of nodes that transmit (forward) the packet, and the extended cluster C_0^+ corresponds to the set of nodes that receive the packet. The only caveat is that, since the source is assumed to always transmit the packet, we must consider only those realizations of the site percolation process in which the origin $\mathbf{0}$ is open. In other words, we must consider the site percolation process conditioned on the event that the origin is open. By extension, the probabilistic forwarding of n coded packets on the lattice \mathbb{Z}^2 corresponds to n independent site percolation processes on \mathbb{Z}^2 , conditioned on the event that the origin is open in all n percolations.

C. Analysis of probabilistic forwarding on a large grid

Our analysis of probabilistic forwarding on the finite grid Γ_m is based on the approximation described next. For the purposes of this discussion, we fix a forwarding probability p . Let $\mathcal{R}_{k,n}^\infty$ denote the set of all nodes that receive at least k of the n coded packets in the probabilistic forwarding protocol on \mathbb{Z}^2 . We use $|\mathcal{R}_{k,n}^\infty \cap \Gamma_m|$ as a proxy for $R_{k,n}^m$, which is the number of nodes receiving at least k out of n packets in the probabilistic forwarding protocol on Γ_m . In general, it is only true that $R_{k,n}^m$ is stochastically dominated⁴ by $|\mathcal{R}_{k,n}^\infty \cap \Gamma_m|$, since a node in $\mathcal{R}_{k,n}^\infty \cap \Gamma_m$ could receive packets from the origin through paths in \mathbb{Z}^2 that do not lie entirely within Γ_m . Nonetheless, we proceed under the assumption that $\mathbb{E}[R_{k,n}^m] \approx \mathbb{E}[|\mathcal{R}_{k,n}^\infty \cap \Gamma_m|]$ for large m . This is vindicated by the fact that our analysis based on this assumption matches the simulation results reasonably well — see Figs. 5 and 6.

Recall also that we want values of the forwarding probability p for which $\mathbb{E}[\frac{1}{m^2} R_{k,n}^m]$ is at least $1 - \delta$, for some (small) $\delta > 0$. Hence, we need $\mathbb{E}[\frac{1}{m^2} |\mathcal{R}_{k,n}^\infty \cap \Gamma_m|] \geq 1 - \delta$. If we would like this to hold for all sufficiently large m , then p must be such that $\mathcal{R}_{k,n}^\infty$ has infinite cardinality. This implies, due to the correspondence between probabilistic forwarding and site percolation on \mathbb{Z}^2 , that p must be such that there exist infinite (open/extended) clusters in the site percolation process. Thus, we must operate in the super-critical region $p > p_c$. It can also be seen from the simulation results in Figs. 5 and 6 that $\tau_{k,n,\delta}$ is minimized when $p_{k,n,\delta}$ is in the super-critical region. We use this to justify considering only $p > p_c$ in our analysis.

⁴A random variable X is stochastically dominated by a random variable Y if $\mathbb{P}(X \geq x) \leq \mathbb{P}(Y \geq x)$ for all $x \in \mathbb{R}$. For non-negative random variables, this implies that $\mathbb{E}[X] \leq \mathbb{E}[Y]$.

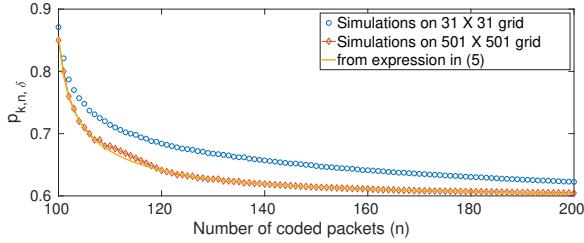


Fig. 5. Comparison of the minimum forwarding probability obtained via simulations on a 31×31 grid and a 501×501 grid, with the results obtained numerically from (5), for $k = 100$ data packets and $\delta = 0.1$.

The following theorem is the main result of this section. The proof, given in [7], is obtained by carefully relating $\mathcal{R}_{k,n}^\infty$ to the set, $C_{k,n}^+$, of all sites in \mathbb{Z}^2 that belong to the IEC containing $\mathbf{0}$ in at least k out of n independent percolations, conditioned on $\mathbf{0}$ being open in all n percolations.

Theorem IV.3. For $p > p_c$, $\lim_{m \rightarrow \infty} \mathbb{E} \left[\frac{1}{m^2} |\mathcal{R}_{k,n}^\infty \cap \Gamma_m| \right] = \sum_{t=k}^n \sum_{j=k}^t \binom{n}{t} \binom{t}{j} (\theta^+(p))^{t+j} (1 - \theta^+(p))^{n-j}$.

From the discussion prior to the theorem, the left-hand side of the equality stated in the theorem is our proxy for $\lim_{m \rightarrow \infty} \mathbb{E} \left[\frac{1}{m^2} R_{k,n}^m \right]$. Thus, for large grids Γ_m , we take $p_{k,n,\delta}$ to be the least value of p for which

$$\sum_{t=k}^n \sum_{j=k}^t \binom{n}{t} \binom{t}{j} (\theta^+(p))^{t+j} (1 - \theta^+(p))^{n-j} \geq 1 - \delta. \quad (5)$$

This can be evaluated numerically using the values of $\theta^+(p)$ plotted in Fig. 4. The results thus obtained are shown in Fig. 5. It is clear that these results match very well with those obtained from simulations on a 501×501 grid.

We next look into estimating the expected total number of transmissions at a given forwarding probability p . Consider the transmission of a single packet on the infinite lattice \mathbb{Z}^2 . The set of nodes transmitting this packet is simply the open cluster C_0 in the percolation framework. Thus, arguing as for packet receptions above, the expected number of transmissions for probabilistic forwarding on a large (but finite) grid Γ_m is well-approximated by $\mathbb{E}[|C_0 \cap \Gamma_m| \mid \mathbf{0} \text{ is open}]$. In [7], we infer from Theorem IV.2 the fact that for $p > p_c$,

$$\lim_{m \rightarrow \infty} \frac{\mathbb{E}[|C_0 \cap \Gamma_m| \mid \mathbf{0} \text{ is open}]}{m^2} = \frac{\theta(p)^2}{p}.$$

Thus, in probabilistic forwarding of a single packet on a large grid Γ_m , the expected number of transmissions, normalized by the grid size m^2 , is approximately $\frac{\theta(p)^2}{p}$. Hence, when we have n coded packets, by linearity of expectation, the expected total number of transmissions, again normalized by the grid size m^2 , is approximately $n \frac{\theta(p)^2}{p}$. In particular, setting $p = p_{k,n,\delta}$, we obtain

$$\frac{1}{m^2} \tau_{k,n,\delta} \approx n \frac{\theta(p_{k,n,\delta})^2}{p_{k,n,\delta}}, \quad (6)$$

provided that $p_{k,n,\delta} > p_c$.

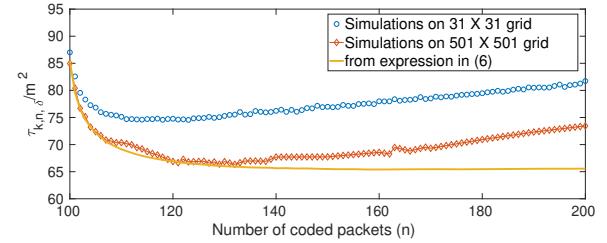


Fig. 6. Comparison of the expected total number of transmissions normalized by the grid size m^2 , obtained via simulations on Γ_{31} and Γ_{501} , with the expression from (6), for $k = 100$ data packets and $\delta = 0.1$.

Fig. 6 compares, for $k = 100$ data packets and $\delta = 0.1$, the values of $\frac{1}{m^2} \tau_{k,n,\delta}$ obtained using (6) with those obtained via simulations on the Γ_{31} and Γ_{501} grids. The curve based on (6) initially tracks the Γ_{501} curve well, but trails off after $n = 130$. Given that the corresponding $p_{k,n,\delta}$ curves are well-matched (Fig. 5), this is perhaps attributable to the fact that the $\theta(p)$ values from Fig. 4 are not very reliable. As even a small change in $\theta(p)$ would significantly affect $n \frac{\theta(p)^2}{p}$, better estimates of $\theta(p)$ may correct the discrepancy observed. Nonetheless, our analysis gives theoretical validation, at least for large grids, for the observed initial decrease in $\tau_{k,n,\delta}$ as a function of n , thus indicating a benefit to introducing some coding into the probabilistic forwarding mechanism on grids.

V. CONCLUSION

The ideas we have used for establishing our results for grids can be extended to other graph structures as well. The key tool for obtaining our results for grids was Theorem IV.2, which was an application of an ergodic theorem for percolation on the infinite \mathbb{Z}^2 lattice. Similar theorems can be obtained for graphs such as the triangular and hexagonal grids through ergodic theorems for percolation on the corresponding infinite lattices.

ACKNOWLEDGEMENTS

The research presented in this paper was supported in part by a Cisco PhD Fellowship awarded to the first author, and in part by the DRDO-IISc “Frontiers” research programme.

REFERENCES

- [1] Y.-C. Tseng, S.-Y. Ni, Y.-S. Chen, and J.-P. Sheu, “The broadcast storm problem in a mobile ad hoc network,” *Wireless Networks*, vol. 8, no. 2/3, pp. 153–167, 2002.
- [2] Y. Sasson, D. Cavin, and A. Schiper, “Probabilistic broadcast for flooding in wireless mobile ad hoc networks,” in *Proc. WCNC 2003*, vol. 2, March 16–20, 2003, pp. 1124–1130.
- [3] Z. J. Haas, J. Y. Halpern, and L. Li, “Gossip-based ad hoc routing,” *IEEE/ACM Trans. Networking*, vol. 14, no. 3, pp. 479–491, 2006.
- [4] B. R. Vinay Kumar, R. Antony, and N. Kashyap, “The effect of introducing redundancy in a probabilistic forwarding protocol,” in *Proc. NCC 2018*, IIT-Hyderabad, Feb 25–28, 2018. [Online]. Available: <http://arxiv.org/abs/1901.02033>
- [5] R. M. Roth, *Introduction to Coding Theory*. Cambridge Univ. Press, 2006.
- [6] R. Vaze, *Random Wireless Networks*. Cambridge Univ. Press, 2015.
- [7] B. R. Vinay Kumar and N. Kashyap, “Probabilistic forwarding of coded packets on networks,” preprint, 2019. [Online]. Available: http://ece.iisc.ac.in/~nkashyap/Papers/prob_fwding_long.pdf
- [8] G. Grimmett, *Percolation*, 2nd ed. Springer-Verlag, 1999.
- [9] C. Newman and L. Schulman, “Infinite clusters in percolation models,” *Journal of Statistical Physics*, vol. 26, no. 3, pp. 613–628, 1981.