A Comparative Analysis of Flooding Methods on Random and Real Network Topologies

Asterios Papamichail[1], Georgios Tsoumanis[1], Spyros Sioutas[2], Konstantinos Oikonomou[1]

[1] Department of Informatics, Ionian University, Corfu, Greece, [2] Computer Engineering and Informatics Department, University of Patras, Patras, Greece

Corresponding author: Asterios Papamichail, aspapa@ionio.gr

The Internet of Things (IoT) is revolutionizing industries by connecting everyday objects, known as smart devices, via the Internet. These devices, embedded with sensors and communication technologies, gather and share data. For the guaranteed gathering of information, the devices share global knowledge with each other, by using dissemination mechanisms in order to broadcast information. This study evaluates four flooding methods for broadcasting information across network nodes, namely: (i) blind flooding; (ii) probabilistic flooding; (iii) m-probabilistic flooding; and (iv) scoped probabilistic flooding, the latter to be introduced here. The evaluation considers random networks that are based on the Burr Type XII distribution and seven real networks. The evaluated flooding methods are studied on three different metrics: (i) coverage achieved; (ii) number of messages exchanged; and (iii) a metric that is based on binomial approximation. The latter is introduced to provide deeper insights into the particulars of the under-evaluation flooding methods. The results show that, under certain conditions, m-probabilistic flooding outperforms probabilistic flooding in terms of coverage, while requiring significantly fewer messages. Additionally, the study revealed that the scoped probabilistic flooding achieves coverage comparable to that of the probabilistic flooding while reducing the number of exchanged messages.

Keywords: dissemination, IoT, probabilistic flooding, real networks

1. INTRODUCTION

The Internet of Things (IoT) describes a network where everyday objects, known as "smart devices," are connected via the Internet. In most cases, these devices are embedded with sensors, processors, and communication technologies, enabling them to gather and share data. The IoT is revolutionizing various sectors, including healthcare, transportation, manufacturing, and energy [1, 2], while encompassing a wide array of devices such as smartphones, vehicles, home appliances, industrial equipment, and even wearable technology. These devices are interconnected, facilitating communication among themselves and with cloud-based systems. This interconnectivity not only allows for the accumulation and analysis of data, providing novel insights and improved efficiency, but also supports the remote automation and control of these devices.

Flooding, a technique for broadcasting messages to every node within a network, is frequently utilized across a range of networking protocols and mechanisms to enable efficient information dissemination. This method is particularly effective in networks consisted of a high density of devices, facilitating rapid and effective information spread. Flooding is integral to a variety of applications, including established routing protocols, data collection, dissemination processes, and device discovery methods. In the realm of the IoT, several mechanisms employ variations of flooding. These mechanisms include, among others, the Constrained Application Protocol (CoAP) [3], Zigbee [4], the protocol for low-power and lossy networks (RPL) [5], and Z-Wave [6]. Each of these applications

Part of this work was originally published in the proceedings of the International Balkan Conference on Communications and Networking (BalkanCom 2023), Istanbul, Turkey, June 5-8, 2023.

leverages the flooding method to enhance network communication and information exchange.

The core concept of flooding, known as *blind flooding*, was introduced in [7]. In a network that utilizes blind flooding, each node is responsible for broadcasting a message to all its neighbors. These neighbors then relay the message further to their respective neighbors. This process continues, propagating the message throughout the network until it either reaches its intended destination or all nodes in the network have received it. While blind flooding is widely used in various applications to ensure reliable information dissemination and to achieve *global coverage*, it can be resource-intensive. Specifically, it may require a substantial number of packet transmissions to cover the entire network. As a result, blind flooding is susceptible to the *broadcast storm problem* as noted in [8]. This issue arises when an excessive number of broadcast messages are sent simultaneously by network nodes, leading to bandwidth overload, network congestion, and potentially severe degradation in network performance, sometimes even rendering the network inoperable.

To mitigate the broadcast storm problem, *probabilistic flooding* has been proposed as an alternative. This approach deviates from the comprehensive broadcast method of blind flooding. In probabilistic flooding, a node does not transmit a message to every neighboring node; instead, it sends the message to a randomly chosen subset of nodes, based on a predetermined forwarding probability. The key benefit of this technique lies in its ability to diminish redundant network traffic. By selectively forwarding packets to only a portion of the network's nodes, it significantly reduces the generation and transmission of duplicate packets. This, in turn, eases network congestion and enhances overall network performance. However, determining the ideal forwarding probability for comprehensive network coverage is a complex task that often necessitates extensive experimentation for optimal results. Furthermore, probabilistic flooding might not be universally applicable or effective across all network types and traffic conditions, especially in networks characterized by high dynamics or congestion levels.

m-Probabilistic flooding, given in [9], is a probabilistic flooding variation where each node forwards the information message to m (randomly selected) neighbor nodes. It was introduced as an alternative for studying the network operations when there is a lack of information regarding the required knowledge (i.e., node degree and largest eigenvalue in [9]'s case). Interestingly, the authors there demonstrated that m-probabilistic flooding could achieve global coverage with a value of m equal to or greater than 4, even in the absence of this critical information.

The fourth considered flooding mechanism is a blend of scoped flooding [10] and probabilistic flooding. Unlike traditional flooding methods that aim for broad network coverage, scoped flooding focuses on limiting the spread of information to specific areas or "scopes" within a network. This method is particularly effective in reducing redundant data transmissions and mitigating network congestion. To have a clearer comparison among the various techniques, the "scoped" ability is added to probabilistic flooding to create the new *scoped probabilistic flooding*.

This study concentrates on assessing the performance of the four distinct flooding variations given above: (i) blind flooding; (ii) probabilistic flooding; (iii) m-probabilistic flooding; and (iv) scoped probabilistic flooding, with initial results featured in [11]. These variations are examined within two types of networks: random networks based on the Burr Type XII distribution [12], and seven real networks. The evaluation focuses on three metrics: coverage, the number of messages, and a specific metric Δ. Δ, originally associated with probabilistic flooding in [9], employs a binomial approximation and is redefined in this study to assess m-probabilistic flooding, specifically its average forwarding probability. Furthermore, an advanced version of the Δ metric for m-probabilistic flooding is proposed, incorporating the consideration of variable forwarding probability.

The evaluation indicated that although probabilistic flooding typically outperforms other methods, there were instances where m-probabilistic flooding achieved better coverage with significantly fewer message exchanges. Furthermore, the study found that the metric Δ, redefined in this paper for the m-probabilistic flooding approach, exhibited behavior similar to the Δ metric previously established for probabilistic flooding in existing literature. Additionally, the introduced scoped probabilistic flooding appeared to achieve comparable coverage to probabilistic flooding but with a reduced number of messages transmitted.

The rest of the paper is organized as follows. Section 2 presents the past related work. In Section 3, the system model is presented along with the definition of Δ that is based on the binomial approximation for probabilistic flooding and the definitions of the two versions of the metric for the m-probabilistic flooding. The results from the evaluation of the random networks are presented in Section 4. Section 5 presents the collected results from the real topologies. Finally, the conclusions are drawn in Section 6.

2. PAST RELATED WORK

There are numerous studies that make use of a probabilistic approach [13] to broadcast information instead of blind flooding [7], in order to avoid the broadcast storm problem [8] and to decrease the overhead of the dissemination. Additionally, probabilistic approaches are used to reduce overhead that is generated by malicious nodes [14]. The probabilistic approaches can be categorized as approaches using fixed forwarding probabilities [15] and approaches that use variable forwarding probabilities depending on node properties [16, 17]. Other approaches [18, 19, 20] use fuzzy logic to select the forwarding probability. Another dissemination mechanism is that of scoped flooding [10, 21], where the information message is transmitted to all nodes of the selected region of the network.

Reina *et al.* [22] studied probabilistic broadcast schemes for wireless ad hoc networks. More specifically, they provided a classification of probabilistic broadcast schemes along with a thorough review on them. Additionally, their evaluation metrics include a broadcast efficiency group that considers the reachability of the schemes, the redundancy which measures the overhead of the broadcast schemes, the time elapsed during the broadcast and the energy consumption. Ruiz *et al.* [23] proposed a new taxonomy of the broadcasting schemes. Additionally, they provided a thorough review of the most noticeable broadcasting protocols.

m-Probabilistic flooding was first introduced by Oikonomou *et al.* [9]. In this paper, the authors evaluated m-probabilistic flooding on geometric random graphs and on Erdős-Rényi networks. Additionally, the authors proposed a spectral approach for selecting the forwarding probability of the probabilistic flooding along with a metric Δ that is based on the binomial approximation and the results showed that the approach covers more than 95% of the network nodes.

3. SYSTEM MODEL

This section presents: (i) definitions of the system model; (ii) descriptions of the evaluated flooding methods; and (iii) the definitions of the metric Δ for all evaluated methods. More specifically, Subsection 3.1 presents general definitions of the system model while Subsection 3.2 describes the evaluated flooding methods. Finally, the definition of the metric Δ which is based on binomial approximation for all the evaluated methods is presented in Subsection 3.3.

3.1 General definitions

Let each network topology of N nodes be represented by an undirected graph $G = (V(G), E(G))$, where $V(G)$ and $E(G)$ are the set of nodes and the set of links between them, respectively. A node of the graph is represented by an integer number $i = 1, 2, ..., N$. Let n_i be the set of neighboring nodes of node i. The node degree (i.e., the number of neighbor nodes) of node i is denoted by $d_i = |n_i|$. Let $\bar{d}$, be the average node degree, given by $\bar{d} = \frac{1}{N}\sum_{i=1}^{N} d_i$. $\mathcal{D}$ denotes the diameter (i.e., the maximum eccentricity) of a graph and D denotes the density of a graph.

Let $\mathcal{A}$ denote the $N \times N$ adjacency matrix of a graph G. If $(\mathcal{A})_{ij}$ is the element of row i and column j, then $(\mathcal{A})_{ij} = 1$, for $(i, j) \in E(G)$ and $(\mathcal{A})_{ij} = 0$, for $(i, j) \notin E(G)$. Since the graph in this paper is assumed to be undirected, $(\mathcal{A})_{ij} = (\mathcal{A})_{ji}$ applies, and the nodes cannot be connected to themselves i.e., $(i, i) \notin E(G)$, and $(\mathcal{A})_{ii} = 0$. Consequently, $\mathcal{A}$ is a real symmetric matrix, i.e., $\mathcal{A} = \mathcal{A}^\top$ and let, $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_N$ denote the eigenvalues of $\mathcal{A}$ in decreasing order.

3.2 The flooding methods

All evaluated methods are initiated by a network's node, to be called hereafter the *initiator*, which is the first node that transmits the information message. When a node receives the information message, it becomes "covered", while if it receives the information message more than once, the information message is discarded. For analysis purposes, it is assumed here that the information messages are transmitted and received in discrete time steps t. The total number of steps that flooding needs to cover all nodes of the network, or reaches a point that it cannot cover any other nodes, is denoted by T.

3.2.1 Blind flooding

Under blind flooding, when a node receives the information message for the first time, forwards the message to all of its neighbors. The information message travels this way until all nodes receive it or until there are no other nodes able to forward it.

3.2.2 Probabilistic flooding

In the context of probabilistic flooding, upon receiving an information message for the first time, a node alters its behavior by transmitting the information selectively to its neighboring nodes based on a predetermined *forwarding probability* denoted as q.

Five different cases of q are explored here, encompassing scenarios beyond the conventional blind flooding which is actually the situation where $q = 1$. These cases include specific values, $q = 0.2, 0.5, 0.7$ and a unique instance represented by $q = 4/\lambda_1$. The first three values of q (i.e., $0.2, 0.5, 0.7$) are straightforward and are employed to evaluate their proximity to blind flooding outcomes.

For the case of $q = 4/\lambda_1$, where λ_1 signifies the largest eigenvalue calculated from the network's adjacency matrix, this specific value has been highlighted in existing literature as the minimum forwarding probability for achieving global coverage under specific conditions [9].

3.2.3 m-Probabilistic flooding

Unlike probabilistic flooding, where an information message is forwarded to every neighbor under a specific probability, in m-probabilistic flooding, upon initial reception of the message, a node transmits it exclusively to m randomly selected neighboring nodes with probability one. If not, the message is disregarded. Three distinct values of m are under consideration: $0.25\bar{d}, 0.50\bar{d}, 0.75\bar{d}$, with $\bar{d}$ denoting the average node degree. The second value of m, i.e. $0.5\bar{d}$, is selected to be the same case as the probabilistic flooding ($q = 0.5$), while the other two are selected arbitrarily.

To assess m-probabilistic flooding in comparison to other probabilistic methods based on q, $\bar{p}$ is considered as the average likelihood of a node being selected among the neighbors of its neighbor (i.e., to receive the message). In m-probabilistic flooding, all m selected neighbors receive the message with certainty; thus, the average probability can be expressed as shown in Eq. *(1)*.

$$\bar{p} = \frac{m}{\bar{d}} \tag{1}$$

3.2.4 Scoped probabilistic flooding

Scoped probabilistic flooding has many similarities to probabilistic flooding, the main differentiation lying into the scoped probabilistic flooding limitation on how many time steps (or hops) it will operate. Scoped probabilistic flooding is evaluated under the same forwarding probabilities q as probabilistic flooding while it is also evaluated on different time step limits.

3.3 The metric

Before delving into the analysis of the metric Δ that is based on binomial approximation, certain definitions are necessary. The collection of nodes covered at a specific step t is denoted as C_t. The set of covered nodes up to step t is represented by $\mathbb{C}_t$ and can be computed as $\mathbb{C}_t = C_0 \cup C_1 \cup \cdots \cup C_t$, with the understanding that $\mathbb{C}_0 \subset \cdots \subset \mathbb{C}_{t-1} \subset \mathbb{C}_t$. Additionally, let $\mathbb{C}'_t$ denote the set of uncovered nodes until time step t, complementing the integral of $\mathbb{C}_t$.

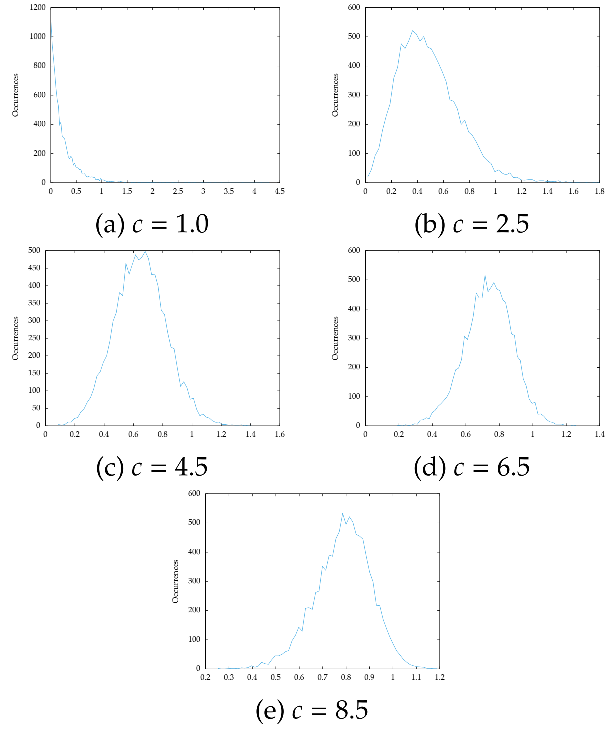

(a) $c = 1.0$ (b) $c = 2.5$ (c) $c = 4.5$ (d) $c = 6.5$ (e) $c = 8.5$

Figure 1 – The histogram of different values of c of the Burr Type XII distribution.

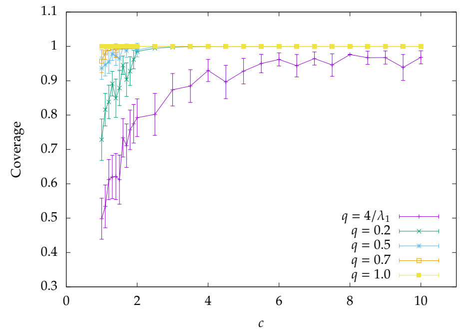

Figure 2 – The coverage of probabilistic flooding for different probabilities q as a function of the shape constant c in 95 % confidence interval.

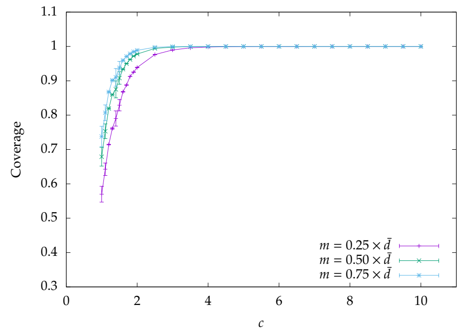

Figure 3 – The coverage of m-probabilistic flooding for different values of m as a function of the shape constant c in 95 % confidence interval.

The metric Δ that characterizes both probabilistic flooding and scoped probabilistic flooding is defined by Eq. *(2)*. It involves summing the count of the number of uncovered neighbors of the nodes that became covered at step t, then multiplying this count by the forwarding probability q.

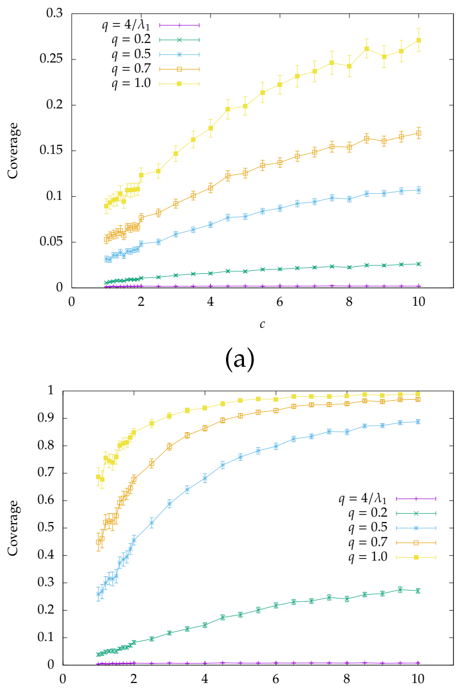

(a)

(b)

Figure 4 – The coverage of scoped probabilistic flooding for different values of q as a function of the shape constant c in 95% confidence interval. Subfigure (a) depicts the results for 2 hops and subfigure (b) presents the results for 3 hops.

This process occurs for all steps $t \in T$, and the result is normalized by the summation of node degrees ($N\bar{d}$). Thus, the metric Δ can be defined as an approximation of the probability that the flooding will continue normalized by the summation of node degrees.

$$\Delta = \frac{1}{N\bar{d}} \sum_{t=0}^{T} \sum_{\forall i \in \mathbb{C}_t'} |C_t \cap n_i| \times q \tag{2}$$

This study presents two versions of the metric Δ for the m-probabilistic flooding. On the one hand, Eq. *(3)* expresses the simple version of the metric for m-probabilistic flooding. This expression is similar to the metric for probabilistic flooding with the difference that the latter, instead of multiplying by the forwarding probability q, it multiplies by the average probability $\bar{p}$ of a node to be selected. On the other hand, Eq. *(4)* expresses the more sophisticated version where it takes into consideration the different nodes' degrees. More specifically, the multiplication is replaced by the summation of the probabilities of each uncovered node to become covered p_j at the next time step, and this probability is given by:

$$p_j = \begin{cases} 1, & \text{if } d_j \leq m, \\ \frac{m}{d_j}, & \textit{otherwise} \end{cases}.$$

$$\Delta = \frac{1}{N\bar{d}} \sum_{t=0}^{T} \sum_{\forall i \in \mathbb{C}_t'} |C_t \cap n_i| \times \bar{p} \tag{3}$$

$$\Delta = \frac{1}{N\bar{d}} \sum_{t=0}^{T} \sum_{\forall i \in \mathbb{C}_t'} \sum_{\forall j \in n_i \cap C_t} p_j \tag{4}$$

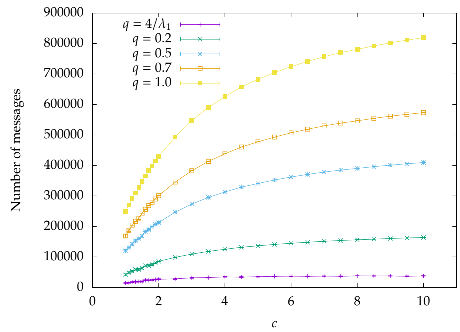

Figure 5 – The number of messages that are transmitted under the probabilistic flooding for different forwarding probabilities q as a function of the shape constant c in 95 % confidence interval.

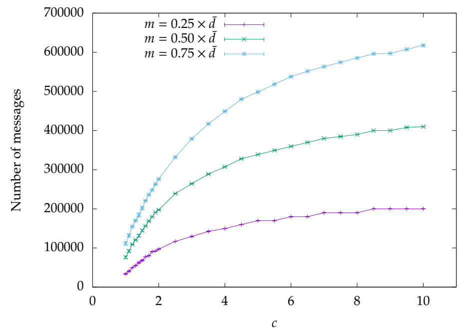

Figure 6 – The number of messages that are transmitted under the m-probabilistic flooding for different values of m as a function of the shape constant c in 95 % confidence interval.

4. RESULTS ON RANDOM NETWORKS

For the evaluation process, simulation scenarios have been implemented in the Python programming language. The networks are managed by the NetworkX Python package [24] while the generation of the randomly generated numbers, and the computation of the largest eigenvalue takes place under the SciPy python package [25].

For studying the different flooding methods, multiple networks have been generated from the Burr Type XII distribution [12], which is continuous, and its probability density function is given by $\frac{x^{c-1}}{(1+x^c)^{k+1}}$, where k and c are shape constants. Fig. 1 presents the histograms of the Burr Type XII distribution for some values of c. The following procedure has been followed to generate the random networks of the Burr Type XII distribution.

First, a sequence of randomly generated numbers is

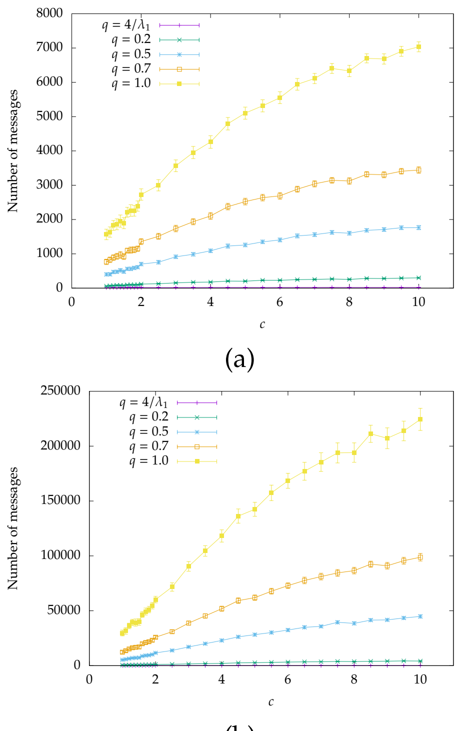

Figure 7 – The number of messages transmitted under scoped probabilistic flooding for different values of q as a function of the shape constant c in 95% confidence interval. Subfigure (a) depicts the results for 2 hops and subfigure (b) presents the results for 3 hops.

produced following the distribution under 27 different values of c, i.e., $\{1.0, 1.1, 1.2, \ldots, 2.0, 2.5, 3.0, \ldots, 10\}$ and $k = 5$ for all sequences. Then, each element of the sequence is multiplied by 100 and then converted to an integer. This is done in order to convert the continuous sequence to a discrete one. Next, all elements of the discrete sequence are incremented by one as a way to ensure that all nodes have at least one neighbor (i.e., $d_i > 0 \forall i \in V(G)$). This sequence is then examined to find if it is graphical or not. If the sequence is graphical then the network is generated by a variation of the Havel-Hakimi algorithm [26]. This variation guarantees the production of a connected graph by picking the nodes with the lowest degrees first, instead of the highest ones. This procedure is repeated for each c under ten different iterations. Table 1 depicts the characteristics of some networks that are generated by the aforementioned procedure for different values of c.

Table 1 – The characteristics of some random networks.

c	$\lvert E(G)\rvert$	$\bar{d}$	λ_1	$4/\lambda_1$	$\mathcal{D}$	D
1.0	130249	26.05	47.625	0.084	5	0.002605
2.0	218852	43.77	54.121	0.074	5	0.004377
2.5	253047	50.61	58.808	0.068	5	0.005061
4.5	333216	66.64	70.543	0.057	5	0.006665
6.5	376384	75.28	77.613	0.052	4	0.007528
8.5	401203	80.24	81.782	0.049	4	0.008025
10.0	413962	82.79	83.986	0.048	4	0.008280

For the evaluation, different values of the forwarding probabilities, the number of nodes, and the number of time steps are considered. For the probabilistic flooding, five different values of the forwarding probabilities q are selected, namely, $q \in \{0.2, 0.5, 0.7, 1.0, 4/\lambda_1\}$, where λ_1 is the largest eigenvalue of the graph's adjacency matrix. The first three values of q are straightforward. When $q = 1.0$ the probabilistic flooding reduces to blind flooding, where the nodes forward the information message to all of its neighbors.

For m-probabilistic flooding, three different numbers of nodes are selected for each network with respect to the average node degree $\bar{d}$, namely $m \in \{0.25\bar{d}, 0.50\bar{d}, 0.75\bar{d}\}$. Here, it is noteworthy that the case of $m = 0.50\bar{d}$ is similar to the case of $q = 0.5$ because the average probability of a node to be informed in the m-probabilistic flooding is 0.5 from *(1)*.

Scoped probabilistic flooding is evaluated under the same values as probabilistic flooding due to their similarities. As already mentioned, the only difference between the two approaches is that scoped probabilistic flooding has a limit on the time steps a message can be forwarded (i.e., hops); thus it is evaluated along with two different values of time steps.

The following sections present the evaluation results, focusing on: (i) the coverage achieved by the flooding methods, (ii) the number of messages sent until the conclusion of each method, and (iii) the metric Δ. The figures provided in the subsequent sections represent the 95% confidence interval of the outcomes derived from ten repetitions for each iteration of the shape constant c. In order to ensure a "fair" experimentation process, all the algorithms start from the same initiator for each network, and the initiator is arbitrarily selected.

4.1 Coverage

The mean coverage of the probabilistic flooding, for different forwarding probabilities q, is depicted in Fig. 2. As expected, blind flooding (i.e., $q = 1$) covers all nodes of all graphs and iterations. In addition, the coverage of all forwarding probabilities improves as the shape constant c increases. For low values of the shape constant, all forwarding probabilities struggle to achieve full coverage. It is noteworthy, that the full coverage of the network is achieved in reverse order of the forwarding probabilities. More specifically, when the forwarding probability is set to 0.7 the full coverage is achieved at $c = 1.5$. Similar to the forwarding probability $q = 0.7$, the forwarding probability $q = 0.5$ achieves full coverage at $c = 2.5$, while the case of $q = 0.2$ covered all nodes at $c = 4.0$. In the case where $q = 4/\lambda_1$, it is shown that probabilistic flooding is not able to cover all nodes but its performance improves as the shape constant increases, as expected [9].

Fig. 3 illustrates the average coverage attained by m-probabilistic flooding across various m values. The coverage in m-probabilistic flooding is directly related to the number of selected neighbors. For instance, when $c = 1.0$, with $m = 0.25\bar{d}$, the average coverage achieved is 0.58, while for $m = 0.75\bar{d}$, the attained coverage rises to 0.75. In general, the performance of m-probabilistic flooding is lower with smaller values of the shape constant, gradually improving as the shape constant increases, until it eventually reaches 1.0 in the reverse order of m.

In Fig. 4, the coverage of scoped probabilistic flooding is depicted for 2 and 3 time steps. As expected, for both the considered time steps, the coverage is proportional to the forwarding probability q for all different values of the shape constant c. For the case of 2 time steps (Fig. 4(a)), the scoped probabilistic flooding is not able to cover all nodes of the network while the achieved coverage of the blind scoped flooding (i.e., $q = 1.0$) does not get over 0.3. It is also noted that the case of $q = 4/\lambda_1$ does not seem to improve as the shape constant increases, with this behavior being similar to the case of 3 hops. This is attributed to the fact that this case needs substantially more time (or hops) to cover the same number of nodes compared to the case of $q = 0.2$, because $4/\lambda_1 \ll 0.2$. The performance of scoped probabilistic flooding is improved substantially for the case of 3 hops as shown by Fig. 4(b). Scoped blind flooding seems to approach the total coverage but does not achieve it. Note that the performance of the case of $q = 0.7$ is very close to the performance of scoped blind flooding for high values of the shape constant c.

Both probabilistic methods (probabilistic flooding and m-probabilistic flooding) demonstrate similar coverage behaviors concerning the shape constant c. They both reach full coverage (1.0) when c is significantly large, with the exception of the case of $q = 4/\lambda_1$. It is interesting to focus on one case of each flooding method in order to compare them. Specifically, the $q = 0.5$ case is selected for probabilistic flooding and the $m = 0.5\bar{d}$ case is selected for m-probabilistic flooding. The latter case is selected because it maintains an equivalent average forwarding probability to the probabilistic flooding. For large values of c (i.e., $c > 2.5$), both methods achieve full coverage. However, when c is small, probabilistic flooding significantly outperforms m-probabilistic flooding. As an illustration, when $c = 1.0$, probabilistic flooding achieves a coverage of 0.93, while m-probabilistic flooding only reaches 0.7.

4.2 Number of messages

Fig. 5 illustrates the average number of messages transmitted in probabilistic flooding for different values of q. As expected, a higher forwarding probability corresponds to an increased message count. Notably, for all cases of q, the number of sent messages rises alongside the shape constant, with the exception of the case $q = 4/\lambda_1$. This behavior can be attributed to the direct relationship between the number of edges and the shape constant c. It is important to highlight that in the case of $q = 4/\lambda_1$, the message count remains relatively constant, in contrast to the results in Fig. 2. This observation is linked to the fact that λ_1 is proportional to the shape constant and to the average degree $\bar{d}$, as well.

Fig. 6 illustrates the average number of messages sent by m-probabilistic flooding for various values of m. The figure demonstrates a consistent pattern that, as m increases, the number of messages sent by m-probabilistic flooding also increases. This behavior is observed across all values of m. Initially, for $c = 1.0$, the message count is relatively low, ranging from $34,438$ to $113,718$ messages. However, as c increases, the number of messages sent by m-probabilistic flooding rises significantly, ranging from $200,000$ to $617,707$ messages.

Fig. 7 presents the average number of transmitted messages that scoped probabilistic flooding was able to transmit. As expected, the number of messages is proportional to the shape constant c and to the forwarding probability. The former is attributed to the fact that the average number of messages increases with the shape constant c. The latter can be explained by the fact that, as q increases, the more messages are transmitted by the method, thus higher coverage is achieved, as shown by Fig. 4. The great difference in the number of messages between the case of 2 hops (Fig. 7(a)) and the case of 3 hops (Fig. 7(b)) is attributed to the fact that the generated networks are of relatively small diameter (i.e., the networks of $c = 1.0$ have a diameter of 5 hops) and the later case is easier to reach all nodes of the network. This is also shown by the fact that, for the case of 3 hops, scoped blind flooding achieved total coverage for large values of c, while the case of 2 hops achieved only 0.28 coverage.

Revisiting the same scenarios for the comparative analysis of probabilistic flooding and m-probabilistic flooding proves intriguing. As anticipated, both flooding methods exhibit a nearly identical average message count across all values of c. It is noteworthy here that scoped probabilistic flooding achieves an outstanding performance with a significant lower number of transmitted messages, compared to the other considered flooding methods.

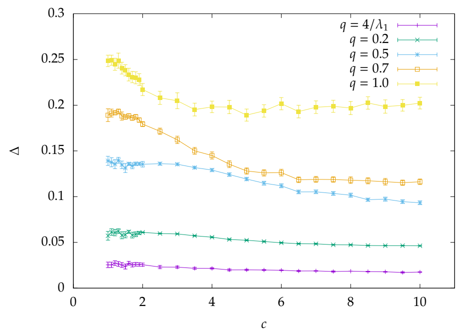

Figure 8 – The metric Δ for probabilistic flooding for different values of q as a function of the shape constant c in 95 % confidence interval.

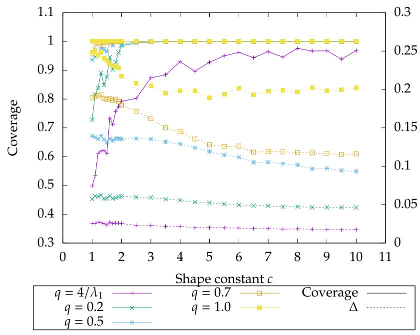

(a)

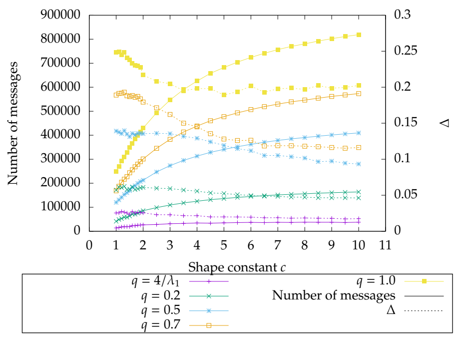

(b)

Figure 9 – The metric Δ in along with the achieved coverage (subfigure (a)) and with the number of transmitted messages (subfigure (b)).

4.3 The binomial approximation

Fig. 8 presents the average value of Δ for various forwarding probabilities q in relation to the shape constant c. Fig. 9 illustrates the behavior of the metric Δ along with the achieved coverage and the number of messages. Across all cases of forwarding probabilities, there is a consistent decrease in Δ as c increases. Furthermore, the results illustrate that Δ is directly proportional to the forwarding probability q and inversely proportional to the coverage and the number of messages.

The simple version of the metric Δ of m-probabilistic flooding is presented in Fig. 10(a) for different values of m as a function of the shape constant c. Similar to the behavior observed in probabilistic flooding, this figure reveals that, increasing m, results in a higher metric

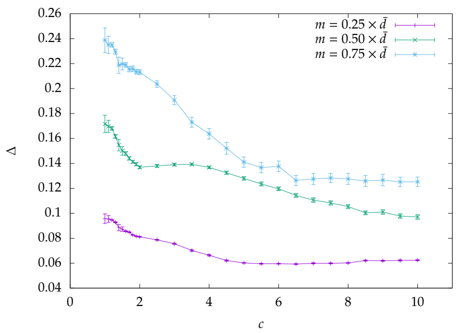

(a) simple version

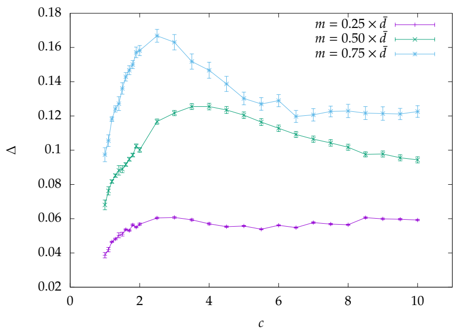

(b) sophisticated version

Figure 10 – The metric Δ for m-probabilistic flooding for different values of m as a function of the shape constant c in 95 % confidence interval.

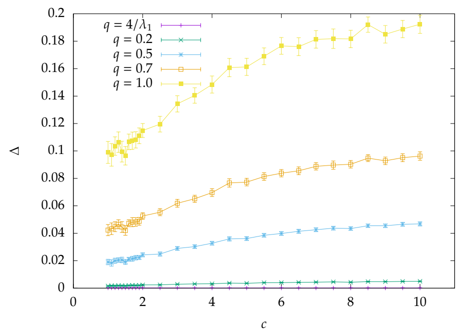

(a)

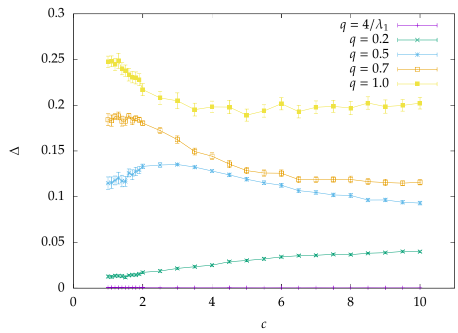

(b)

Figure 11 – The metric Δ for scoped probabilistic flooding for different values of q as a function of the shape constant c in 95% confidence interval. Subfigure (a) depicts the results for 2 hops and subfigure (b) presents the results for 3 hops.

Δ. Additionally, for all m scenarios, Δ decreases as the shape constant c increases. It is interesting that, when considering the same cases of forwarding probabilities (i.e., $q = 0.5$ and $m = 0.5\bar{d}$), both metrics yield similar Δ values on average relative to the shape constant c.

Fig. 10(b) presents the sophisticated version of the metric Δ for different values of m as a function of the shape constant c. Similar to the simple version of the metric Δ, the sophisticated version is proportional to m. In the case of $0.25\bar{d}$ there is an increase for small values of c, i.e., $c < 2$. For the other values of the shape constant c of the case of $0.25\bar{d}$ the sophisticated version of the metric Δ stays relatively stable.

The other two cases have different behaviors. More specifically, the metric Δ for both the other cases ($0.50\bar{d}$ and $0.75\bar{d}$) initially exhibits an increase as the values of c remain small. This trend is then followed by a decrease. It is reasonable here to take into consideration the degree histograms of the key networks, i.e., $c \in \{1.0, 2.5, 6.5\}$ as depicted in Fig. 12(a), 12(b) and 12(c), respectively. The vertical orange line of each figure depicts the mean degree of each network. It is observed that for small values of $c < 2.5$ and as increases, there are more and more nodes that their degree is smaller than the selected m, resulting into $p_j = 1.0$. At $c = 2.5$, the metric Δ reaches the maximum value for the case of $m = 0.75\bar{d}$, and next, the metric Δ decreases until it stabilizes after $c = 6.5$. This behavior is also observed for the case of $m = 0.50\bar{d}$, where the maximum value of the metric Δ is detected when $c = 3.5$ instead of 2.5 of the case of $m = 0.75\bar{d}$ and there is not any visible point that the metric Δ stabilizes for $m = 0.50\bar{d}$.

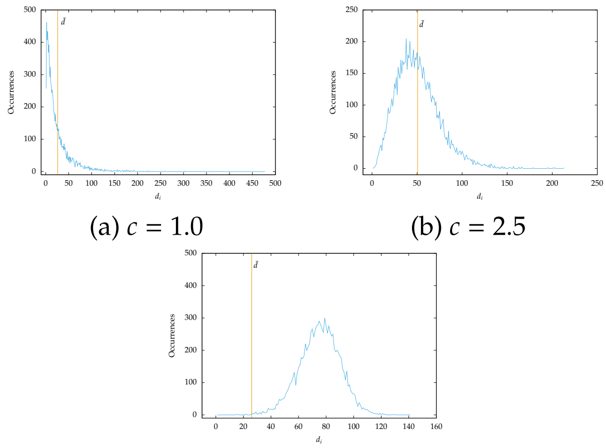

(a) $c = 1.0$ (b) $c = 2.5$ (c) $c = 6.5$

Figure 12 – The degree histograms of randomly generated networks that follow the Burr type XII distribution for different values of c. The vertical orange lines depict the mean degree $\bar{d}$.

Fig. 11 presents the metric Δ for scoped probabilistic flooding as a function of shape constant c and for different values of time steps. For the case of 2 hops (Fig. 11(a)), the metric Δ is proportional to the forwarding probability q and to the shape constant c for the most cases, except the cases of $q = 4/\lambda_1$ and $q = 0.2$. This behavior is attributed to the fact the method was not able to cover enough nodes for each value of the shape constant to be a notable increase in the metric Δ. For the case of 3 hops (Fig. 11(b)), the behavior of the metric Δ for scoped probabilistic flooding, for the three largest values of the forwarding probability, is similar to the behavior of probabilistic flooding. For the case of $q = 0.2$, there is a notable increase in the rate of increment but it does not follow the behavior of the metric Δ of the probabilistic flooding; a similar behavior follows the case of $q = 4/\lambda_1$.

5. RESULTS ON REAL NETWORKS

In order to further evaluate the probabilistic methods, data regarding seven real networks are considered. More specifically, four of them have been obtained by [27], namely *routers-rf*, *WHOIS*, *pgp* and *cegb2802*. Additionally, *facebook* originates from [28] and the other two networks, namely *ca-grqc* and *ca-HepPh* originate from [29]. For all considered real networks, the largest connected component is selected as the network. Table 2 depicts the final seven connected networks with their characteristics/attributes.

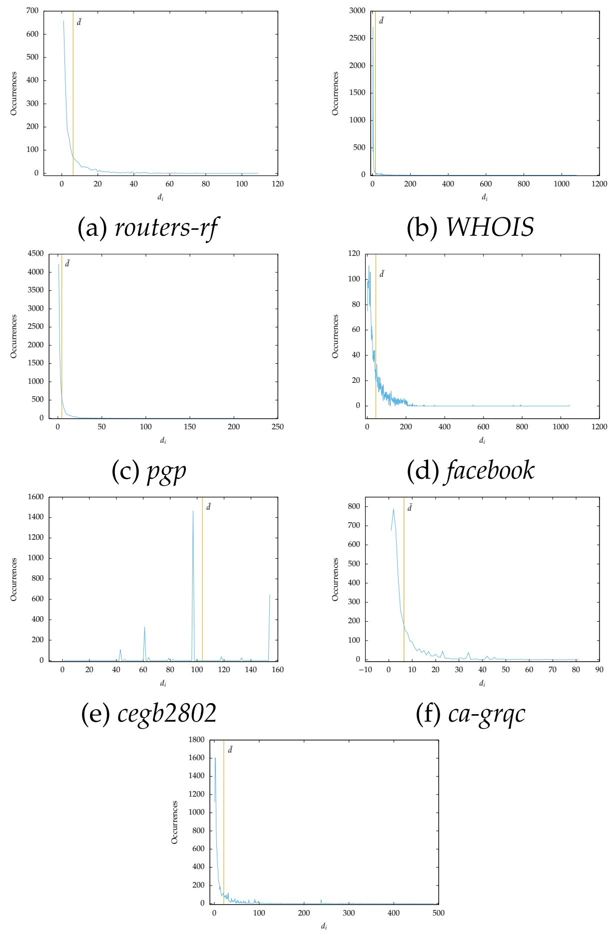

(a) *routers-rf* (b) *WHOIS* (c) *pgp* (d) *facebook* (e) *cegb2802* (f) *ca-grqc* (g) *ca-HepPh*

Figure 13 – The degree histograms of the real networks. The vertical orange lines depict the mean degree $\bar{d}$.

Table 2 – The characteristics of the real networks

	N	$\|E(G)\|$	$\bar{d}$	λ_1	$4/\lambda_1$	$\mathcal{D}$	D
routers-rf	2,113	6,632	6.280	27.671	0.145	12	0.0029
WHOIS	7,476	56,943	15.230	150.859	0.027	8	0.0020
pgp	10,680	24,316	4.550	42.435	0.094	24	0.0004
facebook	4,039	88,234	43.690	162.374	0.025	8	0.0108
cegb2802	2,694	140,028	103.960	116.724	0.034	26	0.0386
ca-grqc	4,158	13,428	6.460	45.617	0.088	17	0.0018
ca-HepPh	11,204	117,649	21.000	244.939	0.016	13	0.0015

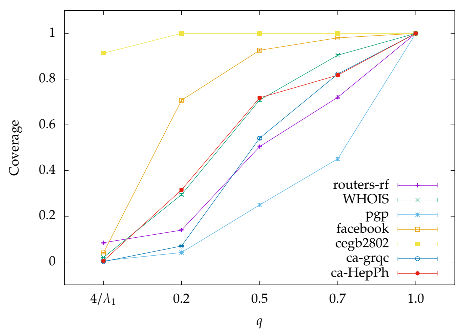

Figure 14 – The mean coverage of probabilistic flooding for five different networks with respect to the forwarding probability q.

The experimentation that took place for these topologies is similar to the previous one. The forwarding probabilities q and m are assigned the same values as the evaluation on random networks. The time step limits of scoped probabilistic flooding are selected with respect to each network's diameter. The time step limits are selected in this manner for the algorithms to be able to reach the same portion of each network. Each experiment is repeated 100 times, with a different initiator for each iteration but the same for all methods, and the results are averaged.

5.1 Coverage

Fig. 14 illustrates how the coverage of probabilistic flooding varies concerning the parameter q. Across all networks, higher forwarding probabilities consistently lead to improved coverage. Notably, for $q = 4/\lambda_1$, the coverage achieved by probabilistic flooding on the *cegb2802* network far surpasses that of the other four networks, each registering coverage levels below 0.1. Additionally, the probabilistic flooding achieves total coverage for the other forwarding probabilities. This discrepancy can be attributed to the *cegb2802* network's D and the proximity of λ_1 to $\bar{d}$. At $q = 0.2$, while probabilistic flooding exhibits better coverage across all topologies, the most striking enhancement occurs within the *facebook* network, where coverage improves from 0.04 (at $q = 4/\lambda_1$) to 0.7. Moreover, probabilistic flooding achieves significant coverage in the *facebook* network for all q but does not reach total coverage until the forwarding probability goes to 1.0. In general, the coverage of probabilistic flooding seems to be proportional to the network's D.

Fig. 15 showcases the mean coverage achieved by m-probabilistic flooding as a function of m. Notably, the coverage remains consistently stable for both the *facebook* and *cegb2802* networks across all values of m. However, a noticeable improvement in coverage is observable across other networks as the value of m increases. It's worth highlighting a connection between the outcomes derived from the Burr distribution and the results observed in the *pgp* network. Specifically, the probabilistic flooding demonstrates superior coverage compared to m-probabilistic flooding, notably in the case where $q = 0.5$ for probabilistic flooding and $m = 0.5\bar{d}$ for m-probabilistic flooding.

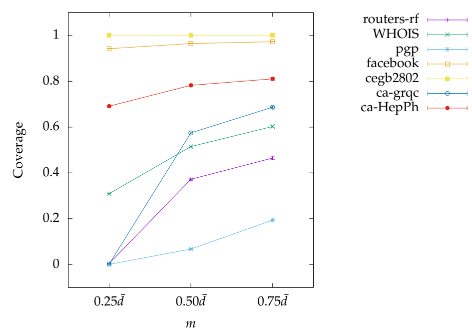

Figure 15 – The mean coverage of m-probabilistic flooding for five real networks with respect to different values of m.

Fig. 16 presents the mean coverage of scoped probabilistic flooding for the real networks with respect to the forwarding probability and the time step limits. It is observed that the coverage for all cases of the time step limits is proportional to the forwarding probability. Additionally, the coverage is proportional to the time step limits. This is attributed to the fact that the methods are restricted to a portion of the network that can be reached. It is noteworthy here that the achieved coverage of scoped probabilistic flooding on the *cegb2802* network, for most time step limits and most forwarding probabilities outperforms the rest. As expected, scoped probabilistic flooding reduces to the probabilistic flooding for large number of time step limits. This is shown in Fig. 16(d) which is mostly identical to the corresponding figure of the probabilistic flooding (Fig. 14).

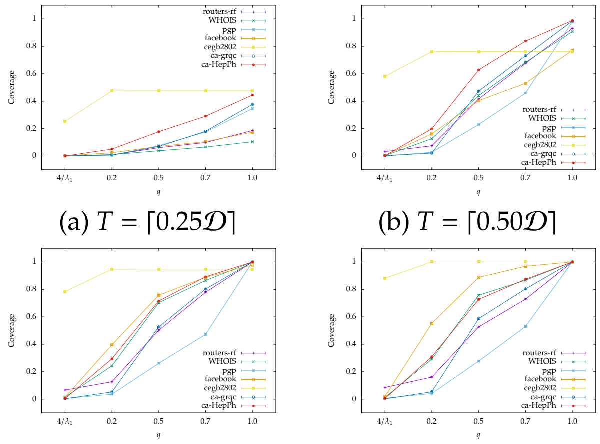

(a) $T = \lceil 0.25\mathcal{D} \rceil$ (b) $T = \lceil 0.50\mathcal{D} \rceil$

(c) $T = \lceil 0.75\mathcal{D} \rceil$ (d) $T = \lceil 1.0\mathcal{D} \rceil$

Figure 16 – The mean coverage of scoped probabilistic flooding for the real networks with respect to the forwarding probability q. The labels of subfigures present the time step limits of scoped probabilistic flooding.

5.2 Number of messages

The mean number of messages sent from probabilistic flooding is presented in Fig. 17. As expected, the number of sent messages rises in tandem with the increment of the forwarding probability. It is noteworthy here to consider the achieved coverage from Fig. 14 and the case of $q = 4/\lambda_1$ where the probabilistic flooding achieves 0.9 coverage in the *cegb2802* network with a small number of messages.

Fig. 18 presents the average count of messages transmitted using m-probabilistic flooding. In both the *facebook* and *cegb2802* networks, in order for probabilistic flooding to achieve an equivalent coverage as m-probabilistic flooding with $m = 0.25\bar{d}$, the forwarding probability needs to be set at 0.5. Moreover, the number of messages transmitted by probabilistic flooding is twice the number of messages that is sent by m-probabilistic flooding.

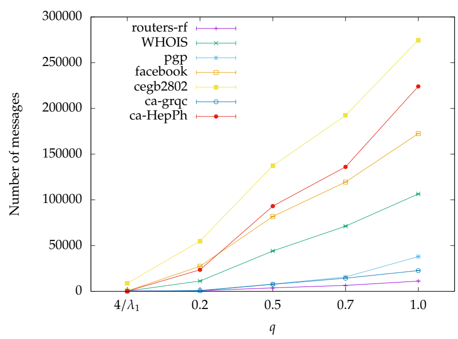

Figure 17 – The mean number of messages transmitted under probabilistic flooding for the real networks on different forwarding probabilities q.

Fig. 19 illustrates the mean number of messages transmitted by scoped probabilistic flooding with respect to the forwarding probability and to the time step limits. The number of messages is proportional to the forwarding probability and to the time step limit. Additionally, it is also proportional to $|E(G)|$. It is noteworthy here that the coverage of scoped probabilistic flooding when the time step limit $T = \lceil 0.75\mathcal{D} \rceil$ is similar to the coverage of probabilistic flooding (Fig. 14) but this is achieved with slightly fewer messages, as shown by the Fig. 17.

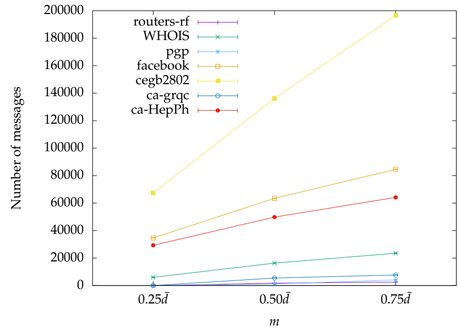

Figure 18 – The mean number of messages transmitted under m-probabilistic flooding for real networks on different values of m.

5.3 The binomial approximation

Fig. 20 illustrates the mean Δ of probabilistic flooding for real networks across varying forwarding probabilities q. The results indicate that there is a proportional relationship between the metric Δ and the forwarding probability q in almost every considered network, except for the *facebook* network. Specifically, in the case of $q = 1.0$, the metric Δ experiences a decrease instead of the expected increase, in contrast to the outcomes observed in the Burr distribution results. This is attributed to the fact that the *facebook* network is more dense than the evaluated random networks, thus, the probabilistic flooding needs fewer steps to cover the network. In addition, although the coverage of probabilistic flooding increases by $0.15 - 0.18$, the Δ doubles for the *facebook* network when $q \in \{0.2, 0.5\}$. Moreover, the metric Δ for the *pgp* network experiences a substantial increase for the case of $q = 1.0$. This is attributed to the fact that this network is loosely connected, its density D is very small as illustrated in Table 2, and that the case of $q = 1.0$ is the only case that total coverage is achieved for this network.

The simple version of the metric Δ for m-probabilistic flooding is presented in Fig. 21(a) and Fig. 21(b) presents the sophisticated version of the metric. As expected, both versions of the metric Δ of m-probabilistic flooding are proportional to m. Despite the fact that, the coverage of m-probabilistic flooding (Fig. 15) for the *facebook* network and for the *cegb2802* network stays the same, the metric Δ increases in both versions with the increasing values of m.

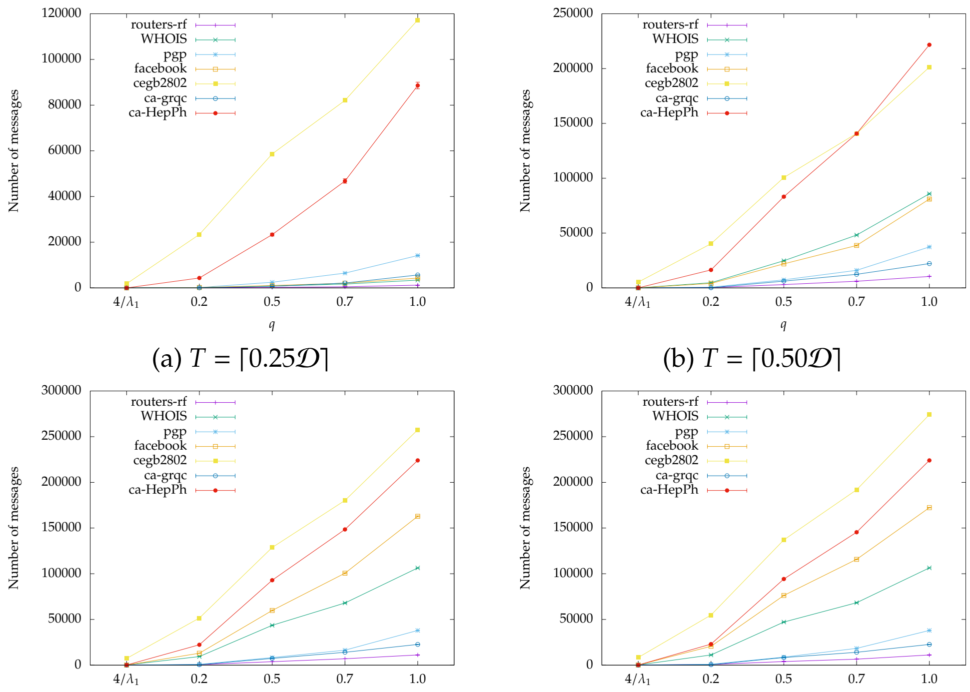

(a) $T = \lceil 0.25\mathcal{D} \rceil$ (b) $T = \lceil 0.50\mathcal{D} \rceil$

(c) $T = \lceil 0.75\mathcal{D} \rceil$ (d) $T = \lceil 1.0\mathcal{D} \rceil$

Figure 19 – The mean number of messages transmitted under scoped probabilistic flooding for real networks with respect to the forwarding probability q. The labels of subfigures present the time step limits of scoped probabilistic flooding.

Fig. 22 presents the metric Δ for scoped probabilistic flooding with respect to the forwarding probability and the time step limits for real network topologies. As expected, from the results obtained for the random topologies, the metric Δ for scoped probabilistic flooding is proportional to the time step limits. Additionally, it is observed that the metric Δ for scoped probabilistic flooding follows similar behavior to the metric Δ of probabilistic flooding as the time step limits increase.

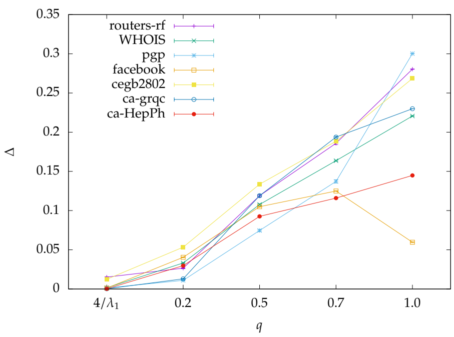

Figure 20 – The mean Δ for the probabilistic flooding for real networks as a function of the different forwarding probabilities q.

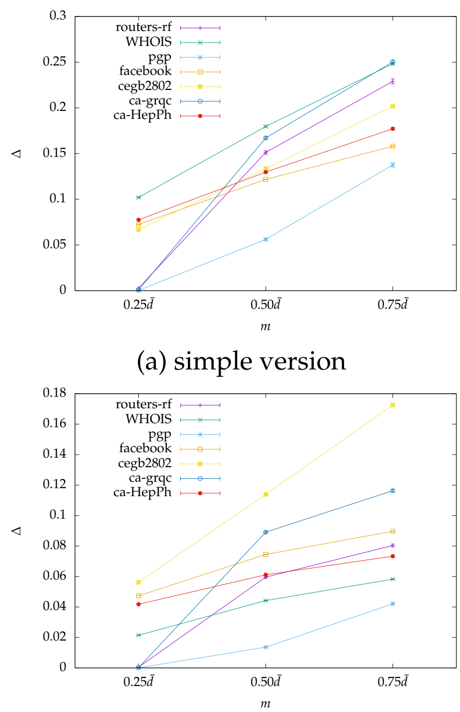

(a) simple version

(b) sophisticated version

Figure 21 – The metric Δ for the m-probabilistic flooding for the real networks as a function of the different values of m.

6. CONCLUSIONS

This paper evaluates the performance of four different flooding methods, namely: (i) blind flooding; (ii) probabilistic flooding; (iii) m-probabilistic flooding; and (iv) scoped probabilistic flooding, and the latter is introduced in this paper. Scoped probabilistic flooding is basically

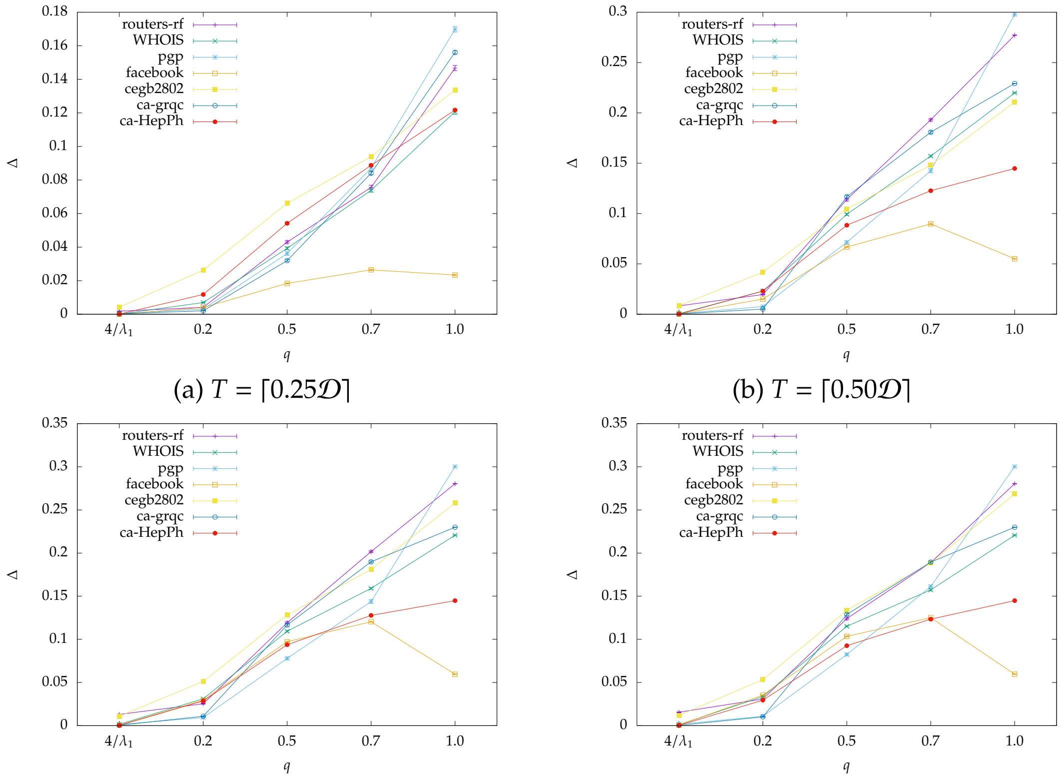

(a) $T = \lceil 0.25\mathcal{D} \rceil$
(b) $T = \lceil 0.50\mathcal{D} \rceil$
(c) $T = \lceil 0.75\mathcal{D} \rceil$
(d) $T = \lceil 1.0\mathcal{D} \rceil$

Figure 22 – The metric Δ of the scoped probabilistic flooding for seven real networks with respect to the forwarding probability q. The labels of subfigures present the time step limits of the scoped probabilistic flooding.

probabilistic flooding with restriction on the number of hops that the message will be transmitted. The flooding methods are evaluated on three metrics: (i) coverage; (ii) number of transmitted messages; and (iii) the metric Δ which is based on binomial approximation. Metric Δ approximates the probability that the flooding will continue normalized by the summation of node degrees. The assessment is carried out on networks structured according to the Burr Type XII distribution for node degree, as well as on real network topologies. The paper also introduces a modified version of the Δ metric for m-probabilistic flooding, displaying similar characteristics to its probabilistic flooding counterpart. This extended Δ metric, specific to m-probabilistic flooding, incorporates individual node degrees to calculate the likelihood of a node becoming covered.

The findings indicate that m-probabilistic flooding, in certain scenarios, outperformed probabilistic flooding in terms of coverage, while requiring significantly fewer message exchanges. Probabilistic flooding, however, generally exhibited superior performance. In addition, the study revealed that scoped probabilistic flooding attains coverage comparable to that of probabilistic flooding but with a reduced number of messages sent. Finally, the metric Δ has similar behavior for random, as well as for real networks, with some differences that are attributed to the networks' density.

REFERENCES

[1] Li Da Xu, Wu He, and Shancang Li. "Internet of things in industries: A survey". In: *IEEE Transactions on industrial informatics* 10.4 (2014), pp. 2233–2243.

[2] Ethirajan Manavalan and Kandasamy Jayakrishna. "A review of Internet of Things (IoT) embedded sustainable supply chain for industry 4.0 requirements". In: *Computers & Industrial Engineering* 127 (2019), pp. 925–953.

[3] Zach Shelby, Klaus Hartke, and Carsten Bormann. *The constrained application protocol (CoAP)*. Tech. rep. 2014.

[4] Sinem Coleri Ergen. "ZigBee/IEEE 802.15. 4 Summary". In: *UC Berkeley, September* 10.17 (2004), p. 11.

[5] Harith Kharrufa, Hayder AA Al-Kashoash, and Andrew H Kemp. "RPL-based routing protocols in IoT applications: a review". In: *IEEE Sensors Journal* 19.15 (2019), pp. 5952–5967.

[6] Carles Gomez and Josep Paradells. "Wireless home automation networks: A survey of architectures and technologies". In: *IEEE communications magazine* 48.6 (2010), pp. 92–101.

[7] Adrian Segall. "Distributed network protocols". In: *IEEE transactions on Information Theory* 29.1 (1983), pp. 23–35.

[8] Sze-Yao Ni, Yu-Chee Tseng, Yuh-Shyan Chen, and Jang-Ping Sheu. "The broadcast storm problem in a mobile ad hoc network". In: *Proceedings of the 5th annual ACM/IEEE international conference on Mobile computing and networking*. 1999, pp. 151–162.

[9] Konstantinos Oikonomou, George Koufoudakis, Sonia Aïssa, and Ioannis Stavrakakis. "Probabilistic Flooding Performance Analysis Exploiting Graph Spectra Properties". In: *IEEE/ACM Transactions on Networking* (2022).

[10] Hung-Chang Hsiao and Chung-Ta King. "Scoped broadcast in dynamic peer-to-peer networks". In: *29th Annual International Computer Software and Applications Conference (COMPSAC'05)*. Vol. 1. IEEE. 2005, pp. 533–538.

[11] Asterios Papamichail, Georgios Tsoumanis, Spyros Sioutas, and Konstantinos Oikonomou. "A Performance Evaluation of Flooding Methods on Real Network Topologies". In: *2023 International Balkan Conference on Communications and Networking (BalkanCom)*. IEEE. 2023, pp. 1–6.

[12] Irving W Burr. "Cumulative frequency functions". In: *The Annals of mathematical statistics* 13.2 (1942), pp. 215–232.

[13] Yoav Sasson, David Cavin, and André Schiper. "Probabilistic broadcast for flooding in wireless mobile ad hoc networks". In: *2003 IEEE Wireless Communications and Networking, 2003. WCNC 2003.* Vol. 2. IEEE. 2003, pp. 1124–1130.

[14] Huaqiang Xu, Haipeng Si, Hang Zhang, Liren Zhang, Yan Leng, Jingjing Wang, and Dengwang Li. "Trust-based probabilistic broadcast scheme for mobile ad hoc networks". In: *IEEE Access* 8 (2020), pp. 21380–21392.

[15] Djahida Ali-Fedila and Mohamed Ould-Khaoua. "Performance evaluation of probabilistic broadcast in low-power and lossy networks". In: *2021 20th International Conference on Ubiquitous Computing and Communications (IUCC/CIT/DSCI/SmartCNS)*. IEEE. 2021, pp. 247–254.

[16] Kannan Shanmugam, Karthik Subburathinam, Arunachalam Velayuthampalayam Palanisamy, et al. "A dynamic probabilistic based broadcasting scheme for manets". In: *The Scientific World Journal* 2016 (2016).

[17] Qi Zhang and Dharma P Agrawal. "Dynamic probabilistic broadcasting in MANETs". In: *Journal of parallel and Distributed Computing* 65.2 (2005), pp. 220–233.

[18] Ankita Srivastava, Arun Prakash, and Rajeev Tripathi. "Fuzzy-based beaconless probabilistic broadcasting for information dissemination in urban VANET". In: *Ad Hoc Networks* 108 (2020), p. 102285.

[19] Anuradha Banerjee and Subhankar Ghosh. "FPR: fuzzy controlled probabilistic rebroadcast in mobile ad hoc network". In: *International Journal of Information Technology* 12.2 (2020), pp. 523–529.

[20] Fatemeh Safari, Herb Kunze, Jason Ernst, and Daniel Gillis. "A novel cross-layer adaptive fuzzy-based ad hoc on-demand distance vector routing protocol for MANETs". In: *IEEE Access* (2023).

[21] Liang Wang, Suzan Bayhan, Jörg Ott, Jussi Kangasharju, Arjuna Sathiaseelan, and Jon Crowcroft. "Pro-diluvian: Understanding scoped-flooding for content discovery in information-centric networking". In: *Proceedings of the 2nd ACM Conference on Information-Centric Networking*. 2015, pp. 9–18.

[22] Daniel Gutiérrez Reina, Sergio L Toral, Princy Johnson, and Federico Barrero. "A survey on probabilistic broadcast schemes for wireless ad hoc networks". In: *Ad Hoc Networks* 25 (2015), pp. 263–292.

[23] Patricia Ruiz and Pascal Bouvry. "Survey on broadcast algorithms for mobile ad hoc networks". In: *ACM computing surveys (CSUR)* 48.1 (2015), pp. 1–35.

[24] Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. "Exploring Network Structure, Dynamics, and Function using NetworkX". In: *Proceedings of the 7th Python in Science Conference*. Ed. by Gaël Varoquaux, Travis Vaught, and Jarrod Millman. Pasadena, CA USA, 2008, pp. 11–15.

[25] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python". In: *Nature Methods* 17 (2020), pp. 261–272. DOI: 10.1038/s41592-019-0686-2.

[26] S Louis Hakimi. "On realizability of a set of integers as degrees of the vertices of a linear graph. I". In: *Journal of the Society for Industrial and Applied Mathematics* 10.3 (1962), pp. 496–506.

[27] Ryan A. Rossi and Nesreen K. Ahmed. "The Network Data Repository with Interactive Graph Analytics and Visualization". In: *AAAI*. 2015. URL: https://networkrepository.com.

[28] Jure Leskovec and Julian Mcauley. "Learning to discover social circles in ego networks". In: *Advances in neural information processing systems* 25 (2012).

[29] Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. "Graph evolution: Densification and shrinking diameters". In: *ACM transactions on Knowledge Discovery from Data (TKDD)* 1.1 (2007), 2–es.

AUTHORS

Asterios Papamichail graduated from the Department of Informatics of the Ionian University, Corfu, Greece, in 2019. In 2021, he received his M.Sc. degree, and he is currently a Ph.D. candidate at the same department. He has a great deal of experience with the Arduino prototyping platform (developed a MAC protocol). He has also investigated the spectral properties of distance matrices with respect to placement problems in SDNs. His research interests include computer networks, facility location and the multi-component application placement problem in the environment of multi-access edge computing.

Georgios Tsoumanis received his undergraduate degree in 2008 from the Department of Teleinformatics and Management, Technological Educational Institute of Epirus, Greece. In 2014 he obtained his MSc in informatics, in the subject area of information systems, from the Department of Informatics, Ionian University, Greece. In 2018 he completed his PhD studies at the same university and in 2021 he completed his postdoctoral research scholarship at University of Ioannina, Greece, funded by the State Scholarship Foundation (I.K.Y.) He currently works as a postdoctoral research associate at the Information

Technologies Institute (ITI) of the Centre for Research and Technology Hellas (CERTH) and as a postdoctoral researcher in Ionian University.

He has been involved in various EC research projects in the smart cities and Internet of Things (IoT) research areas and in various local development projects in several research areas, such as Wireless Sensor Networks (WSNs) and virtual worlds. His research interests include issues related to smart cities, energy consumption in WSN-based IoT, recharging of devices in WSNs, and the placement of services in network environments.

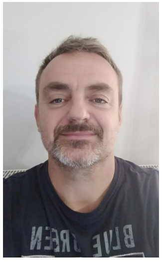

Spyros Sioutas is a full professor of "Data Structures and Software Systems for Big Data Management" in Computer Engineering and Informatics Department (CEID) (School of Engineering, University of Patras) and Head of "Information Systems and Artificial Intelligence" Lab in the Computer Software Division of the same department. His current research interests include: algorithmic data management, database systems, big data systems, large scale machine learning and cloud data engineering, indexing, query processing and query optimization. He has published over 200 papers in various high quality scientific journals and refereed conferences (amongst others SIGMOD, SODA, PODC, SIGKDD, CIKM, ESA, ICALP, CCGRID, ICDT/EDBT, SIGMOD Record, Algorithmica, Theoretical Computer Science, Computer Journal, Data and Knowledge Engineering, Journal of Discrete Algorithms, Distributed and Parallel Databases, Journal of Systems and Software, Knowledge and Information Systems, Information Science, ACM Computing Reviews, TLDKS) and he has more than 2000 citations. He served as editor, chair and invited speaker in more than 40 scientific and prestigious journals, conferences and international technological forums.

He has 24 years working experience as a developer, software tester, database administrator and project manager at Computer Technology Institute (Research Unit 5), MMLab (https://mmlab.ceid.upatras.gr/en/), ISD Lab (http://di.ionio.gr/isdlab/) and ML@Cloud Lab (https://www.ceid.upatras.gr/en/research/labs/laboratory-large-scale-machine-learning-and-cloud-data-engineeringy).

Konstantinos Oikonomou has received his MEng in Computer Engineering and Informatics from University of Patras in 1998. In September 1999 he received his postgraduate degree: M.Sc. in Communication and Signal Processing from the Electrical and Electronic Engineering Department, Imperial College (London). He received his Ph.D. degree in 2004 from the Informatics and Telecommunications Department, University of Athens, Greece. His Ph.D. thesis focuses on medium access control policies in ad hoc networks.

Since 2006 Prof. K. Oikonomou is a faculty member in Computer Networks (a full professor since 2021) at the Department of Informatics of the same university. He has also served as Dean of the Faculty of Information Science and Informatics and Head of the Department of Informatics. Between December 1999 and January 2005 he was employed at Intracom S.A, as a research and development engineer. His current interests involve medium access control in ad hoc networks, performance issues in wireless networks, information dissemination, service discovery, facility location, energy consumption in wireless sensor networks and network cost reduction in cloud computing environments.

Prof. Konstantinos Oikonomou has a great deal of experience with wireless systems and he has been coordinating a number of EC research projects in the area of computer networks, and in various local development projects (e.g., virtual worlds). He is currently member of the editorial board of Computer Networks Journal (Elsevier) and Journal on Future and Evolving Technologies (ITU). He has been a reviewer and TPC member of numerous conferences and journals in the area of computer networks and he holds an award for the best paper from the Hawaii International Conference on System Service (HICSS 2005) and another one from the IEEE Symposium on Computers and Communications (ISCC 2021).