Probabilistic Breadth-First Search – A Method for Evaluation of Network-Wide Broadcast Protocols

Björn Lichtblau
Department of Computer Science
Humboldt-Universität zu Berlin
Berlin, Germany
Email: lichtbla@informatik.hu-berlin.de

Andreas Dittrich
Advanced Learning and Research Institute
Università della Svizzera italiana
Lugano, Switzerland
Email: andreas.dittrich@usi.ch

***Abstract*—In Wireless Mesh Networks (WMNs), Network-Wide Broadcasts (NWBs) are a fundamental operation, required by routing and other mechanisms that distribute information to all nodes in the network. However, due to the characteristics of wireless communication, NWBs are generally problematic. Optimizing them thus is a prime target when improving the overall performance and dependability of WMNs. Most existing optimizations neglect the real nature of WMNs and are based on simple graph models, which provide optimistic assumptions of NWB dissemination. On the other hand, models that fully consider the complex propagation characteristics of NWBs quickly become unsolvable due to their complexity. In this paper, we present the Monte Carlo method Probabilistic Breadth-First Search (PBFS) to approximate the reachability of NWB protocols. PBFS simulates individual NWBs on graphs with probabilistic edge weights, which reflect link qualities of individual wireless links in the WMN, and estimates reachability over a configurable number of simulated runs. This approach is not only more efficient than existing ones, but further provides additional information, such as the distribution of path lengths. Furthermore, it is easily extensible to NWB schemes other than flooding. The applicability of PBFS is validated both theoretically and empirically, in the latter by comparing reachability as calculated by PBFS and measured in a real-world WMN. Validation shows that PBFS quickly converges to the theoretically correct value and approximates the behavior of real-life testbeds very well. The feasibility of PBFS to support research on NWB optimizations or higher level protocols that employ NWBs is demonstrated in two use cases.**

Keywords*—*Wireless mesh networks; Network-wide broadcasts; Probabilistic network graphs; Monte Carlo methods

I. Introduction

Over the last decade, Wireless Mesh Networks (WMNs) have become increasingly diffused in a manifold of applications. A relatively low cost of deployment, ad-hoc and automatic configuration are among their advantages. However, inherent dependability issues due to the wireless communication remain until today. A fundamental operation for WMNs are Network-wide Broadcasts (NWBs), required by many of the auto-configuring protocols, such as routing, for distributing information to all nodes in the network. Current protocols for NWBs lead to the well-known broadcast storm problem, which essentially describes excessive redundant transmissions that happen in a meshed network during basic flooding. Many optimizations for NWBs exist that aim to counter this problem by reducing redundant transmission. However, removing redundancy also hurts reliability. To properly evaluate such trade-offs, researchers need to rely on extensive simulation and experimentation, since analytic methods do not scale to realistic network sizes due to their complexity.

This paper presents the novel Monte Carlo evaluation method *Probabilistic Breadth-First Search* (PBFS) for flooding and similar NWBs. Without reverting to expensive empirical or analytical methods, PBFS can be used to realistically approximate various metrics, such as reachability or average path length. As a network model, PBFS uses a graph with edge weights representing the probability for successful transmissions. It then simulates the dissemination of NWBs within this graph. PBFS has a very low time and space complexity compared to existing approaches and validation shows that it provides a comparable precision. Thus, it promises to significantly advance research on optimizing NWBs. The application of PBFS is demonstrated in two use cases, which give further insight about NWBs in lossy networks and their usage in higher layer network protocols.

The rest of the paper is structured as follows. An overview of related work is provided in Section II. The WMN model and NWB metrics are defined in Section III followed by a description of PBFS in Section IV. PBFS is validated both theoretically and empirically in Section V and its feasibility to support current research work on NWBs demonstrated in Section VI. Section VII concludes this work.

II. Related Work

General problems and challenges of WMNs are presented in [1]. NWBs as a fundamental operation have been the target of numerous optimizations with very different goals, among them OLSR's MPR selection scheme [2], probabilistic flooding and flooding with retransmissions [3]. The prominent broadcast storm problem was pointed out by Tseng et al. [4]. They and others then proposed approaches which focused on removing redundant transmissions which emanate from a simplified, unweighted graph model. Unit Disk Graphs (UDGs) are one way to obtain these unweighted graphs and are often employed for WMNs and NWBs. In an UDG, a link always exists when two nodes have an Euclidian distance of less than 1, otherwise not. Kuhn et al [5] propose an extension called Quasi Unit Disk Graphs (QUDGs), introducing an uncertainty that allows to model the impact of obstacles on wireless transmissions. Few approaches consider link qualities when optimizing NWBs as in [6] or [7], [8], where flooding with different forwarder selection strategies was proposed and evaluated. Evaluation revealed that in reasonably lossy networks, existing optimizations severely decrease NWB reachability when removing redundancy. In this paper, we are interested in calculating the reachability for NWB protocols, given a weighted topology as input. Oikonomou et. al [9] did a similar analysis for probabilistic flooding, using randomly generated graphs and not considering edge weights. Chen et. al [10] define the Flooding Path Probability (FPP), which corresponds to the reachability as defined in Section III. Their provided algorithm to compute the FPP has a reduced but still exponential complexity.

III. Network Model & Broadcast Metrics

A WMN can be modeled as an undirected graph $G = (V, E)$, where each vertex $v \in V$ corresponds to a WMN node and each edge $e \in E$ corresponds to a wireless communication link between two nodes. In QUDGs [5], there is a communication link between two nodes if their Euclidian distance is less than d (with $0 < d < 1$) and none if it is greater, otherwise unspecified. Thus, a link may or may not exist within this uncertainty range. This neglects that links itself are rarely perfect or completely non-existent, but their quality varies in between. Thus, in the model proposed in this section wireless links are not just binary but approximated with a transmission probability $p_{i,j} : E \to [0, 1]$, the probability a transmission from node i to node j succeeds. If $p_{i,j} = 0$, the edge is not included in the graph.

In practice, $p_{i,j}$ usually reflects a measurement-based approximation. Link qualities are never constant but vary over time due to small-scale fading effects, micro-mobility of the environment and interference. Node outages may also happen. The graph model assumes independence between the qualities of different links, any dependence is expected to be already included in the approximation of $p_{i,j}$. For the use cases in Section VI, static networks with no node mobility are assumed. However, the method described in Section IV can also be adjusted to support network topology dynamics reflected in varying values for $p_{i,j}$.

A. Reachability of Network-wide Broadcasts

To evaluate the performance of NWBs within the weighted graph model, we define reachability as the main metric, which captures an important aspect of the dependability of a NWB. We distinguish between the following two reachabilities:

Individual reachability $r_{a \to b}$ – The probability that a node b is reached by a NWB initiated by node a. In practice, $r_{a \to b}$ is averaged over multiple runs as a single run will only give a binary value.

Global reachability r_a – The percentage of nodes that receives a NWB sent by node a. Let R be the number of nodes reached by a NWB initiated from node a. The global reachability is defined as $r_a = \frac{R}{|V|}$ and derived from a single NWB or averaged over multiple NWBs.

So, global and individual reachabilities depend on the given network topology with its transmission probabilities, NWB protocol, NWB source node and, for individual reachability, the destination node. Source independent results can be achieved by considering NWBs from every node in the network and averaging over all of them, which is done in the first use case in Section VI-A. There, the reachability is just denoted as r.

B. Calculating Reachability of Network-wide Broadcasts

There exist purely analytical methods to calculate the reachability metrics based solely on the described weighted graph model. For example, Chen et a. [10] have shown how the individual reachability $r_{a \to b}$ can be obtained by summing up the probabilities of the family of edge sets that contain a path from a to b, as in Equation 1: $2^E_{a \to b}$ is the set of all subsets of edges that contain a path from a to b. Calculating this formula has an exponential complexity due to the iteration of the power set.

$$r_{a \to b} = \sum_{E' \in 2^E_{a \to b}} \left(\prod_{(i,j) \in E'} p_{i,j} \prod_{(k,l) \in E \setminus E'} (1 - p_{k,l}) \right) \quad (1)$$

Another possible approach is to enumerate all given paths between a and b and calculate the probability of not reaching node b over any of these paths. Enumerating all simple paths between two nodes is NP-hard for arbitrary graphs. Furthermore, as these paths are not independent and can contain common subpaths, this approach requires factoring out these common subpaths, which is again NP-hard.

Chen's algorithm for the FPP [10] also computes the individual reachability $r_{a \to b}$. It resorts to computing the joint probability distribution that a packet is received by a subset of the vertex cut in each step of its algorithm. The complexity is $O(N(C\Delta^+ + 2^C \Delta^-))$, where C is the size of the largest vertex cut used by the algorithm, Δ^+ is the maximum out-degree and Δ^- the maximum in-degree of the graph. Therefore, calculating the FPP is efficient if the network has only a small largest vertex cut, which is not true for high connectivity WMNs.

Due to their complexity, these approaches are only feasible in simple networks. To support the reachability evaluation of novel NWB strategies, more efficient methods are needed. Equation 1 will be used in Section V-A, however, to validate the method proposed in Section IV in exemplary small networks.

IV. Probabilistic Breadth-First Search

This section introduces the Monte Carlo method *Probabilistic Breadth-First Search* (PBFS) that models flooding as a slight modification of *Breadth-First Search* (BFS) [11] in a weighted graph as described in Section III. In contrast to regular BFS, in PBFS, for a node i its neighbor j in the graph is only considered if the edge between i and j succeeds a random roll against its transmission probability $p_{i,j}$. Consequently, PBFS can be seen as a flooding process in a graph whose edges are weighted with independent transmission probabilities, abstracting from wireless effects such as contention and collisions. This way, each run of PBFS simulates how a NWB packet would traverse the WMN, without resorting to complex network simulations or testbed experiments, and with considerably lower complexity than a purely analytical evaluation.

```
1: procedure PBFS(G, src)                          ▷ G provides nodes and p_{i,j}.
2:     mark[n] ← False ∀ n ∈ G                     ▷ BFS visited marker.
3:     NoT[n] ← 1 ∀ n ∈ G                          ▷ Number of transmissions setup.
4:     reached ← 0                                 ▷ How many nodes were finally reached.
5:     nodes ← [src]                               ▷ Queue of nodes to be processed.
6:     mark[src] ← True
7:     while nodes ≠ ∅ do
8:         i ← nodes.dequeue()
9:         for j ∈ G.neighbors(i) do
10:            p_success ← 1 − (1 − p_{i,j})^{NoT[i]}
11:            if mark[j] = False ∧ random() < p_success then
12:                nodes.enqueue(j)
13:                mark[j] ← True
14:                reached ← reached + 1
15:            end if
16:        end for
17:    end while
18: end procedure
```

Fig. 1: Demonstration of the basic *Probabilistic Breadth-First Search* (PBFS) algorithm, which mimics flooding in a WMN.

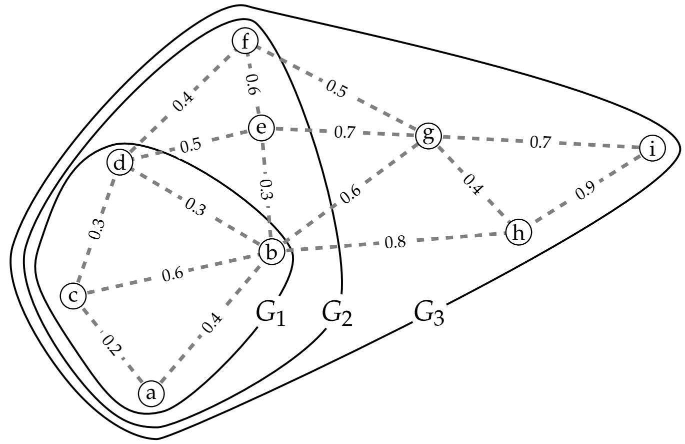

Fig. 2: Three example networks for theoretical validation: $\{a, b, c, d\}$, $\{a, b, c, d, e, f\}$ and $\{a, b, c, d, e, f, g, h, i\}$.

The algorithm is presented in Figure 1. The global reachability r_{src} is the number of nodes that were marked (*reached*), divided by all nodes. In real networks, for example, the probabilistic edge weights $p_{i,j}$ in the graph can be obtained by doing link probing for the *Expected Transmission Count* metric [12]. The PBFS procedure in Figure 1 samples a sufficient number of times to approximate with which probability a destination node b is reached by a NWB from node a (the individual reachability $r_{a \to b}$), Additionally, it approximates over how many hops b is reached, providing a list of tuples with the probability for any path length b is reachable by the flooding process initiated by a. For reasons of brevity, path length distributions have not been included in Figure 1. It should be mentioned that while in this paper, only undirected graphs are considered, PBFS can be used just as well on directed graphs.

PBFS has several advantages over the existing analytical approaches. It inherits the complexity $O(|E| + |V|)$ from BFS, which allows to calculate reachability even for large graphs. PBFS also provides over how many hops a node is reached and with which probability. Furthermore, it does that not only for one node pair but for all reachable nodes from a specific NWB source node. PBFS can work with dynamic values of $p_{i,j}$. Provided a function that calculates $p_{i,j}$ over time or a series of ETX measurements, each run of PBFS could use different values for $p_{i,j}$ as input. Additionally, it is possible to modify PBFS to model modifications of flooding and other NWB approaches, as will be demonstrated in Section VI. For example, PBFS allows for different configurations of transmissions done by each node (see code in Figure 1, Lines 3 and 10).

V. Validation

As PBFS relies on repeated randomized sampling runs, it reflects a Monte-Carlo method to approximate reachability. In a nutshell, more runs increase the accuracy of its approximation. An estimation of a sufficient sampling number for PBFS, its deviation from the exact theoretical reachability and its confidence follows in Sections V-A and V-B.

TABLE I: PBFS deviation (with 99% confidence after 1000 trials) with different number of runs x, implementation runtime included in parentheses for illustration.

Network	Nodes / Edges	Equation 1	PBFS $x = 10^3$	PBFS $x = 10^4$	PBFS $x = 10^5$
$G_1, r_{a \to d}$	4 / 5	0.22824 (0.004 s)	$\pm 5.10\%$ (0.020 s)	$\pm 1.58\%$ (0.199 s)	$\pm 0.47\%$ (1.990 s)
$G_2, r_{a \to f}$	6 / 9	0.19682 (0.116 s)	$\pm 4.71\%$ (0.023 s)	$\pm 1.45\%$ (0.230 s)	$\pm 0.48\%$ (2.275 s)
$G_3, r_{a \to i}$	9 / 16	0.39104 (31.004 s)	$\pm 6.30\%$ (0.032 s)	$\pm 1.90\%$ (0.315 s)	$\pm 0.59\%$ (3.179 s)

A. Theoretical Validation

First, PBFS is evaluated in three exemplary small networks, where the complexity still allows a precise analytical assessment of reachability using Equation 1. These networks are depicted in Figure 2. To save space, all three are depicted within one Figure. The first network G_1 consists of nodes a, b, c and d and connecting links. The second G_2 adds nodes e and f to G_1, while the third network G_3 consists of the whole graph in Figure 2. Edge labels correspond to $p_{i,j}$.

Evaluation results are listed in Table I and demonstrate two phenomena. First, the exponential increase in runtime is visible already for these simple networks. Since these times are based on non-optimized implementations, they have to be taken with a grain of salt. But the effect of combinatorial explosion using Equation 1 is obvious. Second, one can see that PBFS converges relatively fast to the correct reachability value.

Already after only 1000 runs, the relative error is around 5% and below 0.5% after 100000 runs with a confidence of 99%. In fact, with an increasing number of runs x we see the error decreasing by a factor of $\sqrt{x_m / x_n}, m > n$. This is the expected behavior assuming a constant standard deviation and is reflected by a linear decrease of the absolute error in the logarithmic plot in Figure 3. Together with the convergence of reachability it validates PBFS as a proper Monte Carlo method to approximate the behavior of NWBs.

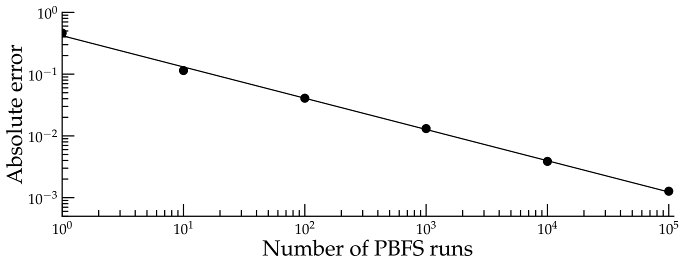

Fig. 3: Absolute PBFS error (y-axis) over number of runs (x-axis). Obtained by comparing PBFS results with Equation 1 for network G_3 in Figure 2. Axes have a logarithmic scale.

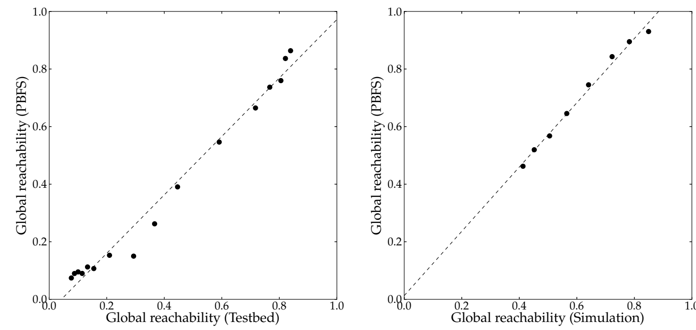

(a) PBFS vs. Testbed (b) PBFS vs. Simulation

Fig. 4: Comparison of global reachability predicted by PBFS (y-axis) and achieved in reality, simulations (x-axis). A point constitutes the reachability for a single scenario.

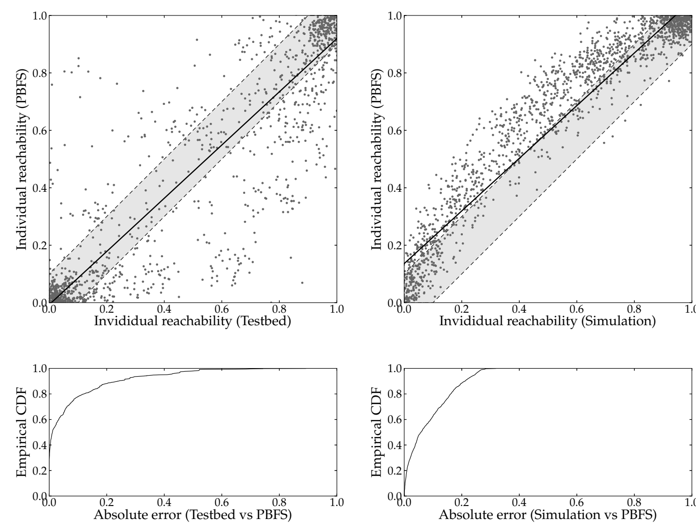

Fig. 5: Comparison of individual node reachabilities predicted by PBFS (y-axis) and achieved in reality or simulations (x-axis) across different scenarios (top). Corresponding CDF plots for the absolute differences shown in the bottom graphs.

B. Experimental Validation

To validate the feasibility of PBFS in realistic WMNs, this section provides an experimental validation, comparing global and individual reachability approximations given by PBFS with results measured in simulations and real-life wireless testbeds. We focus on basic flooding as NWB across different network topologies. Further, hop length distributions for exemplary node pairs are shown and compared. PBFS estimations are expected to have a higher deviation, as it abstracts from short-term effects not reflected in the measured transition probabilities $p_{i,j}$.

1) Experiment Setup: Flooding is implemented within the Click Modular Router framework, which allows to run it either in a wireless testbed or in simulators. For real-word experiments, the DES-Testbed at Freie Universität Berlin is used, which consists of over hundred nodes equipped with IEEE 802.11 hardware that are spread across multiple campus buildings. Due to space limitations in this paper, we refer to [13] for a complete description of the testbed. As simulator, ns-3.14 with extensions for Click is employed. The wireless model is based on the Jakes propagation loss model [14], nodes form an NPART topology [15] consisting of 275 nodes. Different scenarios are obtained by varying the transmission power in the testbed as well as in simulations. Measuring transmission probabilities $p_{i,j}$, which are the input to PBFS, is done by a link prober that sends neighbor discovery packets and calculates the ETX [12]. Flooding packets are sent as link-layer broadcasts with a fixed 1 Mbps data rate and payload sizes of 400 bytes, the same size as the probing packets. In each scenario one node was chosen to start 1000 NWBs.

2) Experiment Results: In Figure 4, the global reachability for flooding is shown, contrasting results obtained by PBFS with measurements from the testbed (Figure 4a) and simulation (Figure 4b). One data point corresponds to reachability in a single scenario. Linear regression results in the dashed line and shows a very strong linear correlation with a correlation coefficient greater than 0.99 and very small p-values in both comparisons. PBFS slightly over-estimates reachability when compared to simulations, an effect that is not shown in testbed results. This bias in simulations hints at a systemic property in the simulation setup that needs further investigation.

Figure 5 shows the individual node reachabilities (top graphs) with their corresponding empirical cumulative distribution functions of the absolute difference between predicted and achieved reachabilities for all nodes (bottom graphs). Again, a data point reflects the reachability of a single node as estimated by PBFS versus the results from testbed or simulation. The shaded area in the top graphs constitutes an absolute error $\pm 10\%$ of the ideal line with slope 1. Linear regression with all data points results in the solid line. While there are outliers when compared with both simulation and testbed results, PBFS reasonably approximates the measured reachability. Simulation shows a much narrower distribution, although the points show a characteristic curve which again hints at some systematic property that warrants further research. In the bottom graphs, corresponding CDF plots for the absolute differences are shown. The CDF plots show that in both setups, the PBFS approximation does not differ more than 0.2 from the measured reachability for 80% of all data points.

Last but not least, PBFS provides the distribution of path length probabilities for each destination node of a NWB. A comparison of these distributions between PBFS and testbed

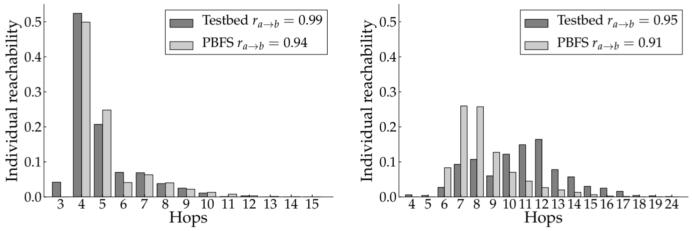

(a) Node t9-155

(b) Node a6-028

Fig. 6: Exemplary path length distributions for NWBs to nodes t9-155 and a6-028.

measurements is shown in Figure 6. In general, PBFS is able to estimate these distributions reasonably accurate as exemplary shown in Figure 6a. This is despite a 5% difference in reachability. However, there are few nodes, especially at higher distance from the NWB source, that show misfits, usually with an overestimated path length as can be in Figure 6b. These inaccuracies, as well as the inaccuracies for global and individual reachabilities, stem from basically two facts: (1) The calculation of link transmission probabilities has an intrinsic error based on its own accuracy and cannot capture short-term effects visible in experiment runs. (2) Transmission probabilities are not independent, especially during flooding (broadcast storm problem). Nonetheless, PBFS provides a valid approximation of reachability for NWBs and it does this with a high efficiency.

VI. USE CASES

This section shows in two different research works how PBFS could be used to facilitate NWB modeling.

A. Flooding with Retransmissions

The efficiency of PBFS allows to use reachability as an optimization goal. This is demonstrated by evaluating the reachability of flooding, modified to use retransmissions. The two evaluated strategies differ in the number of transmissions each node makes:
(1) Globally fixed number of transmissions – In contrast to basic flooding every node repeats a freshly received packet not only once, but a fixed number of times. This number is the same for every node in the network.
(2) Locally optimized number of transmissions – Every node has its own individual number of transmissions. The transmission values are optimized using reachability approximations from PBFS.

A genetic algorithm is used to come to a near-optimal number of transmissions with respect to PBFS reachability in (2), which is implemented within the *Distributed Evolutionary Algorithms in Python* (DEAP) package. The algorithm operates on a list of numbers as individuals which correspond to the number of transmissions for every node. It mutates, combines and selects these individuals according to the parameters given in Table II. It should be noted that this case study was done to demonstrate the capabilities of PBFS and although these parameters were chosen using common sense, there surely is potential for optimization outside the scope of this work.

TABLE II: Parameters of the genetic optimization algorithm

Parameter	Value
Type	Evolutionary $(\mu + \lambda)$ algorithm with $\mu = 100$ and $\lambda = 90$
Individual	List of the number of transmissions (NoT) for every node
Fitness function	r - $\frac{\sum NoT}{1000}$ (reachability slightly muted by the total NoT)
Mutation	Random addition of $\{-3, -2, \ldots, 3\}$ to each NoT with 50% probability
Crossover	Random selection of NoT from the two parents
Offspring	$\lambda = 90$ offspring are created (30% crossover and 70% mutation)
Selection	$\mu = 100$ individuals from both offspring and population (elitist)

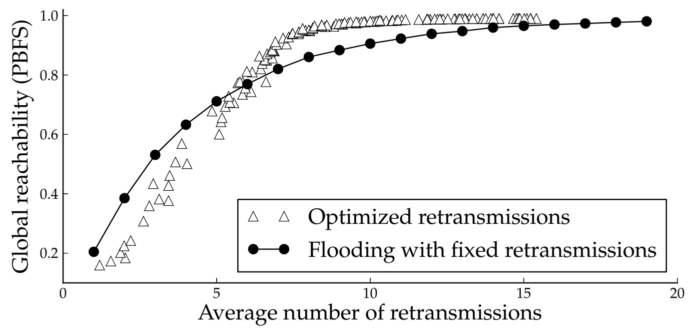

Fig. 7: Comparing reachability of two different strategies for flooding with retransmissions.

Figure 7 shows the evolution of the fittest individuals over 600 generations compared to approach (1) of using a fixed number of transmissions. The corresponding topology was chosen to have very weak links, so that basic flooding achieves only a reachability of around 20%. The curve shows that a high fixed number of transmissions is required for high reachability. For example, 12 transmissions are needed to increase reachability above 90%. In contrast, using locally optimized transmissions for each node arrives at the same reachability with only 7 retransmissions on average. The conclusion is that the optimal number of transmissions can be very different and should not be a globally fixed parameter. Also, optimizing the number of transmissions could not have been achieved without an efficient reachability estimator, such as PBFS.

B. Latency Estimation for Responsiveness Evaluation

PBFS also facilitates evaluation of high level application protocols that rely on NWBs. In [16], this is demonstrated by evaluating the responsiveness of service discovery. Discovery provides a layer in service networks where instances of a specific service can either publish themselves or be enumerated by clients. Current discovery protocols rely heavily on multicast communication, which is realized as NWB. Responsiveness in this case describes the probability of successful discovery within a deadline, even in the presence of faults. It is modeled with a hierarchy of Markov and semi-Markov models. To evaluate the responsiveness of discovery operations, the probability for individual packets to traverse the network within a specific deadline needs to be estimated. This is done using the path length distributions provided by PBFS and combining them with approximated transmission times per hop for specific packet sizes based on the 802.11 standard.

Since the efficiency of PBFS allows to do such evaluation for arbitrary pairs in the network, responsiveness can be

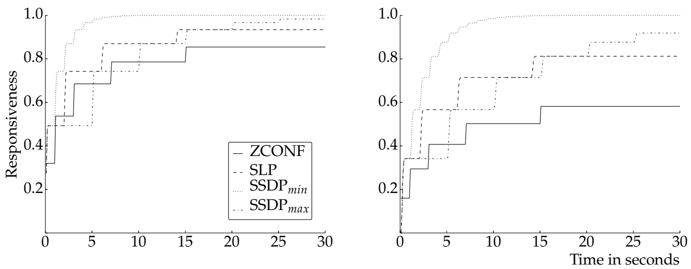

(a) Provider 1, within same cloud (b) Provider 2, maximum distance

Fig. 8: Discovery responsiveness over time for different service discovery protocols when requested from a single client.

evaluated for individual pairs as well as averaged for all clients of a specific service provider. This is done for the three most prevalent discovery protocols in IP networks: SLP, SSDP and Zeroconf. Figure 8 shows the responsiveness over time for two different service providers when requested from the same client. The impact of distance on responsiveness is visible as the closer provider 1 (Figure 8a) can be discovered with considerably higher responsiveness than provider 2 (Figure 8b), which is at a maximum distance from the client. Also, the effect of the individual retry strategies of the protocols is notable: The characteristic steps in the curves happen whenever a discovery request is retried. All details about the methodology and its evaluation can be found in [16].

VII. Conclusion

Network-Wide Broadcasts (NWBs) are a fundamental operation in Wireless Mesh Networks (WMNs). Since they cause a high load in the network, optimizing NWBs remains a major challenge. However, evaluating optimized NWB strategies is either computationally expensive when done purely analytical and with realistic assumptions or time-consuming when resorting to experiments. We propose an efficient Monte Carlo method named *Probabilistic Breadth-First Search* (PBFS) to approximate NWB reachability and diverse related metrics. It works on a network graph where links between nodes are weighted with their transmission probability, a value usually measured during runtime in WMNs. PBFS supports simple flooding and modifications to that NWB model.

Validation of PBFS has been carried out by comparing estimations of PBFS with both analytical solutions and experimental results gathered in simulation and in a real-world WMN testbed. It shows that PBFS approximates the theoretical solutions with a very low error already after few runs and with considerably lower complexity, converging to the correct value with an exponential decrease of the absolute error. This validates PBFS as a proper Monte Carlo method. Despite the abstractions within the graph model, the reachability of NWBs as approximated by PBFS overall reflects the behavior as measured in simulations and reality. This makes PBFS well suited to support development and evaluation of novel strategies for NWBs before deploying them in real WMNs.

The paper describes two case studies that successfully employed PBFS. In the first study, in a NWB strategy modified to use retransmissions, the number of retransmissions has been optimized using reachability estimations given by PBFS. The second study uses PBFS to model the stochastic behavior of NWBs when evaluating the responsiveness of common service discovery protocols, which rely heavily on NWBs. The results in both studies could not have achieved a comparable accuracy without resorting to considerably more complex methods.

References

[1] I. F. Akyildiz and X. Wang, "A survey on wireless mesh networks," *IEEE Communications Magazine*, vol. 43, no. 9, pp. S23–S30, 2005.

[2] J. Hrri, C. Bonnet, and F. Filali, "OLSR and MPR: mutual dependences and performances," in *Challenges in Ad Hoc Networking*, ser. IFIP International Federation for Information Processing, K. Agha, I. Gurin Lassous, and G. Pujolle, Eds. Springer US, 2006, vol. 197, pp. 67–71.

[3] B. Blywis, P. Reinecke, M. Gunes, and K. Wolter, "Gossip routing, percolation, and restart in wireless multi-hop networks," in *2012 IEEE Wireless Communications and Networking Conference (WCNC)*, Apr. 2012, pp. 3019–3023.

[4] 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, pp. 153–167, Mar. 2002.

[5] F. Kuhn, R. Wattenhofer, and A. Zollinger, "Ad hoc networks beyond unit disk graphs," *Wireless Networks*, vol. 14, no. 5, pp. 715–729, October 2008.

[6] T. Zhu, Z. Zhong, T. He, and Z.-L. Zhang, "Exploring link correlation for efficient flooding in wireless sensor networks," in *Proceedings of the 7th USENIX conference on Networked systems design and implementation*, ser. NSDI'10. Berkeley, CA, USA: USENIX Association, 2010, p. 44.

[7] B. Lichtblau and J.-P. Redlich, "Network-wide broadcasts for wireless mesh networks with regard to reliability," in *19th Symposium on Communications and Vehicular Technology in the Benelux (SCVT)*. IEEE Computer Society, November 2012, pp. 1–6.

[8] B. Lichtblau and J.-P. Redlich, "Link quality based forwarder selection strategies for reliable network-wide broadcasts," in *14th International Symposium and Workshops on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*. IEEE Computer Society, June 2013, pp. 1–6.

[9] K. Oikonomou, G. Koufoudakis, and S. Aissa, "Probabilistic flooding coverage analysis in large scale wireless networks," in *19th International Conference on Telecommunications (ICT)*, April 2012, pp. 1–6.

[10] P. Chen, K. H. Johansson, P. Balister, B. Bollobás, and S. Sastry, "Multipath routing metrics for reliable wireless mesh routing topologies," KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden, Tech. Rep. TRITA-EE 2011:033, March 2011.

[11] S. Even, *Graph Algorithms*, 2nd ed. Cambridge University Press, November 2011.

[12] D. S. J. D. Couto, D. Aguayo, J. Bicket, and R. Morris, "A high-throughput path metric for multi-hop wireless routing," in *9th annual international conference on Mobile computing and networking*. San Diego, USA: ACM, 2003, pp. 134–146.

[13] B. Blywis, M. Günes, F. Juraschek, and O. Hahm, "Properties and topology of the des-testbed," Freie Universität Berlin, Computer Systems & Telematics, Takustraße 9, 14195 Berlin, Germany, Tech. Rep. TR-B-11-02, March 2011.

[14] Y. R. Zheng and C. Xiao, "Simulation models with correct statistical properties for rayleigh fading channels," *IEEE Transactions on Communications*, vol. 51, no. 6, pp. 920–928, June 2003.

[15] B. Milic and M. Malek, "NPART - node placement algorithm for realistic topologies in wireless multihop network simulation," in *2nd International Conference on Simulation Tools and Techniques (Simutools)*, ser. Simutools '09. ACM, March 2009, pp. 9:1–9:10.

[16] A. Dittrich, B. Lichtblau, R. Rezende, and M. Malek, "Modeling responsiveness of decentralized service discovery in wireless mesh networks," in *MMB & DFT*, ser. Lecture Notes in Computer Science, K. Fischbach and U. R. Krieger, Eds. Springer International Publishing Switzerland, March 2014, vol. 8376, pp. 88–102.