

Efficient and Resilient Overlay Topologies over Ad Hoc Networks

Sandrine Calomme and Guy Leduc

Research Unit in Networking
Electrical Engineering and Computer Science Department
University of Liège, Belgium

Abstract. We discuss what kind of overlay topology should be proactively built before an overlay routing protocol enters a route search process on top of it.

The basic overlay structures we study are the *K-Nearest Neighbours* overlay topologies, connecting every overlay node to its K nearest peers.

We introduce a family of optimizations, based on a pruning rule. As flooding is a key component of many route discovery mechanisms in MANETs, our performance study focusses on the delivery percentage, bandwidth consumption and time duration of flooding on the overlay. We also consider the overlay path stretch and the overlay nodes degree as respective indicators for the data transfer transmission time and overlay resilience.

We finally recommend to optimize the K-Nearest Neighbours overlay topologies with the most selective pruning rule and, if necessary, to set a minimal bound on the overlay node degree for improving resilience.

1 Introduction

Ad hoc networks are formed without the use of any existing network infrastructure nor centralized administration. The devices in contact can have different hardware capabilities, software, application needs, and mobility pattern. Plenty of multi-hop routing protocols have been proposed for MANETs and, in such heterogeneous networks, the best one may be different for each set of communicating nodes. The preferred routing solution could also change along time, because of mobility and varying network conditions. Consequently, the requirement of choosing a routing protocol and imposing it to all ad hoc devices in order to form a MANET is a limitation. To overcome this restriction, we propose to copy the layered approach of Internet [1]: agree only on a few unspecialized protocols at the physical, data link and routing layers, imposed by their proved qualities or de facto, and over this basic architecture, develop plenty of more specialized solutions, from routing to application, thanks to the overlay technique. Overlay routing could promote the deployment of ad hoc networks, offering a very flexible ground for a variety of applications using the overlay routing protocol that best fits their specific needs.

In this paper, we discuss what kind of overlay topology should be proactively built before an overlay routing protocol enters a route search process on top

of it. As flooding is a key component of many route discovery mechanisms in MANETs, our study focusses on the bandwidth consumption and time duration of flooding on the overlay. We also consider the overlay path stretch and the overlay nodes degree as respective indicators for the data transfer transmission time and overlay resilience.

The interference level is not directly addressed. We let the task of reducing interferences to the underlay topology control algorithm and assume that reducing the number of packets emitted per flood is an efficient way to pace collisions due to the overlay use. We also do not present how the studied topologies can be built or maintained. For example, the reader will not find any test on mobile networks, which does not mean that we assume a static network. The overlay topology control protocol itself will be presented in a follow-up paper. Its design guideline will be to maintain, in a mobile context, the overlay topology as close as possible to a target overlay topology, chosen in accordance with the conclusions drawn in Section 6.

2 Related Work

A major part of the current literature about overlays addresses peer-to-peer applications. Although this work was not conducted for P2P networks, it could probably be exploited in unstructured peer-to-peer middleware. We did not explore this open issue but can compare in some points our work to what have been done for ad hoc P2P networking. In several works, for example ORION [2] and [3], it is assumed that all nodes run the proposed protocol. The use of an overlay allows to get rid of this restriction. In [4], the Gnutella protocol is optimized for ad hoc networks. XL-Gnutella proactively builds an overlay on top of which queries can be efficiently disseminated by the underlying routing protocol. This is similar to our objective of building an overlay for the propagation of overlay routing requests. In Section 5.3, we compare the topologies we recommend for overlay routing to the XL-Gnutella ones.

Topology control (TC) consists of selecting a subset of edges in a graph representing the communication links between network nodes, with the purpose of maintaining some global graph property (e.g., connectivity), while reducing energy consumption and/or interference [5]. Similarly, our problem requires to select a subset of paths between overlay nodes, with the purpose of maintaining their connectivity, while reducing the number of messages they induce in the whole network when they flood an overlay message. Hence, the roots of this work can be found in the TC literature. Mechanisms for building the presented overlay topologies are inspired from two topology control protocols : k-Neigh [6] and XTC [7]. We first present overlay topologies where each overlay node must be aware of a minimal number of the closest other overlay nodes, a process identical to the one used by k-Neigh, but without discarding asymmetric neighbours. We then propose optimizations of the obtained topologies based on an XTC-like criterion.

3 Study Overview

In the sequel, all concepts related to the whole ad hoc network, and not only to overlay nodes, will be identified by the term “underlay”.

We consider a connected underlay and assume that a routing protocol that builds the shortest symmetric paths is available to all nodes. Overlay nodes are randomly and uniformly distributed on the set of ad hoc nodes. The proportion of overlay nodes is called the overlay density.

3.1 Fundamental Properties

The overlay topologies we discuss are strongly connected, i.e. there exists a path on the overlay graph between any pair of overlay nodes, at least with a high probability.

They can be built by a fully distributed algorithm. We also take care of locality: The topology can be built even if each overlay node is allowed to exchange only a few messages with a limited number of nearest overlay nodes. As we do not make any assumption about the underlay routing protocol type, locality is an important feature. With reactive on-demand protocols, like AODV, the control traffic necessary for building a data path between overlay neighbours increases exponentially with the number of hops that separates them.

3.2 Performance Criteria

The objective of the overlay creation and maintenance is to offer a logical communication structure between the overlay nodes which allows the deployment of efficient overlay routing protocols. From this angle of view, the quality of an overlay is strongly linked to desired properties of overlay routing protocols. We translate this in terms of the following objectives.

1. Bandwidth: as routing control traffic is often generated by flooding, the bandwidth necessary to send a message from one overlay nodes to all other ones by using a simple flooding procedure must be as low as possible.
2. Diffusion time: in order to quickly compute valid routes, the overlay control traffic must be flooded rapidly.
3. Delivery: in order to find routes, the overlay control traffic must be received by all overlay nodes.
4. Stretch: the average cost of the shortest overlay path between any pair of overlay nodes must be as close as possible to its value in the underlay. Its maximal cost must also be kept reasonable. In this paper, we use the hop metric. Other metrics, as for example the path delay, could be considered.

3.3 Flooding Technique

As flooding is a key component of many route discovery mechanisms in MANETs, the above performance criteria mainly focus on the flooding of a message on the overlay. We assume that, once the overlay is built, each overlay node knows the

hop distance to every neighbour it has selected. In order to spare bandwidth, an overlay node employs the following flooding technique:

1. For all overlay neighbours located only one hop away, it emits a single overlay message, which is actually broadcast in the underlay with a *Time To Live* (TTL) field set to one.
2. For every overlay neighbour located further away, an individual overlay message is created, which will be unicast to it by the underlay routing protocol.

3.4 Simulations Description

All simulations in this paper, except for Section 4.4, were realized with ns-2.29.

The ad hoc nodes are randomly and uniformly distributed on a square field. We vary their number from 50 to 250. Overlay nodes are randomly chosen in the set of ad hoc nodes. All experiments were made for overlay densities ranging from 10 to 90%. For the sake of brevity, we only present graphics for the 50% overlay density. Analysis is identical for all densities.

For a given set of ad hoc nodes, the more efficient underlay topology control (TC) algorithm is used, the more traffic is needed for the construction, use and maintenance of overlays built on top of its resulting logical topology [8]. Hence, in order to test the overlay topologies in a stringent environment, we employ the logical topologies obtained after the use of an ideal homogeneous underlay TC technique which assigns the same value r to each node's radio transmission range, r being the minimal value that makes the underlay connected (i.e. the critical radio transmission range).

The underlay routing protocol used is AODV [9]. The performance criteria are only evaluated on strongly connected overlay topologies. The overlay topologies are calculated offline and provided as input to the ns simulator. A source node emits 23 overlay messages of 64 bytes, at the rate of one message per second. The performance study ignores the period elapsed during the transmission of the first 3 messages. Their flooding necessitates the building of AODV paths between the overlay neighbour pairs. Consequently, the AODV traffic is heavier at the beginning of the simulations and the diffusion time of the first overlay messages is higher than for the following messages. When there is no congestion, the latter must be forwarded on AODV paths that are already up. Each point on the graphics is a mean calculated on 20 trials.

4 Building Topologies That Fulfill the Locality and Connectivity Properties

4.1 Ropt: The Critical Neighbourhood Range

One simple way to give preference to nearest neighbours, and thus to respect the locality principle, is to fix the maximal hop distance between overlay neighbours, the *neighbourhood range*. For any underlay and subset of overlay nodes, one can compute the critical neighbourhood range, that is the minimal neighbourhood




Fig. 1. Example of the Ropt and Kopt overlay topologies for the same underlay topology

range R_C such that the overlay is connected [8]. We denote *Ropt* (R optimal) a topology obtained when each overlay node considers as a neighbour any overlay node that is located at a distance less than or equal to R_C .

4.2 Kopt: The Critical Neighbourhood Cardinality

Another simple way to respect locality is to fix the maximal number of overlay neighbours. For any underlay and subset of overlay nodes, one can compute the critical number of overlay neighbours, that is the minimal neighbourhood cardinality K_C such that the overlay is connected. We denote *Kopt* (K optimal) a topology obtained when each overlay node considers as a neighbour its K_C nearest neighbours, the distance metric being the number of hops. Let k_i be the number of overlay nodes located at i hops from a given overlay node U . If there exists an integer j such that $\sum_{i=1}^{i=j} k_i < K$ and $\sum_{i=1}^{i=j+1} k_i > K$, the required number of overlay neighbours is randomly picked in the set of overlay nodes located at distance $j + 1$ from U ¹.

4.3 Ropt and Kopt Delivery Percentage

Figure 1 shows an example of the Ropt and Kopt overlay topologies for the same underlay. There are 500 nodes and the overlay density equals 50 %. The 250 overlay nodes are represented with (red) squares. The remaining nodes, represented with (blue) circles, are drawn if and only if they are on the shortest path between a pair of overlay neighbours. For this particular underlay and assignment of overlay nodes, the critical neighbourhood range equals 3 and the

¹ We evaluated some more sophisticated policies, but none provided significantly better performance.

critical neighbourhood cardinality equals 4. This figure also illustrates that the Ropt overlay topologies are much denser than the Kopt ones. This is confirmed on Figure 2(a) which shows their average overlay nodes degree. The high overlay nodes degree of Ropt topologies explains their very weak delivery percentage for flooded messages. Severe congestion problems arise for a moderate amount of overlay nodes. Figure 2(b) illustrates this problem for an overlay density equal to 0.5.




Fig. 2. Average overlay nodes degree and overlay flooding delivery percentage for Ropt and Kopt topologies

4.4 KNN: The Minimal Number of Overlay Neighbours Needed for Connectivity

Problem statement. *Kopt* topologies provide better delivery percentages but are difficult to build in practice. There is no analytical function that gives the optimal number of overlay neighbours needed for connectivity. One could imagine a distributed algorithm that determines K_C . For example, the algorithm employed in [10] for electing the best radio transmission range could be adapted. However, this would require the exchange of a lot of information in the whole network. We reject this solution because of its high bandwidth demand.

In the TC field, it has been demonstrated that for any protocol which preserves worst-case connectivity of the ad hoc network, there exists a placement of n nodes such that the maximum physical node degree in the controlled topology equals $n - 1$ [5]. In other words, there is no given number of physical neighbours k , with $k < n - 1$, that implies the connectivity of every network composed of n nodes. However, it has also been shown that setting the minimum number of physical neighbours to 9 is sufficient to obtain connected networks with high probability for ad hoc networks with the number of nodes ranging from 50 to 500 [6].

Similarly, an extensive set of simulations allowed us to determine empirically a parameter K that assures with a high probability the overlay connectivity for a wide range of ad hoc network sizes and overlay densities. We denote KNN (K-Nearest Neighbours) a topology obtained when each overlay node considers as neighbours its K nearest overlay nodes. We now describe the experiment we conducted in order to obtain the K value necessary for the simulations described on Section 3.4.

Testbed. We model the ad hoc network by a graph. Vertices represent the ad hoc nodes that we randomly and uniformly distribute on a unitary square field. We vary their number from 50 to 1000 and the overlay density from 10 to 90%.

For a given set of ad hoc nodes and communication links, the traffic needed for the construction, use and maintenance of overlays is higher on top of sparse logical underlay topologies, that is when an efficient underlay TC algorithm is used [8]. Hence, in order to test the overlay topologies in a stringent environment, we employ the logical topologies obtained after the use of an ideal homogeneous TC technique which assigns the same value r to each node’s radio transmission range, r being the minimal value that makes the underlay connected.

Reduction and extension rules for building the KNN overlay graph.

Let L_U^K denote the set of K nearest overlay neighbours of U . Overlay nodes U and V are K -symmetric neighbours if and only if $U \in L_V^K$ and $V \in L_U^K$. Figure 3 shows an example with $K = 1$.

Many MANET routing protocols assume bidirectionnal links. Moreover, using unidirectional links in route searches only provides an incremental benefit because of the high overhead needed to handle them [11]. We thus fix as an objective to build overlay topologies where the neighbourhood relation is symmetric.




Fig. 3. With $K = 1$, V is a K -asymmetric neighbour of U . V and W are K -symmetric neighbours.

Let L_U denote the set of overlay neighbours selected by overlay node U . For each pair of overlay nodes U and V , there could be two rules to ensure symmetry of the overlay topology:

1. Reduction rule: $V \in L_U$ iff. $U \in L_K^V$ AND $V \in L_K^U$,
2. Extension rule: $V \in L_U$ iff. $U \in L_K^V$ OR $V \in L_K^U$ (graph symmetric closure)




Fig. 4. Result of the reduction and extension rules on the same example topology, with asymmetric neighbours

With the reduction rule, only the symmetric K -neighbours of a node are included in its neighbourhood. With the extension rule, asymmetric K -nearest neighbours are also considered. For a given value K , the topology obtained with the extension rule is a super-graph of the topology obtained with the reduction rule. Its connectivity probability is thus higher. An example is given on Figure 4.

Results. Figure 5 shows the evolution of the percentage of overlays that are connected, for 200 tests with 500 nodes, as a function of the number of nearest overlay nodes (K) for both rules.

The same experiment has been conducted for nodes ranging from 50 to 1000. For a given overlay density, the curves obtained for the different underlay sizes were very close from each other. In other words, we observed that the percentage of connected overlays was more influenced by the overlay density than by the number of nodes.

The three lowest curves are obtained with the reduction rule and the three highest with the extension rule. With both rules, there is a phase where the connectivity probability is very low and a phase where it is very high. The transition from the low-probability phase to the high-probability one arrives sooner and is sharper with the extension rule.

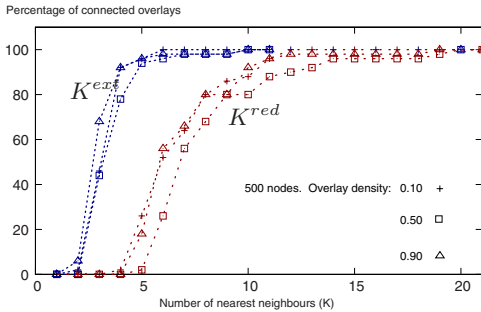


Fig. 5. Percentage of connected overlays as a function of the number of neighbours, with the reduction and extension rules, and for 500 nodes

Ov. density	K_{95}^{ext}	R_{95}^{ext}	K_{95}^{red}	R_{95}^{red}
0.1	5	5	11	8
0.3	7	3	17	5
0.5	8	3	19	4
0.7	8	2	25	4
0.9	6	2	15	3

Fig. 6. Neighbourhood cardinality needed for a connectivity probability equal to 0.95 for the extension and reduction discovery rules (1000 nodes)

Table on figure 6 shows the minimum number of nearest overlay neighbours that must be considered for obtaining 190 connected overlay topologies over 200, for 1000 nodes and different overlay densities. We respectively denoted K_{95}^{red} and K_{95}^{ext} this value for overlays built with the reduction and with the extension rule. All results show that the value of K_{95}^{ext} is far less than K_{95}^{red} . The maximal value of K_{95}^{ext} on our whole set of experiments equals 8, while the maximal value of K_{95}^{red} reached 30.

This table also shows the neighbourhood range (resp. R_{95}^{ext} and R_{95}^{red}) that must be admitted in order to allow the corresponding K_{95}^{red} and K_{95}^{ext} number of

overlay neighbours. For each overlay density, the needed neighbourhood range is one to three hops longer, which is not negligible. Assume that the overlay nodes discover their neighbours by sending hello packets. In this case, the bandwidth consumption for building the overlay rapidly grows with the distance at which these packets must be diffused.

The number of overlay neighbours K needed to obtain a high probability of connectivity for a KNN overlay is much lower with the extension rule, that is if we do include the K -asymmetric neighbours, than with the reduction rule. Moreover the K_{95} value with the extension rule is more reliable than with the reduction one because of the sharper transition from the disconnected to the connected phase. Finally, the discovery uses less bandwidth when accepting K -asymmetric neighbours.

Our conclusion is that the extension rule should be used. It however has a side-effect: Though the nearest neighbour lists have a limited size, a given overlay node could be included in the neighbourhood of a larger number of overlay nodes, due to the symmetric closure. Hence, there is no bound on the overlay nodes degree in the KNN extended topology. Nevertheless, as will be discussed in Section 5, a sufficiently selective pruning criterion moderates a lot this drawback.

In this testbed, with the extension rule, the number of nearest overlay neighbours needed for ensuring the connectivity of 95% of the overlay graphs for any overlay density equals 8.


We would like to point out that the solution we propose is not restricted to the simple underlay model used in these simulations, which are only presented as illustrations of the principles exposed. The important information they bring is not the particular value of $K = 8$ but how it can be determined and why it is preferable to use the symmetric closure and thus let the overlay node degree unbounded.

5 Optimizing the Topologies for Overlay Routing

KNN topologies are connected with a high probability. However, as Ropt topologies, they are too dense. Their delivery percentage of flooded overlay messages is low. In this section, we explore methods for eliminating edges while preserving the connectivity property.

5.1 Shortest Path Pruning

Consider figure 7. The overlay nodes (U , V and W) are grey-shaded. Thick arrows represent the flooding of an overlay message from U and thin ones the corresponding underlay packets. In fig. 7(a), the Kopt overlay topology is used; it is composed of the three edges (U, V) , (V, W) and (U, W) . The flooding of the overlay message on the Kopt topology generates 6 packets on the underlay. However, as illustrated in fig. 7(b), the propagation from U to V , followed by the forwarding from V to W would have been sufficient for all overlay nodes to receive the messages and would have generated only 3 packets. The longest edge of the triangle is unnecessary.



(a) Kopt topology without SPP

(b) Kopt topology with SPP

Fig. 7. Motivation for the Shortest Path Pruning. Thick arrow = overlay message, thin arrow = packet.

We thus introduce the following *Shortest Path Optimization*. Consider three overlay nodes U , V and W , and a distance metric d . The distance metric can be the hop count, the path average delay or any other real positive and symmetric function. Assume that the edge (U, W) is the longest: $d(U, V) \leq d(U, W)$ and $d(V, W) \leq d(U, W)$. The Shortest Path Optimization sets aside the edge (U, W) if and only if $d(U, V) + d(V, W) \leq d(U, W)$. It preserves the connectivity of any overlay graph because an overlay edge is suppressed if and only if an alternative path exists on the overlay.

5.2 Maximal Pruning

Shortest Path Pruning improves the delivery percentage of flooded messages on KNN topologies. However, this pruning method is not sufficiently selective. It can be generalized by setting aside any overlay edge (U, W) such that $d(U, V) + d(V, W) \leq \alpha d(U, W)$, with $\alpha \geq 1$. Connectivity is still preserved.

The higher value is assigned to α , the more edges are pruned. We call this parameter the *pruning selectivity*. Maximal Pruning is reached when any edge (U, W) is suppressed as soon as there exists two shorter edges (U, V) and (V, W) . This behaviour is already obtained for $\alpha = 2$: (U, W) being the longest edge, the inequality $d(U, V) + d(V, W) \leq 2d(U, W)$ is always satisfied.

Let us make the distinction between the one-hop overlay neighbours, or *broadcast neighbours*, and the overlay neighbours located farther, the *unicast neighbours*. The emission of only one broadcast packet is sufficient for an overlay flooded message to reach all the broadcast neighbours. Thus, keeping all broadcast neighbours does not increase the bandwidth consumed per overlay flooding. On the other hand, it densifies the final overlay, without increasing the number of unicast neighbours of any overlay node. The consequence is a lower diffusion time and stretch. It also improves the overlay resilience. We thus modify a little the generalized rule in order to maintain as neighbours every pair of overlay nodes located at one hop from each other.

Therefore we finally define the following generic pruning rule.

Consider three edges $E_1 = (U, V)$, $E_2 = (V, W)$ and $E_3 = (U, W)$, a distance metric d , and a real number α with $1 \leq \alpha \leq 2$. Assume E_3 is the longest edge: $d(E_1) \leq d(E_3)$ and $d(E_2) \leq d(E_3)$.

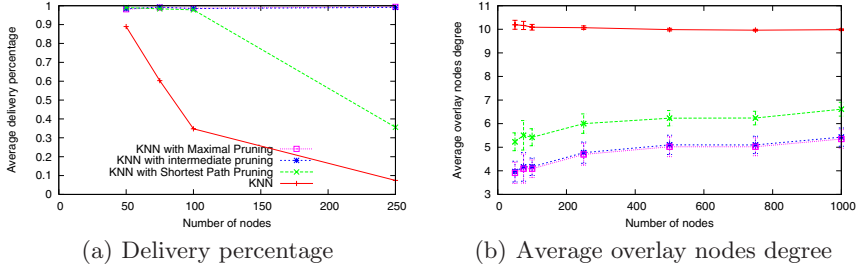


Fig. 8. The delivery percentage and average overlay nodes degree of KNN overlay topologies pruned with various selectivity factors. A common key for both figures is given on the left one.

Edge E_3 is pruned if and only if:

1. E_3 is longer than one hop, and
2. $d(E_1) + d(E_2) \leq \alpha d(E_3)$.

Figure 8 shows the delivery percentage and the overlay node degree for various pruning selectivity (α) on KNN overlay graphs. The distance metric used is the hop count. For the intermediate pruning selectivity, parameter α is set to 1.5. The delivery percentage increases with the selectivity of the pruning method. It is correlated with the average number of overlay neighbours. Flooding an overlay message consumes much bandwidth. Congestion is avoided on sparse overlay graphs.

The average overlay nodes degree of KNN overlay topologies with Maximal Pruning is above 4, with a tight 95%-confidence interval. Maximal Pruning thus preserves some resilience on KNN overlay topologies. Note that resilience is also provided by the underlay topology and routing protocol. The underlay often offers several different paths between each pair of overlay nodes, and a new route can be built when a path between two overlay neighbours breaks.

5.3 Final Performance Study

A brief comparison with XL-Gnutella. We do not criticize the XL-Gnutella protocol, which is intended to be used for P2P data search, not for overlay routing applications. The point here is to show the utility of our own work in the context of overlay routing.

XL-Gnutella is an optimization of the Gnutella protocol for ad hoc networks. To remain fully compatible with the legacy Gnutella protocol, an overlay edge selection algorithm maintains the number of neighbouring peers between 4 and 8.

The delivery percentage of flooded messages on the XL-Gnutella and KNN topologies are compared on Figure 9(a). Recall that the underlays we use for our simulations are very sparse. In this environment, forcing every overlay node to reject neighbours once the overlay node degree has reached the highest water




Fig. 9. XL-Gnutella overlay topologies are intended to be used in a P2P networking context, not for overlay routing

mark of 8 leads a lower connectivity probability than for KNN overlay topologies, for which such restriction does not exist. For the same reason, some overlay edges are longer in XL-Gnutella than KNN topologies. This increases a lot the bandwidth required per overlay message flooding (Figure 9(b)). We also expect, when the underlying routing protocol is reactive, the discovery of XL-Gnutella topologies to consume much more bandwidth than the discovery of KNN topologies, again because some overlay neighbours are selected very far away. In the XL-Gnutella paper, authors use a proactive routing protocol, OLSR, and a cross-layer architecture that allows the P2P middleware to be aware of every overlay node identity and distance, with a low bandwidth consumption. They mention that experiences were also successful with AODV, but that results are better with OLSR.

Comparison of Kopt and KNN with Maximal Pruning. The performance of flooding a message on KNN and Kopt with Maximal Pruning topologies are compared on Figure 10. These are similar, which indicates that the use of the empirical value $K = 8$ before optimization, common for all simulations, instead of the exact minimal number of nearest neighbours needed for overlay connectivity K_C , which value must be determined for each simulation, is not a handicap.

We can also observe that the flooding of an overlay message, which can collect and propagate interesting information for the overlay routing applications, does not consume much more bandwidth than the flooding of a packet on the underlay (exactly 1 packet per node). Note also the reasonable value of the overlay path stretch.

Improving resilience. One could use an intermediate value for α instead of Maximal Pruning, for the purpose of improving the overlay topology resilience. Performance obtained on the KNN topologies pruned with $\alpha = 1.5$ and $\alpha = 2$ for instance are very close (their delivery percentage is compared on fig. 8). However, the gain in resilience is difficult to quantify.

Setting a minimum overlay node degree is another way to increase the redundancy of the overlay, is easier to evaluate and simple to implement. In some cases, it is even required. This is the case, for example, when one wants to deploy

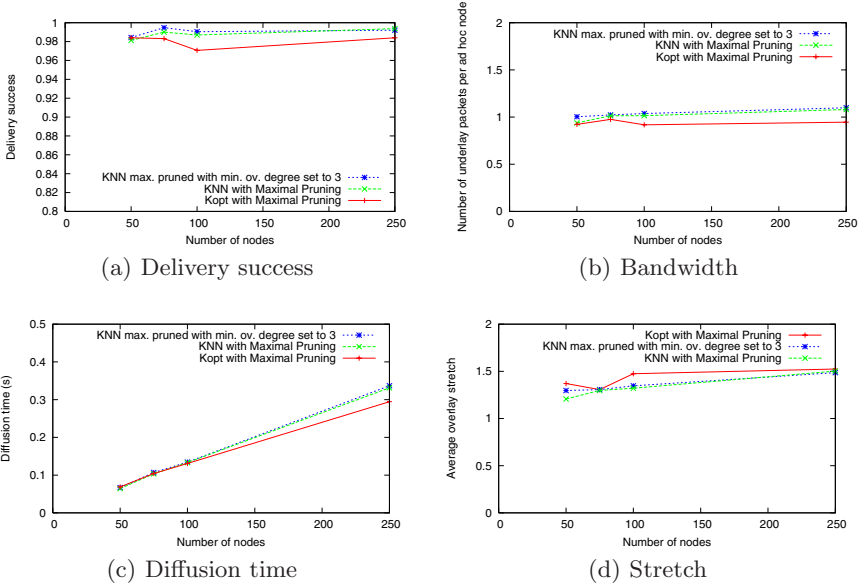


Fig. 10. After pruning, flooding a message on KNN overlay topologies provides similar performance results than on Kopt ones. Setting the minimal overlay degree to 3 does not modify significantly the performance obtained on KNN with Maximal Pruning.

multipath routing on the overlay. A minimal number K_{min} of overlay neighbours is easily guaranteed by reading the nearest overlay nodes list in increasing order of distance and beginning to apply the pruning rule only at the $K_{min} + 1$ element. On Figure 10, we also compare the performance obtained with Maximal Pruning on KNN topologies when applying the pruning rule to the 3 nearest overlay nodes and when systematically keeping them in the final neighbourhood.

6 Conclusions

Overlay routing is well-suited to ad hoc networks. In an ad hoc network, there is no centralized administration. Overlay routing would ease the test and the introduction of new routing protocols, without preliminary agreement between the whole nodes set. Furthermore, an ad hoc network is often composed of several groups of users with specific routing needs. The network conditions, for example the available bandwidth and the mobility level, can vary a lot. Overlay routing would allow each group of users to employ a customized routing protocol for their common application, or to adapt it to the network conditions.

In this paper, we discussed what kind of overlay topology should be proactively built before an overlay routing protocol enters a route search process on top of it.

The basic overlay structures we studied are the *K-Nearest Neighbours* overlay topologies, connecting every overlay node to its K nearest peers. These overlays

can be established with respect to the locality principle, whatever the overlay routing protocol type. This feature is necessary for providing a sustainable building and maintenance cost of the overlays. Parameter K must be empirically tuned. In order to obtain an overlay where the neighbourhood relation is symmetric, the symmetric closure of the K -nearest neighbour graph is preferable to its reduction. The extension method is expected to consume less overlay topology control traffic and is also more reliable, as the corresponding K_{ext} value depends less on the number of ad hoc devices and overlay density. The extension rule lets the overlay node degree unlimited. However, an optimization of the resulting overlay topology cancels this drawback.

We introduced a family of optimization rules of the K -Nearest Neighbours topologies, based on a pruning rule. As flooding is a key component of many route discovery mechanisms in MANETs, our performance study focusses on the delivery percentage, bandwidth consumption and time duration of flooding on the overlay. Simulations illustrated the gain in performance when flooding a message on pruned topologies. The most selective rule, Maximal Pruning, suppresses any overlay edge such that there exists an alternative path in the overlay graph, while preserving from pruning any pair of overlay neighbours that are in the direct communication range of each other. We also considered the overlay path stretch and the overlay nodes degree as respective indicators for the data transfer transmission time and overlay resilience. Maximal Pruning does not increase a lot the path stretch, but can have an undesired effect on the overlay resilience. It can be easily adapted such as to impose a minimal overlay node degree K_{min} . For reasonable values of K_{min} , the performance remains very close to the one obtained with the primary Maximal Pruning rule.

The overlay topology control protocol itself will be presented in a follow-up paper. Its design guideline will be to maintain, in a mobile context, the overlay topology as close as possible to the target K -Nearest Neighbours overlay topology with Maximal Pruning and a minimal bound on the overlay node degree.

Acknowledgements

This work has been partially supported by the European Union under the ANA FET project (FP6-IST-27489).

References

1. Clark, D.: The design philosophy of the DARPA Internet protocols. *Computer Communication Review* 18(4), 106–114 (1988)
2. Klemm, A., Lindemann, C., Waldhorst, O.: A special-purpose peer-to-peer file sharing system for mobile ad hoc networks. In: *Proc. of IEEE Semiannual Vehicular Technology Conference (VTC2003-Fall)*, IEEE Computer Society Press, Los Alamitos (2003)
3. Duran, A., Shen, C.C.: Mobile ad hoc p2p file sharing. In: *Proc. of IEEE Wireless Communications and Networking Conference (WCNC'04)*, IEEE Computer Society Press, Los Alamitos (2004)

4. Conti, M., Gregori, E., Turi, G.: A cross-layer optimization of gnutella for mobile ad hoc networks. In: Proc. of ACM MobiHoc 05, pp. 343–354. ACM Press, New York (2005)
5. Santi, P.: Topology control in wireless ad hoc and sensor networks. *ACM Comp. Surveys* 37(2), 164–194 (2005)
6. Blough, D., Leoncini, M., Resta, G., Santi, P.: The k-neigh protocol for symmetric topology control in ad hoc networks. In: Proc. of ACM MobiHoc 03, pp. 141–152. ACM Press, New York (2003)
7. Wattenhofer, R., Zollinger, A.: XTC: A practical topology control algorithm for ad-hoc networks. In: Proc. of 4th International Workshop on Algorithms for Wireless, Mobile, Ad Hoc and Sensor Networks (WMAN) (2004)
8. Calomme, S., Leduc, G.: The critical neighbourhood range for asymptotic overlay connectivity in ad hoc networks. *Ad Hoc & Sensor Wireless Networks journal* 2(2) (2006)
9. Perkins, C., Royer, E.M.: Ad hoc on-demand distance vector routing. In: Proc. IEEE Workshop on Mobile Computing Systems and Applications(WMCSA'99), IEEE Computer Society Press, Los Alamitos (1999)
10. Narayanaswamy, S., Kawadia, V., Sreenivas, R., Kumar, P.R.: Power control in ad-hoc networks: Theory, architecture, algorithm and implementation of the compow protocol. In: Proc. of European Wireless 2002. Next Generation Wireless Networks: Technologies, Protocols, Services and Applications, Florence, Italy, pp. 156–162 (February 2002)
11. Marina, M.K., Das, S.R.: Routing performance in the presence of unidirectional links in multihop wireless networks. In: *MobiHoc '02: Proceedings of the 3rd ACM international symposium on Mobile ad hoc networking & computing*, pp. 12–23. ACM Press, New York, NY, USA (2002)