



HAL
open science

TopoFormer: An Efficient Link-Set Prediction Architecture for Ad Hoc Network Topology Generation

Félix Marcoccia, Victor Fagoo, Gilles Monzat de Saint Julien, Cédric Adjih,
Thomas Watteyne, Paul Mühlethaler

► To cite this version:

Félix Marcoccia, Victor Fagoo, Gilles Monzat de Saint Julien, Cédric Adjih, Thomas Watteyne, et al.. TopoFormer: An Efficient Link-Set Prediction Architecture for Ad Hoc Network Topology Generation. 2025. ⟨hal-04988831v2⟩

HAL Id: hal-04988831

<https://hal.science/hal-04988831v2>

Preprint submitted on 10 Sep 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

TopoFormer: An Efficient Link-Set Prediction Architecture for Ad Hoc Network Topology Generation

Félix Marcoccia
Inria Paris, Thales

Victor Fagoo
Thales

Gilles Monzat de Saint Julien
Thales

Cédric Adjih
Inria Saclay

Thomas Watteyne
Inria Paris

Paul Mühlethaler
Inria Paris

Abstract

In this paper, we present TopoFormer, a powerful architecture for predicting links between communication nodes in mobile networks. The goal is to imitate, in real time, the results of a costly combinatorial algorithm that generates topologies for networks with directional antennas. These antennas offer excellent performance but require complex, interdependent steering decisions in real time. Our Transformer-based architecture is enhanced with efficient components that add useful inductive biases, making it suitable for environments where scaling is limited. A key contribution is the introduction of directional density encodings, which help the attention mechanism better separate nodes in dense clusters. Equipped with our modules, a single Transformer block of dimension 12 achieves over 95 % accuracy, reducing the gap to optimality by half compared to a plain 1-block Transformer while requiring only 12 % more computation. Using two blocks, the model comes close to perfect accuracy.

1 Introduction

1.1 Link Topology Generation for Mobile Ad Hoc Networks

We address the generation of link topologies in wireless mobile networks using directional antennas. We tackle real-time link topology generation in MANETs with directional antennas, selecting links that form a bipartite backbone for antenna steering under sectorization, range, and interference constraints. The bipartiteness corresponds to the need to allocate emission and reception slots to each node, alternately. Fig. 1 illustrates our problem. While protocols like OLSR [1] and AODV [2] adapt well with omnidirectional antennas, they incur high interference; directional links boost throughput [3] but require complex, interdependent decisions. Existing approaches, UAV placement, topology control, adaptive beam/power tuning [4, 5, 6, 7, 8], or greedy link selection [9], are often sub-optimal or mobility-dependent. Combinatorial optimization produces quality static solutions [10, 11], but is too slow for mobility; instead, we learn to replicate these patterns from feasible topologies for fast and robust generation.

The challenge is exacerbated by the limited computational resources of mobile nodes, which must also perform signal decoding, neighbor position prediction, and topology updates. Thus, the topology predictor must be efficient to preserve resources for other network tasks.

Our approach is to replace a computationally expensive algorithm with an efficient neural network, capable of producing similar outputs in real time. We therefore enhance a Transformer, well-suited to point clouds [12, 13]—with a set of computationally efficient modules that provide useful structural inductive biases, partially avoiding the need to widen layers, which would increase computation quadratically. Embedded architectures, notably small planes or UAVs, which generally do not allow for massive parallel computation, could typically benefit from such modules.

While prior works on neural networks for wireless systems mainly address routing and performance prediction [14, 15, 16], we focus on efficient link topology generation.




Figure 1: Given a set of aerial nodes, a neural network is used to generate a link topology. We then assign the nodes a binary emission-reception scheduling.

1.2 Message Passing and Its Expressivity

In graph learning, rich node representations can be built via message passing [17, 18, 19], graph convolutions [20, 21, 22], or attention [23, 24, 25]. When edges are missing or must be predicted, GNNs require heuristic edges to enable message passing, such as the latent-space k -NN used in [26]. However, this scheme is equivalent to local-masked attention with a fixed query vector, limiting its ability to capture long-range dependencies and non-local patterns, an essential capability for edge prediction [12, 13].

While GNNs suffer from oversmoothing due to the message-passing bottleneck [27], attention-based models face similar issues [28]: as layers accumulate, node representations become indistinguishable. This is critical in small architectures with few, thin layers.

Handcrafted node features can help but are often unsuitable here as they typically rely on predefined edges. Laplacian positional encodings [29] require a predefined edge set, Graphormer [12] encodes hop-based distances, and Graph-BERT [13] orders features by shortest-path hops. Without edges, such encodings degrade to simple Euclidean distances, scalar proximity measures that many models already recover.

We address this by introducing **Directional Density Encodings (DDE)**: a node’s DDE quantifies the density of neighbors along each axis of its local reference frame. In 2D, this yields a four-dimensional geometric signature capturing direction-weighted density, providing a strong relative positional cue even in edge-free settings and improving multiple link prediction.

Furthermore, we incorporate cross-attentive modulation (CAM) tokens [30] to assist the message passing mechanism in adapting its behavior to the node layout, while also enabling a form of global control.

These additional modules and features allow TopoFormer to hold particularly well-suited inductive biases for the prediction of full link topologies for any node layout.

1.3 Towards more Expressive Link-Set Prediction

Pairwise link prediction generally assumes that link candidates’ embeddings can be aligned pairwise. In [31], link prediction is formulated as a graph algorithm problem: deciding the existence of an edge is associated with evaluating a function over node pairs, which a message-passing mechanism can be trained to approximate.

Although expressive and capable of good generalization, this paradigm assumes that a single decoding function, generally an inner product, can capture the many different relationships that underlie the existence of a link, which we assume could be improved using a mechanism that can model the different relations that can qualify a pair of nodes to form a link. The feature aggregator must then learn to precisely align nodes pairwise in the latent space to predict links, while having to capture complex patterns and global properties to ensure that these predicted links are correct. In the context of mobile networks, the heterogeneous nature of the topologies combined with the need for small and shallow neural architectures tend to challenge the capabilities of the feature aggregator.

As a simple way to enhance the capabilities of the prediction layer, we equip TopoFormer with a multichannel factorization method that can be interpreted as a flexible generalization of RESCAL [32] or Tucker [33] for one-dimensional link prediction. In our case, static relation-specific matrices are replaced by learned parameters that perform factorization over an arbitrary number of dimensions. It allows us to enhance the expressivity of the prediction layer while computationally scaling linearly with the number of channels.

2 Problem Description

Our work is motivated by a network topology problem that consists in finding an optimal set of links given a set of nodes described by their positions. It corresponds to a typical combinatorial problem that optimizes the link topology of a wireless network with directional antennas. The topologies feature a small (<32) number of nodes. It is defined by several constraints and properties:

- Each node may have at most four links.
- There are strict link sectorization constraints regarding the placement of antennas.
- Interference between links must be minimized, which implies avoiding acute angles between emitting nodes and their unintended receiving neighbor nodes.
- The link topology must be connected.

The problem cannot be trivially linearized and exhaustive methods are too heavy to be carried out in any reasonable amount of time, even to create a dataset. We therefore rely on a heuristic algorithm (Alg. 1) that follows:

Algorithm 1 Greedy Network Topology Algorithm

Require: V ▷ set of nodes
Require: L ▷ set of all possible links
Require: θ ▷ interference threshold
Ensure: $G = (V, E)$ ▷ connected graph

- 1: $E \leftarrow \emptyset$
- 2: **while** the graph (V, E) is not connected **do**
- 3: **while** $L \neq \emptyset$ **do**
- 4: Select the best link $l \in L$ according to a heuristic (e.g., link length or throughput)
- 5: **if** interference of l with E is below θ **then**
- 6: $E \leftarrow E \cup \{l\}$ ▷ add the link offering best throughput
- 7: **else**
- 8: $L \leftarrow L \setminus \{l\}$ ▷ abandon the link
- 9: **return** (V, E)

The topologies obtained achieve about 95 % of the throughput obtained using an exact and exhaustive algorithm, they allow for a **1.7** times higher theoretical throughput than if using a standard omnidirectional protocol topology for a 16-node-network, and a **3.7** factor of improvement for a 32-node-network.

The algorithm is $\mathcal{O}(m^2)$ complex if the length of the link is used as a heuristic and at least $\mathcal{O}(m^3)$ complex if the link throughput is used instead, m being the number of edges, even when the heuristics are used. It is not suitable for real-time use, but can be used to generate a dataset. Imitating such graphs with a neural network demands both global control capabilities in order to provide plausible large-scale connectivity patterns and the ability to model the different kinds of relationship between nodes that can lead to the presence of a link. Accurately reproducing the patterns and construction rules of such dataset graphs would enable finding high-performance complex topologies in real time.

3 Our Architecture

In this section, we detail the neural architecture that we adopt in order to imitate the results of our costly algorithm.

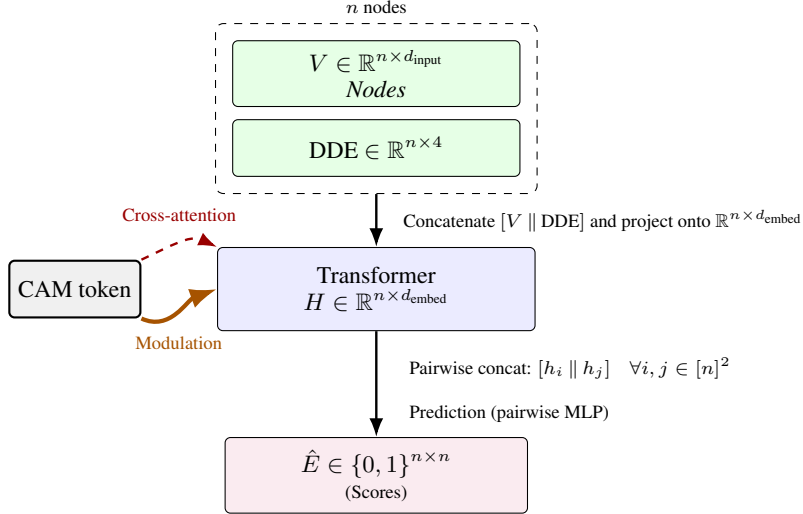


Figure 2: The architecture of TopoFormer.

3.1 Overview of Our Architecture

Inspired by the effectiveness of attention-based models to deal with point clouds [12, 13], we adopt the Transformer [23] as the backbone of our architecture. We implement it with no sequence-based encodings [23], in order to keep permutation invariance, nor spectral [29] positional encoding, since the edges are the object we seek to predict.

Given the limited width of the layers in TopoFormer, a trade-off arises between using a small number of attention heads or assigning a low dimensionality to each head. We opt for a larger number of lightweight (1- or 2-dimensional) attention heads. Our empirical results presented in Table 3 justify this choice.

The nodes are enriched with directional density encodings, which we detail in Sec. 3.2.

To enhance its adaptability to heterogeneous topologies, we augment it with CAM tokens [30], which provide dynamic modulation of the model’s behavior based on the overall node layout. CAM tokens are only $\mathcal{O}(TN)$ complex, T and N being respectively the number of CAM tokens and the number of nodes. They allow the model to efficiently adapt its behavior to the heterogeneous nature of network topologies.

Finally, TopoFormer is equipped with multichannel prediction layers, as proposed in Sec. 3.3

Fig. 2 shows the architecture of our framework.

3.2 Directional Density Encoding

We equip TopoFormer with directional density encodings in order for the nodes to be easily distinguishable by attention mechanism, especially when they are located within dense clusters. Fig. 3 shows an illustration of the computation of DDEs for an arbitrary node.

Let v be a node with a 2D position $\mathbf{v} = (x_v, y_v) \in \mathbb{R}^2$. We define a 4-dimensional attribute vector for v as follows:

$$\mathbf{d}(v) = \left[D_{\text{right}}(v), D_{\text{left}}(v), D_{\text{up}}(v), D_{\text{down}}(v) \right],$$

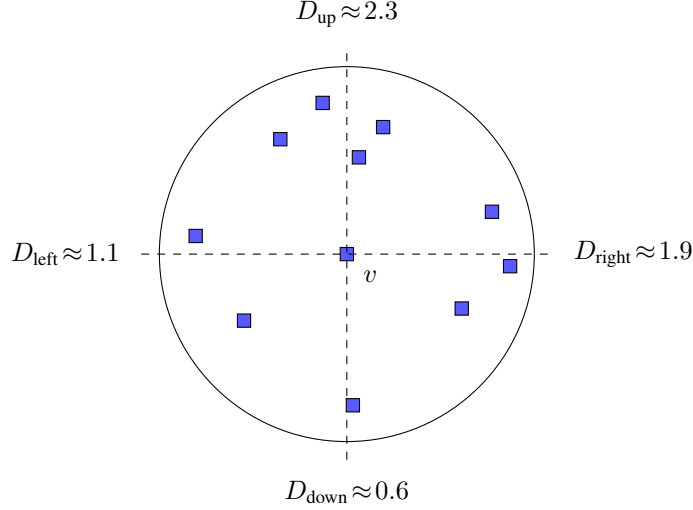


Figure 3: Illustration of the computation of DDEs for a node v .

where each component is given by:

$$D_{\text{right}}(v) = \sum_{\substack{u \in V \\ u \neq v}} \exp\left(-\frac{\|\mathbf{x}_u - \mathbf{x}_v\|^2}{2\sigma^2}\right) \max(0, \cos \theta_{uv}), \quad (1)$$

$$D_{\text{left}}(v) = \sum_{\substack{u \in V \\ u \neq v}} \exp\left(-\frac{\|\mathbf{x}_u - \mathbf{x}_v\|^2}{2\sigma^2}\right) \max(0, -\cos \theta_{uv}), \quad (2)$$

$$D_{\text{up}}(v) = \sum_{\substack{u \in V \\ u \neq v}} \exp\left(-\frac{\|\mathbf{x}_u - \mathbf{x}_v\|^2}{2\sigma^2}\right) \max(0, \sin \theta_{uv}), \quad (3)$$

$$D_{\text{down}}(v) = \sum_{\substack{u \in V \\ u \neq v}} \exp\left(-\frac{\|\mathbf{x}_u - \mathbf{x}_v\|^2}{2\sigma^2}\right) \max(0, -\sin \theta_{uv}), \quad (4)$$

with

$$\theta_{uv} = \arctan 2\left((\mathbf{x}_u - \mathbf{x}_v)_y, (\mathbf{x}_u - \mathbf{x}_v)_x\right),$$

One can also derive a learnable and continuous generalization of DDEs to an arbitrary number of dimensions by introducing H learnable directional vectors, $\{\mathbf{q}_h\}_{h=1}^H \subset \mathbb{R}^d$, which act as heads similar to attention mechanisms. For a given node v with position $\mathbf{x}_v \in \mathbb{R}^d$, the h -th directional density component is computed as:

$$D_h(v) = \sum_{\substack{u \in V \\ u \neq v}} \phi(\|\mathbf{x}_u - \mathbf{x}_v\|) \cdot \mathbf{q}_h^\top (\mathbf{x}_u - \mathbf{x}_v) \quad (5)$$

where:

- $\mathbf{q}_h \in \mathbb{R}^d$ is a learnable direction vector (the h -th head),
- $\mathbf{x}_v, \mathbf{x}_u$ are the latent positions of nodes v and u respectively,
- $\phi: \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is a distance-based weighting function, such as $\phi(r) = \exp(-r^2/2\sigma^2)$.

The resulting vector $\mathbf{d}(v) = [D_1(v), \dots, D_H(v)] \in \mathbb{R}^H$ encodes the anisotropic density of neighboring nodes around v along each learned direction. Unlike standard attention mechanisms, this




Figure 4: Learnable channel values are multiplied with the node-pair embeddings. The result is summed over the channels and scaled before being passed to a sigmoid. " \cdot " denotes either a Hadamard product or a matrix multiplication, depending on the model version, between each channel with each node pair.

formulation does not use a softmax normalization, allowing each head to accumulate density information instead of normalizing it, thereby preserving both directionality and local concentration, allowing them to “count” easily.

Throughout this paper, we prefer the non-learnable version of DDEs, as it is much more computationally efficient.

In practice, because the Gaussian weight $\phi(r) = \exp(-r^2/2\sigma^2)$ decays rapidly, only nodes lying within a fixed radius $r \approx 3\sigma$ contribute non-negligibly to the directional density encoding (DDE). To avoid inspecting all $\sim N^2$ pairs, we first insert the N node positions into a kd -tree [34], a binary space-partitioning structure that recursively splits the point set along coordinate axes, yielding a balanced search tree of depth $\mathcal{O}(\log N)$.

A radius query on this tree visits only the buckets intersecting the hypersphere of radius r around a query node and therefore returns, on average, a constant number k of nearby neighbours, independent of N for bounded density. With k treated as a small constant, the overall complexity is $\mathcal{O}(N \log N)$, and the aggregation of the four DDE components adds only an $\mathcal{O}(Nk)$ linear pass over the reported neighbours. Hence the truncated scheme retains nearly all the relevant mass while reducing the theoretical cost from quadratic to near-linear.

3.3 Multichannel scoring mechanisms.

To model the diversity of structural factors that can lead to link formation, ranging from local motifs to more global topological cues, we explore two bilinear scoring strategies inspired by relational representation learning. Fig. 4 illustrates the principle of such multichannel prediction layers.

First, we introduce a diagonal bilinear formulation using learnable channels $\{c_i\}_{i=1}^C \subset \mathbb{R}^d$ that interact with the pair embedding $p \in \mathbb{R}^d$ via element-wise (Hadamard) multiplication. Each channel produces a score through a shared projection, the score of an edge prediction \hat{e} corresponding to the node pair p is hence given by:

$$\hat{e} = \sum_{i=1}^C w^\top \text{LN}(c_i \odot p),$$

where w is a common weight vector and LN denotes a LayerNorm applied prior to aggregation. This mechanism corresponds to a multichannel, low-rank-factorization-like of a diagonal bilinear form:

$$\hat{e} = \sum_{i=1}^C \sum_{k=1}^d w_k c_{i,k} p_k.$$

While this Hadamard-channels formulation is not strictly more expressive than the neurons of a simple linear layer, it offers several practical advantages. In particular, it enforces a channel-wise separation of parameters that improves optimization, keeps channel-pair Hadamard products in structured form longer before summation, which stabilizes training and preserves richer interactions.

Following the same general idea, we also implement a dense bilinear variant faithful to a RESCAL [32] formulation, where K relation-specific matrices $R^{(k)} \in \mathbb{R}^{d \times d}$ interact with the pair embedding

following the equation:

$$\hat{e} = \sum_{k=1}^K p^\top R^{(k)} p.$$

This dense formulation enables the model to capture complex cross-dimensional interactions and more nuanced structural dependencies.

However, throughout the paper, we adopt the diagonal multichannel formulation as our default, since it offers nearly equivalent performance while being significantly more efficient and stable to train. These learned channels could also be parametrized by a feature aggregator to dynamically adapt to different or more complex network settings.

4 Results

We conduct our experiments with an Intel Xeon(R) E5-2650v3 at 2.30 GHz CPU and a Tesla T4 GPU. The models are trained using standard binary cross-entropy. The models are equipped with one Transformer block, unless stated otherwise. The prediction layer consists in the concatenation of node pairs fed into a small MLP. In addition to attention-based baselines, we benchmark a graph neural networks in which virtual edges connect each node to the neighbors in its communication range. The convolutional variant can also be seen as a 1-layer DCGNN [26]. The embedding size of the models is $d = 12$, which corresponds to the lowest number of parameters allowing us to achieve a satisfactory accuracy. CAM-enhanced ablations use a single CAM token.

The models are trained using adaptive learning rates using AdamW [35] optimizer for about 5000 epochs depending on the observed convergence. They are then trained using decreasingly high learning rates (from 10^{-6} to 10^{-3}) until validation performance does not improve anymore. The model achieving the highest accuracy is marked in bold.

The model described as TopoFormer is a Transformer equipped with CAM tokens, DDEs and a multichannel prediction layer. The uncertainty is computed as the standard deviation over bootstrapped measurements. The 2-block TopoFormer also follows a second disjoint CAM block iteration, in addition to the supplementary Transformer block.

4.1 Accuracy

Table 1: Accuracy of different ablations and models on our test set. The FLOPs column indicates the relative percentage of computational cost compared to that of the standard 1-block Transformer.

Method	Accuracy (%)	FLOPs
Graph Convolutional Network [36]	87.4±.4	-14%
Graph Attention Network [24]	88.1±.4	-8%
Graph Transformer [25]	89.9±.3	+8%
Transformer	90.9±.2	Reference
Transformer w/ CAM	92.9±.2	+7%
Transformer w/ DDE	93.7±.2	+3%
TopoFormer w/ Multichannel Hadamard	95.1±.2	+12%
TopoFormer w/ Multichannel Factorization	95.2±.2	+14%
2-block Transformer	99.0±.1	+89%
TopoFormer w/ 2 blocks	99.9±.1	+100%

Table 1 shows that TopoFormer outperforms the baseline Transformer, and significantly outclasses a range of GNN models. The results underline that DDEs have a high impact, and that the rest of the architecture enhancements also bring about an improvement for a low computational overhead. The bilinear multichannel prediction offers the best accuracy but seems less computation-efficient than the Hadamard-based version. Attention-based models perform better than locality-based graph models. Learnable DDEs were less computationally efficient than their handcrafted counterpart, their results are not showcased for clarity. Please note that, in a Transformer model, increasing the embedding dimension from 12 to 14, the next non-prime number (required for multi-head attention), would raise the number of FLOPs required by roughly 27%.

Table 2: Comparison of the results of our Hadamard multichannel prediction layer with the pairwise linear prediction, both without any backbone network. *Variance* is used as a proxy for the ability to clearly separate 0s and 1s.

Prediction Layer	Accuracy (%)	Variance
Pairwise Linear Layer	79.42±.5	0.038
Large Pairwise Linear Layer	82.89 ±.5	0.057
Multichannel Hadamard	86.91±.5	0.086

Table 3: Exploration of the head number/head size trade-off for our model.

Ablation	Accuracy (%)
One 12D head	94.07±.3
Two 6D heads	94.46±.3
Three 4D heads	94.81±.3
Six 2D heads	95.16±.2




Figure 5: A real example of a neural-network-generated 16-node topology.

Fig. 5 shows an example of a generated topology using our architecture.

In Table 2, we isolate the effect of the prediction layers by measuring the accuracy on our prediction task without using any complex backbone network, with the nodes simply being fed to a small node-wise MLP. The *Large Pairwise Linear Layer* is made k times larger than the *Pairwise Linear Layer*, k being the number of channels in the *Multichannel Hadamard* in order to verify that the sheer number of parameters is not enough to explain the performance increase. The Hadamard-based multichannel layer appears to possess some intrinsic expressivity that allows it to discriminate links better than simple linear pairwise matching.

Table 3 shows that equipping the backbone Transformer with numerous but small attention heads yields better accuracy than fewer but larger attention heads.

4.2 Throughput and Application

In order to be operational, the generated topologies should be post-processed in two ways. First, a sanity check is necessary to ensure that the topology respects the physical constraints that qualify admissible solutions. Second, we build a simple, straightforward communication scheduling assignment. In our case, since the links are designed to avoid interference, using directional antennas, we can aim for a simple emission-reception two-slot scheduling, where half of the nodes emits while the other half receives, alternately.

The pruning part of the algorithm is $\mathcal{O}(S * m)$ complex, with S being the average number of supernumerary links and m the number of links. The parity assignment part is $\mathcal{O}(N + m)$ complex. Usually $S \ll m$ and $S < n$, which allows this algorithm to run in less than 300 milliseconds.

Algorithm 2 Building a Bipartite Communication Network with TX/RX Parity

Require: $G_{\text{pred}} = (V, E_{\text{pred}})$ ▷ predicted links
Ensure: $G = (V, E)$ ▷ pruned, connected graph
Ensure: $p : V \rightarrow \{\text{TX}, \text{RX}\}$ ▷ parity labels

- 1: $E \leftarrow E_{\text{pred}}$
- 2: **for all** $v \in V$ **do**
- 3: Partition v 's neighbors into $\{\text{NE}, \text{NW}, \text{SE}, \text{SW}\}$
- 4: **for all** each sector q **do**
- 5: **while** $|\{(v, u) \in E : u \in q\}| > 1$ **do**
- 6: pick candidate (v, u) in q
- 7: **if** removing (v, u) keeps $(V, E \setminus \{(v, u)\})$ connected **then**
- 8: $E \leftarrow E \setminus \{(v, u)\}$
- 9: **else**
- 10: try another candidate ▷ — 2-coloring + removal of monochromatic edges —
- 11: Initialize $\forall v : p(v) \leftarrow \text{UNCOLORED}$
- 12: $toRemove \leftarrow \emptyset$
- 13: **for all** each component of (V, E) **do**
- 14: pick an uncolored root r , set $p(r) \leftarrow \text{TX}$
- 15: push r onto stack S
- 16: **while** $S \neq \emptyset$ **do**
- 17: $u \leftarrow \text{pop}(S)$
- 18: **for all** $(u, v) \in E$ **do**
- 19: **if** $p(v) = \text{UNCOLORED}$ **then**
- 20: $p(v) \leftarrow \text{opposite}(p(u))$
- 21: push v onto S
- 22: **else if** $p(v) = p(u)$ **then**
- 23: $toRemove \leftarrow toRemove \cup \{(u, v)\}$
- 24: $E \leftarrow E \setminus toRemove$ ▷ remove ineligible links
- 25: **return** (V, E, p)

Table 4: Avg instantaneous throughput upper bound (Mbps)

	OLSR-like Omni.	Transformer	TopoFormer	Target
Throughput ↑	47.24	69.36	74.77	79.33

In Table 4, we report theoretical throughputs as upper bounds on the simultaneous information exchange in the network. Topologies are obtained through the post-processing algorithm 2, which derives a parity-based communication scheme, and throughput is computed following geometrical arguments from [3, 37], mainly reflecting link density and interference. The omnidirectional baseline uses OLSR-like topologies with a communication range of about one quarter of the network diameter. Our approach halves the gap between Transformer and target throughput, and significantly outperforms the omnidirectional case by leveraging directional antennas. Since the target algorithm achieves 95 % of optimal throughput, TopoFormer reaches ~ 90 % with one block (and nearly yields the same throughput as the target's with two blocks).

5 Conclusion

We presented TopoFormer, a Transformer-based architecture to infer link sets from node sets, mimicking a complex network algorithm for steering directional antennas in mobile ad hoc networks. Efficient modules enhance expressivity without enlarging the model, which only requires minimal post-processing allowing for a fast scheduling enabling near-target performance.

References

- [1] P. Jacquet, P. Muhlethaler, T. Clausen, A. Laouiti, A. Qayyum, and L. Viennot. Optimized link state routing protocol for ad hoc networks. In *Proceedings. IEEE International Multi Topic Conference, 2001. IEEE INMIC 2001. Technology for the 21st Century.*, pages 62–68, 2001. 1
- [2] C.E. Perkins and E.M. Royer. Ad-hoc on-demand distance vector routing. In *Proceedings WMCSA '99. Second IEEE Workshop on Mobile Computing Systems and Applications*, pages 90–100, 1999. 1
- [3] Su Yi, Yong Pei, and Shivkumar Kalyanaraman. On the capacity improvement of ad hoc wireless networks using directional antennas. In *Proceedings of the 4th ACM international symposium on Mobile ad hoc networking & computing, MobiHoc '03*, pages 108–116, New York, NY, USA, June 2003. Association for Computing Machinery. 1, 9
- [4] Antonio Guillen-Perez and Maria-Dolores Cano. Flying ad hoc networks: A new domain for network communications. *Sensors*, 18(10):3571, oct 2018. 1
- [5] Zhenming Wang, Qin Zhang, and Hai Li. UAV mobile control strategy based on potential field method. In Tao Lei and Dehai Zhang, editors, *Sixth International Conference on Information Science, Electrical, and Automation Engineering (ISEAE 2024)*, volume 13275, page 132753L. International Society for Optics and Photonics, SPIE, 2024. 1
- [6] Zhuochuan Huang and Chien-Chung Shen. A comparison study of omnidirectional and directional mac protocols for ad hoc networks. volume 1, pages 57–61 vol.1, 12 2002. 1
- [7] Zhuochuan Huang, Chien-Chung Shen, Chavalit Srisathapornphat, and Chaiporn Jaikaeo. Topology control for ad hoc networks with directional antennas. 12 2002. 1
- [8] Qing Li and Imrich Chlamtac. A framework for topology control protocol using directional antennas in wireless networks. In *Proceedings of the IEEE International Conference on Communications*, volume 5, pages 2958–2962. IEEE, 2005. 1
- [9] Lichun Bao and J.J. Garcia-Luna-Aceves. Transmission scheduling in ad hoc networks with directional antennas. In *Proceedings of the 8th Annual International Conference on Mobile Computing and Networking, MobiCom '02*, page 48–58, New York, NY, USA, 2002. Association for Computing Machinery. 1
- [10] Wei Feng, Yong Li, Depeng Jin, Li Su, and Sheng Chen. Millimetre-Wave Backhaul for 5G Networks: Challenges and Solutions. *Sensors*, 16(6):892, June 2016. Number: 6 Publisher: Multidisciplinary Digital Publishing Institute. 1
- [11] Amal Benhamiche, Wesley da Silva Coelho, and Nancy Perrot. Routing and Resource Assignment Problems in Future 5G Radio Access Networks. June 2019. 1
- [12] Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu. Do transformers really perform bad for graph representation?, 2021. 1, 2, 4
- [13] Jiawei Zhang, Haopeng Zhang, Congying Xia, and Li Sun. Graph-bert: Only attention is needed for learning graph representations, 2020. 1, 2, 4
- [14] Krzysztof Rusek, José Suárez-Varela, Paul Almasan, Pere Barlet-Ros, and Albert Cabellos-Aparicio. RouteNet: Leveraging Graph Neural Networks for Network Modeling and Optimization in SDN. *IEEE Journal on Selected Areas in Communications*, 38(10):2260–2270, October 2020. Conference Name: IEEE Journal on Selected Areas in Communications. 1
- [15] Abdelhadi Azzouni, Raouf Boutaba, and Guy Pujolle. Neuroute: Predictive dynamic routing for software-defined networks, 2017. 1
- [16] Miquel Ferriol-Galmés, Jordi Paillisse, José Suárez-Varela, Krzysztof Rusek, Shihan Xiao, Xiang Shi, Xiang Cheng, Pere Barlet-Ros, and Albert Cabellos-Aparicio. Routenet-fermi: Network modeling with graph neural networks. *IEEE/ACM Transactions on Networking*, 31(6):3080–3095, December 2023. 1
- [17] Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural message passing for quantum chemistry, 2017. 2
- [18] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks?, 2019. 2
- [19] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated graph sequence neural networks, 2017. 2

- [20] Thomas N. Kipf and Max Welling. Variational graph auto-encoders, 2016. 2
- [21] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering, 2017. 2
- [22] William L. Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs, 2018. 2
- [23] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention Is All You Need, December 2017. arXiv:1706.03762 [cs]. 2, 4
- [24] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks, 2018. 2, 7
- [25] Vijay Prakash Dwivedi and Xavier Bresson. A generalization of transformer networks to graphs, 2021. 2, 7
- [26] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, and Justin M. Solomon. Dynamic graph cnn for learning on point clouds, 2019. 2, 7
- [27] T. Konstantin Rusch, Michael M. Bronstein, and Siddhartha Mishra. A survey on oversmoothing in graph neural networks, 2023. 2
- [28] Xinyi Wu, Amir Ajorlou, Zihui Wu, and Ali Jadbabaie. Demystifying oversmoothing in attention-based graph neural networks, 2024. 2
- [29] Vijay Prakash Dwivedi, Chaitanya K. Joshi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. Benchmarking graph neural networks, 2022. 2, 4
- [30] Félix Marcoccia, Cédric Adjih, and Paul Mühlethaler. Netdiff: Deep graph denoising diffusion for ad hoc network topology generation, 2024. 2, 4
- [31] Andrew Dudzik and Petar Veličković. Graph neural networks are dynamic programmers, 2022. 2
- [32] Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. A three-way model for collective learning on multi-relational data. In *Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML'11*, page 809–816, Madison, WI, USA, 2011. Omnipress. 3, 6
- [33] Ivana Balazevic, Carl Allen, and Timothy Hospedales. Tucker: Tensor factorization for knowledge graph completion. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, 2019. 3
- [34] Jon Louis Bentley. Multidimensional binary search trees used for associative searching. In *Communications of the ACM*, volume 18, pages 509–517, 1975. 6
- [35] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. 7
- [36] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks, 2017. 7
- [37] P. Gupta and P.R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, March 2000. 9