Structural Vulnerability of Greedy Routing in Hyperbolic Embedded Networks

Dániel Ficzere[‡&], András Majdán[‡], Gergely Hollósi[‡], Pál Varga[‡], József Bíró[‡]

[‡]*Dept. of Telecommunications and Artificial Intelligence*, *Faculty of Electrical Engineering and Informatics*
Budapest University of Technology and Economics, Műegyetem rkp. 3., H-1111 Budapest, Hungary
[&]*IIoT and AI Division, AITIA International Inc.*, 48-50 Czetz J. Str. H-1039 Budapest, Hungary
corr.: ficzere@tmit.bme.hu

***Abstract*—Greedy geometric routing enables scalable, stateless communication in IoT and edge networks by relying solely on local coordinate information. However, this efficiency is underpinned by a structural backbone—the Hyperbolic Minimum Spanning Tree (HMST)—which introduces critical implications for performance and resilience. We analyze the HMST as a locally inferable routing scaffold emerging from hyperbolic embeddings, using Edge Load Centrality (ELC) to assess its importance. We show that HMST edges carry a disproportionate routing load and that removing them causes up to 50% failure in greedy routing—far exceeding the impact of random removals. While this reveals a vulnerability, the HMST's local reconstructability also enables autonomous recovery of the network. We discuss the implications for resilient, self-organizing network management and coordinate-aware security.**

***Index Terms*—Geometric routing, Distributed routing, IoT, Network resilience, Self-organizing networks, Energy-efficient routing, Decentralized systems, Topology-aware management**

I. Introduction

Greedy geometric routing is a scalable alternative to traditional table-based protocols, enabling efficient communication in distributed networks such as the Internet of Things (IoT), edge computing systems, and decentralized architectures. By using only local coordinate information to forward packets, it eliminates the need for global topology knowledge or routing tables, reducing overhead and complexity [1]. However, this apparent simplicity conceals nontrivial structural dependencies that affect both performance and resilience.

Hyperbolic embeddings have emerged as a compelling framework for supporting greedy routing. By assigning each node polar coordinates in a negatively curved space, hyperbolic geometry captures both hierarchical (radial) and similarity-based (angular) relationships in real-world networks [2], [3]. This enables highly navigable, low-dimensional representations that support stateless routing [4], [5]. These embeddings are increasingly used beyond packet networks, including data center architectures, knowledge graphs, and large-scale semantic systems [4], [6].

Despite these advantages, the latent structures induced by hyperbolic embeddings remain underexplored. In particular, we focus on the Hyperbolic Minimum Spanning Tree (HMST)—a locally identifiable subgraph formed by linking each node to its nearest neighbor with a smaller radial coordinate. While implicit in greedy navigation, this structure plays a critical role in routing performance. Unlike classical spanning trees, the HMST is emergent from coordinate proximity and requires no global state for its construction [7].

In this paper, we investigate the HMST as a hidden backbone that supports efficient greedy routing but also introduces systemic vulnerability. Using Edge Load Centrality (ELC), we show that HMST edges carry a disproportionate share of routing traffic. When HMST edges are removed, greedy success rates degrade by up to 50%, compared to minimal degradation under random link removals. These findings underscore the dual nature of the HMST: it is both a key enabler of stateless efficiency and a high-value attack surface.

Finally, we discuss the implications for network and service management. While the local identifiability of the HMST creates vulnerabilities, it also enables autonomous reconstruction and self-healing. These insights suggest that geometry-aware management strategies must account not only for performance but also for the structural exposures induced by coordinate-based routing frameworks.

II. Background and Related Works

The increasing scale, heterogeneity, and dynamism of modern communication networks—especially in IoT, edge computing, and autonomous systems—pose critical challenges to traditional routing paradigms. Table-based routing protocols often struggle to cope with scalability, update overhead, and convergence delays in such distributed, state-constrained environments. As a result, there is growing interest in stateless routing mechanisms, particularly those based on geometric embeddings [7]–[9]. A prominent example is greedy geometric routing, where each node forwards packets to its neighbor that is closest to the destination according to a geometric distance function. This strategy is highly scalable and requires only local information, but its success hinges on how well the underlying network topology can be embedded in a metric space that supports such local decisions [7], [9].

Hyperbolic geometry has emerged as a powerful tool for embedding real-world complex networks. Due to its exponential expansion, hyperbolic space naturally models hierarchical and tree-like structures found in communication, social, and biological networks. In contrast to Euclidean embeddings, which often require high dimensions to preserve routing

978-3-903176-75-1

properties, hyperbolic embeddings enable low-dimensional (typically 2D) representations that support near-optimal greedy routing. In a hyperbolic embedding, each node is assigned polar coordinates: a radial coordinate reflecting node centrality or hierarchy, and an angular coordinate capturing similarity. The hyperbolic distance function—significantly different from its Euclidean counterpart—is used to guide greedy routing decisions. As shown by Kleinberg [10], every connected graph can be greedily embedded in the hyperbolic plane, providing theoretical guarantees for routing success under appropriate coordinate assignments. Further foundational work [2] demonstrated that hyperbolic mapping can sustain Internet-like connectivity with minimal overhead. Stateless routing using hyperbolic coordinates avoids the need for routing tables and periodic updates, making it particularly suitable for dynamic and large-scale networks. The concept of geohyperbolic routing, where coordinates can be inferred from geographic and hierarchical information, supports efficient and scalable forwarding in real-world settings.

The utility of hyperbolic embeddings extends beyond packet routing. In data center architectures [5], they underpin scalable, latency-aware communication structures. In knowledge graph embeddings and large language model alignment, hyperbolic spaces capture latent hierarchies and semantic similarities with high fidelity [4]. The embedding also provides a geometric foundation for content-addressable storage and distributed hash tables, as shown in Kleinberg's routing scheme [10], where keys are hashed to hyperbolic coordinates and data is retrieved via greedy navigation. These embedding-based methods offer compelling advantages for network and service management. As discussed in the book chapter by Karyotis and Stai, hyperbolic models facilitate big data analytics, support the prediction of evolving network dynamics, and enable efficient topological optimization [11]. This positions hyperbolic embedding as not only a routing tool but a broader framework for dynamic, autonomic network management.

Despite their strengths, most prior studies focus on routing performance rather than structural analysis of routing backbones. While greedy routing appears stateless, it implicitly depends on latent structures formed by coordinate proximity. Our work introduces the notion of the HMST as a central routing scaffold emergent from local coordinates. Unlike synthetic or globally constructed spanning trees, the HMST can be inferred locally, making it both a potential enabler of resilience and a vector for fully distributed attacks. We build on the geometric routing literature by shifting focus from success ratios to structural load centrality, revealing the network's implicit vulnerabilities and opportunities for self-organizing optimization.

Our earlier work proposed the integration of greedy navigational cores and hyperbolic trees into a unified architecture for low-complexity routing in communication networks [12]. While that approach emphasized memory efficiency and architectural scalability, the present study focuses specifically on the structural role and load distribution of HMST in greedy navigation. The concept of HMST, although introduced as part of broader design goals, had not been isolated and studied in the context of routing efficiency or vulnerability until now.

III. Methods

A. Network Models and Hyperbolic Embedding

We consider both synthetic and real-world hyperbolic complex networks, embedded in the native representation of the hyperbolic plane. Each node is assigned polar coordinates (r, θ), where r is the radial coordinate representing node centrality, and θ is the angular coordinate representing similarity. A reference point defines the origin for radial coordinates, and a reference direction defines the angular zero. The hyperbolic distance $d_H(u, v)$ between nodes u and v with coordinates (r_u, θ_u) and (r_v, θ_v) is given by:

$$\cosh(d_H(u, v)) = \cosh(r_u)\cosh(r_v) - \sinh(r_u)\sinh(r_v)\cos(\Delta\theta), \tag{1}$$

where $\Delta\theta = |\theta_u - \theta_v|$. Equation (III-A) is the hyperbolic law of cosines, where the negative curvature of the space leads to exponential expansion and the angular term enters through $\cos(\Delta\phi)$, capturing similarity between nodes.

The main motivation for using hyperbolic geometry is its ability to produce low-dimensional embeddings (typically 2D) that support scalable and efficient greedy routing, unlike Euclidean embeddings which often require higher dimensions to achieve comparable navigability.

We use two sets of synthetic networks for evaluation. In both cases, $N = 1000$ nodes are distributed in a hyperbolic disk of radius R. Radial coordinates r follow the probability density: $P(r) = \frac{\alpha \sinh(\alpha r)}{\cosh(\alpha R) - 1}$, where α is a tunable parameter controlling the density profile and the decay of the node degree distribution [13]. The angular coordinates θ are drawn uniformly from $[0, 2\pi)$.

In the first experiment, we fix $\alpha = 1$, which yields a uniform node density over the entire disk. We vary R in the set $10.0, 10.5, 11.0, 11.5, 12.0$ to generate networks with different average degrees $\bar{k}$. For each value of R, 20 network instances are generated, and results are averaged.

In the second setup, we maintain $\bar{k} \approx 11$ by jointly tuning α and R. Smaller α values lead to a denser, more inclusive core (more nodes with smaller r), while larger α values yield sparser cores. Again, 20 network instances are generated per setting for statistical reliability.

In addition to synthetic networks, we include the Autonomous System (AS) level topology of the Internet as a real-world example. The hyperbolic coordinates of its nodes are obtained using standard embedding techniques [2].

An illustrative synthetic network and its corresponding HMST are shown in Fig. 1, highlighting the structural relationship between the full network and its locally definable backbone. In Fig. 2 the highest ranked 1000 nodes can be seen from the embedded AS network topology.

B. HMST Definition, Routing Models, and Evaluation Metrics

We define the HMST as a substructure consisting of directed edges based on local geometric rules. For each node u, among

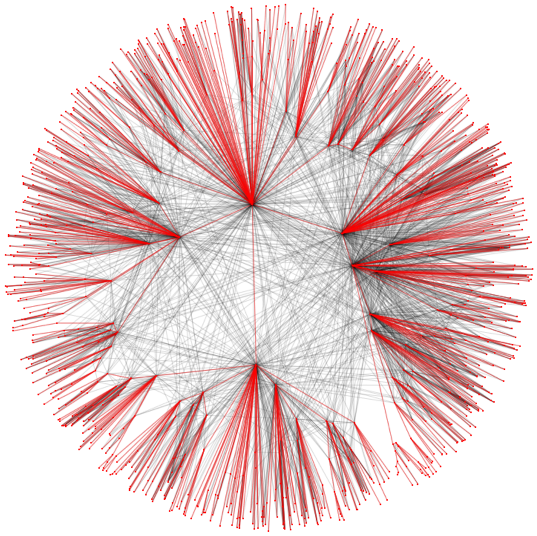

Fig. 1. An example synthetic network and its HMST links (red lines).

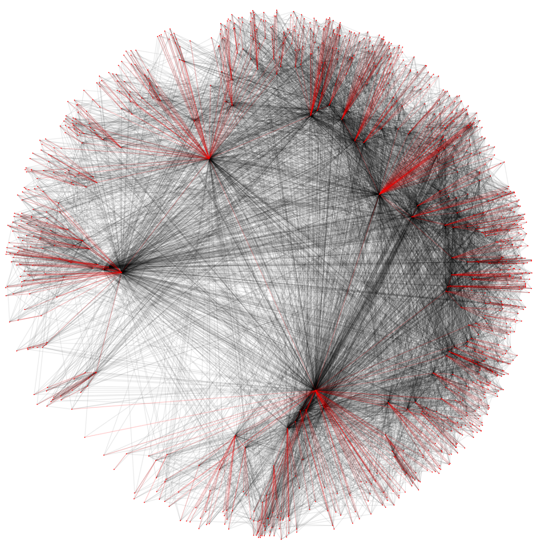

Fig. 2. Top 1000 nodes in the network and the corresponding HMST subgraph (red edges).

its neighbors $v \in \mathcal{N}(u)$ with $r_v < r_u$, we select the neighbor v^* that minimizes the hyperbolic distance:

$$v^* = \arg \min_{v \in \mathcal{N}(u), r_v < r_u} d_H(u, v). \tag{2}$$

The edge (u, v^*) is then included in the HMST. This procedure ensures that each node links to one more central neighbor (if available), forming a spanning tree rooted at the node with the smallest radial coordinate.

We evaluate two routing models. In greedy routing, packets are forwarded to the neighbor with the smallest hyperbolic distance to the destination. A route is successful if it reaches the target without revisiting any node. As a baseline, shortest path routing computes the minimum-hop path between source and destination using the network topology.

For both routing models, we compute ELC, defined as the number of distinct source-destination paths that traverse a given edge. We distinguish between ELC values for HMST edges and for non-HMST edges.

To assess the vulnerability and structural role of HMST, we perform edge removal experiments. First, we remove all HMST edges from the network and measure the resulting Greedy Success Ratio (GSR), defined as the fraction of successful greedy paths. Then we remove an equal number of randomly selected edges and compare the resulting GSR. The routing degradation is expressed as $1 - \text{GSR}$.

All reported values are averages over 20 independent realizations per parameter setting. This ensures statistical robustness and highlights consistent patterns in how the HMST structure impacts routing performance and edge centrality.

In experiments involving edge removals—whether targeting HMST edges or selecting them at random—it is possible for the network to fragment into disconnected components. However, we consistently observed that such fragmentation results in the formation of a dominant giant component accompanied by several much smaller components. The presence of this giant component is important for interpreting routing performance, particularly under greedy routing, where fragmentation inherently limits reachability between disconnected regions of the network.

To ensure fair and focused evaluation of routing performance, we isolate and analyze only the giant component in each experimental scenario. Specifically, we measure the degradation of greedy routing success solely within this largest connected subgraph. This design choice allows us to quantify the direct impact of edge removals on greedy routing efficiency, independently of the indirect impact caused by topological disconnection. As the smaller fragments are negligible in size and traffic potential compared to the giant component, we believe this approach offers a more accurate comparison between targeted (HMST) and random edge removal scenarios.

IV. RESULTS

We begin by presenting the results obtained from synthetic networks with varying average degrees, corresponding to different link densities, under greedy routing. Figure 3 displays box plots comparing the distributions of ELC values for HMST edges and non-HMST edges across five network densities. A clear and consistent pattern emerges: in all five cases, the ELC values for HMST edges are substantially higher than those for non-HMST edges. This observation strongly suggests that the HMST forms a central structural backbone for greedy routing, concentrating a significant portion of the routing load. Notably, the vertical scale in Figure 3 is logarithmic, further underscoring the magnitude of the differences between the two distributions. Another important trend is visible in how the ELC values for HMST edges evolve with network sparsity. As the average degree decreases, the ELC distributions for HMST edges shift upward, indicating that in sparser networks, the load concentrates even more heavily on HMST edges. While it is expected that average ELC values rise in sparser networks due to fewer available edges, the extent to which this increase is localized on HMST edges is non-trivial and highlights their critical role in the greedy routing infrastructure.

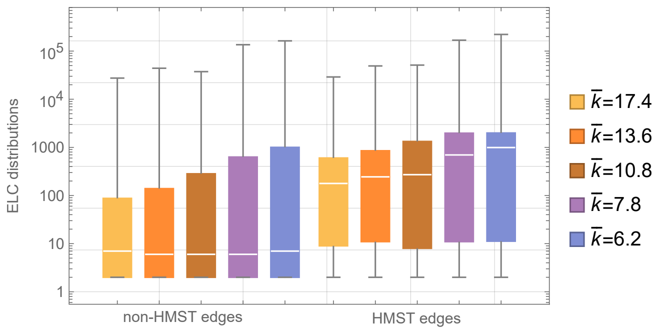

Fig. 3. Distribution of Edge Load Centrality values in case of greedy routing for HMST and non-HMST edges, compared across synthetic networks at five different density levels. Observe that in all five cases, the ELC values for HMST edges are substantially higher than those for non-HMST edges

We now turn to the results obtained from networks with varying levels of core connectivity, while maintaining a constant average degree of approximately $\bar{k} \approx 11$. The variation in core-connectivity is controlled by the parameter α: lower values of α result in a denser core with a larger number of core nodes, whereas higher values produce a sparser and smaller core. This parameter also affects the overall degree distribution of the network. In real-world embedded networks, values of α typically range between 0.5 and 1. Specifically, for the Internet AS-level topology, $\alpha \approx 0.7$ is typically observed. Figure 4 presents the ELC distributions for HMST and non-HMST edges across five different α values. The results are consistent with the earlier findings: HMST edges consistently dominate in terms of ELC values, regardless of the underlying core-connectivity level. This further reinforces the conclusion that HMST edges carry a disproportionate amount of the routing load under greedy navigation, even when the network's hierarchical structure varies.

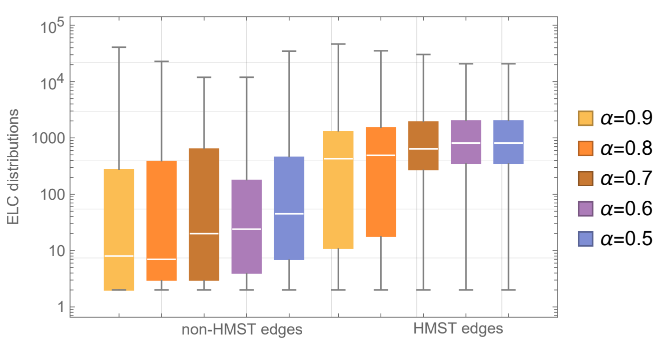

Fig. 4. Distribution of Edge Load Centrality values in case of greedy routing for HMST and non-HMST edges, compared across synthetic networks at five different core-connectivity levels. The results are consistent with the earlier findings: HMST edges consistently dominate in terms of ELC values, regardless of the underlying core-connectivity level.

We also examine the role of HMST edges in the real-world Internet AS-level topology. Specifically, we conduct experiments on two subsets of the network: the top 1000 and top 2000 providers, identified as the nodes with the smallest radial coordinates in the hyperbolic embedding. Figure 5 presents the ELC distributions for HMST and non-HMST edges within these subsets. The results are consistent with previous observations: HMST edges again dominate in terms of ELC values, further confirming their disproportionate structural and functional importance in supporting greedy routing.

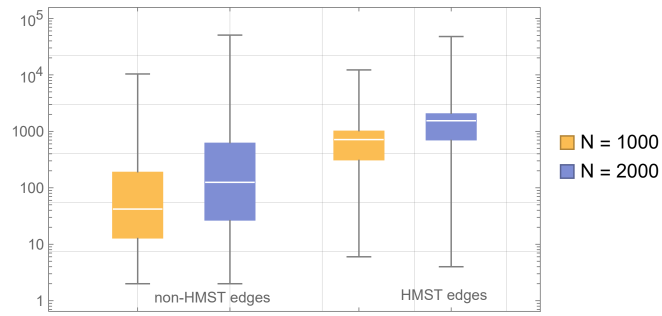

Fig. 5. Distribution of Edge Load Centrality values under greedy routing for the top 1000 and top 2000 ranked nodes (by radial coordinate) in the Internet AS-level topology. The results are consistent with previous observations: HMST edges again dominate in terms of ELC values.

We now continue the analysis by examining the degradation of the greedy routing success ratio when networks are pruned either by removing HMST edges or an equal number of randomly selected edges. Figure 6 shows the results across five network sets with varying average degrees. Across all cases, HMST-based pruning leads to significantly higher degradation in routing success compared to random edge removal. As expected, for both types of removal, the degradation lessens as the network becomes denser. This is largely because average ELC values tend to decrease with increasing average degree, distributing the routing load more evenly across a greater number of edges. Importantly, while random edge removal causes negligible impact in dense networks (with degradation falling below 1%), HMST-based pruning still results in a pronounced decline in routing success—exceeding 10% in some cases. These findings underscore the centrality of HMST edges in supporting effective greedy routing and highlight the serious performance consequences that arise from their absence. Similar patterns are observed when varying the core-connectivity level through changes in the parameter α. As shown in Figure 7, the degradation of greedy success ratio tends to decrease with increasing α when HMST edges are removed. However, this degradation remains significantly higher than that observed under random edge removals. In the latter case, the impact on routing success is limited to just a few percent, whereas HMST-based removals continue to result in substantial losses, reinforcing the structural importance of HMST edges across varying hierarchical profiles.

We also conducted similar experiments using shortest path routing to compute ELC values, where each edge's centrality reflects the number of shortest paths it supports. Fig. 8 and

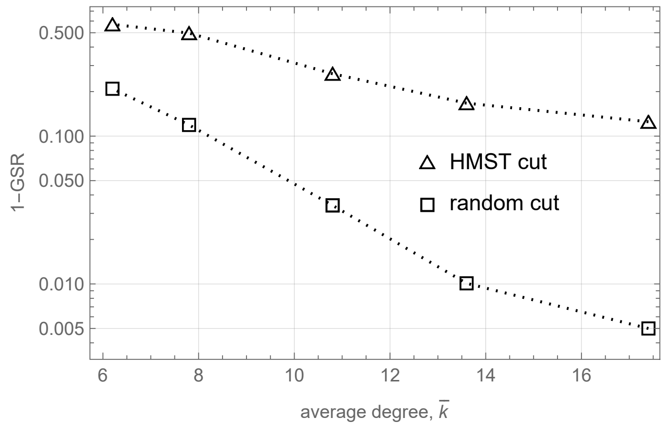

Fig. 6. Degradation of Greedy Routing Success Ratio (1 – GSR) as a function of average node degree in synthetic networks, comparing HMST-based and random edge removals.

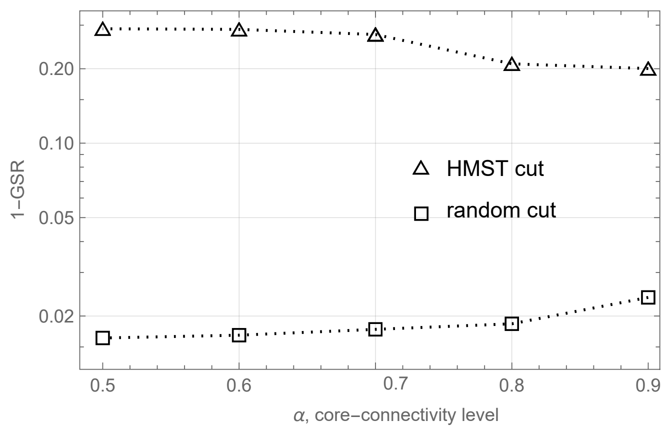

Fig. 7. Greedy Routing Success Ratio degradation (1 – GSR) versus core-connectivity level under HMST-targeted and random edge removal scenarios.

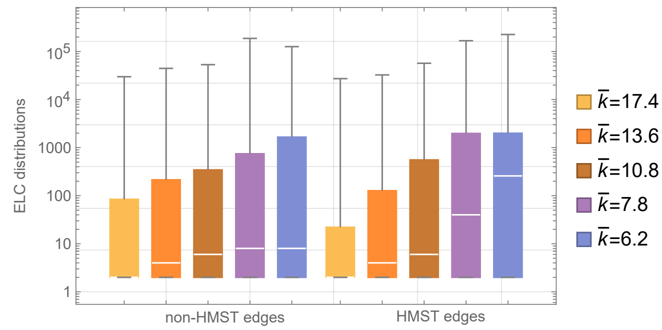

Fig. 8. Comparison of Edge Load Centrality distributions for HMST and non-HMST edges under shortest path routing across five synthetic network densities. This indicates that HMST edges are not structurally central in the shortest path routing framework.

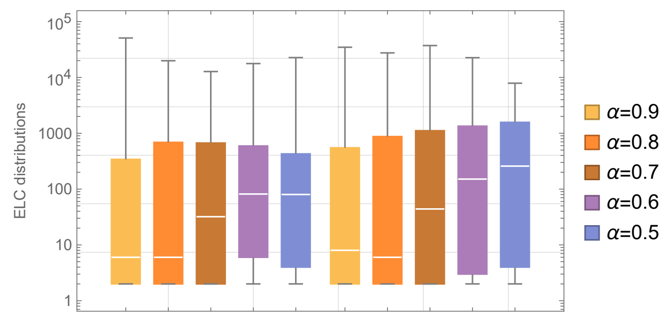

Fig. 9. Comparison of Edge Load Centrality distributions for HMST and non-HMST edges under shortest path routing, across synthetic networks at varying core-connectivity levels. This indicates that HMST edges are not structurally central in the shortest path routing framework.

Fig. 9 show the resulting ELC distributions for HMST and non-HMST edges under the same synthetic network configurations. Unlike the results observed for greedy routing, the difference in ELC values between HMST and non-HMST edges is much less pronounced. This indicates that HMST edges are not structurally central in the shortest path routing framework. These findings confirm that HMST plays a distinctly influential role in greedy routing, while its impact is significantly reduced in shortest path-based routing strategies.

V. Discussion (Management Implications)

The experimental findings presented in this study demonstrate that the Hyperbolic Minimum Spanning Tree plays a central role in enabling efficient greedy routing. Across a range of synthetic and real-world networks, HMST edges consistently exhibit substantially higher ELC values compared to non-HMST edges. This pattern highlights the structural centrality of HMST in supporting the majority of routing paths.

Ongoing theoretical investigations by the authors aim to further explain this phenomenon. Preliminary observations suggest that for an HMST edge $u \to v$ where $r_u > r_v$, node v is often closer in hyperbolic space to a large number of potential destination nodes. This implies that HMST edges frequently serve as effective next-hop choices in greedy routing. Compared to other neighbors, the endpoint of an HMST edge tends to offer broader reachability across the network, enhancing the probability of successful routing. These findings suggest that the HMST is not just a navigational aid but a geometrically grounded backbone that maximizes local forwarding efficiency. While the locality and determinism of HMST edge identification are operational advantages, they also introduce a significant vulnerability. HMST edges can be disclosed purely through local coordinate analysis, without requiring global topology knowledge. This makes the HMST a prominent attack surface, especially in distributed and adversarial environments.

A fully distributed, self-organizing attack model can be envisioned in which adversaries embedded in network nodes independently identify HMST edges using only their own coordinates and those of their neighbors. These agents can then target and disrupt key links in a stealthy, asynchronous fashion. Because each adversary operates independently, such attacks are difficult to detect and may degrade the network gradually, ultimately leading to severe routing failures. Moreover, the visibility of hyperbolic coordinates themselves becomes a risk factor. Since these coordinates enable the identification of HMST edges, they must be considered sensitive information and protected accordingly.

From a network management perspective, protecting HMST edges and the coordinates that reveal them is critical. Several strategies can be considered: limiting access to raw hyperbolic coordinates through access control; applying lightweight randomization or coordinate rotation to make HMST edge prediction more difficult; introducing redundancy or priority monitoring for high-ELC edges, especially those aligned with the HMST; and continuously monitoring routing performance and edge-level traffic for asymmetric degradation patterns that may signal targeted attacks. These strategies align with the goals of secure and dependable networking and promote a more resilient routing infrastructure.

Beyond its critical role in normal operations, the HMST offers an inherent advantage for recovery and self-healing. In scenarios involving large-scale disruptions, catastrophic failures, or network fragmentation, the HMST can be reconstructed in a fully distributed manner. Since each node requires only local coordinate information to identify its HMST edge, the recovery process can be executed without centralized orchestration. This property makes HMST not only a potential vulnerability but also a powerful instrument for restoring connectivity [14]. It supports the design of autonomous, coordinate-driven recovery mechanisms capable of rapidly restoring connectivity and reestablishing a navigable backbone in degraded networks.

The HMST emerges from this study as a dual-use structure. It is a high-value subgraph essential for efficient greedy routing, a critical attack surface due to its local discoverability, and a robust tool for self-organizing recovery. These insights have substantial implications for future network and service management frameworks, particularly in decentralized, coordinate-driven environments. They call for a holistic approach that recognizes the operational, security, and resilience roles of geometric structures like the HMST.

VI. Conclusion

Geometric embeddings, and in particular hyperbolic embeddings, are powerful tools for the design and management of scalable, decentralized networks. Their primary advantage lies in enabling stateless routing, where packet forwarding decisions rely solely on node coordinates rather than global routing tables. Greedy routing remains highly relevant for today's IoT and edge systems, where stateless forwarding and lightweight local decisions align with resource-constrained and decentralized environments. Hyperbolic space is especially suited for this task, as it naturally captures the hierarchical and heterogeneous structure of real-world networks. The embedding process leverages global topological information and encodes it locally in each node's coordinates, which not only support routing but also reveal deeper structural insights. One such insight is the Hyperbolic Minimum Spanning Tree (HMST), a latent scaffold that can be identified purely through local geometric rules.

This duality—where coordinates empower both functionality and exposure—has important implications for network and service management. On the one hand, the local identifiability of structures like the HMST opens pathways for distributed, self-organizing attack models, as malicious nodes could exploit coordinate information to gradually dismantle essential routing backbones. On the other hand, these same geometric properties enable distributed identification of critical substructures, which supports autonomous recovery, connectivity restoration, and routing optimization. As such, hyperbolic embeddings introduce both opportunities and vulnerabilities into the management landscape. Future architectures must embrace this dual nature, combining the efficiencies of geometric design with robust mechanisms for securing and regulating coordinate-based infrastructure.

References

[1] Moyses M Lima, Horacio ABF Oliveira, Eduardo F Nakamura, Leandro N Balico, and Antonio AF Loureiro. Greedy routing and data aggregation in wireless sensor networks. In *2013 IEEE Symposium on Computers and Communications (ISCC)*, pages 000342–000347. IEEE, 2013.

[2] Marián Boguná, Fragkiskos Papadopoulos, and Dmitri Krioukov. Sustaining the internet with hyperbolic mapping. *Nature communications*, 1(1):62, 2010.

[3] Thomas Bläsius, Tobias Friedrich, Anton Krohmer, and Sören Laue. Efficient embedding of scale-free graphs in the hyperbolic plane. *IEEE/ACM transactions on Networking*, 26(2):920–933, 2018.

[4] Ines Chami, Adva Wolf, Da-Cheng Juan, Frederic Sala, Sujith Ravi, and Christopher Ré. Low-dimensional hyperbolic knowledge graph embeddings. *arXiv preprint arXiv:2005.00545*, 2020.

[5] Marton Csernai, Attila Korosi, Balazs Sonkoly, and Gergely Biczok. Poincare: A hyperbolic data center architecture. In *2012 32nd International Conference on Distributed Computing Systems Workshops*, pages 8–16. IEEE, 2012.

[6] Peng Cui, Xiao Wang, Jian Pei, and Wenwu Zhu. A survey on network embedding. *IEEE transactions on knowledge and data engineering*, 31(5):833–852, 2018.

[7] Ivan Voitalov, Rodrigo Aldecoa, Lan Wang, and Dmitri Krioukov. Geohyperbolic routing and addressing schemes. *ACM SIGCOMM Computer Communication Review*, 47(3):11–18, 2017.

[8] Nitul Dutta, Hiren Kumar Deva Sarma, Rajendrasinh Jadeja, Krishna Delvadia, and Gheorghita Ghinea. Routing schemes used in content delivery. In *Information Centric Networks (ICN) Architecture & Current Trends*, pages 73–93. Springer, 2021.

[9] Weihong Yang, Yang Qin, and Bingbing Wu. A hyperbolic routing scheme for information-centric internet of things with edge computing. *Wireless Networks*, 27:4567–4579, 2021.

[10] Robert Kleinberg. Geographic routing using hyperbolic space. In *IEEE INFOCOM 2007-26th IEEE International Conference on Computer Communications*, pages 1902–1909. IEEE, 2007.

[11] Vasileios Karyotis and Eleni Stai. Hyperbolic big data analytics for dynamic network management and optimization. In *Big Data and Computational Intelligence in Networking*, pages 177–208. CRC Press, 2017.

[12] András Majdán, Lejla Pasic, Zalán Heszberger, Rolland Vida, and József Bíró. Navigable architectures for complex communication networks. In *2024 IEEE Virtual Conference on Communications (VCC)*, pages 1–5. IEEE, 2024.

[13] Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguná. Hyperbolic geometry of complex networks. *Physical Review E*, 82(3):036106, 2010.

[14] András Majdán, Lejla Pasic, Dániel Ficzere, Gergely Hollósi, Zalán Heszberger, Rolland Vida, and József Bíró. Iot connectivity management by hyperbolic trees. In *NOMS 2025-2025 IEEE Network Operations and Management Symposium*, pages 1–7. IEEE, 2025.