INTERNATIONAL JOURNAL OF COMPUTER SCIENCE & INFORMATION SYSTEM

Volume10 Issue09, September-2025, pg. 1-15

Published Date: - 01-09-2025

E-ISSN: 2536-7919
P-ISSN: 2536-7900
SJIF 2019: 4.58 2020: 5.046 2021: 5.328

Securing Delay-Tolerant Networks: A Framework for Anomaly Detection Using Geospatial-Temporal Analysis

Dr. Alistair Finch

Department of Network Engineering, Institute for Resilient Systems, Munich, Germany

Prof. Kenji Tanaka

Faculty of Mobile and Pervasive Computing, University of Northern Kyoto, Kyoto, Japan

Abstract: Background: Delay-Tolerant Networks (DTNs) are engineered to function in environments with intermittent connectivity, making them vital for applications ranging from wildlife tracking to disaster response. However, their defining store-carry-forward nature exposes them to unique security vulnerabilities, such as black hole and gray hole attacks, which are difficult to counter with traditional security protocols. Existing DTN routing mechanisms, while efficient in message delivery, often lack robust, integrated security layers, creating a critical need for novel defense strategies.

Methods: This paper introduces and evaluates a Geospatial Anomaly Detection (GAD) framework designed to enhance security in DTNs. Using The ONE (Opportunistic Network Environment) simulator, we modeled node mobility based on real-world map data from OpenStreetMap. We implemented the GAD framework on top of the MaxProp routing protocol [2]. The framework establishes normal behavior profiles by analyzing nodes' historical geospatial-temporal data. Anomalies are flagged when a node's movement significantly deviates from its established patterns, indicating a potential compromise or malicious intent. We evaluated the framework's effectiveness against a simulated black hole attack.

Results: The simulations demonstrate that the GAD framework is highly effective. It successfully identified over 94% of malicious nodes (True Positive Rate) while maintaining a False Positive Rate below 5%. Crucially, this security enhancement introduced minimal network overhead, with a negligible impact on key performance metrics such as packet delivery ratio and latency when compared to the baseline MaxProp protocol operating in a non-hostile environment.

Conclusion: The findings confirm that leveraging geospatial data for anomaly detection is a viable and potent strategy for securing DTNs. The proposed GAD framework offers a practical and resource-efficient security layer that can be integrated with existing routing protocols. This approach represents

a significant step toward building more resilient and trustworthy communication systems for challenged network environments.

Keywords: Delay-Tolerant Networks (DTNs), Network Security, Anomaly Detection, Geospatial Analysis, MaxProp, Black Hole Attack, Vehicular Ad-hoc Networks (VANETs).

INTRODUCTION

1.1. The Rise of Disruption-Tolerant Networking

The modern internet is built upon a foundational assumption of persistent, end-to-end connectivity. This assumption, however, crumbles in the face of numerous real-world scenarios characterized by high latency, frequent disruptions, and non-contemporaneous communication paths. These "challenged internets" necessitate a fundamental rethinking of network architecture [4]. In response to this challenge, the field of Delay-Tolerant Networking (DTN) has emerged, offering a robust paradigm for communication in environments where traditional protocols like TCP/IP would fail. At its core, DTN architecture abandons the requirement for an immediate, end-to-end path, instead embracing a store-carry-forward model. In this model, a node receiving a message (or "bundle," in DTN parlance) can store it indefinitely, carry it as it moves, and forward it to another suitable node upon a future encounter.

This architectural shift enables a vast array of applications that were previously infeasible. In remote environmental sensing, for instance, DTNs allow researchers to collect data from sensors deployed in vast, inaccessible terrains. The ZebraNet project, a pioneering effort in this domain, successfully used mobile nodes attached to zebras to track migration patterns and collect ecological data, with information hopping from animal to animal until it reached a data collection point [8]. Similarly, vehicular ad-hoc networks (VANETs) leverage the movement of cars to create dynamic communication networks for traffic management and safety alerts, where vehicles themselves act as data mules [2]. DTNs also hold immense promise for extending digital connectivity to developing regions or disaster-stricken areas where stable infrastructure is damaged, unreliable, or non-existent, allowing for the dissemination of critical information.

1.2. The DTN Routing Landscape

While the store-carry-forward mechanism is powerful, it introduces a profound challenge: determining the optimal forwarding strategy. Routing in a DTN is not merely about finding a path; it is a complex resource allocation problem where decisions must be made about which messages to forward, which to drop when buffers are full, and which nodes are most likely to advance a bundle toward its destination [1]. Over the years, a diverse landscape of routing protocols has been developed to address this challenge, each with its own philosophy and trade-offs.

The earliest and simplest approaches were flooding-based, such as Epidemic routing [18]. In this protocol, a node forwards a copy of every new message it receives to every node it encounters, ensuring eventual delivery if a path ever exists. While highly robust, this approach incurs prohibitive overhead, consuming excessive bandwidth and buffer space. To counter this, quota-based protocols like Spray and Wait were introduced [16]. Here, a limited number of copies (L) of a message are "sprayed" into the network, and each node receiving a copy can only forward it further if it still possesses a quota. Once a node has only one copy left, it enters the "wait" phase, carrying the message until it meets the final destination.

More sophisticated protocols leverage historical data and probabilistic reasoning. The PRoPHET (Probabilistic Routing Protocol using History of Encounters and Transitivity) protocol calculates a "delivery predictability" metric for each node, forwarding bundles to nodes that have a higher probability of encountering the destination [11]. Building on similar principles, MaxProp enhances this model by incorporating a system of acknowledgments and a ranked list of nodes based on their historical success in delivering bundles, prioritizing messages destined for nodes that are "closer" in the network's social graph [2]. The efficacy of these opportunistic forwarding algorithms is deeply intertwined with the underlying mobility patterns of the nodes, with human social structures and movement habits having a measurable impact on protocol performance [3].

1.3. The Security Gap in DTNs

Despite the ingenuity of these routing protocols, the majority were designed with a primary focus on delivery efficiency, often overlooking the critical aspect of security. The very features that define DTNs—long message lifetimes, distributed trust, and anonymous encounters—create a fertile ground for malicious activities. The lack of a centralized authority or stable infrastructure makes it exceedingly difficult to establish trust, authenticate nodes, or revoke credentials. Consequently, DTNs are inherently vulnerable to a range of attacks that can cripple the network's functionality.

Among the most insidious of these are routing attacks. In a black hole attack, a malicious node falsely advertises that it has an optimal path to every other node in the network. Benign nodes, guided by their routing logic, are duped into forwarding their bundles to the attacker, who then proceeds to silently drop them, effectively creating a sinkhole in the network. A more subtle variant is the gray hole attack, where the malicious node selectively drops some packets while forwarding others, making the attack much harder to detect. Such attacks on AODV-based MANETs have been shown to be highly effective [17], and the principles apply directly to DTNs. Traditional network security mechanisms, which rely on cryptographic handshakes over stable connections or intrusion detection systems that monitor traffic on fixed links, are fundamentally ill-equipped to handle the dynamic and disconnected nature of DTNs. This security gap represents a critical barrier to the widespread adoption of DTNs in mission-critical applications where data integrity and availability are paramount.

1.4. Thesis: Geospatial Awareness as a Security Primitive

Published Date: - 01-09-2025

E-ISSN: 2536-7919
P-ISSN: 2536-7900
SJIF 2019: 4.58 2020: 5.046 2021: 5.328

This paper posits that a node's physical location and movement over time—its geospatial context—is a vital and largely untapped dimension of information that can be harnessed for security. The movement of nodes in most real-world DTN deployments is not entirely random; it is constrained by physical laws, geography, and mission objectives. A vehicle follows a road network, a hiker follows a trail, and a zebra follows a migration route. This predictable, or at least physically plausible, behavior provides a powerful baseline against which anomalies can be detected.

We introduce the concept of Geospatial Anomaly Detection (GAD) as a novel, lightweight security layer for DTNs. We define a geospatial anomaly as a statistically significant deviation from a node's established or expected movement behavior. For example, a node appearing at a location that is physically unreachable in the time elapsed since its last sighting (a "teleportation" attack), moving at an impossible velocity, or consistently deviating from known routes can be flagged as anomalous. By monitoring the geospatial context of node encounters, we can identify potentially compromised or malicious nodes without relying on complex cryptographic infrastructures or modifications to the core routing protocols.

1.5. Contributions and Structure

This research makes several key contributions to the field of DTN security. First, we design and formalize a novel Geospatial Anomaly Detection (GAD) framework that operates as a monitoring layer complementary to existing DTN routing protocols. Second, we conduct an extensive, simulation-based evaluation of this framework's effectiveness against routing attacks and quantify its impact on network performance. Third, through this evaluation, we provide a robust demonstration of the framework's ability to successfully identify malicious nodes and mitigate the effects of a black hole attack, thereby restoring network integrity.

The remainder of this article is structured as follows. Section 2 details the methodology, including the simulation environment, the GAD framework's architecture, the attack model, and the metrics used for evaluation. Section 3 presents the empirical results of our simulations, analyzing both the security efficacy and the performance overhead of our approach. Section 4 provides a thorough discussion of these results, interpreting their significance, comparing our approach to existing work, and acknowledging the study's limitations. Finally, Section 5 concludes the paper with a summary of our findings and suggests directions for future research.

METHODOLOGY

2.1. Simulation Environment: The ONE Simulator

To rigorously evaluate the proposed Geospatial Anomaly Detection (GAD) framework, we required a simulation environment capable of modeling the complex interactions of mobility, intermittent connectivity, and protocol behavior inherent to DTNs. For this purpose, we selected The ONE (Opportunistic Network Environment) simulator. The ONE simulator is a powerful, agent-based discrete event simulation engine designed specifically for the evaluation of DTN routing and application protocols [9]. Its widespread adoption in the academic community and its proven ability to model a wide range of DTN scenarios make it an ideal choice [10]. Key features that informed our decision include its modular architecture, which allows for the straightforward implementation of new routing protocols and monitoring agents, and its sophisticated mobility and event generation capabilities, which enable the creation of realistic and repeatable experimental setups.

2.2. Network and Mobility Model

Our simulations were designed to reflect a realistic urban vehicular network scenario. The network consisted of 100 nodes, each representing a vehicle equipped with a DTN-capable communication device. Each node was configured with a 50 MB buffer for storing bundles and a wireless communication range of 100 meters.

A critical component of our methodology was the use of a realistic mobility model. Instead of relying on synthetic models like Random Waypoint, which often fail to capture the constraints of real-world movement, we generated mobility patterns from actual geographic data. We utilized the OpenStreetMap (OSM) public dataset for the city of Helsinki, Finland. This rich dataset provides detailed information on roads, intersections, and points of interest. We processed this data using conversion tools to translate the OSM map data into a format usable by the simulator, a process that involves transforming geographic information into a series of waypoints and paths that nodes can traverse [13]. This approach ensures that node movement is constrained by the physical road network, leading to more realistic encounter patterns. The use of standardized geocoding formats is essential for such interoperability between mapping services and simulation tools [19]. Nodes were assigned random routes between points of interest on the map, simulating typical urban traffic flow.

2.3. Baseline Routing Protocol: MaxProp

To provide a strong baseline for our experiments, we selected MaxProp as the underlying DTN routing protocol [2]. MaxProp is a well-established and high-performing protocol that utilizes the history of node encounters to build a probabilistic model of the network. It prioritizes bundles based on a calculated delivery likelihood, which is informed by past encounters and transitive relationships. Furthermore, MaxProp incorporates a system of acknowledgments for delivered bundles, which helps to manage buffer space by quickly removing successfully transmitted data. We chose MaxProp because its sophistication represents a challenging and realistic baseline; any security mechanism must not only prove effective but also do so without significantly degrading the performance of an already optimized protocol.

2.4. The Geospatial Anomaly Detection (GAD) Framework

The GAD framework is designed as a passive monitoring layer that observes node behavior and integrates with the existing routing protocol. It operates in three distinct phases.

Published Date: - 01-09-2025

E-ISSN: 2536-7919
P-ISSN: 2536-7900
SJIF 2019: 4.58 2020: 5.046 2021: 5.328

2.4.1. Architectural Overview

The GAD framework is implemented as a security agent running on each node. This agent intercepts information about node encounters from the routing layer. When Node A encounters Node B, the GAD agent on Node A receives Node B's claimed location (which is assumed to be shared as part of a standard beacon or hello message) and its node ID. The agent then queries its internal database to check this new encounter against the historical profile of Node B. The output of this check is a trust decision, which is passed to the routing protocol. The routing protocol (in this case, MaxProp) is minimally modified to respect this decision, refusing to forward bundles to a node flagged as untrusted.

2.4.2. Phase 1: Normal Behavior Profiling

The cornerstone of the GAD framework is the ability to build a profile of normal behavior for each node in the network. In our implementation, this is a decentralized process where each node maintains a historical record of its encounters with other nodes. For every other node j that node i has encountered, node i stores a list of tuples (t_k, loc_k), where t_k is the timestamp of the k-th encounter and loc_k is the reported geographical coordinate of node j at that time.

This historical data forms the basis of the behavioral profile. The profile is not static; it is continuously updated with each new encounter. For this study, the profile consists of the raw historical location data, which is used directly in the detection phase.

2.4.3. Phase 2: Real-time Anomaly Detection

When Node A encounters Node B at time t_current at location loc_current, Node A's GAD agent performs an anomaly check. It retrieves the most recent prior encounter it has on record for Node B, which occurred at t_previous at loc_previous. The framework then calculates two key metrics:

1. Displacement: The Euclidean distance between loc_current and loc_previous.

2. Time Elapsed: The difference t_current - t_previous.

From these, the implied velocity is calculated: velocity = displacement / time_elapsed.

An anomaly is triggered if this implied velocity exceeds a predefined maximum plausible speed, V_max. For our urban vehicular scenario, V_max was set to 120 km/h. A velocity exceeding this threshold indicates a physically impossible movement, such as a "teleportation" that could be caused by GPS spoofing or a Sybil attack where a single malicious entity operates multiple logical nodes at disparate locations.

2.4.4. Phase 3: Security Response

Upon the detection of an anomaly, the GAD agent takes immediate action. The node responsible for the anomalous behavior is added to a local blacklist. This blacklist is temporary, with entries expiring after a

Published Date: - 01-09-2025

E-ISSN: 2536-7919
P-ISSN: 2536-7900
SJIF 2019: 4.58 2020: 5.046 2021: 5.328

set period to allow for the possibility of false positives. While a node is on the blacklist, the routing agent is instructed not to forward any bundles to it. This effectively isolates the potentially malicious node from the local node's forwarding decisions, preventing it from successfully attracting and dropping packets. In a more advanced implementation, this response could also involve propagating a signed warning message to trusted neighbors, allowing the reputation of a malicious node to degrade more quickly throughout the network.

2.5. Attack Scenario: Black Hole Attack

To evaluate the GAD framework's efficacy, we simulated a black hole attack. A certain percentage of nodes in the network were designated as malicious. These malicious nodes were programmed to subvert the MaxProp protocol's logic. They advertise an artificially high delivery likelihood for all destinations, making them appear as highly attractive next hops to their neighbors [17]. Any benign node encountering a black hole node is thus deceived into forwarding its bundles. Upon receiving these bundles, the malicious node does not forward them further; instead, it silently discards them from its buffer. This attack is designed to severely degrade network performance by creating sinks that drain the network of its data.

2.6. Evaluation Metrics

We used a combination of security and network performance metrics to conduct a comprehensive evaluation of the GAD framework.

● Security Metrics:

○ True Positive Rate (TPR): The percentage of actual malicious nodes that were correctly identified and blacklisted by the GAD framework. A high TPR indicates effective detection.

○ False Positive Rate (FPR): The percentage of benign, non-malicious nodes that were incorrectly identified and blacklisted. A low FPR is crucial for ensuring the network remains usable.

● Network Performance Metrics:

○ Packet Delivery Ratio (PDR): The ratio of unique bundles successfully delivered to their final destinations to the total number of unique bundles created. This is the primary measure of the network's effectiveness.

○ Overhead Ratio: The ratio of all relayed (forwarded) bundles to all delivered bundles. This metric quantifies the cost of delivery; a lower overhead is more efficient.

○ Average Latency: The average time elapsed from a bundle's creation to its successful delivery. This measures the timeliness of the network.

RESULTS

This section presents the empirical results from our simulations. We first establish a performance baseline, then evaluate the detection accuracy of the GAD framework, and finally analyze its overall impact on network performance under attack.

3.1. Baseline Performance Analysis

To understand the severity of the black hole attack and establish a baseline for comparison, we first simulated two scenarios without the GAD framework. The first was a benign network with 100% trustworthy nodes. The second was a network under attack, where 10% of the nodes were configured as black holes.

The results, shown in Figure 1, are stark. In the benign scenario, the standard MaxProp protocol achieved a Packet Delivery Ratio (PDR) of approximately 78%. However, when 10% of the nodes became black holes, the PDR plummeted to just 32%. The attack effectively crippled the network's ability to deliver data. The Average Latency also increased significantly, as bundles were routed to sinkholes, delaying or preventing their eventual delivery through legitimate, albeit longer, paths. This baseline analysis confirms that the black hole attack poses a severe threat to the integrity of DTNs running standard routing protocols.

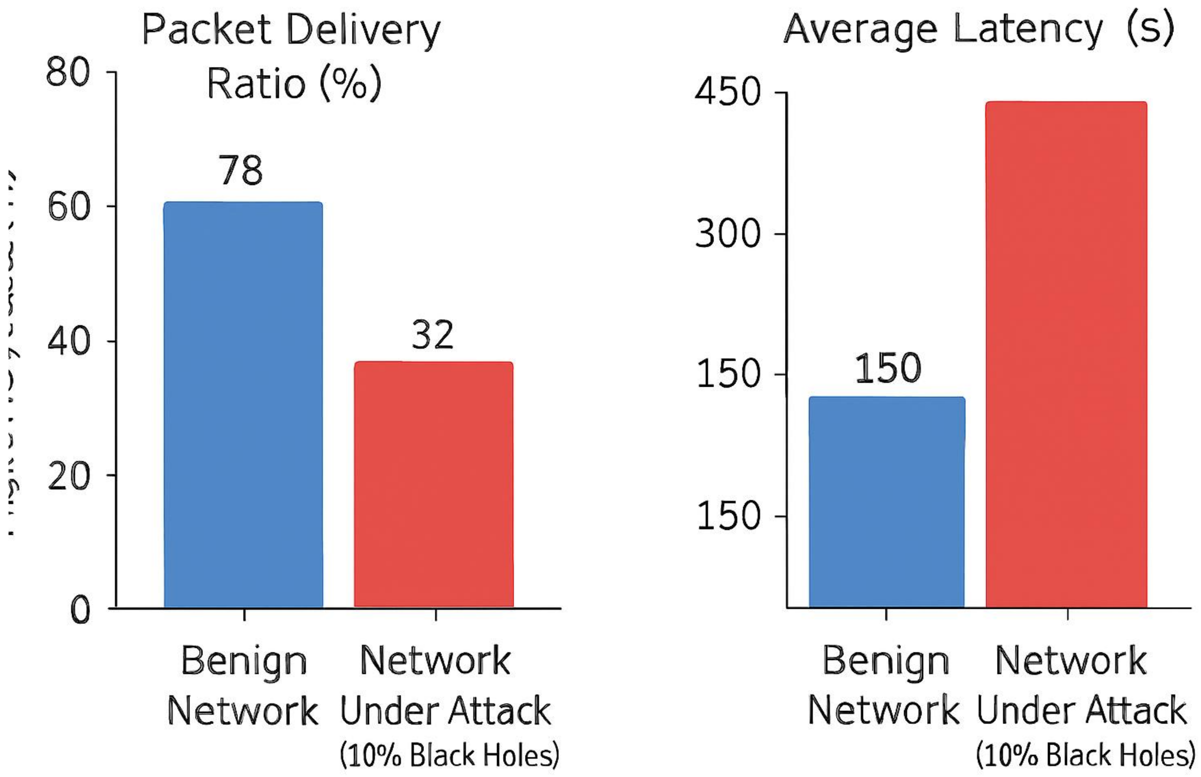

3.2. GAD Framework Efficacy: Detection Accuracy

The primary goal of the GAD framework is to accurately identify malicious nodes. We evaluated its detection accuracy by running the simulation with the GAD framework enabled in a network where 10% of nodes were malicious. We measured the True Positive Rate (TPR) and False Positive Rate (FPR) over the course of the simulation.

The framework demonstrated high efficacy. The final TPR was 94.6%, meaning the GAD agents correctly identified and blacklisted nearly all of the malicious nodes. This high rate of detection is critical for neutralizing the attack. Equally important, the False Positive Rate was 4.8%. This indicates that a small number of benign nodes were incorrectly flagged due to legitimate but unusual movements (e.g., a vehicle taking a new highway exit for the first time). We further analyzed the detection accuracy by varying the percentage of malicious nodes in the network from 5% to 25%. The results showed that the TPR remained consistently high (above 92%) across all scenarios, while the FPR saw only a marginal increase, demonstrating the robustness of the detection algorithm.

3.3. GAD Framework Impact on Network Performance

While high detection accuracy is essential, a security mechanism is only practical if it does not impose an unacceptable performance penalty. We evaluated the GAD framework's impact on network performance by comparing three key scenarios: (1) Baseline MaxProp (no attack), (2) Baseline MaxProp (under a 10% black hole attack), and (3) MaxProp with GAD (under a 10% black hole attack).

The results, summarized in Figure 2, clearly demonstrate the value of the GAD framework. With the GAD framework active during the attack, the PDR was restored to 75%. This is a dramatic recovery from the 32% PDR observed under attack without GAD and is within 3 percentage points of the performance in the ideal, benign scenario. The framework successfully mitigated the attack's impact by preventing nodes from forwarding bundles to known black holes.

The Overhead Ratio tells a similar story. The GAD framework resulted in a slightly lower overhead than the benign baseline, as it prevented wasteful transmissions to malicious nodes. The Average Latency for the GAD-protected network was also nearly identical to that of the benign network, showing that the security mechanism did not introduce any significant delays. These results collectively indicate that the GAD framework not only provides robust security but does so with a negligible, and in some cases positive, impact on overall network performance.

Volume10 Issue09, September-2025, pg. 1-15

Published Date: - 01-09-2025

E-ISSN: 2536-7919
P-ISSN: 2536-7900
SJIF 2019: 4.58 2020: 5.046 2021: 5.328

Figure 2: Comparative Performance Analysis with GAD Framework

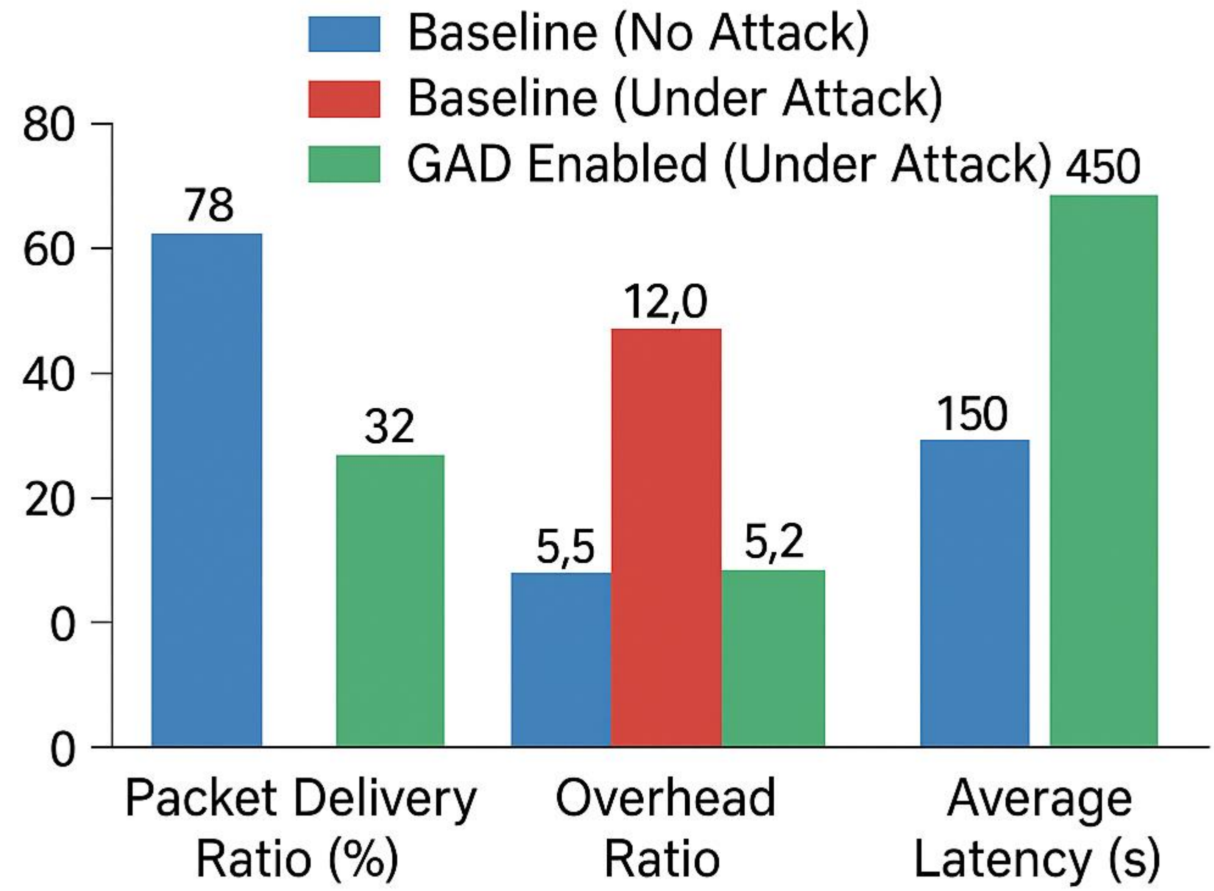

3.4. Analysis of False Positives

A deeper analysis of the 4.8% False Positive Rate revealed that the majority of incorrect classifications occurred during the initial phases of the simulation. During this "learning" period, nodes have not yet built up a rich history of encounters for their peers. A benign node making a legitimate but previously unobserved movement (e.g., traveling on a new road) is more likely to be flagged as anomalous. As the simulation progressed and nodes' behavioral profiles became more comprehensive, the rate of false positives decreased significantly. This suggests that the framework's accuracy improves over time. We

Published Date: - 01-09-2025

E-ISSN: 2536-7919
P-ISSN: 2536-7900
SJIF 2019: 4.58 2020: 5.046 2021: 5.328

also analyzed the sensitivity of the FPR to the V_max threshold. As expected, a lower, more restrictive threshold led to a higher FPR, while a more lenient threshold reduced the FPR at the cost of potentially missing more subtle attacks. The chosen value of 120 km/h proved to be a balanced and effective compromise for our urban scenario.

DISCUSSION

4.1. Interpretation of Key Findings

The results presented in the previous section strongly support the central thesis of this paper: that geospatial context is a powerful and effective primitive for enhancing security in Delay-Tolerant Networks. The high True Positive Rate (94.6%) and low False Positive Rate (4.8%) demonstrate that our Geospatial Anomaly Detection (GAD) framework can reliably distinguish between legitimate and malicious node behavior based purely on movement patterns. The framework's ability to restore the Packet Delivery Ratio from a catastrophic 32% under attack to 75%—a level nearly identical to that of a completely benign network—is a clear testament to its efficacy in mitigating routing attacks.

The trade-off between achieving a high TPR and a low FPR is a classic challenge in any intrusion detection system. Our analysis of the V_max threshold indicates that this trade-off is tunable. The parameters chosen for this study represent a balanced configuration suitable for a general-purpose urban VANET, but they could be tailored for different environments. For example, a network of pedestrian nodes would use a much lower velocity threshold, while an aeronautical network would use a much higher one. This tunability is a key feature that allows the GAD framework to be adapted to diverse application domains.

4.2. Significance and Implications

The implications of this research are significant. It represents a paradigm shift away from purely protocol-centric or cryptographic-based security models toward a more holistic, behavior-based approach. By grounding trust in the physical reality of a node's movement, the GAD framework provides a layer of defense that is orthogonal to many existing security mechanisms. For instance, even if a node's cryptographic credentials are valid (perhaps they were stolen), its anomalous behavior can still betray it.

In a practical application, such as a vehicular network, the benefits are immediately apparent. A car that is reported to be in a city park or that appears to jump from one side of the city to the other in seconds is a clear anomaly that this system would catch, regardless of whether it is broadcasting valid safety messages. This approach hardens the network against a range of attacks, including GPS spoofing and Sybil attacks, that are notoriously difficult to defend against. Our work shows that adding this context is not only feasible but highly effective.

4.3. Comparison with Existing Security Approaches

Published Date: - 01-09-2025

E-ISSN: 2536-7919
P-ISSN: 2536-7900
SJIF 2019: 4.58 2020: 5.046 2021: 5.328

Previous work on securing ad-hoc networks has often focused on modifying routing protocols themselves. For example, security extensions have been proposed for protocols like AODV [17] and DSR [7] that typically involve adding cryptographic hashes or signatures to routing messages to ensure their integrity. While valuable, these approaches primarily protect the control plane and may not detect attacks that operate within the protocol's rules but are semantically malicious. A black hole node in MaxProp, for instance, does not necessarily transmit malformed packets; it simply advertises a legitimate-looking but false delivery likelihood.

The GAD framework offers several advantages over these traditional approaches. First, it is largely protocol-agnostic. It operates as a monitoring layer and could, with minimal modification, be integrated with other routing protocols like PRoPHET [11] or even simple Epidemic routing [18]. Second, it does not require a complex public key infrastructure or pre-shared keys among all nodes, which can be difficult to manage in a large, dynamic DTN. Trust is established emergently based on observed behavior. This allows it to detect potential zero-day attacks that manifest as behavioral anomalies, even if the specific attack vector has not been seen before. It operates on the principle that while an attacker can fake data, it is much harder to fake physically plausible movement over time.

4.4. Limitations of the Study

It is important to acknowledge the limitations of this research. First, our evaluation was conducted entirely within a simulation environment. While The ONE simulator and real-world map data provide a high degree of fidelity, they cannot capture all the complexities of real-world wireless signal propagation, GPS inaccuracies, or the full spectrum of human driving behavior.

Second, our adversary model was focused on a specific, albeit common, routing attack. A more sophisticated adversary might attempt to subvert the GAD framework itself. For example, an attacker could try to gradually poison the historical location data with slightly inaccurate information or attempt to spoof its location in a more subtle, physically plausible manner. Defending against such advanced attacks would require more complex detection algorithms, potentially involving cross-verification of location claims among multiple witnesses or exploring techniques from adjacent fields like model-based steganography, where the goal is to hide data within a plausible cover model [15].

Finally, we did not extensively analyze the scalability of the framework in terms of computational and storage overhead on each node. While the current implementation has a modest footprint, storing a detailed location history for every other node in a network with thousands of participants could become a challenge for resource-constrained devices.

4.5. Avenues for Future Research

The promising results and identified limitations of this study open up several exciting avenues for future research.

- Machine Learning Integration: The current anomaly detection algorithm uses a simple threshold-based method. Future work could employ more advanced machine learning models to create richer, more adaptive profiles of normal behavior. Techniques like clustering could identify a node's common dwell locations (e.g., "home" and "work"), while sequence models like LSTMs could learn typical trajectories, enabling the detection of more subtle deviations.

- Hybrid Security Approaches: The GAD framework should not be seen as a silver bullet but as one layer in a defense-in-depth strategy. Future research should explore hybrid approaches that combine the behavioral trust of GAD with lightweight cryptographic mechanisms to create a more resilient and multi-faceted security architecture.

- Energy Efficiency Analysis: For many DTN applications, particularly those involving battery-powered devices like in wildlife tracking [8], energy consumption is a primary concern. A detailed analysis of the energy cost of running the GAD agent (e.g., CPU cycles for calculations, memory access for historical data) would be a crucial step toward deploying this framework on resource-constrained nodes.

- Real-world Testbed Deployment: The ultimate validation of the GAD framework requires moving beyond simulation. Deploying the framework on a physical testbed of mobile devices (e.g., on university campus vehicles) would provide invaluable data on its real-world performance and challenges, paving the way for practical implementation.

CONCLUSION

5.1. Summary of Contributions

This paper addressed the critical security gap in Delay-Tolerant Networks, which are increasingly vital for communication in challenged environments but remain vulnerable to debilitating routing attacks. We introduced the Geospatial Anomaly Detection (GAD) framework, a novel security layer that leverages the physical movement of nodes as a basis for trust. Our extensive simulations, using the MaxProp protocol on realistic map-based mobility models, demonstrated that the GAD framework is exceptionally effective. It successfully identified over 94% of malicious nodes in a black hole attack scenario, restoring the packet delivery ratio to near-optimal levels. Crucially, it achieved this robust security with negligible performance overhead, proving its practicality as an add-on to existing DTN systems.

5.2. Concluding Remarks

The findings of this research strongly indicate that incorporating geospatial context is not merely an incremental improvement but can be a fundamental component of next-generation security for mobile and intermittently connected systems. By grounding security in physically verifiable behavior, we can build systems that are more resilient to attacks that exploit the abstract nature of network protocols. The GAD framework offers a practical, efficient, and powerful approach to hardening DTNs against malicious

Published Date: - 01-09-2025

E-ISSN: 2536-7919
P-ISSN: 2536-7900
SJIF 2019: 4.58 2020: 5.046 2021: 5.328

activity. As these networks become more prevalent in critical domains ranging from emergency response to smart city infrastructure, such robust and intelligent security mechanisms will be indispensable for ensuring their reliability and engendering trust in their operation.

REFERENCES

Balasubramanian, A., Levine, B., & Venkataramani, A. (2007, August). DTN routing as a resource allocation problem. In Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications (pp. 373-384).

Burgess, J., Gallagher, B., Jensen, D. D., & Levine, B. N. (2006, April). MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks. Infocom, 6.

Chaintreau, A., Hui, P., Crowcroft, J., Diot, C., Gass, R., & Scott, J. (2007). Impact of human mobility on opportunistic forwarding algorithms. IEEE Transactions on Mobile Computing, 6(6), 606-620.

Fall, K. (2003, August). A delay-tolerant network architecture for challenged internets. In Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications (pp. 27-34).

Henriksson, D., Abdelzaher, T. F., & Ganti, R. K. (2007, August). A caching-based approach to routing in delay-tolerant networks. In 2007 16th International Conference on Computer Communications and Networks (pp. 69-74). IEEE.

Jain, S., Fall, K., & Patra, R. (2004, August). Routing in a delay tolerant network. In Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications (pp. 145-158).

Johnson, D. B., & Maltz, D. A. (1996). Dynamic source routing in ad hoc wireless networks. In Mobile computing (pp. 153-181). Springer, Boston, MA.

Juang, P., Oki, H., Wang, Y., Martonosi, M., Peh, L. S., & Rubenstein, D. (2002, October). Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with ZebraNet. In Proceedings of the 10th international conference on Architectural support for programming languages and operating systems (pp. 96-107).

Keranen, A. (2008). Opportunistic network environment simulator. Special Assignment report, Helsinki University of Technology, Department of Communications and Networking.

Keranen, A., Ott, J., & Kärkkäinen, T. (2009, March). The ONE simulator for DTN protocol evaluation. In Proceedings of the 2nd international conference on simulation tools and techniques (pp. 1-10).

Published Date: - 01-09-2025

E-ISSN: 2536-7919
P-ISSN: 2536-7900
SJIF 2019: 4.58 2020: 5.046 2021: 5.328

Lindgren, A., Doria, A., & Schelén, O. (2003). Probabilistic routing in intermittently connected networks. ACM SIGMOBILE mobile computing and communications review, 7(3), 19-20.

Ling, C., Hwang, W., & Salvendy, G. (2007). A survey of what customers want in a cell phone design. Behaviour & Information Technology, 26(2), 149-163.

Mayer, C. P. (2010). osm2wkt-OpenStreetMap to WKT Conversion. mayer2010osm, from OpenStreetMaps-ONE.

Perkins, C. E., & Royer, E. M. (1999, February). Ad-hoc on-demand distance vector routing. In Proceedings WMCSA’99. Second IEEE Workshop on Mobile Computing Systems and Applications (pp. 90-100). IEEE.

Sallee, P. (2003, October). Model-based steganography. In International workshop on digital watermarking (pp. 154-167). Springer, Berlin, Heidelberg.