Network Hexagons Under Attack: Secure Crowdsourcing of Georeferenced Data

Okemawo Obadofin
Carnegie Mellon University, Africa
Carnegie Mellon University

João Barros
Carnegie Mellon University, Africa
Carnegie Mellon University

***Abstract*—A critical requirement for modern-day Intelligent Transportation Systems (ITS) is the ability to collect georeferenced data from connected vehicles and mobile devices in a safe, secure and anonymous way. The Nexagon protocol, which builds on the IETF Locator/ID Separation Protocol (LISP) and the Hierarchical Hexagonal Clustering (H3) geo-spatial indexing system, offers a promising framework for dynamic, privacy-preserving data aggregation. Seeking to address the critical security and privacy vulnerabilities that persist in its current specification, we apply the STRIDE and LINDDUN threat modelling frameworks and demonstrate, among other findings, that the Nexagon protocol is susceptible to user re-identification, session linkage, and sparse-region attacks. To address these challenges, we propose an enhanced security architecture that combines public key infrastructure (PKI) with ephemeral pseudonym certificates. Our solution guarantees user and device anonymity through randomized key rotation and adaptive geospatial resolution, thereby effectively mitigating re-identification and surveillance risks in sparse environments. A prototype implementation over a microservice-based overlay network validates the approach and underscores its readiness for real-world deployment. Our results show that it is possible to achieve the required level of security without increasing latency by more than 25% or reducing the throughput by more than 7%.**

***Index Terms*—Secure Crowdsourcing, Mobile Networks, Privacy-Preserving Protocols, Threat Modeling for Geo-Privacy, Mobile Edge Security.**

I. Introduction

Distributed machine learning (ML) running on connected vehicles and mobile devices promises to deliver significant gains in transportation efficiency and sustainability, most notably by improving route planning and traffic optimization in real time [1] with multiple layers of contextual information [2]. As illustrated in Figure 1, connected vehicles are now serving as edge devices capable of acquiring massive amounts of georeferenced data for a variety of use cases, from road defects to crash detection. The accuracy and granularity of such data ultimately determines the overall safety and performance of intelligent transportation systems [3].

A key challenge in this class of data-driven solutions is how to collect real-time data from millions of connected vehicles and mobile edge devices, while addressing the very significant security and privacy concerns of all those who own and operate them in real-world situations. In other words, to ensure that no attacker is able to track or compromise any of the mobile nodes or their data, the underlying intelligent transportation ecosystem requires secure data aggregation systems that are implemented at scale [4], [5].

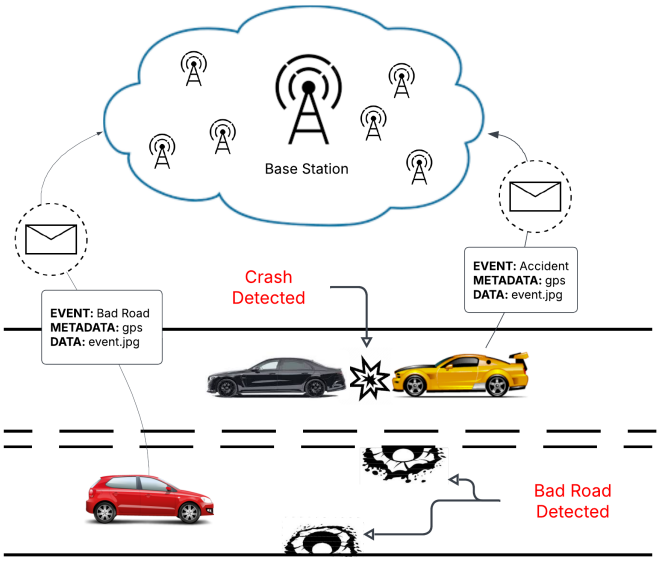

Fig. 1. A traffic scene showing a network of connected vehicles capturing and reporting road conditions and vehicle crashes in real time.

In light of these challenges, the Internet Engineering Task Force (IETF) is currently discussing and standardizing the use of network hexagons (”nexagons”) [6] for secure mobile data collection. This so called *Nexagon Protocol* builds on the Locator/ID Separation Protocol (LISP) [7] and enables a network of mobile nodes to map traffic signs, vehicle routes, construction works, natural hazards and other road conditions in real time. The protocol relies on geographically distributed agents to transmit sensor-derived attributes, while ensuring timely updates across the entire network.

The Nexagon protocol lacks detailed guidelines for practical authentication mechanisms necessary to support its security requirements. Therefore, our research focuses on critically examining the protocol’s architecture to identify aspects that undermine the security and privacy of connected vehicles and their users. Our goal is to conduct a systematic threat

This publication was developed as part of a program managed by Carnegie Mellon University Africa and supported by the MasterCard Foundation. The views expressed in this document are solely those of the authors and do not necessarily reflect those of the Carnegie Mellon University Africa or the MasterCard Foundation. This work was also partially funded by the Bill and Melinda Gates Foundation Carnegie through Upanzi Network at CMU-Africa.

analysis and implement a real-world prototype with effective mitigations for identified vulnerabilities that informs future studies.

Our research addresses critical gaps in the Nexagon protocol by providing: (1) a comprehensive analysis of security and privacy threats within the Nexagon protocol, along with corresponding mitigations; (2) a PKI-based authentication system employing pseudonym certificates and secure key rotation to guarantee secure discovery and enhanced anonymity; and (3) a prototype Nexagon implementation along with performance metrics that demonstrates the feasibility and overhead of PKI-based authentication enhancements.

The remainder of the paper is structured as follows: Section II offers an overview of related research on privacy-preserving protocols within ITS context, leading into our contributions. Section III elaborates on the Nexagon architecture, detailing its association with the LISP and the Hierarchical Hexagonal Grid System (H3). Section IV conducts a security and threat analysis of the Nexagon protocol, identifying potential vulnerabilities, and proposing mitigations. Section V details our implementation strategy for the Nexagon protocol, followed by a comprehensive evaluation of our security attachments. Finally, Section VI concludes with a summary of our findings and suggests areas for future research.

II. Related Work

The heavy dependence of ITS applications on geo-referenced data introduces inherent privacy risks [8] for participating users. Providing user anonymity in such systems is essential to maintain public trust and encourage widespread adoption. Research on k-anonymity [9] introduced a mathematically robust solution to address privacy risks by obscuring user identities. The anonymity model ensures that each individual's data is indistinguishable from at least k-1 other individuals in the dataset. This is achieved by introducing dummy users that enhance anonymity and provide statistical guarantees against re-identification. In contrast, differential privacy (DP) [10] adopts a different strategy by adding randomness to the data through the introduction of noise, ensuring that the output of queries reveals minimal information about any individual while maintaining overall dataset utility.

While privacy-preserving techniques like k-anonymity show promising results, significant challenges persist when applied to real-time, dynamic environments like ITS. For instance, [10] used k-anonymity to obscure the actual location for users of location based services (LBS) by introducing dummy requests from users with spoofed locations. This approach ensures user anonymity but leads to inefficiencies in environments with infrastructural resource constraints. In ITS, where timely updates on routes and traffic conditions are critical, this method results in significant bandwidth consumption without proportional benefits. Additionally, mobile clients using this scheme for spoofing face performance degradation, as duplicate requests are required to maintain anonymity.

In contrast, differential privacy introduces controlled noise into the data, ensuring that individual user data remain protected without requiring any redundancy. However, the introduction of noise can obscure critical details, resulting in less precise traffic predictions, route optimization, or fleet management insights [11].

Several other research efforts have explored real-world solutions aimed at preserving privacy in mobile crowdsourcing systems. The Methods outlined by [12]–[15], safeguard user data and maintain anonymity while still allowing for efficient data aggregation and task allocation. However, these solutions either fall short when applied to dynamic and real-time environments or are too technically challenging to implement and maintain. Existing approaches to guarantee anonymity do not meet the operational needs of ITS, highlighting the need for a more efficient, scalable, and adaptable solution. The Nexagon protocol [6] addresses these challenges by introducing a solution designed for real-time data exchange with security extensions tailored to distributed connected vehicles and other kinds of mobile edge devices.

III. Network Hexagons: An Overview

The Nexagon protocol establishes a distributed network of mobile edge devices designed to support real-time streaming of geo-referenced data while ensuring user privacy. The protocol leverages the advanced addressing semantics of the Locator/ID Separation Protocol (LISP) [7] for efficient client management and employs the H3 hierarchical spatial indexing system to effectively localize clients. LISP addresses the challenges of the modern internet by separating the concerns of uniquely identifying a device and its routing context within the network. H3 was developed to improve the user experience in applications that rely on location-based services [16], [17].

As a key component of the Nexagon network, LISP provides a foundational architecture for localizing clients within hexagons. By separating global and local scopes, LISP mirrors traditional network designs through its use of Routing Locators (RLOCs) and End-Point Identifiers (EIDs). RLOCs enable communication between LISP-enabled edge routers, offering wide-area network access, while EIDs facilitate interactions among clients within the Nexagon overlay. To ensure efficient data forwarding and seamless network discovery, edge routers maintain mapping caches that store the associations between EIDs and RLOCs.

The H3 package, developed by Uber Technologies, was designed to enhance location-based services for user applications. The library converts the GPS coordinates of clients into unique identifiers that correspond to a specific hexagon within a defined grid. Each hexagon, also known as a "H3 tile", groups devices with similar GPS coordinates in the same tile. The granularity of these tiles is determined by a parameter called resolution, which defines the level of detail for a given geographical area [17]. As illustrated in Figure 2, the sectional map visually demonstrates how H3 tiles establish boundaries, enabling vehicles to be grouped based on their geographic location.

In addition to mobile edge devices or clients, the Nexagon specification outlines several key components that enable and

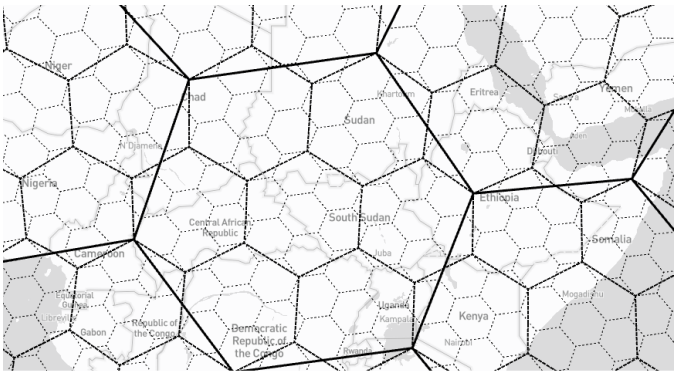

Fig. 2. Sectional map with hexagonal tiles of varying granularity superimposed showing how vehicles and mobile devices can be indexed.

support the services it provides for edge devices. Figure 3 provides a comprehensive visual representation of all the participating elements within this scheme, highlighting their interactions with other components and their specific role in enabling secure, scalable, and privacy-preserving data aggregation.

- **Authentication Nodes:** Handles client association and onboarding into the Nexagon network by securely providing the necessary credentials through a specified method.
- **Geo-Mapping Nodes:** Maintains a spatial database that maps geo-referenced updates to corresponding client Endpoint Identifiers (EIDs).
- **H3 Aggregation Nodes:** Processes data collected within each H3 tile, enabling localized analysis and supporting a wide range of services for users. Aggregated results are stored in data lakes for future analytics and insights.

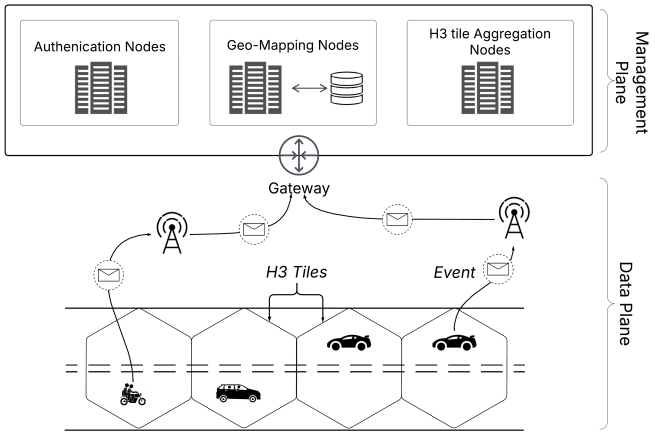

Fig. 3. An Overview of the Nexagon protocol architecture.

IV. Security and Privacy Analysis

As it stands, the Nexagon protocol specification does not define the specific mechanism to be used for client authentication. To ensure the network is robustly protected against rogue actors, it is critical to conduct a comprehensive investigation of potential attack vectors and develop strong security mechanisms capable of mitigating these threats effectively. This analysis is aimed specifically for safeguarding the integrity and privacy of mobile edge devices. To achieve this, we apply the STRIDE [18] and LINDDUN [19] threat modeling frameworks. Together, these complementary techniques enable us to systematically identify privacy risks within the Nexagon protocol.

First, we decompose the Nexagon system into Data Flow Diagrams (DFDs), providing a visual representation of the interactions and information flows between its components. Next, we map threat categories to corresponding DFD elements, identifying components that are susceptible to security and privacy risks. Subsequently, attack tree templates from the STRIDE and LINDDUN frameworks are leveraged to uncover potential attack scenarios and develop mitigation strategies. Finally, we summarize our findings in a comprehensive table that outlines the identified threats, proposed mitigations, and privacy-enhancing solutions to be considered within our implementation setup.

A. System Decomposition

To evaluate the security and privacy risks of the Nexagon protocol in a systematic way, we decompose its architecture into foundational elements using a DFD. Therefore, each element is represented as an entity, process, data store, or data flow. As shown in Figure 4, the resulting trust boundaries are defined across three different zones: (1) *Untrusted External* (client processes and edge infrastructure), (2) *Semi-Trusted "Demilitarized Zone" (DMZ)* (aggregation processes and edge routers), and (3) *Trusted Management Plane* (core authentication processes). Whereas external entities like connected vehicles and other mobile edge devices are assigned to the Untrusted External Zone, edge routers are scoped to the Semi-Trusted DMZ, because they mediate the traffic between untrusted clients and trusted internal processes. The Authentication process operates within the Trusted Management Plane because it handles sensitive tasks such as issuing pseudonym credentials and managing cryptographic secrets. Data flows such as geo-referenced updates are scoped across trust boundaries.

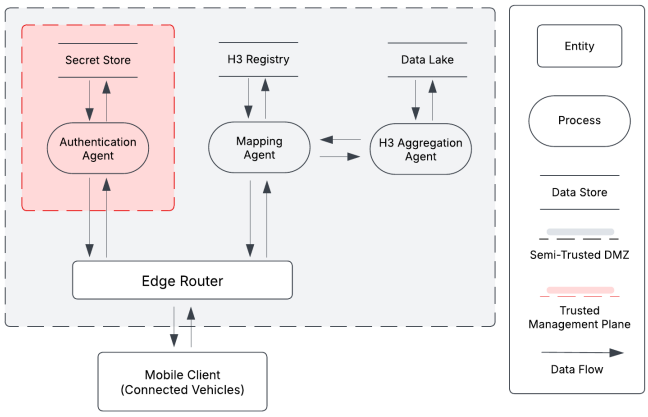

Fig. 4. Data flow diagram for the Nexagon protocol.

B. Mapping Threat Categories

Table I provides a comprehensive mapping of security and privacy threat categories to their respective DFD elements in the Nexagon architecture. Each intersection signifies the applicability of a specific threat category to a particular DFD element: *entity, process, data store, or data flow*. This mapping highlights how different components of the system are exposed to varying risks and serves as a foundation for targeted mitigation strategies.

For instance, properties like *Linkability* and *Disclosure* span across multiple DFD elements, reflecting their pervasive nature in systems handling sensitive data. Similarly, threat categories like Spoofing and Tampering emphasize vulnerabilities in entities and processes, while the impact of Denial of Service is more prevalent to processes and data flows due to their role in maintaining system availability. Repudiation and Information Disclosure were excluded from the STRIDE mapping because they are already addressed within the LINDDUN framework. Similarly, Content Unawareness and Non-Compliance were excluded as they do not impact the privacy of users in Nexagons.

TABLE I
MAPPED THREAT CATEGORIES TO DFD ELEMENTS

	Threat Category	DFD Element			
		Entity	Process	Data Store	Data Flow
LINDDUN	Linkability	X	X	X	X
	Identifiability	X	X	X	X
	Non-repudiation		X		X
	Detectability	X			X
	Disclosure	X		X	X
STRIDE	Spoofing	X	X		
	Tampering		X	X	
	Denial of Service		X		X

C. Threats, Mitigations, and Privacy Enhancing Solutions

In Table II, we present a comprehensive list of threats discovered during our systematic threat modeling process. For each identified threat, we used insights from attack trees provided by the LINDDUN framework to develop attack scenarios and propose mitigations based on evaluated privacy-enhancing solutions. This systematic approach ensures that both security vulnerabilities and privacy risks are addressed holistically. We summarize our findings by presenting two critical threat models that highlight vulnerabilities with high-risk factors. We address the vulnerabilities identified during our prototype implementation, as outlined in the following section. For brevity, we include a table summarizing attack scenarios, providing descriptions for other critical attacks not explicitly detailed.

1) Narrow-region attacks in hierarchical clustering: While hexagons offer an efficient approach to localizing data from mobile clients, privacy concerns arise when H3 tiles are overlaid on road networks in sparsely populated or remote areas. In such regions, the low density of clients can inadvertently expose individuals to privacy risks. As illustrated in Figure 5, a lone client moves from Point **A** to Point **B** along a defined road with overlaid hexagonal tiles. As the client provides periodic updates, its position within the grid can be easily tracked. Each time the client enter a new hexagon, its presence can be easily noticed, making it a trivial task to infer the client's direction and ultimately predict the final destination.

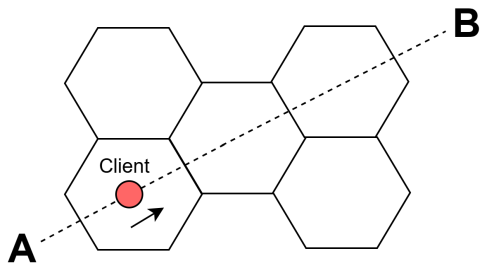

Fig. 5. A mock scenario illustrating a lone clients roaming within sparsely populated hexagons.

One approach to mitigate this risk is to take advantage of the H3 modules that enable us to adjust how dense or coarse the localization of mobile edge devices should be. A parameter specified as the resolution ranging from 1 to 15 is set to control this feature. By dynamically altering the resolution in remote areas, we can achieve an effect similar to k anonymity, while maintaining the overall utility of geo-referenced data gathered. Varying the resolution in such scenarios to accommodate more clients will pose a difficult task for adversaries to decipher.

2) Spoofed Control Plane Agent: The Nexagons management agents serve as vital abstractions that oversee the network's overall state and manage core services handling packet delivery, data pre-processing, and authentication for mobile edge clients. Ensuring that clients communicate exclusively with legitimate agents is essential to maintain trust among entities. Malicious actors may attempt to disrupt the service discovery process or compromise client privacy during device initialization by impersonating a legitimate authentication server. To combat this threat, a trusted platform module (TPM) [24] can provide essential protection for onboarding mobile edge clients by securely generating, storing, and managing cryptographic keys and credentials. When a client attempts to authenticate in a nexagon, the TPM

TABLE II
THREATS, AFFECTED COMPONENTS, ATTACK DESCRIPTIONS, AND MITIGATIONS IN NETWORK HEXAGONS

Threat	**Affected Component(s)**	**Attack Description**	**Mitigation**	**Risk Level**
Session Linkage	–Mobile Client –Authentication Agent	An attacker links users by connecting credentials across different processes and tracking actions across sessions through static identifiers or similar patterns.	–Use pseudonymized EIDs that are rotated dynamically [20]. –Add dummy traffic to prevent timing-based correlation.	High
Request Profiling	–Mobile Client	An attacker links user behavior across processes to identify patterns (for example request frequency).	–Add noise to request patterns to obscure user behavior [21]. –Route requests through mix networks to anonymize traffic.	Medium
Sparse Region Attack	–Mobile Client	An attacker exploits sparsely populated or remote areas by leveraging the low density of clients within static hexagonal grids.	–Expand hexagonal regions dynamically in sparse areas. –Ensure at least k clients are indistinguishable in any region [22].	High
User Re-identification	–Mobile Client –Mapping Agent	An attacker identifies users by correlating pseudonyms with external datasets (for example IP addresses).	–Authenticate users without revealing identity. –Encrypt all metadata in communications and storage or use an onion router.	Medium
Forged Logs or Audit Entries	–Mobile Client –Authentication Agent –Mapping Agent –Aggregation Agent	A malicious actor modifies logs to deny responsibility for specific actions.	–Use blockchain or append-only storage for audit trails. –Cryptographically sign logs to ensure integrity.	Low
Client Request/Response Replay	–Authentication Agent –Mapping Agent	An attacker replays valid authentication requests to bypass non-repudiation mechanisms.	–Include nonces in requests to prevent replay attacks. –Use mutual TLS (mTLS) for bidirectional verification.	Medium
User Data Leakage and Eavesdropping	–Mobile Client –Authentication Agent	Attackers exploit weak encryption or poor access controls to intercept unencrypted transmissions (including geo-referenced data), exposing sensitive user information.	–Encrypt all data in transit and at rest. –Enforce strict role-based access controls (RBAC).	Medium
Spoofed Agent	–Mobile Client –Mapping Agent	An attacker impersonates a legitimate client or authentication agent.	–Use hardware-backed attestation for device verification. –Require both parties to authenticate each other using certificates [23].	High

High - Very strong likelihood of occurring and has a critical effect on the Nexagons, and **Not** currently addressed by well-known schemes
Medium - Very strong likelihood of occurring, has a critical effect on the Nexagons, and currently addressed by well-known schemes
Low- Very weak likelihood of occurring and has a critical effect on the Nexagons

provides verifiable information, such as a server name to be resolved and secure cryptographic keys necessary to solve a challenge, ensuring that only legitimate clients can access management services. Previous research has demonstrated that software implementations of TPM, such as firmware-based TPM (fTPM) offer robust security guarantees comparable to dedicated TPM hardware and have been deployed in millions of mobile devices [25].

V. IMPLEMENTATION AND PERFORMANCE ANALYSIS

The Nexagon [6] protocol specification does not explicitly define which layer of the internet stack the protocol should be implemented. However, since it leverages LISP, a Layer 3 protocol, it's development naturally aligns with the network layer. However, this approach requires significant changes to existing standards and configurations. Deploying the protocol as an overlay offers significantly more flexibility when integrating with existing infrastructure, promoting faster integrations and seamless compatibility with current systems [26], [27]. Based on this observation, we now discuss in detail our approach for protocol's deployment as an overlay, outlining development processes for integrating mitigations for threats, and some architectural abstractions along with core technologies utilized.

A. Deployment Architecture

Leveraging the microservice architecture, each component was developed as a isolated services, enabling it to operate independently while interacting through specified APIs. Core functionalities were encapsulated within distinct containers, each representing a Nexagon agent. In summary, we complemented the Nexagon components with the following security enhancements:

1) Authentication: We designed the authentication process to operate as a root Certificate Authority (CA), serving as the single source of truth. However, an alternative deployment strategy could adopt a hierarchical CA structure with intermediate CAs to provide flexibility and scalability, particularly when accommodating existing vendors and cloud service

providers. Intermediate CAs will act as delegated entities authorized by the root CA to issue certificates, which enhances security and operational efficiency [28]. While adhering to the core principles of PKI [29], we augmented our solution with two essential mechanisms integrated into its operations.

- Anonymity through randomness: To enhance security, we periodically rotate the keys used by the agent for signing, forcing all agents to re-authenticate periodically. This offers an additional layer of security compared to a traditional PKI deployments that maintain certificate validity for extended periods without requiring renewal or re-validation. A similar approach is applied to the EIDs generated by the client agent discussed in the following section. Additionally, a variable hexagonal resolution is employed after each key rotation for clients who are sampled by the management plane to be at risk of being located in a sparse region. This approach mitigates the privacy risks associated with the hierarchical clustering we previously highlighted.
- Pseudonym certificates for enhanced privacy: The use of pseudonym certificates mitigates the potential for unwanted surveillance and limits the exposure of sensitive data. This means that even if a certificate is compromised, its impact is minimized, as it cannot be linked back to the client. Pseudonym certificates differ from standard X.509 certificates in that they omit or anonymize identifying fields, such as the Subject Name or Distinguished Name, to enhance user privacy and prevent direct correlation with a specific identity, while still enabling authentication and secure communication. Successful generation of certificates is backed by the software TPM of a legitimate client during initialization in a Nexagon. The TPM generates a fresh key pair and uses its private key to sign a request that includes the newly generated public key and an encrypted representation of its long-term identity. This request is sent to the CA, which verifies it with the TPM's attestation. After validation, the CA issues a pseudonym certificate tied to the new key pair, containing metadata like permissions and expiration [30].

Although other cryptographic methods such as group signatures [31] exist that guarantee privacy and accountability, pseudonym certificates offer better applicative advantages in Nexagons due to their scalability, flexibility, and enhanced privacy features. With pseudonym certificates individual-level anonymity is guaranteed through dynamic key rotation and omission of identifying fields, preventing re-identification and linkage attacks.

2) Mobile Client: The mobile agent is capable of generating unique EIDs, which is essential to publish events while maintaining anonymity. After coming online, the client is configured to first resolve a pre-configured server name via DNS. It then contacts the authentication agent to obtain the necessary credentials. Another significant addition is the ability for the agent to periodically swap EID when forwarding data across the network. The frequency of certificate swaps, combined with the randomness of these unique identifiers, effectively preserves the privacy of clients and their users. Figure 6 presents a sequence diagram that illustrates the steps undertaken by a mobile client during the initial onboarding process. This process is essential to ensure that the events published are acknowledged by the network.

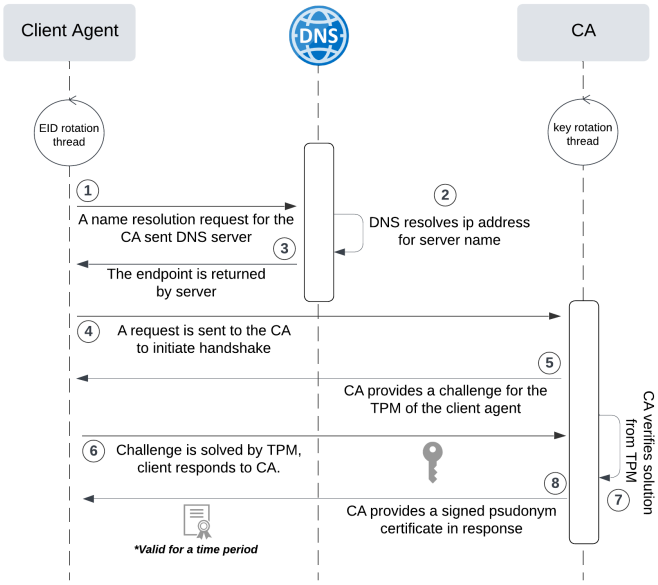

Fig. 6. Sequence diagram illustrating the interaction between the client agent and the Certificate Authority (CA) during the initial onboarding phase.

3) Mapping and Aggregation: The mapping agent is equipped to store the current mappings of active clients as a key-value (KV) data store. We developed endpoints to enable clients to publish events when a detection occurs. The agent is also configured to forward events directly to the aggregation endpoints configured during authentication. The aggregation agent implements a streaming pipeline designed to process data transmitted by mapping agents when clients upload events.

4) Environment: We setup a test environment to assess the operation and performance of the Nexagon protocol. We used two virtual machines (VM), which were configured identically, each equipped with 4 GB of RAM and 2 CPU cores. This setup allowed us to run tests to mimic events from edge devices with and without the proposed security extensions. In the first VM, we ran a composite deployment of the entire management plane, with each component deployed as a microservice. The second VM hosted the client process, which was engineered to publish events that mimicked requests from mobile edge device. A comprehensive documentation of our architecture and codebase can be accessed in our repository [32] upon request, providing further insight into our design and implementation choices.

B. Performance Analysis

In this section, our evaluations are presented together with the key results obtained during our mock tests.

1) Evaluation Methodology: We conducted load tests with varying numbers of user reqests made to compare the performance impact of our authentication extension. Initial tests established a baseline by processing client requests using a pre-shared key. Subsequent tests employed using our authentication mechanism as outlined previously. This approach enabled us to quantify the performance overhead introduced by the security enhancements. A summary of the key metrics captured are presented in Table III. In addition, Table IV provides a detailed breakdown of response times and their percentile distributions.

2) Metrics: The results indicate a latency increase of 10% to 25%. we also observed a slight decrease in throughput that ranged from 3% to 7%. These results are well within the acceptable limits for a real-world deployment.

TABLE III
SUMMARY RESULTS FOR KEY METRICS AND CONFIDENCE INTERVALS

	Without Extension	**With Extension**
Average Latency (ms)	306 (±73)	384 (±45)
Throughput (req/sec)	260 (±19)	250 (±10)
CPU Utilization (%)	42 (±6)	57 (±3)

TABLE IV
DISTRIBUTION OF LATENCY AT DIFFERENT PERCENTILES

	Latency Percentiles (milliseconds)	
	Without Extension	**With Extension**
50th Percentile	276	330
80th Percentile	290	373
90th Percentile	330	416
95th Percentile	400	460

VI. CONCLUSION

Aiming at a safe, secure and efficient solution for crowdsourcing geo-referenced data in intelligent transportation systems, we delivered a detailed vulnerability analysis, mitigation strategies and a set of security extensions for IETF's Nexagon protocol. Our prototype implementation takes advantage of concepts from the addressing capabilities of LISP and the indexing structure of H3. Our results show that a PKI with pseudo-random ephemeral certificates and identifiers delivers a robust solution that protects user privacy and the integrity of the Nexagon network without compromising latency and throughput. Future research could explore the deployment of the Nexagon protocol in real-world vehicular networks to evaluate its performance under dynamic conditions. In addition, advanced mechanisms such as federated learning for decentralized and privacy-preserving data processing can also be explored to further improve the feature set of the protocol and inform future studies.

REFERENCES

[1] L. Zhu, F. R. Yu, Y. Wang, B. Ning, and T. Tang, "Big data analytics in intelligent transportation systems: A survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 1, pp. 383–398, 2019.

[2] J. Liu, H. Shen, H. S. Narman, W. Chung, and Z. Lin, "A survey of mobile crowdsensing techniques: A critical component for the internet of things," *ACM Trans. Cyber-Phys. Syst.*, vol. 2, no. 2, Jun. 2018.

[3] J. Dold and J. Groopman, "The future of geospatial intelligence," *Geo-spatial Information Science*, vol. 20, no. 4, pp. 151–162, 2017.

[4] G. Abdelkader, K. Elgazzar, and A. Khamis, "Connected vehicles: Technology review, state of the art, challenges and opportunities," *Sensors*, vol. 21, no. 22, 2021.

[5] D. Hahn, A. Munir, and V. Behzadan, "Security and privacy issues in intelligent transportation systems: Classification and challenges," *IEEE Intelligent Transportation Systems Magazine*, vol. 13, no. 1, pp. 181–196, 2021.

[6] IETF Working Group, *draft-ietf-lisp-nexagon-54*, Internet Engineering Task Force, Sep. 2024.

[7] A. Rodriguez-Natal, M. Portoles-Comeras, V. Ermagan, *et al.*, "Lisp: A southbound sdn protocol?" *IEEE Communications Magazine*, vol. 53, no. 7, pp. 201–207, 2015.

[8] A. Palia and R. Tandon, "Optimizing noise level for perturbing geo-location data," *Advances in Intelligent Systems and Computing*, pp. 63–73, 2018.

[9] L. Sweeney, "K-anonimity: A model for protecting privacy," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, no. 05, pp. 557–570, 2002.

[10] C. Dwork, "Differential privacy," in *Proceedings of the 33rd International Conference on Automata, Languages and Programming - Volume Part II*, ser. ICALP'06, Venice, Italy: Springer-Verlag, 2006, pp. 1–12, ISBN: 3540359079. DOI: 10.1007/11787006_1. [Online]. Available: https://doi.org/10.1007/11787006_1.

[11] A. Machanavajjhala, X. He, and M. Hay, "Differential privacy in the wild: A tutorial on current practices & open challenges," *Proceedings of the VLDB Endowment*, vol. 9, pp. 1611–1614, Sep. 2016.

[12] L. Xiao, D. Jiang, D. Xu, W. Su, N. An, and D. Wang, "Secure mobile crowdsensing based on deep learning," *China Communications*, vol. 15, no. 10, pp. 1–11, 2018.

[13] J. Ni, K. Zhang, Q. Xia, X. Lin, and X. S. Shen, "Enabling strong privacy preservation and accurate task allocation for mobile crowdsensing," *IEEE Transactions on Mobile Computing*, vol. 19, no. 6, pp. 1317–1331, 2020.

[14] A. J. Perez and S. Zeadally, "Secure and privacy-preserving crowdsensing using smart contracts: Issues and solutions," *Computer Science Review*, vol. 43, p. 100 450, 2022.

[15] J. Xiong, R. Ma, L. Chen, Y. Tian, L. Lin, and B. Jin, "Achieving incentive, security, and scalable privacy protection in mobile crowdsensing services," *Wireless Communications and Mobile Computing*, vol. 2018, no. 1, p. 8 959 635, 2018.

[16] V. Uher, P. Gajdoš, V. Snášel, Y.-C. Lai, and M. Radecký, "Hierarchical hexagonal clustering and indexing," *Symmetry*, vol. 11, no. 6, 2019.

[17] H. Documentation. "H3 documentation." Accessed: 2024-08-27. (2024), [Online]. Available: https://h3geo.org/docs/.

[18] R. Scandariato, K. Wuyts, and W. Joosen, "A descriptive study of microsoft's threat modeling technique," *Requirements Engineering*, vol. 20, Jun. 2013.

[19] M. Deng, K. Wuyts, R. Scandariato, B. Preneel, and W. Joosen, "A privacy threat analysis framework: Supporting the elicitation and fulfillment of privacy requirements," *Re-*

quirements Engineering, vol. 16, pp. 3–32, 2011. [Online]. Available: https://api.semanticscholar.org/CorpusID:856424.

[20] M. J. Freedman and R. Morris, "Tarzan: A peer-to-peer anonymizing network layer," in *Proceedings of the 9th ACM Conference on Computer and Communications Security*, ser. CCS '02, Washington, DC, USA: Association for Computing Machinery, 2002, pp. 193–206, ISBN: 1581136129. DOI: 10.1145/586110.586137. [Online]. Available: https://doi.org/10.1145/586110.586137.

[21] S. Chakravarty, *Traffic analysis attacks and defenses in low latency anonymous communication*. Columbia University, 2014.

[22] C.-Y. Chow, "Spatial cloaking algorithms for location privacy," *Encyclopedia of Geographical Information Science, Springer, USA*, 2008.

[23] A. Haj-Hassan, Y. Imine, A. Gallais, and B. Quoitin, "Zero-touch mutual authentication scheme for 6tisch industrial iot networks," in *2022 International Wireless Communications and Mobile Computing (IWCMC)*, 2022, pp. 354–359. DOI: 10.1109/IWCMC55113.2022.9824568.

[24] M. Achemlal, S. Gharout, and C. Gaber, "Trusted platform module as an enabler for security in cloud computing," pp. 1–6, 2011.

[25] H. Raj, S. Saroiu, A. Wolman, *et al.*, "{Ftpm}: A {software-only} implementation of a {tpm} chip," in *25th USENIX Security Symposium (USENIX Security 16)*, 2016, pp. 841–856.

[26] T. Muhammad, "Overlay network technologies in sdn: Evaluating performance and scalability of vxlan and geneve," *International Journal of Computer Science and Technology (IJCST)*, vol. 5, no. 1, pp. 39–75, 2021.

[27] S. Fahmy and M. Kwon, "Characterizing overlay multicast networks and their costs," *IEEE/ACM Transactions on Networking (TON)*, vol. 15, pp. 373–386, 2007.

[28] K. Jeziorski, D. McFadden, and R. Lennon, "Verification of certificate authorities (cas) and integration with cloud providers for enhanced security," in *2023 Cyber Research Conference - Ireland (Cyber-RCI)*, 2023, pp. 1–8. DOI: 10.1109/Cyber-RCI59474.2023.10671564.

[29] M. Schukat and P. Cortijo, "Public key infrastructures and digital certificates for the internet of things," pp. 1–5, 2015.

[30] F. Haidar, M. Makassikis, M. Sall, H. Bakhti, A. Kaiser, and B. Lonc, "Experimentation and assessment of pseudonym certificate management and misbehavior detection in c-its," *IEEE Open Journal of Intelligent Transportation Systems*, vol. 2, pp. 128–139, 2021. DOI: 10.1109/OJITS.2021.3085366.

[31] S. Rong-Hua, "An efficient secure group signature scheme," in *2002 IEEE Region 10 Conference on Computers, Communications, Control and Power Engineering. TENCOM '02. Proceedings.*, vol. 1, 2002, 109–112 vol.1. DOI: 10.1109/TENCON.2002.1181226.

[32] O. A. Obadofin, *Nexagon*, https://github.com/okemawo/Nexagon, GitHub repository, 2024.