

RHyTHM: A Randomized Hybrid Scheme To Hide in the Mobile Crowd

Mohammad Khodaei, Andreas Messing, and Panos Papadimitratos
Networked Systems Security Group, KTH Royal Institute of Technology, Stockholm, Sweden
{*khodaei, amessing, papadim*}@kth.se

Abstract—Any on-demand pseudonym acquisition strategy is problematic should the connectivity to the credential management infrastructure be intermittent. If a vehicle runs out of pseudonyms with no connectivity to refill its pseudonym pool, one solution is the *on-the-fly* generation of pseudonyms, e.g., leveraging anonymous authentication. However, such a vehicle would stand out in the crowd: one can simply distinguish pseudonyms, thus signed messages, based on the pseudonym issuer signature, link them and track the vehicle. To address this challenge, we propose a randomized hybrid scheme, RHyTHM, to enable vehicles to remain operational when disconnected without compromising privacy: vehicles with valid pseudonyms help others to enhance their privacy by randomly joining them in using *on-the-fly self-certified* pseudonyms along with aligned lifetimes. This way, the privacy of disconnected users is enhanced with a reasonable computational overhead.

Index Terms—Security, Privacy, Vehicular PKI

I. INTRODUCTION

In Vehicular Communication (VC) systems, vehicles beacon Cooperative Awareness Messages (CAMs) and Decentralized Environmental Notification Messages (DENMs) periodically at a high rate in order to provide cooperative awareness. Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I) (V2X) communication is protected with the help of public key cryptography: a set of short-term anonymous credentials, i.e., *pseudonyms*, are provided to each vehicle by the Vehicular Public-Key Infrastructure (VPKI), e.g., [1]. Thus, vehicles switch from one pseudonym to another for message unlinkability as pseudonyms are inherently unlinkable.

One can provide vehicles with valid pseudonyms for a long period, e.g., 25 years [2]. However, extensive preloading with millions of pseudonyms per vehicle for such a long period is computationally costly, inefficient in utilization and cumbersome in revocation [3]. On the contrary, several proposals suggest more frequent Vehicle-to-VPKI interactions, namely *on-demand* schemes, e.g., [1], [4]. This strategy is more efficient in terms of pseudonym utilization and revocation and more effective in fending off misbehavior. However, the more frequent the interaction with the VPKI, the more dependent vehicles are on connectivity. This may hurt Vehicle-to-VPKI connectivity on intermittent coverage of sparsely-deployed Roadside Units (RSUs), or highly overloaded existing cellular infrastructure. Thus, it is necessary to ensure that any vehicle at any time can continue its operation securely without harming privacy, even if the VPKI is not reachable or available for other reasons, e.g., during a Denial of Service (DoS) attack [1].

Obviously, signing CAMs with the private keys corresponding to expired pseudonyms or the long-term certificate, is insecure and harm user privacy, i.e., messages are trivially linkable. *On-the-fly* generation of pseudonyms, the *hybrid* scheme [5], using other anonymous authentication primitives, i.e., group signatures [6], is a promising alternative. Each vehicle is equipped with a group public key, common among all the group members, along with a distinct group signing key. In order to generate on-the-fly pseudonyms, each vehicle generates a pair of public/private keys and signs the public key using the group signing key instead of having a pseudonym signed by the corresponding Certification Authority (CA). This essentially eliminates the need to request pseudonyms from the VPKI entities, especially when the latter is unreachable. This provides authenticity, integrity, accountability, and non-repudiation. Furthermore, a node can be evicted from the system if it deviated from the system security policy.

If only a few vehicles use their self-certified pseudonyms while the rest of the vehicles rely on the VPKI-provided pseudonyms, the *baseline* scheme, they would “*stand out in a crowd*”: one can simply distinguish the pseudonyms, thus the pseudonymously signed messages, based on the pseudonym issuer’s signature. Moreover, the self-certified pseudonyms lifetimes are not aligned with each other and the global system time, i.e., the VPKI clock. As a result, all the vehicles in a region will be transmitting under pseudonyms which are distinguishable based on their timing information [1].

To address this challenge, we propose a cooperative and adaptive scheme, RHyTHM, to mitigate this privacy issue: a vehicle with no valid VPKI-provided pseudonyms initiates RHyTHM protocol by setting a flag in the upcoming CAMs. Neighboring vehicles with VPKI-provided pseudonyms randomly opt in to utilize their self-certified pseudonyms with the probability of r in upcoming pseudonym updates. RHyTHM enhances the privacy of users running out of pseudonyms at the cost of reasonable processing overhead for neighboring vehicles. This ensures the operation of every legitimate vehicle without harming user privacy even if the infrastructure fails to provide them credentials.

In the rest of the paper, we describe our system and adversarial model (Sec. II), present our scheme (Sec. III) and security and privacy analysis (Sec. IV). We provide the performance evaluation of our scheme (Sec. V) before conclusion (Sec. VI).

II. SYSTEM AND ADVERSARIAL MODEL

We assume a VPKI with distinct entities and roles [1]: the Long Term CA (LTCA) is responsible for vehicles registration in a domain [3]; the Pseudonym CA (PCA) issues pseudonyms for the registered vehicles; and the Resolution Authority (RA) is able to initiate a process to resolve a pseudonym of a misbehaving vehicle. Furthermore, a Group Manager (GM) enables any legitimate vehicle to sign a message on behalf of the group without disclosing its actual identity. Upon registration of a vehicle by the LTCA in the bootstrapping phase, each vehicle is provided with an *anonymous ticket*, with which the GM registers the vehicle, thus authorizes it to anonymously operate in some circumstances, e.g., the VPKI is unreachable.

We consider external and internal adversaries that try to harm or abuse RHyTHM. External adversaries could sign messages with fake private keys. Internal adversaries could initiate RHyTHM protocol continuously for two purposes: (i) to be provided with multiple simultaneously valid pseudonyms, thus performing Sybil-based [7] attacks; (ii) to compromise the availability of neighboring vehicles by incurring extra workload towards DoS attacks. Moreover, a global observer, e.g., an *honest-but-curious* VPKI entity [1], might be tempted to link the VPKI-provided pseudonyms to the self-certified ones to infer user sensitive information towards harming user privacy.

III. RHYTHM OPERATION

In order to achieve full unlinkability, we assume that a universally fixed interval, Γ , is specified in a region [3] and all pseudonyms in that region are issued with the lifetime aligned with the global system time, i.e., the VPKI clock. As a result of this policy, at any point in time, all the vehicles transmit using pseudonyms indistinguishable, from one another, thanks to this time alignment. This essentially eliminates any distinction among pseudonym sets of different vehicles, thus achieving user privacy protection. We refer readers to [1] for further details. The On-Board Unit (OBU) “decides” when to trigger the pseudonym acquisition process based on various parameters [8]. This can happen even within the lifetime of the last single valid pseudonym should the connectivity to the VPKI entities be reliable. However, if the VPKI entities are out of reach for any reason, the OBU initiates the RHyTHM protocol to use its self-certified pseudonyms during the next pseudonym update. If the OBU has no valid VPKI-provided pseudonyms, it initiates RHyTHM protocol with its self-certified pseudonym. Table I summarizes notation used in the protocol.

The vehicle, V , generates multiple Elliptic Curve Digital Signature Algorithm (ECDSA) key pairs and aligns the validity intervals with the known VPKI clock (steps 1–7). The OBU does not need to be fully synchronized with the VPKI clock; it simply aligns the pseudonyms lifetimes, τ_p , in the continuation of its last valid VPKI-provided pseudonym. In case of having VPKI-provided pseudonyms from a distant past and being unable to be synchronized by any other means, the OBU aligns the self-certified pseudonyms based on the pseudonym information, piggybacked in neighbors CAMs. This eliminates any distinction among self-certified pseudonyms (signed by the

TABLE I
NOTATION USED IN THE PROTOCOLS

(K_v^i, k_v^i)	pseudonymous public/private keys, corresponding to current pseudonym
gsk_v	group signing key
t_{now}, t_s, t_e	fresh/current, starting, and ending timestamps
$Sign(key, msg)$	signing a message with private key or group signing key
$(msg)_{\sigma_v}, (msg)_{\Sigma_{k_v^i}}$	a signed message with k_v^i or gsk_v

RHYTHM Initiation Protocol

```

1: procedure RHYTHMINIT( $t_s, t_e$ )
2:   for  $i:=1$  to  $n$  do
3:     Begin
4:       Generate( $K_v^i, k_v^i$ )
5:        $\zeta \leftarrow (K_v^i, t_s^i, t_e^i)$ 
6:        $(K_v^i)_{\Sigma_{k_v^i}} \leftarrow Sign(gsk_v, \zeta)$ 
7:     End
8:      $Flag_{rhythm} \leftarrow True$ 
9:      $CAM \leftarrow \{Fields, Flag_{rhythm}, t_{now}\}$ 
10:     $(CAM)_{\sigma_{k_v^i}} \leftarrow Sign(CAM, K_v^i)$ 
11:  end procedure

```




Fig. 1. RHyTHM overview: if $b = True$, the vehicle will utilize its self-certified pseudonym; otherwise, it relies on its VPKI-provided pseudonym.

group signing key, gsk_v); moreover, the anonymity set becomes equal to the number of vehicles with self-certified pseudonyms. Finally, V signs them using the gsk_v . It then piggybacks CAMs to explicitly inform its neighbors of initializing the RHyTHM protocol for the next pseudonym update (steps 8–10). Upon reception of a RHyTHM initiation query, the neighboring vehicles check if the VPKI entities are indeed out of reach. Having had the same viewpoint on the VPKI reachability, they explicitly set the RHyTHM flag in the upcoming CAMs to inform their neighbors, thus epidemically distributing the message. This ensures the distribution of the RHyTHM initiation query.

Fig. 1 illustrates five vehicles, out of which V_5 runs out of pseudonyms. It initiates the RHyTHM protocol by setting the RHyTHM flag in the upcoming CAMs. Neighboring vehicles, i.e., V_1 – V_4 , randomly opt in to utilize their self-certified pseudonyms with probability r in the first pseudonym update. V_2 and V_3 “decide” to switch to utilize their self-certified pseudonyms, thus, they generate a pair of keys, align the validity interval with the global system time, and sign them with gsk . For the second pseudonym update in Γ , only V_3 “opted in” to use its VPKI-provided pseudonym while the rest of vehicles “decided” to utilize their self-certified pseudonyms. V_5 is the only vehicle that uses its self-certified pseudonyms during the entire Γ period while other vehicles randomly opt in to use either of the two. As vehicles randomly switch between the two sets, it is hard to link two pseudonyms to the same vehicle, or identify a vehicle that uses solely self-certified pseudonyms within a Γ period. Once access to the VPKI entities is restored, V_5 refills its pseudonym pool; however, the user privacy is enhanced if it keeps switching between the two sets. In other words, if a vehicle solely relies on its VPKI-provided pseudonyms, the probability of linking two successive pseudonyms, belonging to itself, will be increased.

The exact threshold for how far to distribute the RHyTHM initiation query depends on different factors, e.g., the number of nearby VPKI-disconnected vehicles. The more vehicles without valid VPKI-provided pseudonyms, the less is the needed support from the rest of vehicles. Clearly, initiating the RHyTHM protocol with a high probability of r to switch to self-certified pseudonyms and assist few vehicles is inefficient: it imposes extra overhead on the entire system. However, to enhance the privacy of a few users, it is sufficient to receive a small “contribution” from other vehicles (it becomes clear later). Moreover, if the number of disconnected nodes without valid pseudonyms is much higher than the number of nodes with VPKI-provided pseudonyms, all the nodes “should” switch to self-certified pseudonyms, issued with aligned lifetimes. Dynamically determining an optimal r remains as future work.

IV. SECURITY & PRIVACY ANALYSIS

Non-repudiation, authentication and integrity: The RHyTHM initiation is signed by a currently valid pseudonym, thus we achieve authentication and integrity. Digital signatures and pseudonyms ensure non-repudiation, thus, each entity can be held accountable for its actions.

Thwarting Sybil-based misbehavior: An internal adversary could be equipped with two valid pseudonyms when RHyTHM is active. We rely on the *Hardware Security Module* to ensure that all outgoing signatures are signed under one private key of a single valid (VPKI-provided or self-certified) pseudonym. To mitigate generation of multiple self-certified pseudonyms, one can employ group signature schemes with such a feature [5].

Revocation: If a vehicle deviates from the security policies, it will be evicted from the system based on the underlying VPKI operations. More precisely, the RA interacts with the PCA, the GM, and the LTCA to resolve, and possibly revoke, a misbehaving vehicle, thus, distributing the revocation list.

Thwarting clogging DoS attack: RHyTHM initiation flag, integrated in CAMs, is epidemically broadcasted. Upon reception of a CAM with RHyTHM initiation request, if vehicles can confirm a connection to the VPKI, they simply ignore it (or choose a low value of r). Moreover, RHyTHM only lasts while the VPKI entities are out of reach, i.e., vehicles switch back to utilizing their VPKI-provided pseudonyms at the end of Γ period (if there is no more RHyTHM initiation request).

Honest-but-curious VPKI entities: Due to the separation of duty, no single VPKI entity is able to fully de-anonymize a user or link pseudonyms over a long period of time. RHyTHM improves privacy protection for vehicles with valid VPKI-provided pseudonyms that participate in RHyTHM: pseudonyms used for secure communication are partially linkable by the PCA and partially by the GM within a Γ interval. Communication with self-certified pseudonyms for vehicles without VPKI-provided ones is linkable by the GM.

Privacy: RHyTHM increases user privacy in compared to the baseline scheme. We consider here a suitable privacy metric: the probability of linking two (successive) pseudonyms belonging to the same vehicle. After each pseudonym changing process, an observer might be tempted to link two pseudonyms within

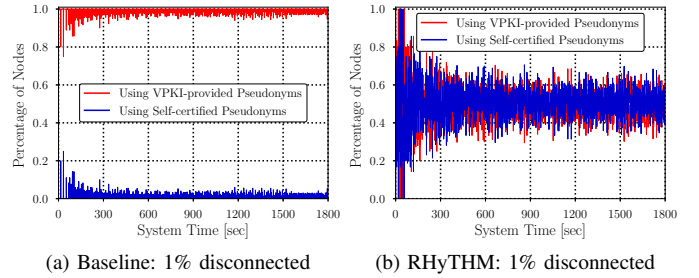


Fig. 2. 1% of nodes run out of pseudonyms ($\tau_P = 60$ sec, $r = 0.5$)

a region at a specific time window. Note that the RHyTHM initiation query is signed either by a VPKI-provided or a self-certified pseudonym(s), belonging to the disconnected node(s); thus, one can simply link that pseudonym(s) to the self-certified ones. Fig. 2 shows the percentage of nodes using their VPKI-provided or self-certified pseudonyms for an actual mobility trace (www.vehicularlab.uni.lu/), during the rush hour (7-7:30 am). Fig. 2.a illustrates that 1% of nodes cannot access the VPKI to refill their pseudonyms pool, e.g., due to sparse deployment of the RSUs. As a result, there is a huge difference between their anonymity set size, thus harming user privacy. We define the anonymity set as the set of vehicles using indistinguishable pseudonyms at any given point in time. Fig. 2.b shows how RHyTHM could enhance user privacy: nodes with valid VPKI-provided pseudonyms randomly and independently switch to utilizing their self-certified pseudonyms to help other vehicles protect their privacy. Thus, the anonymity set size of the two groups is balanced. This does not harm the privacy of users from the larger set since they change their set randomly for each pseudonym update (it becomes clear next).

Assuming there are N vehicles equipped with VPKI-provided pseudonyms and M vehicles run out of pseudonyms, thus using their self-certified pseudonyms. The probability of switching to self-certified pseudonyms is r . Using the baseline scheme, the probability of linking two VPKI-provided pseudonyms belonging to the same vehicle is $\frac{1}{N}$. However, by using RHyTHM, the probability of linking two VPKI-provided pseudonyms belonging to the same vehicle becomes $\frac{(1-r)}{N-(r \times N)} = \frac{1}{N}$. If a vehicle with a VPKI-provided pseudonym decides to utilize its self-certified pseudonym in the next pseudonym update, the probability of linking those two pseudonyms becomes $\frac{r}{M+(r \times N)} = \frac{1}{N+\frac{M}{r}}$. Since $\frac{1}{N+\frac{M}{r}} < \frac{1}{N}$, if $M > 0$, the probability of linking decreases, thus enhancing user privacy. Thereby, employing RHyTHM does not compromise the privacy of users. If a vehicle decides to utilize its self-certified pseudonym, the probability of linking decreases at the cost of extra computation overhead. Simply put, by switching back and forth between utilizing VPKI-provided and self-certified pseudonyms, the probability of linking two pseudonyms, belonging to the same vehicle, decreases exactly because an adversary cannot know which anonymity set they belong to. If a fraction of vehicles join RHyTHM, the probability of linking pseudonyms for those who always utilizes their VPKI-provided pseudonyms increases exactly because an adversary should link a VPKI-provided pseudonym to a VPKI-provided one since the probability is higher (i.e., $\frac{1}{N} > \frac{1}{N+\frac{M}{r}}$).




Fig. 3. (a) Comparing the probability of linking two pseudonyms using baseline and RHyTHM schemes ($N = 100, r = 0.2$). (b) Probability of linking two VPKI-provided pseudonyms ($N = 100, r = 0.5$).

Fig 3.a compares the probability of linking two pseudonyms using the baseline and the RHyTHM schemes: by employing RHyTHM, the probability of linking self-certified pseudonyms of vehicles that must use it significantly decreases, becomes 0.05, when $M = 1$ and $r = 0.2$, i.e., 20 vehicles switch to their self-certified pseudonyms. Moreover, the probability of linking pseudonyms of vehicles that opt in to participate in RHyTHM decreases slightly. When the majority of vehicles run out of pseudonyms, vehicles are highly encouraged to switch to use their self-certified pseudonyms in order to enhance their privacy. Vehicles with VPKI-provided pseudonyms could simply ignore RHyTHM and always use their pseudonyms. We define K ($0 \leq K \leq N$) as the number of vehicles, equipped with VPKI-provided pseudonyms, but never join the RHyTHM protocol. Fig. 3.b shows that the probability of linking two VPKI-provided pseudonyms on average, becomes:

$$Pr = \frac{K}{[K + (N - K) \times (1 - r)]^2} + \frac{N - r \times (N - K) - K}{[K + (N - K) \times (1 - r)]^2} \times (1 - r)$$

The first term is the probability of linking two successive pseudonyms belonging to a vehicle not using RHyTHM. It is the probability of the pseudonym being in K set ($\frac{K}{[K + (N - K) \times (1 - r)]}$), multiplied by the probability of linking it to its successive pseudonym ($\frac{1}{[K + (N - K) \times (1 - r)]}$). The denominator is the size of the entire VPKI-provided pseudonym set. The second term is for the rest of the vehicles using RHyTHM: the probability of a pseudonym belonging to a vehicle using RHyTHM ($\frac{N - r \times (N - K) - K}{[K + (N - K) \times (1 - r)]}$), multiplied by the probability of linking it to its successive pseudonym ($\frac{(1 - r)}{[K + (N - K) \times (1 - r)]}$).

If $K = 0$, i.e., all the vehicles use RHyTHM, or $K = N$, i.e., the baseline scheme that vehicles with valid VPKI-provided pseudonyms always use their pseudonyms, then the probability of linking, on average, becomes: $\frac{1}{N}$. Fig. 3.b illustrates that using RHyTHM increases the uncertainty as one cannot simply predict the destination set after each pseudonym update. The probability of linking two successive pseudonyms for vehicles using RHyTHM is always less than the probability of linking for vehicles always using their VPKI-provided pseudonyms.

V. PERFORMANCE EVALUATION

We emulate a large neighborhood with 7 Nexcom boxes (Dual-core 1.66 GHz, 1GB memory) from PRESERVE project (www.preserve-project.eu) to evaluate the performance of our scheme. Our implementation is in C, and we use OpenSSL and

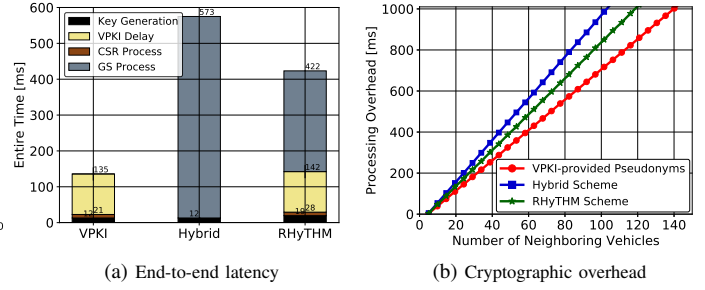


Fig. 4. (a) End-to-end latency to acquire 10 pseudonyms, averaged over 500 runs. (b) Processing overhead as a function of the neighborhood size ($\tau_P = 30$ sec, ratio of received messages: up to 60 beacon/sec, $r = 0.5$)

an implementation (github: IAIK/pairings_in_c) of short-group signature [6] with security level of 112 bits for cryptographic operations and primitives. The average signing and verification latency for group signature is 56 ms and 82.5 ms, respectively; thus, the extra computation overhead, when $r = 0.2$, for every vehicle in the system is around 1.6 sec per τ_P .

Fig. 4.a shows the end-to-end latencies for obtaining 10 pseudonyms using the baseline and the RHyTHM schemes. As the figure shows, employing RHyTHM results in 287 ms extra overhead, mainly for generating the public/private key pairs and signing them with the gsk , for the vehicles equipped with valid VPKI-provided pseudonyms. This overhead pays off as their privacy is improved compared to only using VPKI-provided pseudonyms (beyond assisting vehicles in need). As illustrated in Fig. 4.b, the total number of neighboring vehicles that an OBU could face, if all the vehicles utilize their self-certified and VPKI-provided pseudonyms, is 100 and 140, respectively. By employing RHyTHM with $r = 0.5$, an OBU could verify the signature of CAMs from one up to 120 neighbors.

VI. CONCLUSION AND FUTURE WORK

We presented RHyTHM as a privacy-preserving scheme to help vehicles operate and protect their privacy even if they run out of pseudonyms. As future work, we plan to investigate the provision of incentives for the vehicles to participants in RHyTHM and the optimal probability of switching to utilizing self-certified pseudonyms in different circumstances.

REFERENCES

- [1] M. Khodaei, *et al.*, "SECMACE: Scalable and Robust Identity and Credential Management Infrastructure in Vehicular Communication Systems," in *the IEEE TITS*, Mar. 2018. Online: <https://arxiv.org/abs/1707.05518>.
- [2] V. Kumar *et al.*, "Binary Hash Tree based Certificate Access Management for Connected Vehicles," in *ACM WiSec*, Boston, USA, July 2017.
- [3] M. Khodaei *et al.*, "The Key to Intelligent Transportation: Identity and Credential Management in Vehicular Communication Systems," *IEEE VT Magazine*, vol. 10, no. 4, pp. 63–69, Dec. 2015.
- [4] F. Schaub *et al.*, "V-tokens for Conditional Pseudonymity in VANETS," in *IEEE WCNC*, Sydney, Australia, pp. 1–6, Apr. 2010.
- [5] G. Calandriello *et al.*, "On the Performance of Secure Vehicular Communication Systems," *IEEE TDSC*, vol. 8, no. 6, pp. 898–912, Nov. 2011.
- [6] D. Boneh *et al.*, "Short Group Signatures," in *Advances in Cryptology CRYPTO*. Springer, 2004.
- [7] J. R. Douceur, "The Sybil Attack," in *ACM Peer-to-peer Systems*, London, UK, Mar. 2002.
- [8] M. Khodaei *et al.*, "Evaluating On-demand Pseudonym Acquisition Policies in Vehicular Communication Systems," in *Proceedings of the IoV-Vol*, Paderborn, Germany, pp. 7–12, July 2016.