Timing vs Covert Detours: A Benchmarking Study in Path Validation

Dorine Chagnon
Gérard Chalhoub
dorine.chagnon@doctorant.uca.fr
gerard.chalhoub@uca.fr
LIMOS, Université Clermont-Auvergne, CNRS, Mines Saint-Etienne
Aubière, France

Kevin Thiry-Atighehchi
kevin.thiry@ecole-air.fr
CRéA, École de l'Air et de l'Espace
Salon-de-Provence, France

Abstract

Path-validation protocols authenticate that packets traverse a specified sequence of routers, but they do not prove the absence of covert detours. Most designs embed timestamps in path-validation proofs and recommend checking a trusted delay against a threshold, yet there is no systematic analysis of how well such thresholding detects detours under realistic network dynamics.

In this work-in-progress, we take first steps towards a benchmarking framework for timing-based detour detection. We build a synthetic dataset of one-way delays from ns-3 simulations with configurable detour length and cross-traffic intensity, extract delay-based features at the flow level, and compare simple threshold-based rules against supervised classifiers and deep autoencoders. Our preliminary results suggest that (*i*) simple delay thresholds are fragile once queueing effects make detoured and non-detoured paths overlap in delay, and (*ii*) supervised models such as logistic regression and SVMs maintain high F1-scores across the regimes we study, while unsupervised autoencoders often outperform simple thresholds but exhibit larger performance drops in the most challenging traffic regimes.

CCS Concepts

• **Security and privacy** → *Security protocols.*

Keywords

Path Validation, Detour Attacks, Wormhole Attack, ns-3, Delay Analysis, Machine Learning

ACM Reference Format:
Dorine Chagnon, Gérard Chalhoub, and Kevin Thiry-Atighehchi. 2026. Timing vs Covert Detours: A Benchmarking Study in Path Validation. In *The 41st ACM/SIGAPP Symposium on Applied Computing (SAC '26), March 23–27, 2026, Thessaloniki, Greece.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3748522.3779984

Path-validation protocols let endpoints enforce and verify a data-plane path by requiring routers along the path to cooperatively construct a cryptographic proof that packets traversed a claimed sequence of hops, and their design space has been surveyed in recent overviews [3, 4]. These mechanisms complement control-plane protections such as RPKI and BGPsec, which validate origins and AS paths but do not bind the concrete forwarding path that packets follow [9, 12, 5, 10]. However, traffic may still experience covert detours and wormhole-like behavior: packets can be steered through tunnels or service-function chains that preserve the advertised sequence of routers while significantly changing geography, latency, or jurisdiction [7, 6, 13, 11].

Several path-validation proposals include timestamps and suggest checking end-to-end delay against a trusted bound [3, 14, 8], and distance-bounding protocols study delay-based guarantees [1, 2], but we are not aware of any empirical evaluation of timing-based detour detection under realistic queueing and cross-traffic. This work-in-progress takes first steps towards filling this gap. Our contributions are: (i) we build a synthetic dataset of one-way delays from ns-3 simulations with configurable detour length and cross-traffic intensity; (ii) we derive simple flow-level delay features and compare threshold-based rules with supervised classifiers and deep autoencoders; and (iii) we identify regimes where delay thresholds become fragile due to overlap between direct and detoured delays, while supervised models remain robust and unsupervised autoencoders behave as a promising but more variable alternative.

Preliminary Results: Detour Length Influence

This section presents the preliminary results of our analysis on one varying parameter: the detour length. The delay analysis is performed on 6 different models: 2 baseline models based on simple statistical criteria defined by Arfaoui *et al.*([2, 1]), 2 supervised models (Support Vector Machines (`SVM`) and Logistic Regression (`REGLOG`)) and 2 autoencoders (Denoising Autoencoder (`DAE`) and Variational Autoencoder (`VAE`)). Figure 1 show the performance, as F_1-scores, of the different models as a function of the detour path length with the baseline models (first line, green triangles), the supervised models (second line, red squares), and the autoencoders (third line, blue dots). The F_1-scores are averaged over multiple datasets, with grey confidence intervals. Delay-timestamp plots, in Figure 2, display direct-flow delays (blue crosses) and detour-flow delays (red crosses) for a representative dataset, illustrating typical delay patterns under each scenario. The simulation takes place in a network with a direct path length of 4 routers and a parasite rate of 60 Mbps. 10% of the flows are deviated towards the detour path, the rest takes the direct path.

SAC '26, Thessaloniki, Greece

ACM ISBN 979-8-4007-2294-3/2026/03
https://doi.org/10.1145/3748522.3779984

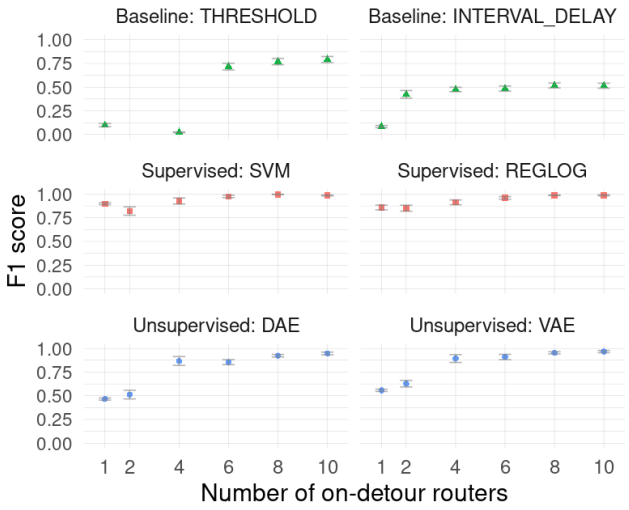

Figure 1: F_1-Score of models as a function of the detour length with a path length of 4 routers and a parasite rate of 60 Mbps.

Globally, models perform better with a longer detour path. It is intuitive that the greater the detour, the greater the difference between direct and detour delays. This separation is present in Figure 2d for a detour of 10 routers. The baseline models do not manage to make good classification for detour length of 1 and the THRESHOLD also struggle for detour length of 2 and 4. THRESHOLD does not manage to classify a single detour flow as deviated and its F_1-Score is null for a detour of 2 routers. Figure 2 shows that for detours of length 1, 2 or 4 routers, the delays overlap, which is why the baseline models find it difficult to classify them. Supervised models have good performance with F_1-Scores above 0.85 for every detour length. They are both better than the baseline models. Regarding unsupervised models, the DAE and the VAE have better performance than the baseline models except for a detour length of 2 routers where the INTERVAL_DELAY model have a greater score than DAE. The difference between the direct and the detour delays increase with each on-detour router added and the classification by models becomes easier.

Conclusion and perspectives

In this poster paper, we presented an evaluation methodology that allows to benchmark detour attacks detection method. We evaluated detection methods that rely on the end-to-end delay in order to distinguish between traffic that was detoured and traffic that was not. These methods are based on supervised and unsupervised machine learning models, in addition, we included two non-ML detection methods as baselines. We relied on the famous ns-3 simulator to generate datasets that represent a wide range of scenarios by varying the detour length.Results in terms of F_1-Score show the superiority of ML-based models compared to non ML methods. Also, supervised models yielded better overall results than unsupervised models.

In this study we did not consider the case where traffic can travel over multiple paths. Also, we did not take into account the fact

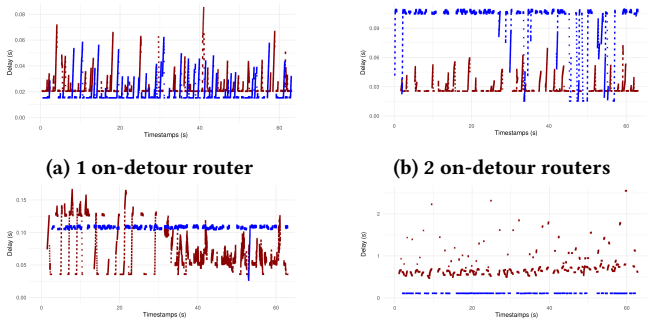

(a) 1 on-detour router **(b) 2 on-detour routers**

(c) 4 on-detour routers **(d) 10 on-detour routers**

Figure 2: Evolution of the variation of delays as a function of the detour length for a path length of 4 routers and a parasite rate of 60 Mbps.

that the propagation duration between two routers depends on the length of the link. Many other aspects will also be studied in our future work including the impact of the path length or the network congestion on the efficiency of the detection methods and the robustness of the methods when faced with previously unseen delay values.

Acknowledgments

The authors acknowledge the support of the Chaire de confiance numérique (12LIMO11LIMOS on FCA) of the Clermont Auvergne University Foundation.

References

[1] Ghada Arfaoui, Gildas Avoine, Olivier Gimenez, and Jacques Traoré. 2021. How distance-bounding can detect internet traffic hijacking. In *Cryptology and Network Security, CANS 2021, Proceedings.*

[2] Ghada Arfaoui, Gildas Avoine, Olivier Gimenez, and Jacques Traoré. 2022. ICRP: internet-friendly cryptographic relay-detection protocol. *Cryptogr.*

[3] Kai Bu, Avery Laird, Yutian Yang, Linfeng Cheng, Jiaqing Luo, Yingjiu Li, and Kui Ren. 2020. Unveiling the Mystery of Internet Packet Forwarding: A Survey of Network Path Validation. *ACM Comput. Surv.*

[4] Dorine Chagnon, Kevin Thiry-Atighehchi, and Gérard Chalhoub. 2025. A survey on path validation: towards digital sovereignty. *Computer Networks.*

[5] Taejoong Chung et al. 2019. RPKI is coming of age: a longitudinal study of RPKI deployment and invalid route origins. In *Proceedings of the Internet Measurement Conference.*

[6] Benoit Donnet, Matthew Luckie, Pascal Mérindol, and Jean-Jacques Pansiot. 2012. Revealing MPLS tunnels obscured from traceroute. *ACM SIGCOMM Computer Communication Review.*

[7] Anne Edmundson, Roya Ensafi, Nick Feamster, and Jennifer Rexford. 2016. Characterizing and avoiding routing detours through surveillance states. (2016).

[8] Anxiao He, Jiandong Fu, Kai Bu, Ruiqi Zhou, Chenlu Miao, and Kui Ren. 2024. Symphony: Path Validation at Scale. *Network and Distributed System Security Symposium 2024.*

[9] Matt Lepinski and Stephen Kent. 2012. An Infrastructure to Support Secure Internet Routing. RFC 6480. (2012).

[10] Qi Li, Jiajia Liu, Yih-Chun Hu, Miao Xu, and Jianping Wu. 2019. Bgp with bgpsec: attacks and countermeasures. *IEEE Network.*

[11] Paul Quinn and Uri Elzur. 2018. Network service header (nsh). RFC 8300. (2018).

[12] Job Snijders, Ben Maddison, Matt Lepinski, Derrick Kong, and Stephen Kent. 2024. A Profile for Route Origin Authorizations (ROAs). RFC 9582. (2024).

[13] Yves Vanaubel, Jean-Romain Luttringer, Pascal Mérindol, Jean-Jacques Pansiot, and Benoit Donnet. 2019. TNT, Watch me Explode: a light in the dark for revealing MPLS tunnels. In *IFIP/IEEE Traffic Measurement and Analysis (TMA).*

[14] Fan Yang, Ke Xu, Qi Li, Rongxing Lu, Bo Wu, Tong Zhang, Yi Zhao, and Meng Shen. 2020. I Know If the Journey Changes: Flexible Source and Path Validation. In *2020 IEEE/ACM 28th International Symposium on Quality of Service (IWQoS).*