Institutional Repository - Research Portal
Dépôt Institutionnel - Portail de la Recherche
researchportal.unamur.be
University of Namur

RESEARCH OUTPUTS / RÉSULTATS DE RECHERCHE

Merging Traffic to Improve Privacy and Performance

Dejaeghere, Jules

Publication date:
2024

Document Version
Publisher's PDF, also known as Version of record

Link to publication

Citation for pulished version (HARVARD):
Dejaeghere, J 2024, 'Merging Traffic to Improve Privacy and Performance', 2024 Cyberwal in Galaxia, Redu, Belgium, 2/12/24 - 6/12/24.

General rights
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy
If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Download date: 27. Apr. 2025

Merging Traffic to Improve Privacy and Performance

Jules Dejaeghere

University of Namur

Research setting

Let's assume we have a **distributed network** of machines **operated by volunteers**.

- Many different operators (e.g., NPO, individuals, universities)
- No central administrator
- Many different physical and network locations

How can we leverage such distributed network to provide privacy enhancing services?

Examples of privacy enhancing services: anonymous web browsing, anonymous asynchronous messaging.

Existing works already studied the privacy impact of blending different latency traffic [1, 2].

Update framework for distributed networks

We recently suggested a framework to securely update members of a distributed network [3]. The framework has the following features:

- Expressiveness
- Low overhead compared to native
- Zero downtime
- No human intervention required
- Authenticity checked
- Safe execution
- Fast distribution

Current limitations

Most low-latency privacy enhancing services are vulnerable to traffic analysis. An external attacker can use the timing and size of the exchanged data to infer the activity of the en user, as in Figure 1.

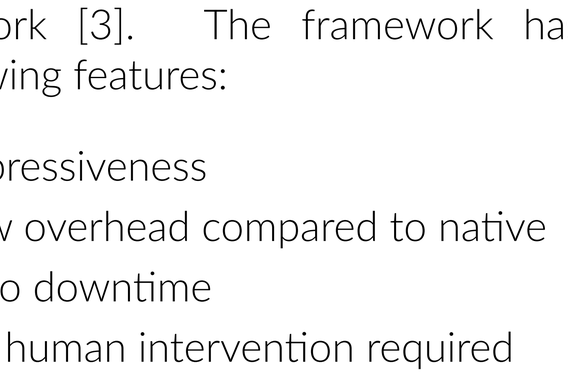

Figure 1. Encrypted traffic can be fingerprinted, based on size and time.

CyberExcellence project, grant 2110186 funded by the SPW Recherche

Traffic analysis: deanonymizing encrypted exchanges online

By observing the **size** and the **timing** of a data flow, an attacker could be able to infer what the user is doing. **Different online activities** are likely to generate **different traffic patterns** (e.g., streaming a video versus checking a webmail service).

Figure 2 depicts a few common countermeasures against traffic analysis.

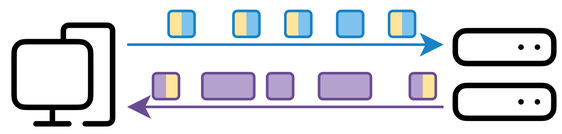

(a) Using padding to hide traffic patterns.

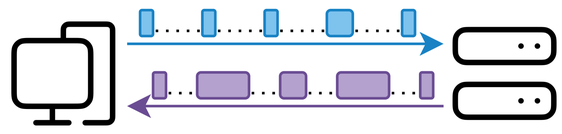

(b) Using delay to hide traffic patterns.

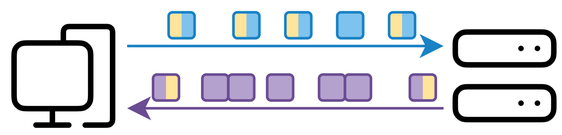

(c) Using fixed-size sending unit to hide traffic patterns.

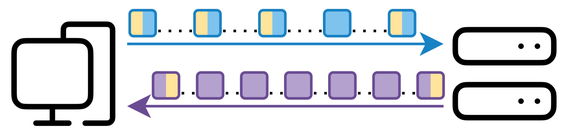

(d) Using the techniques a, b, and c to hide traffic patterns.

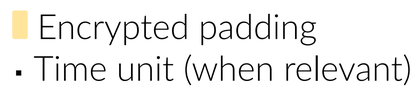

Figure 2. Common techniques to prevent traffic analysis.

Downsides of traffic analysis defenses

- Add dummy packet on the network link
- Delay sending some legitimate packets

Merging traffic types

Multiple applications could have their **traffic merged together** to make traffic analysis harder for an adversary and optimize the use of resources. Not all applications require low-latency from the network:

- Anonymous remailer with Mixminion [4] (high-latency)
- Web browsing with Tor [5] (low-latency)

Figure 3 depicts an example of merging different traffic types.

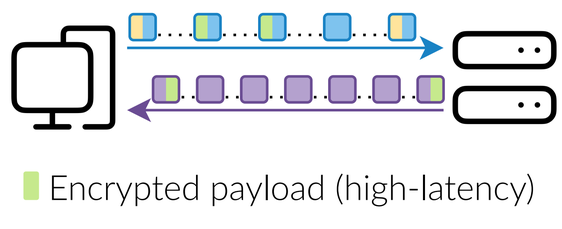

Figure 3. Merging different traffic types.

Tor as a physical infrastructure for many applications

The Tor Project already operates a distributed network of volunteer-operated relays. The relays are distributed around the globe, as show in Figure 4.

Such network could benefit from our framework [3] to reprogram the nodes. This way, nodes could support multiple applications and multiplex their traffic.

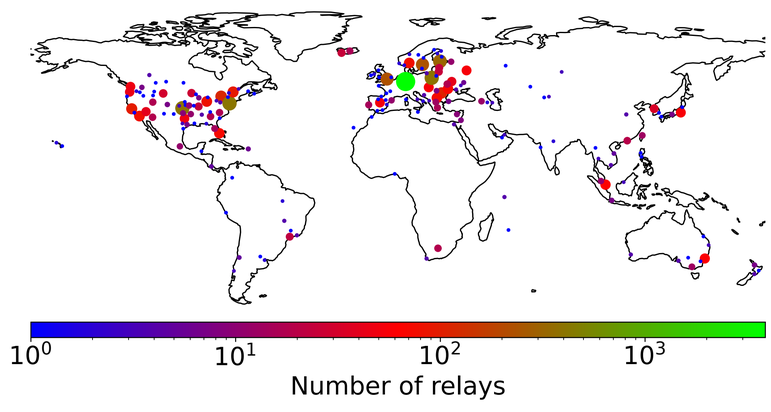

Figure 4. Location of the 8135 Tor relays on 2024-10-29.

Key benefits

- Use the full capacity of the network link
- Require less padding
- Blend multiple traffic types together

Open questions

- How to schedule the different types of traffic?
- What network parameters should we use (padding, sending size)?
- How to safely enable third-party applications on the network?
- What governance model should we adopt?

References

[1] Iness Ben Guirat, Debajyoti Das, and Claudia Diaz. Blending Different Latency Traffic With Beta Mixing. *Proceedings on Privacy Enhancing Technologies*, 2024(2):464–478, April 2024.

[2] Roger Dingledine, Andrei Serjantov, and Paul Syverson. Blending Different Latency Traffic with Alpha-mixing. In George Danezis and Philippe Golle, editors, *Privacy Enhancing Technologies*, Lecture Notes in Computer Science, Berlin, Heidelberg, 2006. Springer.

[3] Florentin Rochet, Jules Dejaeghere, and Tariq Elahi. Towards Flexible Anonymous Networks. In *23rd Workshop on Privacy in the Electronic Society (WPES '24)*, Salt Lake City, UT, USA, October 2024. ACM.

[4] George Danezis, Roger Dingledine, and Nick Mathewson. Mixminion: design of a type III anonymous remailer protocol. In *2003 Symposium on Security and Privacy*, Berkeley, CA, USA, 2003. IEEE Computer Society.

[5] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The Second-Generation Onion Router. In *13th USENIX Security Symposium (USENIX Security 04)*, 2004.

This poster includes GeoLite2 Data created by MaxMind and vector icons by Dazzle UI on SVG Repo.