Network-Physical Layer Mapping for Enhanced Internet Resilience and Security

~ Revealing submarine cables via which traceroutes travel

Presenter: Riya Ponraj (University of Oregon)

Advisors: Ram Durairajan (University of Oregon, Link Oregon) , Yu Wang (University of Oregon), Alexander Gamero-Garrido (UC Davis) Amreesh Phokeer (Internet Society

Pulse Internet Measurement Forum, 2025
Spain

ripo@uoregon.edu

Why This Problem Matters Globally

99% of international traffic flows through submarine cables

Outages disrupt banking, cloud services, education, and national security

Critical Lack of Visibility

- Policymakers, users, and researchers **cannot see** which cables their traffic depends on
- We lack **transparent, data-driven** visibility
- Operators rarely publish accurate **physical-layer routing information**

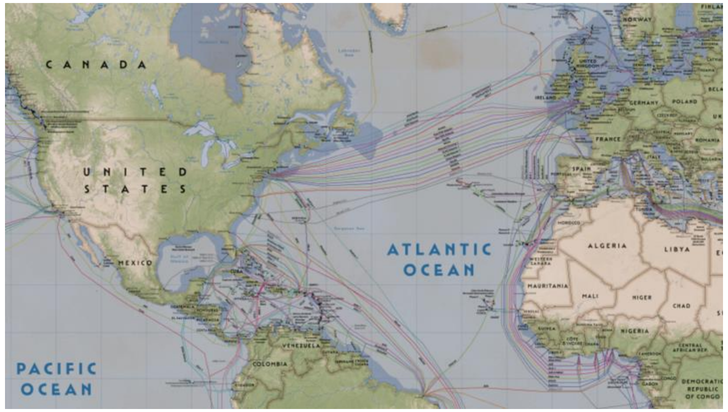

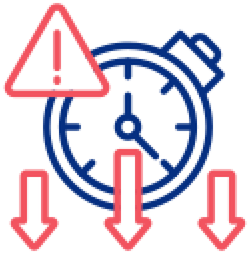

Motivation: Connecting the Logical Paths and Physical Cables

Traceroutes crossing oceans (Logical layer)

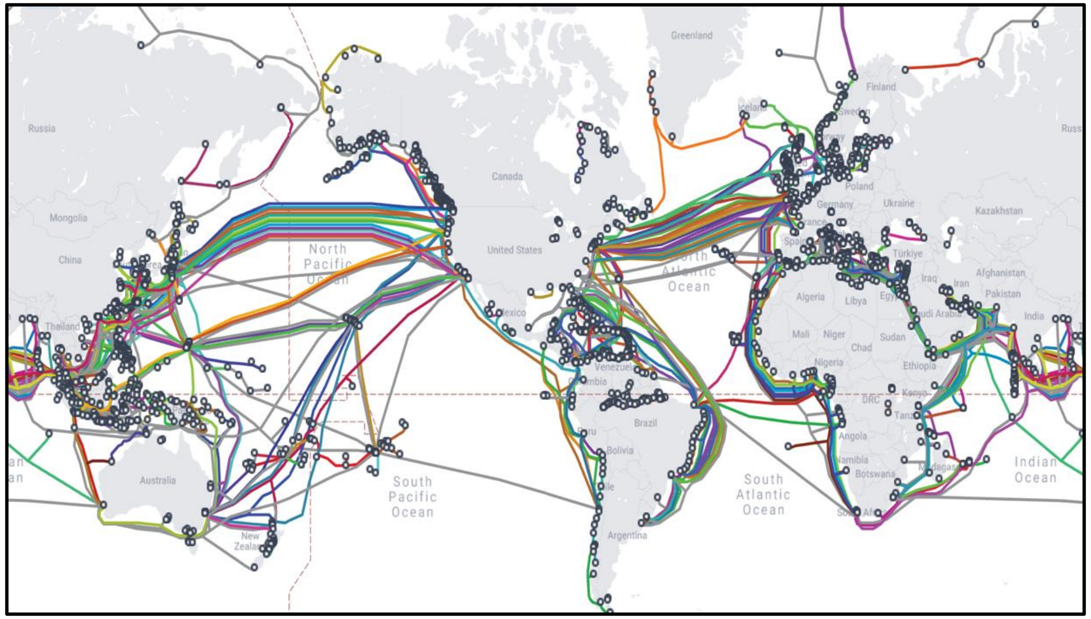

Submarine cable map (Physical Layer)

Problem:
Which submarine cable a path uses ? How traffic changes during outages?

Why?
⚡ Outages and failures
🗺️ Geopolitical dependencies
⚙️ Performance and reliability

Goal:
A principled approach for mapping traceroutes to submarine cables that goes beyond **heuristics**

Why Existing Approaches Fall Short

- **IP Geolocation**: Inaccurate and Inconsistent sometimes
- **Heuristic Inference** : Works well only in some cases
- **Operator Maps**
  - Not public and complete

InterTubes | iGDB | Nautilus | Calypso

- Rely on heuristic rules
- Fail to generalize to real-world complexity
- Struggle under multi-level noise

Our approach

Why it is hard

Traceroute gives
- ✓ IP address
- ✓ RTT
- ✓ Hop number

Can We *Learn* Submarine Cable Fingerprints directly From RTT Patterns?

But not

x Where the physical cable lies
x Which landing station is used
x When paths change due to failure

Key intuition:

- When traffic crosses an ocean, it shows **characteristic RTT spikes**.
- These patterns repeat across thousands of paths.
- A learning model can pick up these patterns on its own.

So, we need a method that learns purely from the structure of the data.

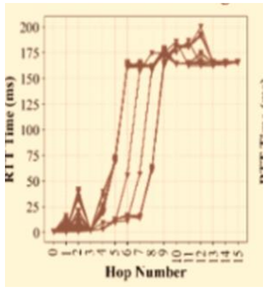

What do we do?

- Without using *any* labels:
- Learn **representations** of paths based solely on their hop-by-hop RTT behavior
- Discover **clusters** of paths that correspond to distinct cable systems
- Reveal structural similarities between oceans (e.g., Asia–Europe vs US–EU)
- Learn **embeddings** (representations) of cable segments
- Generalize across **new regions, new datasets**

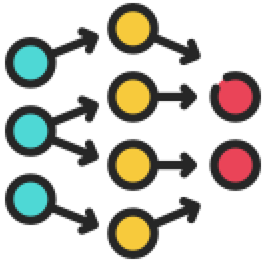

Insights

- Each cable has unique RTT pattern.
- The model learns these patterns directly from measurements
- Same cable → similar RTT signature → tight cluster
- Same corridor -> Nearby clusters
- Different cables → distinct signatures → separate clusters
- Shifts in clusters corresponded to outages, reroutes, or congestion events

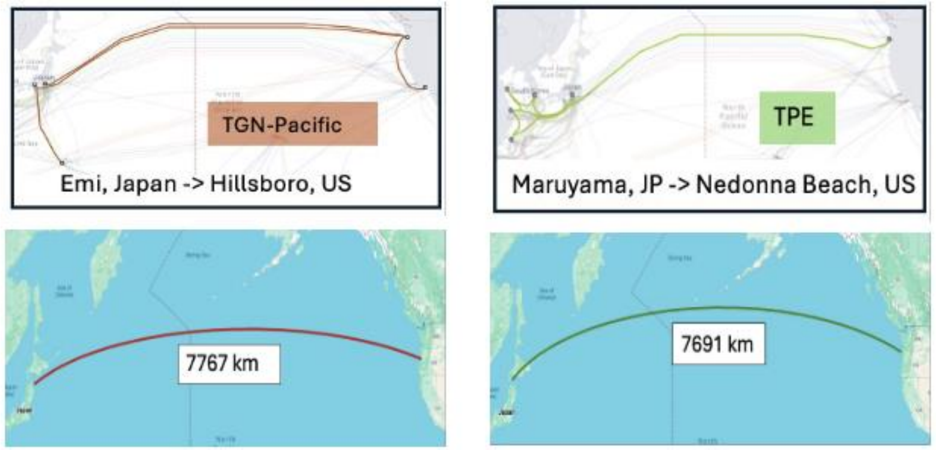

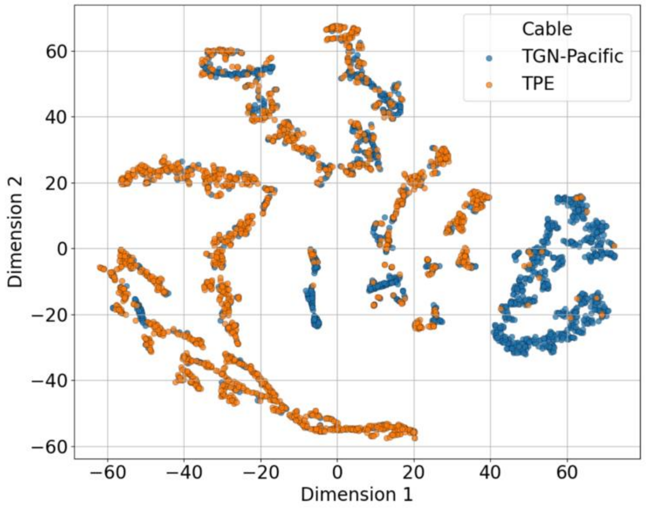

Why This Matters for Resilience

Resilience Planning

- Identify which international routes are critical for a country
- Detect single points of failure
- Understand the impact of cable outages in real time
- Improve preparedness for natural disasters or geopolitical risks

Next Steps

- Expand it to more cables.
- Modeling it for different corridors
- Analyzing different types of cables :
    - Parallel, isolated, trunk and branch
- Collaborate with operators for validation.

ripo@uoregon.edu

Conclusion

- **Shows submarine cable usage can be inferred without labels**
- **Provides transparency into physical-layer dependencies**
- **Demonstrates learning the structure of the Internet directly from the Internet itself**

Thank you!
Happy to discuss more.

Contact:
ripo@uoregon.edu