NC STATE UNIVERSITY

Server, Client, or Relay? Dual-Role Detection of Circumvention Relays

Sultan Almutairi, Khaled Harfoush, Yannis Viniotis

North Carolina State University

Motivation

- Single IP address architecture:
  - Many circumvention tools use a single relay IP.
  - The relay accepts client traffic and forwards it outward.
- Their Defense
  - Obfuscating the client–proxy link.
  - Adopting probe-resistant techniques.

What remains observable

- These defenses protect the link, not relay behavior
- Censors still see traffic generated by proxy especially flow metadata at scale.

Client and Server Traffic

We classify roles per flow

- **Client role:** Endpoint initiates a connection using an ephemeral src port.
- **Server role**: Endpoint receives the connection on a well-known dst port
- **Exclude ambiguous flows** (peer-to-peer or server-server)

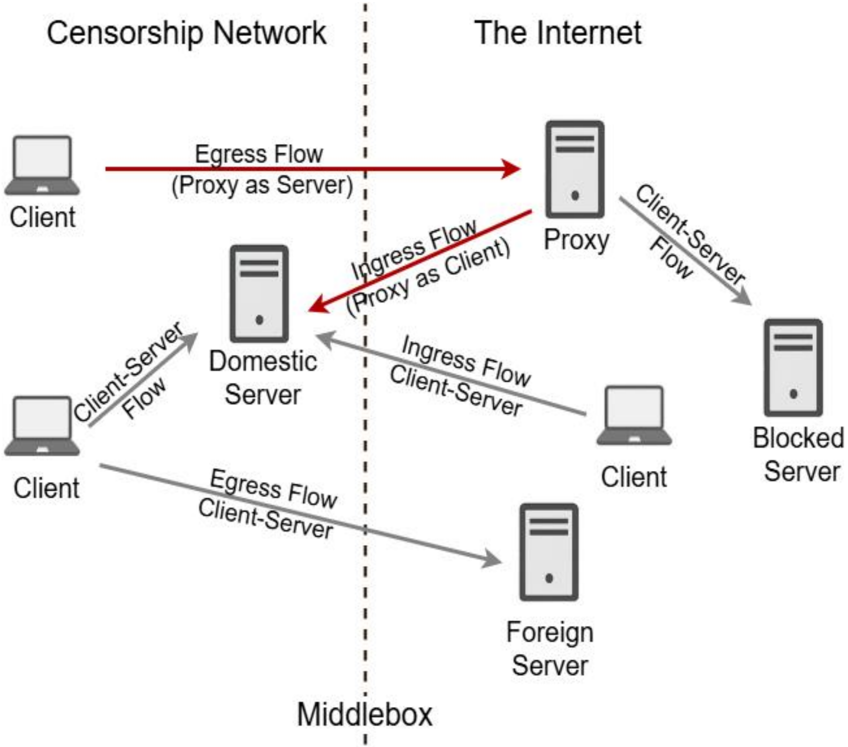

Dual-role behavioral fingerprint

- **Dual Role:**
  - **In practice:** Any endpoint exhibits persistent role.
  - **Relay** with a single IP exhibits **both roles**.

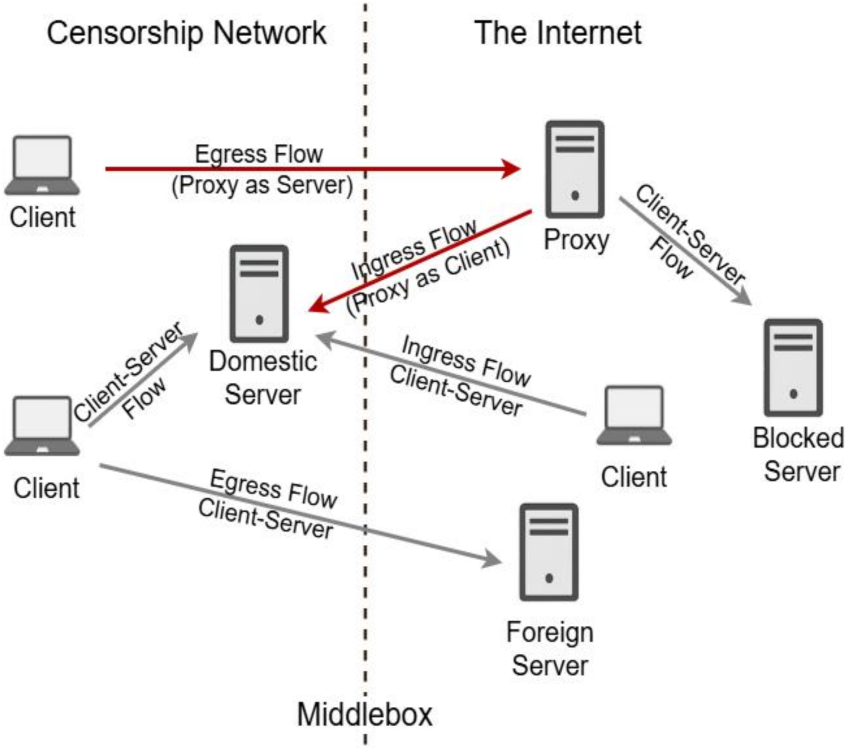

Research Question

Can the dual-role behavioral fingerprint provide a distinguishing feature that exposes relay servers even when traffic is encrypted or obfuscated?

Design

- Threat model
  - State-level monitor at scale (normal filtering)
  - Domestic vs non-domestic classification.
  - Low-cost filtering using VPS-dense ASNs.
  - DPI sees metadata (e.g domain, direction flow)

Design

- Three-stage pipeline
  - Stage 1: Candidate Selection
  - Stage 2: Dual-Role Detection
  - Stage 3: Suspicion Scoring and Classification (RSS)

Design

- **Stage 1**: Candidate Selection
  - Focus on non-domestic servers in VPS-dense ASNs
- **Stage 2**: Dual-Role Detection
  - Define Dual-Role Instance (DRI) within Observation Window (W):
    - Domestic client (c) →external server (r), then this server→destination (d) on ports 80 and 443
  - Discard candidates with zero DRIs.
- **Stage 3**: Suspicion Scoring and Classification (RSS)
  - Score relays by destination types in DRIs.
    - Weight user-facing (high) vs infrastructure domains (Low).
  - Classify relay if $\mathrm{RSS}(r) > \tau$.

Evaluation

- **Dataset**: WIDE backbone traces, 17 TB, April 9, 2025.
- **Goal**: Evaluate the dual-role detection heuristic.
- **Flow Extraction Process:**
  - Extract unidirectional 5-tuple (client_ip, server_ip, client_port, server_port, protocol)
  - Enrich endpoints with Geo-IP
- **Two types of traffic**
  - Ground-truth relays : Use TShark protocol filters (OpenVPN, WireGuard, and SOCKS) to get IP relay candidates.
  - Benign baseline (general servers):
    - Any flow has server_port on 443 as benign server

Evaluation

- Flow Classification:
  - **Foreign servers**: endpoints located outside Japan (Geo-IP)
  - **Egress Flow**: Japan-based client → foreign server
  - **Ingress Flow**: foreign client → Japan-based server
- What we test (per foreign server R)
  - R appears as a server in at least one Egress Flow
  - R later appears as a client in an Ingress Flow to Japan (dst port 80/443)
  - If both occur, R exhibits the dual-role behavioral fingerprint
- Apply the same test to both types of traffic to compute TP/FN and FP/TN.

Evaluation

Table 1: Summary of Dual-Role Detection Results

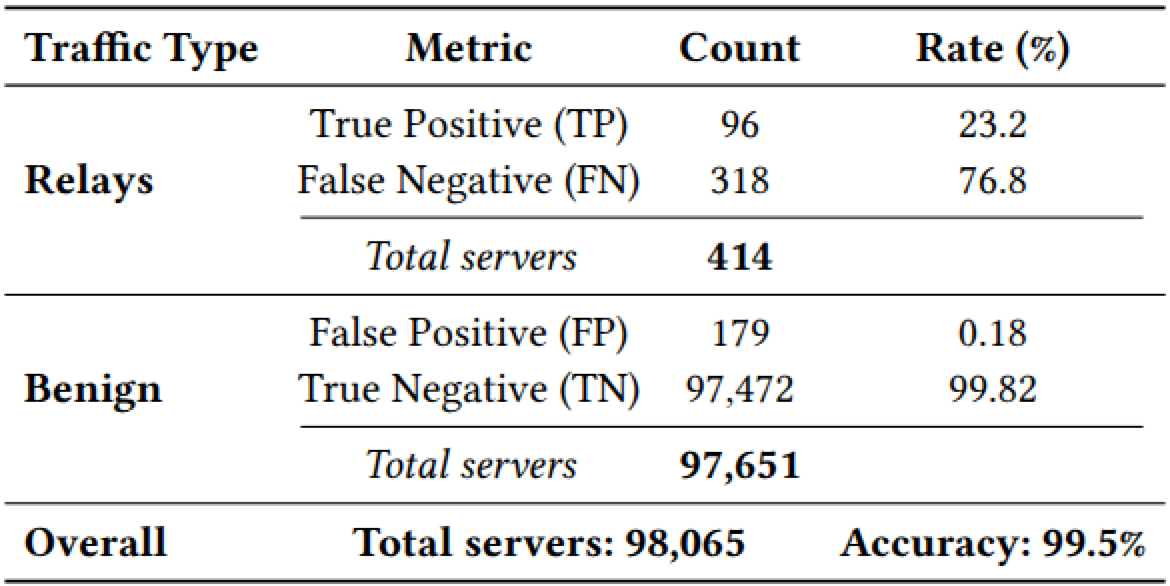

Traffic Type	Metric	Count	Rate (%)
Relays	True Positive (TP)	96	23.2
	False Negative (FN)	318	76.8
	Total servers	**414**	
Benign	False Positive (FP)	179	0.18
	True Negative (TN)	97,472	99.82
	Total servers	**97,651**	
Overall	**Total servers: 98,065**		**Accuracy: 99.5%**

Takeaways

why it matters?

- Single-IP relays can expose a **dual-role behavioral fingerprint**
- Obfuscation and probe resistance do not remove this architecture signal
- **Practical use:** a cost-sensitive censor can use it as a low-cost filter at scale

Any Questions?