SMURF: Privacy-Preserving Multi-Party Graph Mining for AML

Dhruvan Gnanadhandayuthapani[1] **Supervisor**: Dr. Zekeriya Erkin[1]

[1]EEMCS, Delft University of Technology, the Netherlands

1. Introduction

Problem & Challenge:

- **Scale:** Money laundering hides $800B-$2T annually (2-5% of GDP); <1% is seized; Shifting the burden onto the taxpayer.
- **Pattern:** Criminals use multi-bank scatter-gather ("smurfing") patterns. It looks normal to a single bank.
- **Barrier:** Centralized detection is ideal, but blocked by privacy laws (GDPR, FADP). Existing work supports only 2 parties.

Research Question: *"How can scatter-gather patterns be detected across multi-party financial networks without revealing sensitive transaction-level data?*

Contributions: 1) Multi-Party Framework, 2) Rust Implementation, 3) Experimental Evaluation

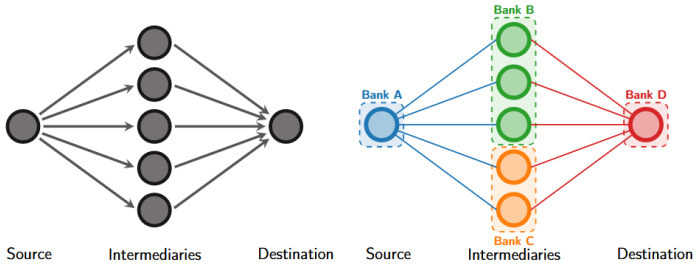

Figure 1: The scatter-gather pattern (a) and its partitioned variant across four banks (b).

2. Background

- **Setup:** n financial institutions $(B_i) + 1$ Central Coordinator (CC).
- **Additive Homomorphic Encryption (AHE):** Allows the CC to sum encrypted transaction counts privately.
- **Threshold Decryption:** Requires $\geq m$ banks to cooperate to decrypt final totals. Prevents unilateral decryption.
- **Oblivious Pseudorandom Functions (OPRF):** Blindly aligns cross-bank source-destination pairs using anonymous tokens (τ_{st}).

3. SMURF Framework

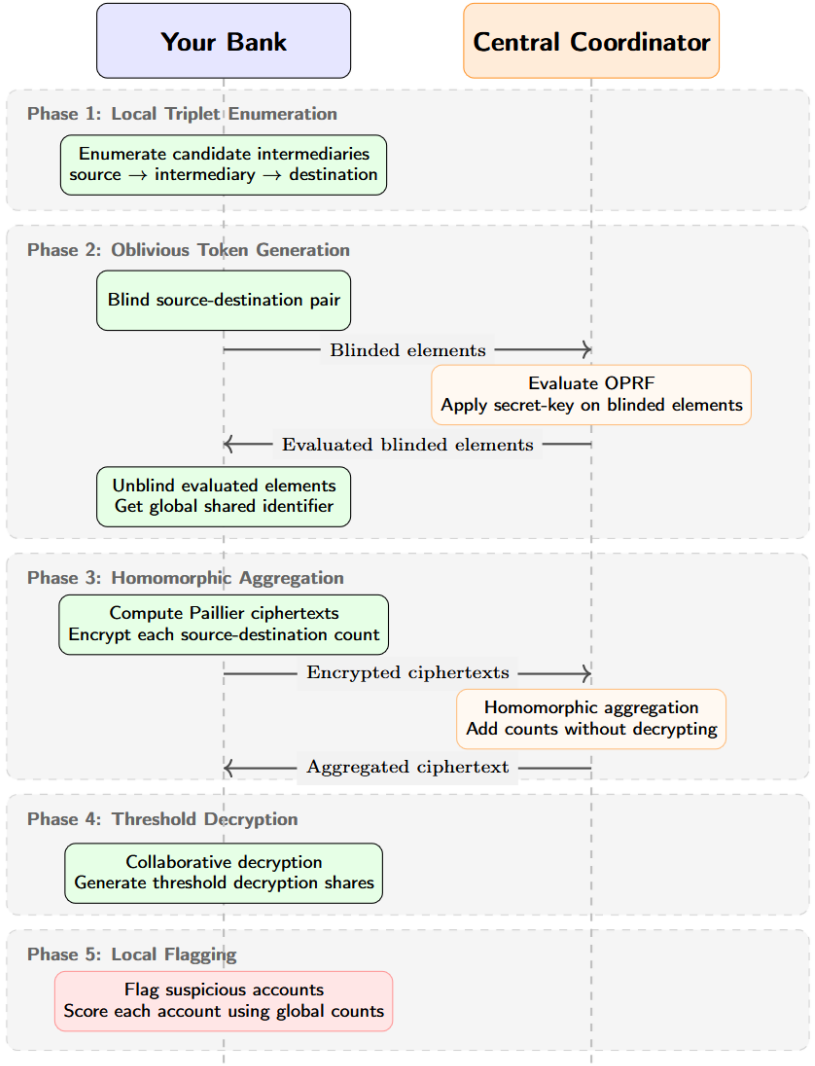

Figure 7: Overview of SMURF (Scatter-gather Mining Using Randomized Functions)

4. Results

- **Hardware:** 4th gen. AMD EPYC (64 cores, 128GiB of RAM from AWS).
- **Data:** AMLWorld. Small (~7M LI/~5M HI) and Medium (~31M transactions).

Table 2: Classification Performance Across AMLWorld at the Threshold With Optimal F_1

Dataset	Intensity	Optimal θ	Accuracy	Precision	Recall	F_1-Score	PR-AUC
Small	Low (LI)	8	99.99%	100.00%	69.09%	81.72%	0.4423
	High (HI)	6	99.97%	88.75%	76.96%	82.44%	0.6874
Medium	Low (LI)	9	99.98%	86.47%	53.36%	65.99%	0.5172
	High (HI)	8	99.94%	88.14%	59.71%	71.19%	0.6671

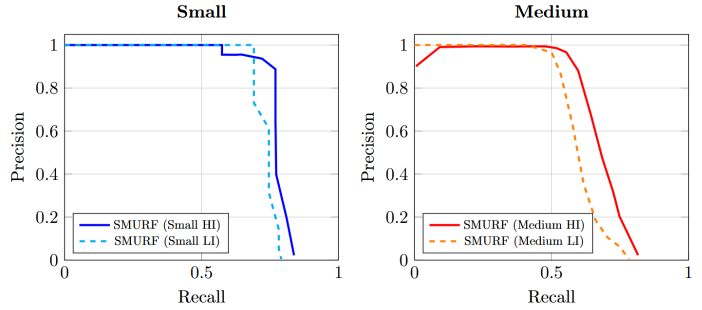

Figure 3: Precision-Recall Curves on Small and Medium AMLWorld datasets for $1 \leq \theta \leq 20$.

5. Conclusions

- **High Performance:** SMURF achieves >80% F_1-Score. Processes approx. 30M transactions in 1.6 hours using a 64-core processor.
- **Linear Complexity:** Traversal scales linearly with local vertices. OPRF and homomorphic encryption scale linearly per bank based on dataset size.
- **Cryptographic Overhead:** Paillier encryption introduces heavy compute and communication costs.
- **Security Assumption:** Currently relies on a semi-honest model; needs Zero-Knowledge Proofs (ZKPs) to prevent malicious data injection.
- **Structural Limitation:** Currently restricted to detecting 2-hop scatter-gather patterns.
- **Metadata:** Omitting transaction amounts and timestamps limits stronger detection signals.