DarkSIM

A similarity-based time-series analytic framework for darknet traffic

Max Gao
Esteban Carisimo
Ricky Mok
kc Claffy
et al.

Background and Motivation

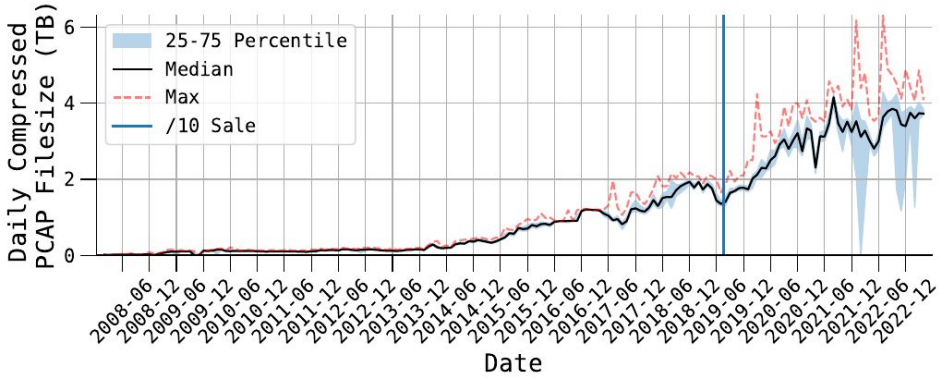

- Darknets provide insight into global scanning, outage, and backscatter event
- Challenging to extract insight from today's traffic using existing methods due to traffic growth
  - Traditional packet analysis approaches lack flexibility
  - Classical statistical approaches (e.g., change-point detection) rely on assumptions
  - Representation learning techniques are black boxes
- Can we design an approach that addresses these limitations?

Overview of *DarkSIM*

- DarkSIM detects events based on pattern similarity in time-series segments
  - Uses Dynamic Time-Warping (DTW) to compute similarity scores

- Statistical analysis of scores identifies three types of anomalies
  - **One-Off**
  - **Repeated**
  - **Concurrent** (pictured)

- Packet processing handled by Corsaro3
  - corsarotagger: geolocation and AS-tagging of packets
  - corsarotrace: time-series generation

- Parallel score computation handled by Python libraries

- Computations run on SDSC's Expanse cluster

Overview of *DarkSIM*

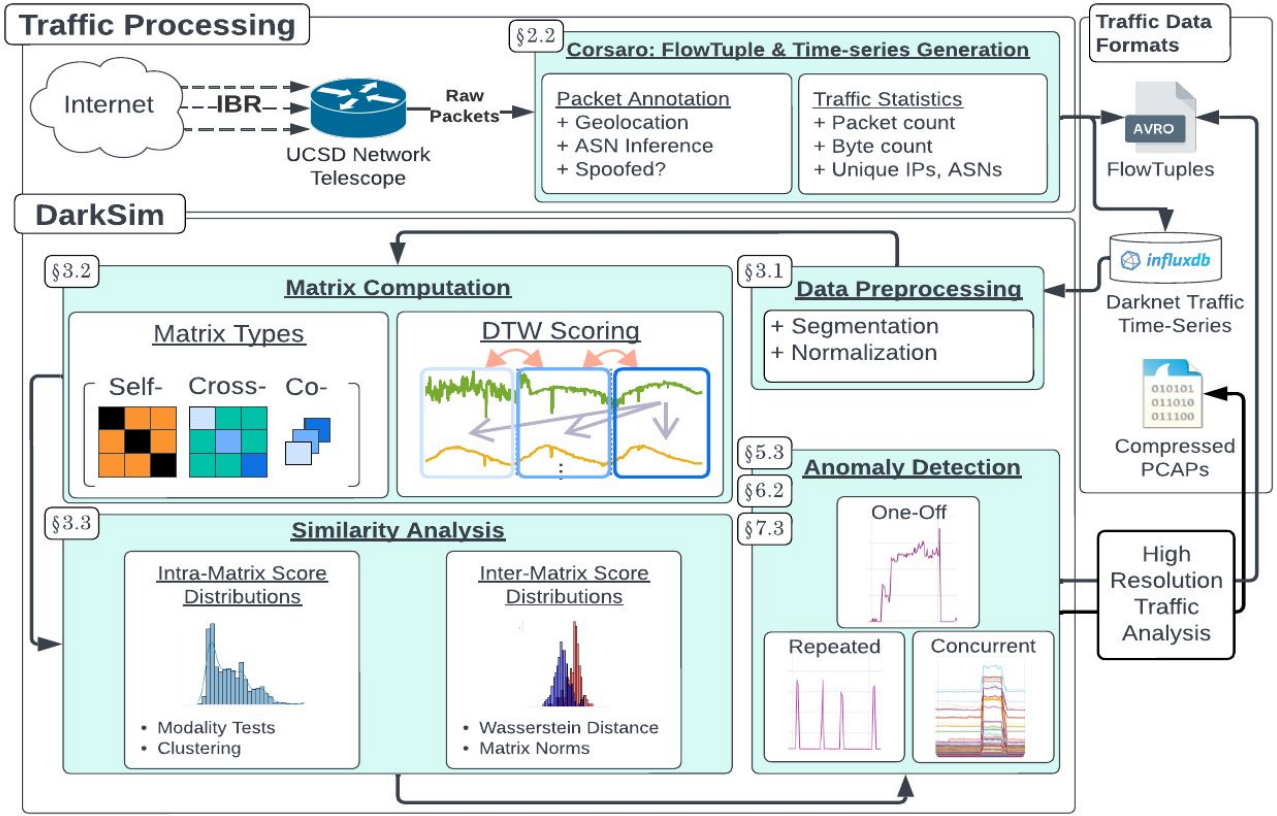

Baseline Method Comparison: *DarkGLASSO*

- Why *DarkGLASSO*?
  - Similar fundamental approach
  - Application of a well-studied statistical inference algorithm to darknet traffic
- Experiment Setup
  - Analyzed unique sender count time-series
    (Top-128 TCP ports by volume, June 2023)
  - Ablated method parameters
    - Segment Length *b* (Both)
    - L1-Penalty λ (*DarkGLASSO*)
    - Warp-Width *w* (*DarkSIM*)
- Comparisons
  - Detection sensitivity
  - Detection accuracy
  - Computational performance

Parameter	Values
b	15, 15, 30, 60, 180, 720, 1440
λ	0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9
w	0, 0.5, 1.0

Baseline Method Comparison - Results

	DarkSIM	*DarkGLASSO*
Accuracy	- Detection sensitivity best for day-long lengths - True positives among most-severe detections (across all lengths) - Misses 28% of DarkGLASSO's detections	- Detection sensitivity best at 3-hour lengths - Unreliable detections (at most 73% at a 6-hour length) - Misses 96% of DarkSIM's detections
Performance	- Larger warp-widths increase computation time - Wall-time per matrix (b=5): 0.015s - Wall-time per matrix (b=1440): 1.7s	- Larger λ decreases computation time - Wall-time per matrix (b=5): 0.175s - Wall-time per matrix (b=1440): 0.21s

Baseline Method Comparison - Results

	DarkSim Co-SM			*DarkGLASSO* Σ_-^{-1}			*DarkGLASSO* R_-^{-1}		
b	A	B	C	A	B	C	A	B	C
5	15	0	0	0	5	10	1	4	0
15	15	0	0	0	4	11	3	4	0
30	15	0	0	1	5	9	2	2	11
60	15	0	0	0	10	5	5	6	4
180	15	0	0	0	6	9	9	1	5
360	15	0	0	1	6	8	10	1	4
720	15	0	0	0	8	7	6	2	7
1440	15	0	0	4	2	9	8	1	6

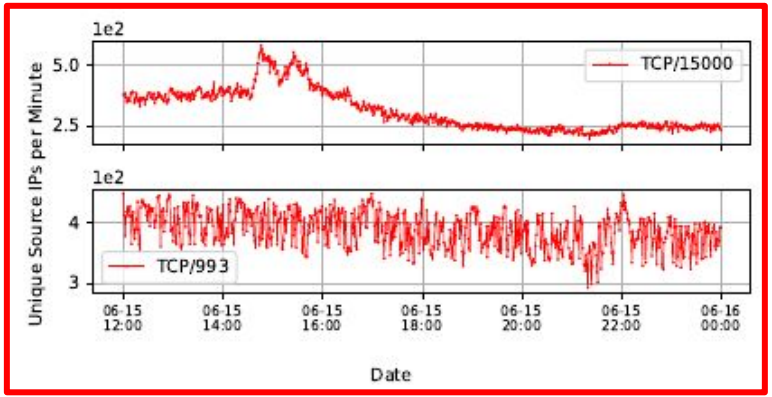

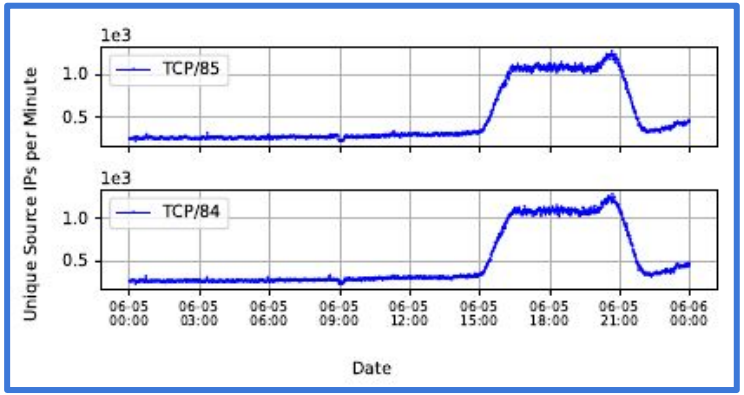

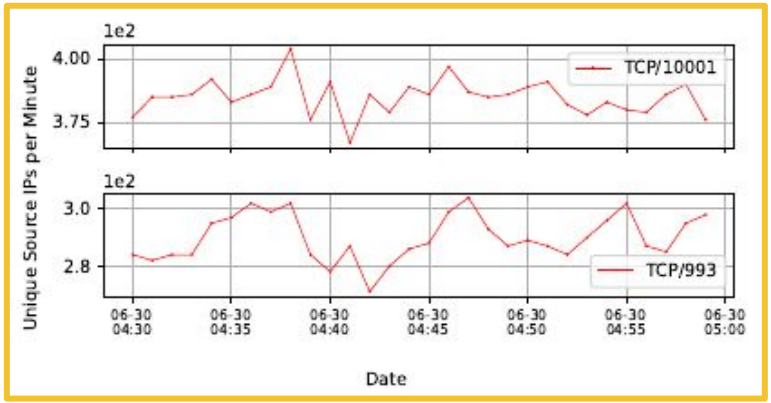

Baseline Method Comparison - Results

N	b	Co-SM	Σ^{-1}	R^{-1}
8640	5	131.13	312.15	1514.74
2880	15	75.57	108.97	426.35
1440	30	61.80	51.94	231.01
720	60	54.71	26.38	127.17
240	180	49.12	9.03	45.77
120	360	48.61	4.82	24.52
60	720	48.94	2.54	13.52
30	1440	51.48	1.46	6.31

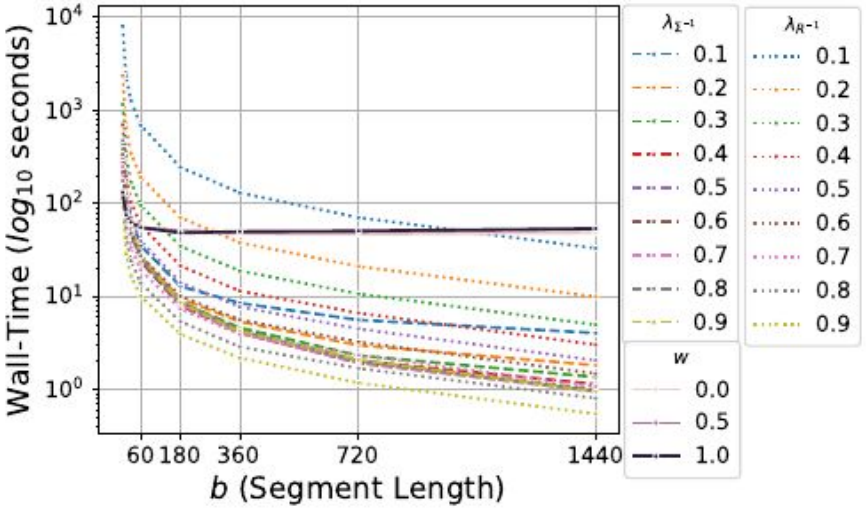

Case Study: Post-Disclosure Events

- Microsoft's February 2023 Patch Tuesday
  - Released on Feb. 2
  - RCEs reported for 14 ports in 3 applications
- Analyzed unique source IP count time series for **concurrent anomalies**
  - For the time frame between Feb. 1 to Mar. 31
- Detected notable changes on 4 days
  - Feb. 23, Feb. 24, Mar. 3, Mar. 15

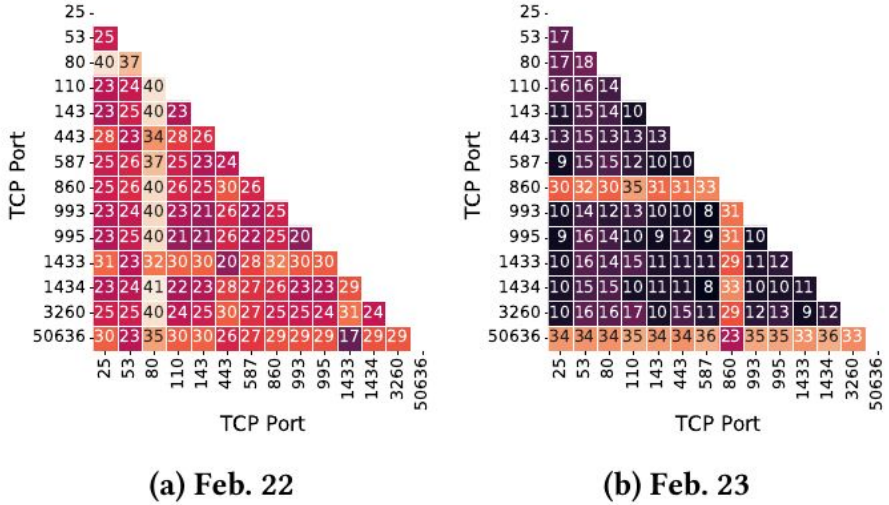

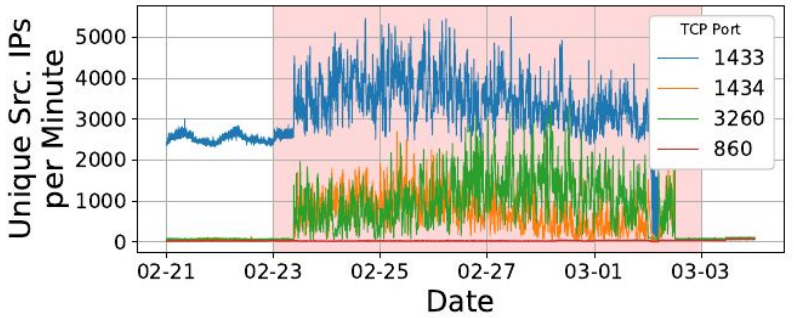

Case Study: Post-Disclosure Events

- Identified 8 Origin ASes responsible for the event lasting from Feb. 24 to Mar. 3
  - Chinese telecoms (4837, 4847, 4134, 9808) and siblings
- Investigated changes before and during the event onset of each AS' sender behavior
  - Increase in packets sent per source IP
  - Increase in unique destination /16s targeted per source IP
  - Non-sequential scan strategy

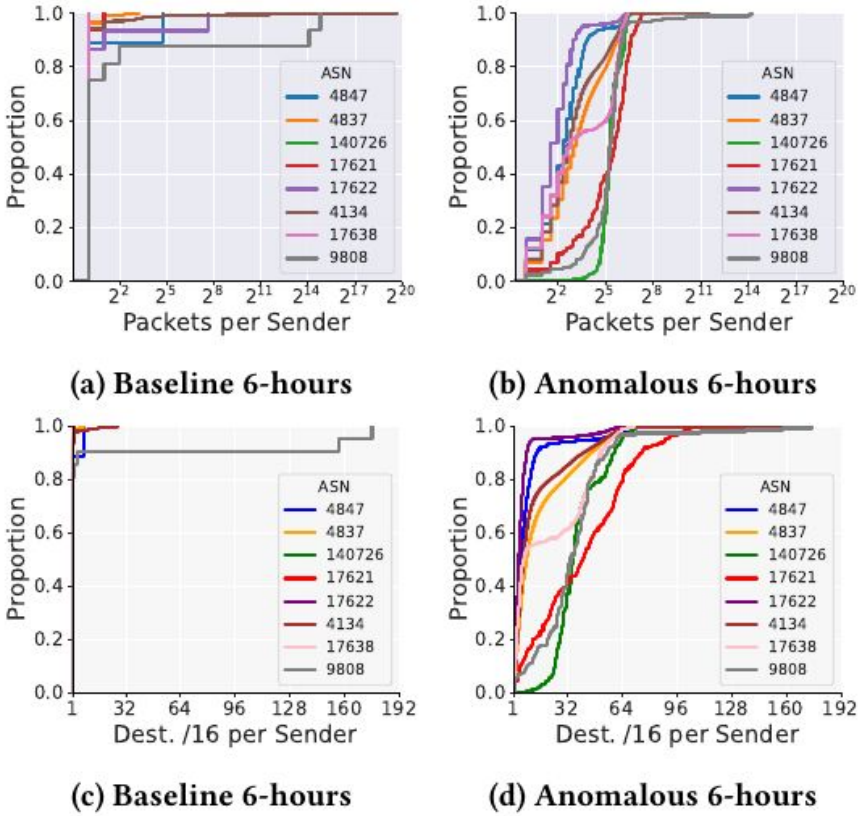

Case Study: Post-Disclosure Events

- Use of reference pattern to discover **repeated** anomalies

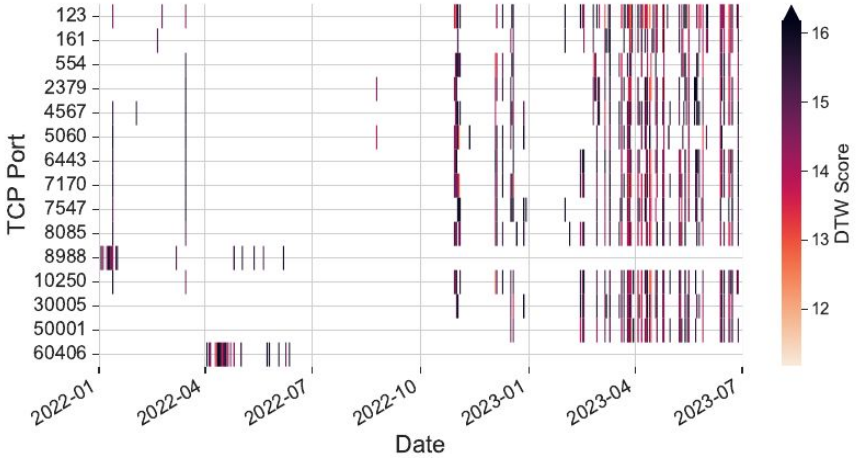

- Notable targeted ports and their inferred services

2022 (Jan. 1 - Dec. 31)		2023 (Jan. 1 - June 31)	
Service	Port - Count	Service	Port - Count
	60406 - 23	NTP	123 - 54
	8988 - 17	Kubernetes	10250 - 40
SIP	5060 - 15	NSRP	7170 - 39
NSRP	7170 - 13	IBM Cloud Orch.	50001 - 39
CPE Mgmt.	8085 - 12	TR-069	30005 - 39
NTP	123 - 12	Kubernetes	2379 - 38
Verizon Routers	4567 - 12	Kubernetes	6443 - 38
TR-069	7547 - 12	Verizon Routers	4567 - 38
Kubernetes	10250 - 11	SNMP	161 - 37
RSTP	554 - 11	CPE Mgmt.	8085 - 36

Summarizing Our Findings

- *DarkSIM* detects a variety of events in our experiments
  - Coordinated scanning campaigns
  - Known-scanner campaigns
  - Potential outages

- More semantically accurate detections compared to state-of-art method *DarkGLASSO*
  - Evaluated under purely “unsupervised” setting
  - Based on a subset of most-severe detections we manually investigated

Future Work

- Design benchmarks for automatic darknet event detection methods
  - Accuracy and performance tradeoffs between novel approaches
  - Strengths and weaknesses of each approach

- Extend current framework capabilities
  - Improve framework's event labeling capabilities
  - Correlate/cross-validate darknet event signals with telescope-exogenous datasets