Attacker Behavior Modeling using Temporal Data Analysis

Azqa Nadeem
PhD candidate @ Cyber Analytics Lab
Department of Intelligent Systems
Delft University of Technology

TU Delft

07 June 2022

Cyber Analytics Lab

Dynamic observables

- Program execution → observable data
  - Software logs
  - Network traffic
  - Intrusion alerts
  - …

Pros:
- Shows exactly what happened (not what could've happened)
- (Relatively) obfuscation free

Human/Automated threat actor

- Proxy to attacker intent

TU Delft

Temporal traces

- Temporal events → insightful patterns

Challenges:
- Curse of dimensionality
- Visualization?
- Distance measure?
- Performance
- …

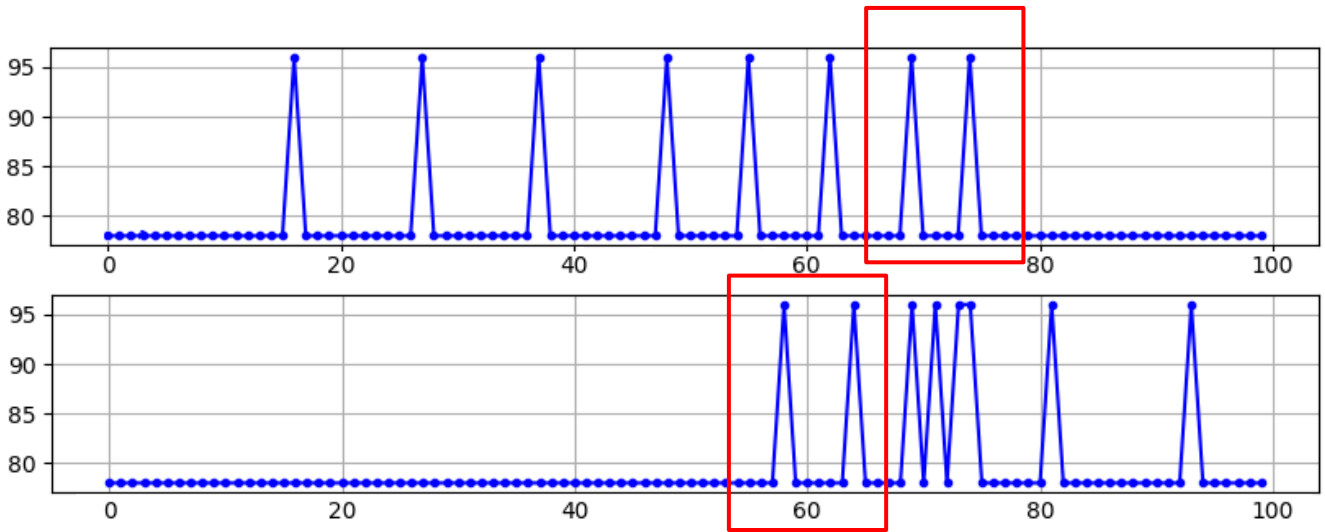

TU Delft

Attacker behavior modeling

- Attacker strategy extraction
  - Data: Intrusion alerts
  - Method: Probabilistic Deterministic Finite Automata

- Malware behavioural profiling
  - Data: Network traffic
  - Method: Hierarchical density-based clustering

Use Case I: Attacker Strategy Extraction

Alert management is expensive

- Security analysts handle >1M intrusion alerts/day* → Alert fatigue

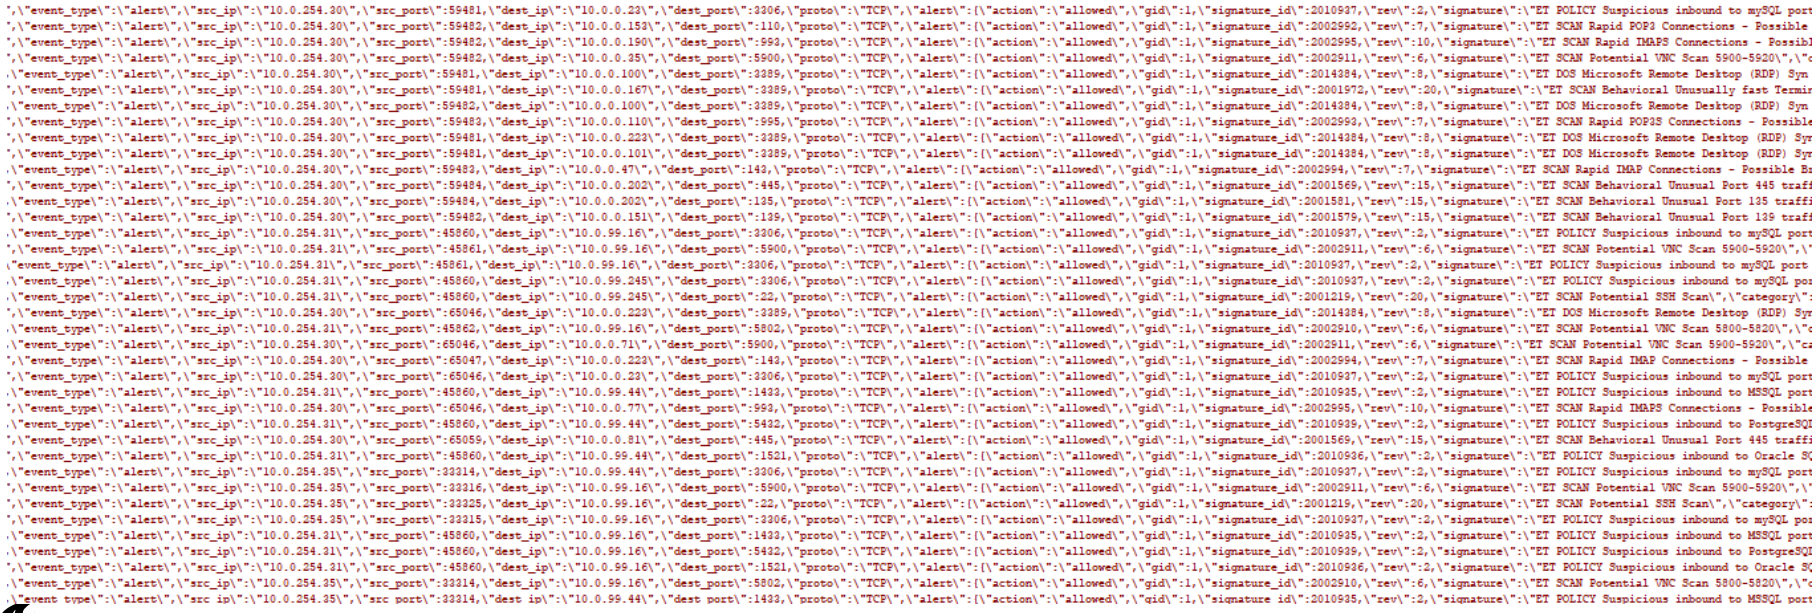

TU Delft

https://www.imperva.com/blog/27-percent-of-it-professionals-receive-more-than-1-million-security-alerts-daily/

Threat assessment

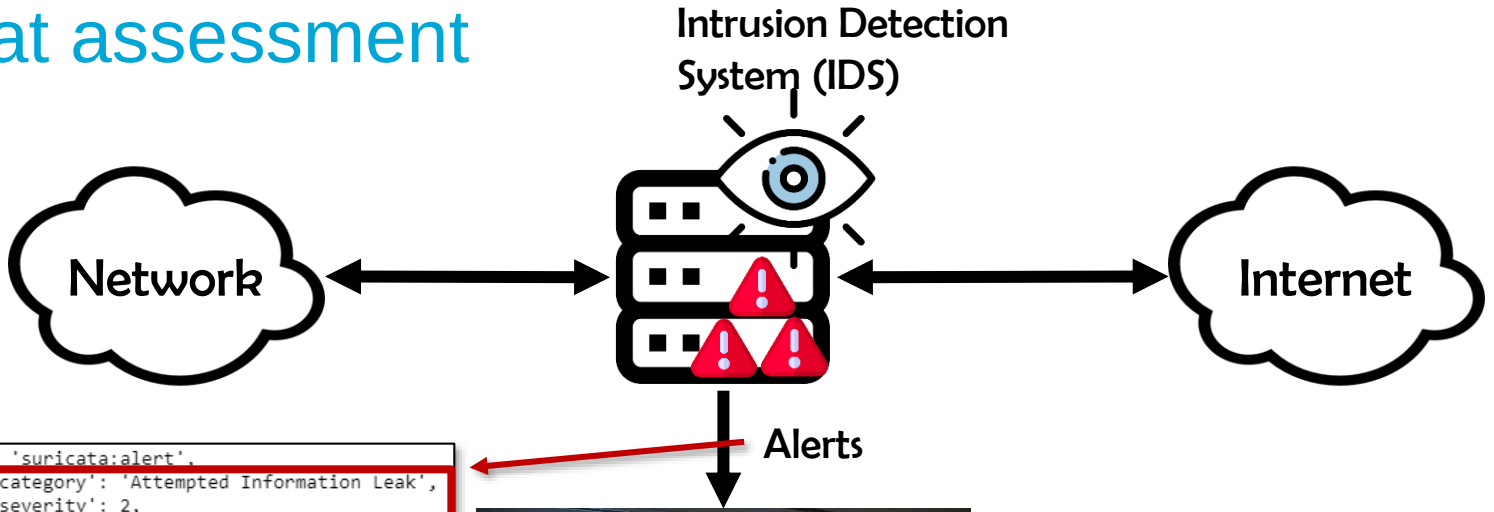


```
{   '_sourcetype': 'suricata:alert',
    'alert': {   'category': 'Attempted Information Leak',
                 'severity': 2,
                 'signature': 'ET POLICY Python-urllib\\/ '
                              'Suspicious User Agent'},
    'dest_ip': '169.254.169.254',
    'dest_port': 80,
    'src_ip': '10.0.0.20',
    'src_port': 56952,
    'timestamp': '2018-11-03T13:51:58.205548+0000'}
```


Security Operations Center (SOCs)

True threat or False alarm?
What's happening?
Attacker strategy?
Multiple attackers?
...

RIT
TU Delft

* https://itsec.group/blog-post-csoc-guide.html

Attack graphs

- "Attack Graphs" display attacker strategies

Network topology
+
Vulnerability reports

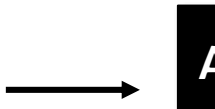

Attack Graph generator

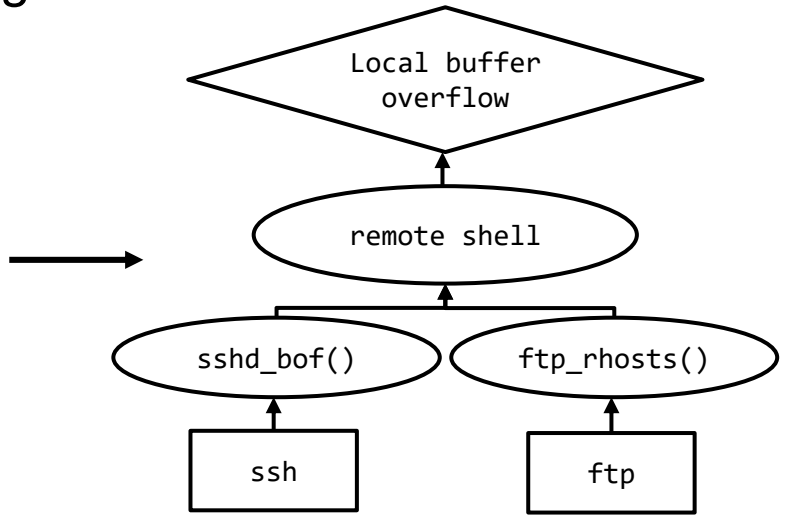

Static & hypothetical view of the network!

"Measuring Security Risk of Networks Using Attack Graphs". Steven Noel et al. (Int. J. Next Gener. Comput. 2010)

Alert-driven Attack graphs

- How to discover and display attacker strategies?
- SAGE: IntruSion alert-driven Attack Graph Extractor

RIT
TU Delft

Alert-driven attack graphs using S-PDFA. Azqa Nadeem, Sicco Verwer, Stephen Moskal, Shanchieh Jay Yang. In IEEE Transactions on Dependable and Secure Computing, 2021.

Anatomy of an Alert-driven Attack Graph

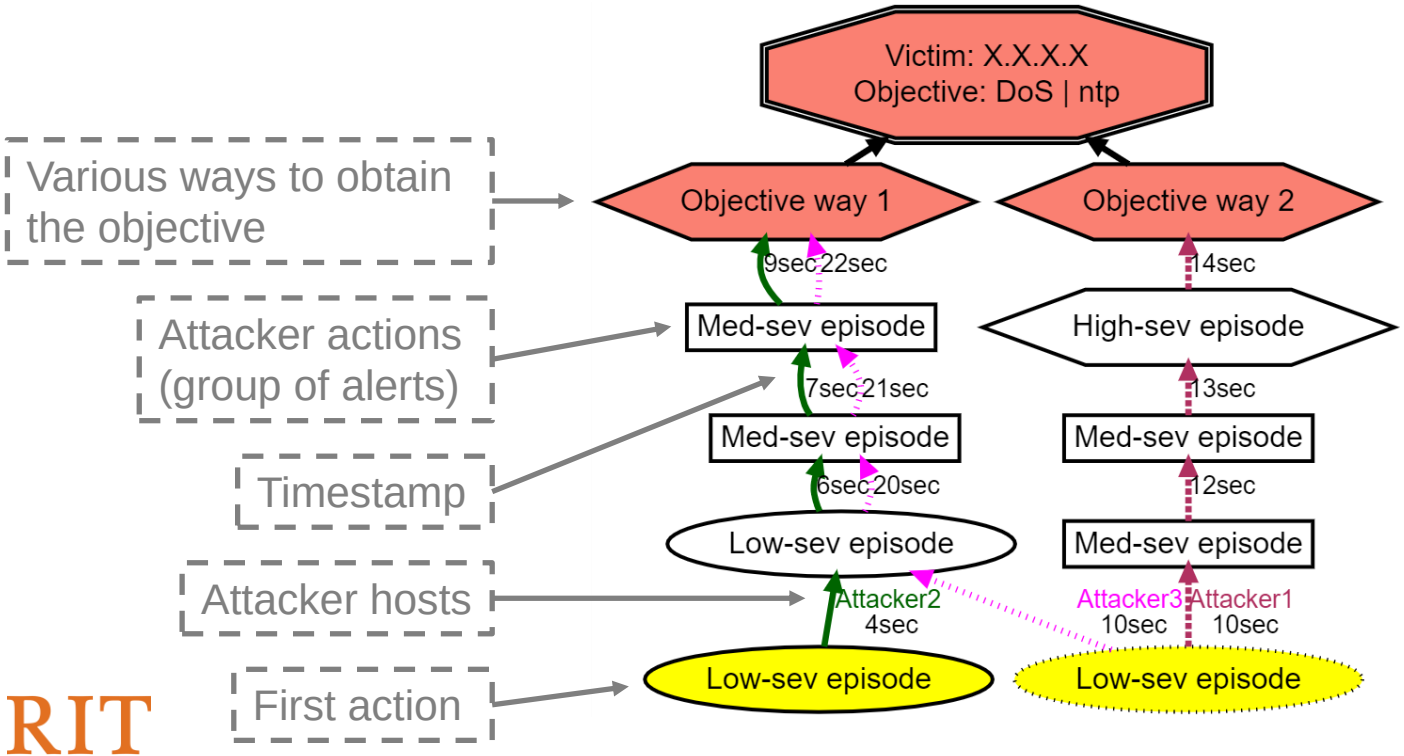

The anatomy of Alert-driven Attack Graphs. https://cyber-analytics.nl/blogposts/2022-01-03-attack-graphs/

Key design challenges

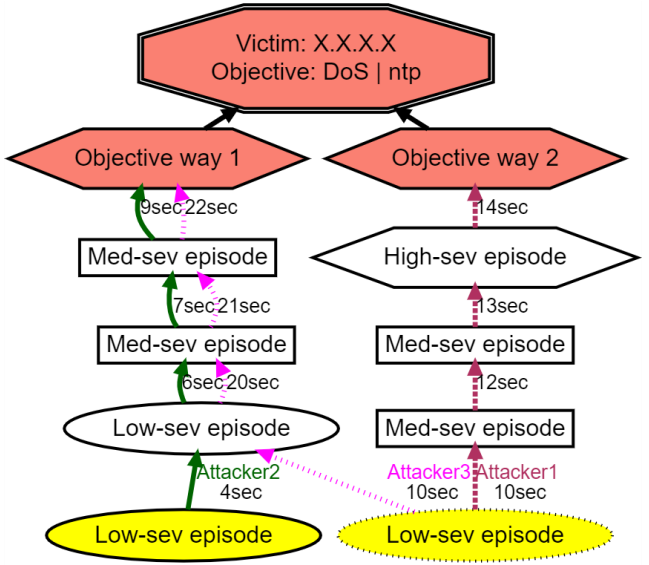

Key design challenges

1. Alert-type imbalance

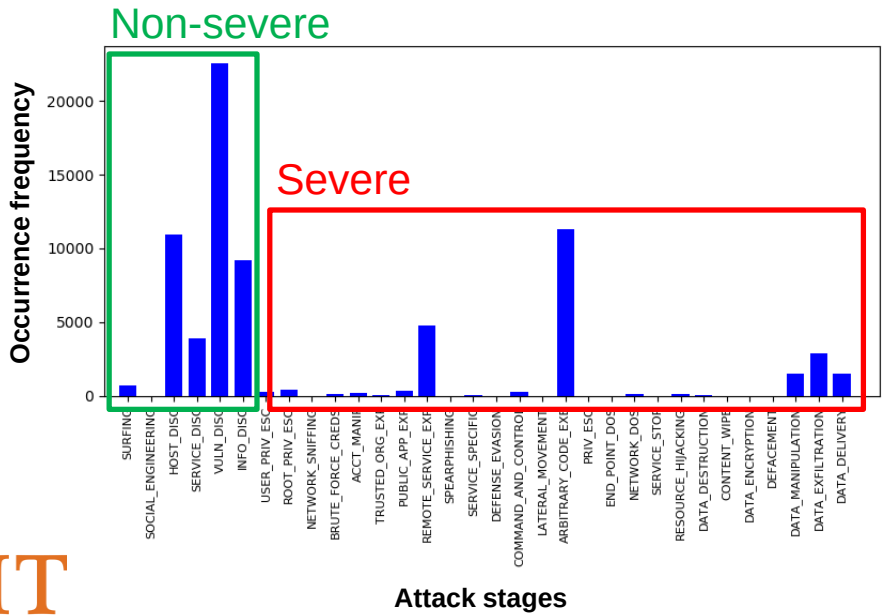

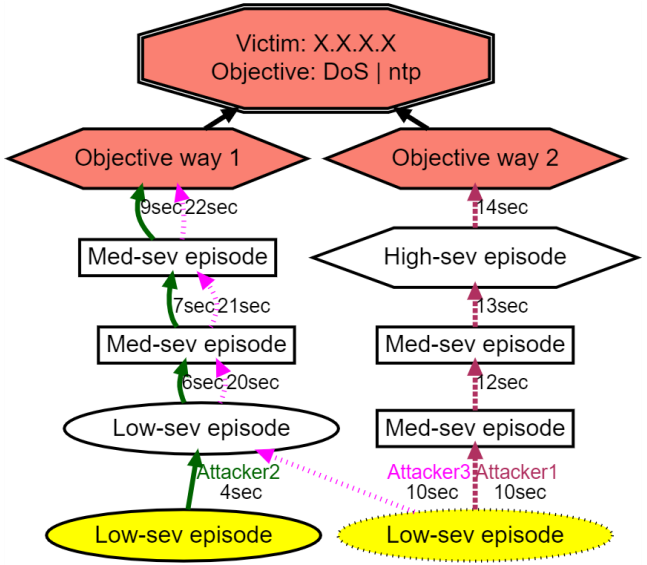

Key design challenges

1. Alert-type imbalance
2. Context matters

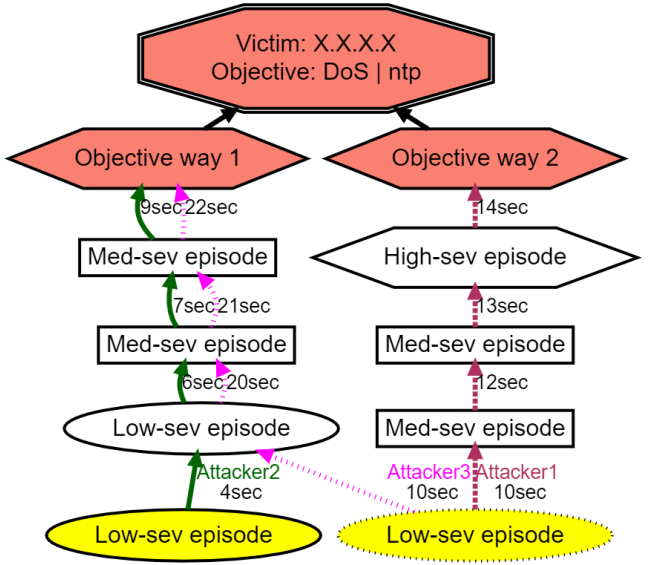

$Scan_1$, $Scan_2$, ACE, Exfiltration, ...
vs.
$Scan_1$, $Scan_1$, $Scan_2$, $Scan_2$, …

RIT
TU Delft

Key design challenges

1. Alert-type imbalance
2. Context matters
3. Explainable approach

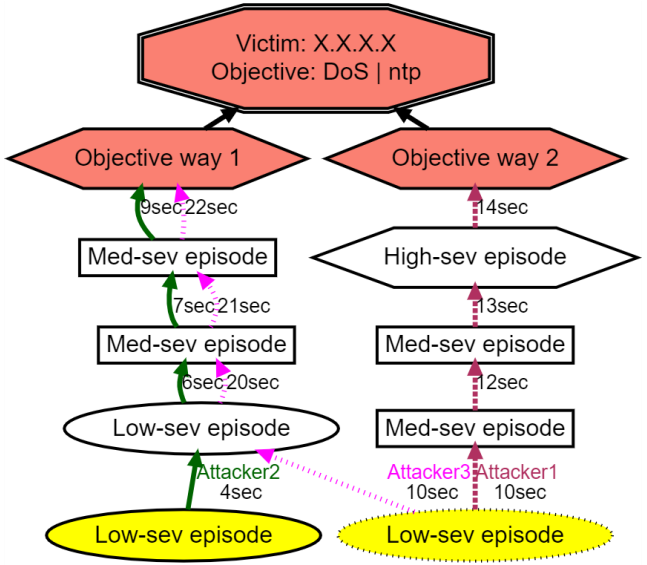

RIT
TU Delft

SAGE: IntruSion alert-driven Attack Graph Extractor

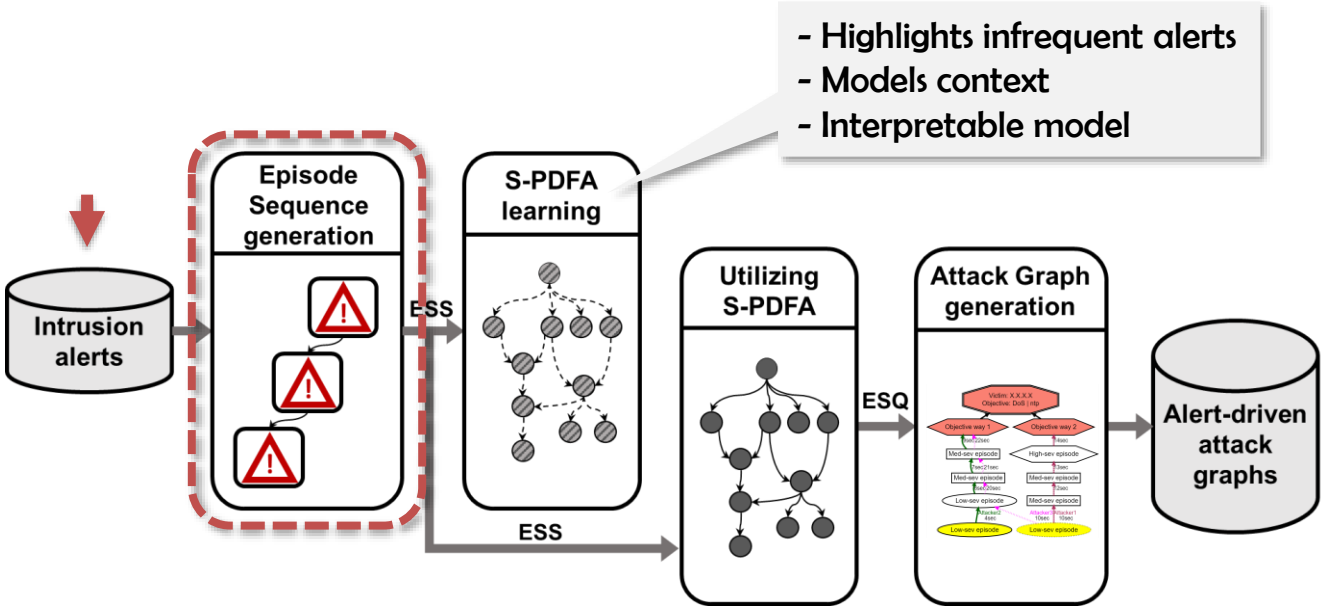

RIT
TU Delft

Alerts → Episode/Action sequences


```
{    '_sourcetype': 'suricata:alert',
     'alert': {    'category': 'Attempted Information Leak',
                   'severity': 2,
                   'signature': 'ET POLICY Python-urllib\\/
                                 'Suspicious User Agent'},
     'dest_ip': '169.254.169.254',
     'dest_port': 80,
     'src_ip': '10.0.0.20',
     'src_port': 56952,
     'timestamp': '2018-11-03T13:51:58.205548+0000'}}
```


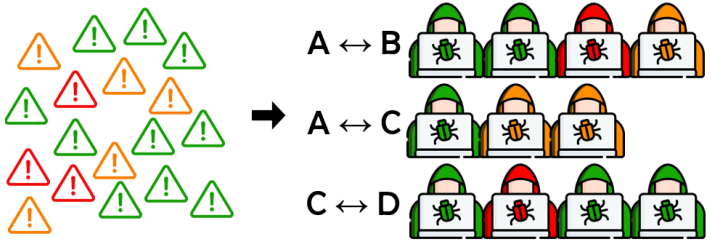

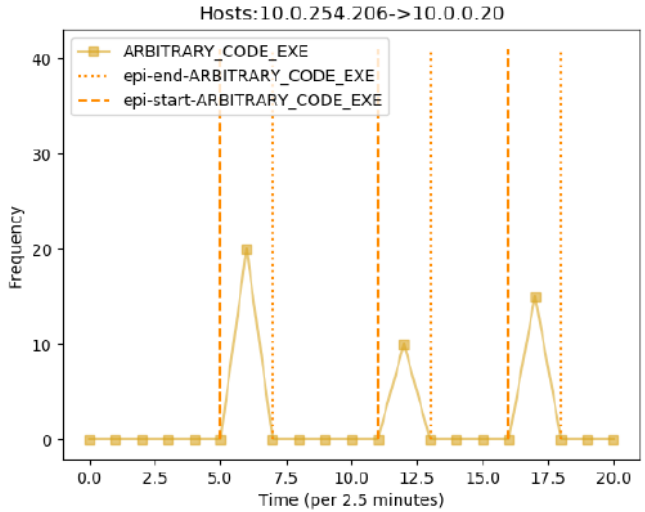

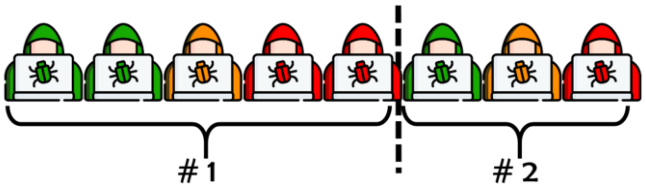

RIT TU Delft

Suffix Tree

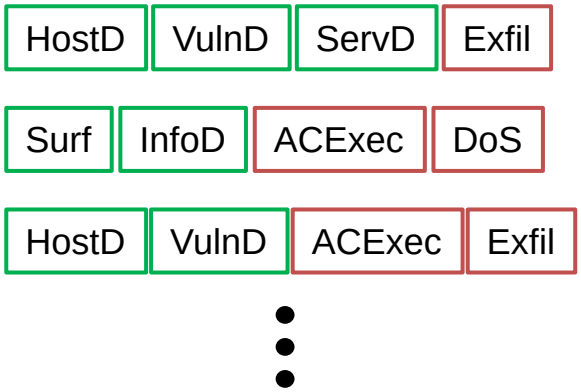

… all sub-sequences!

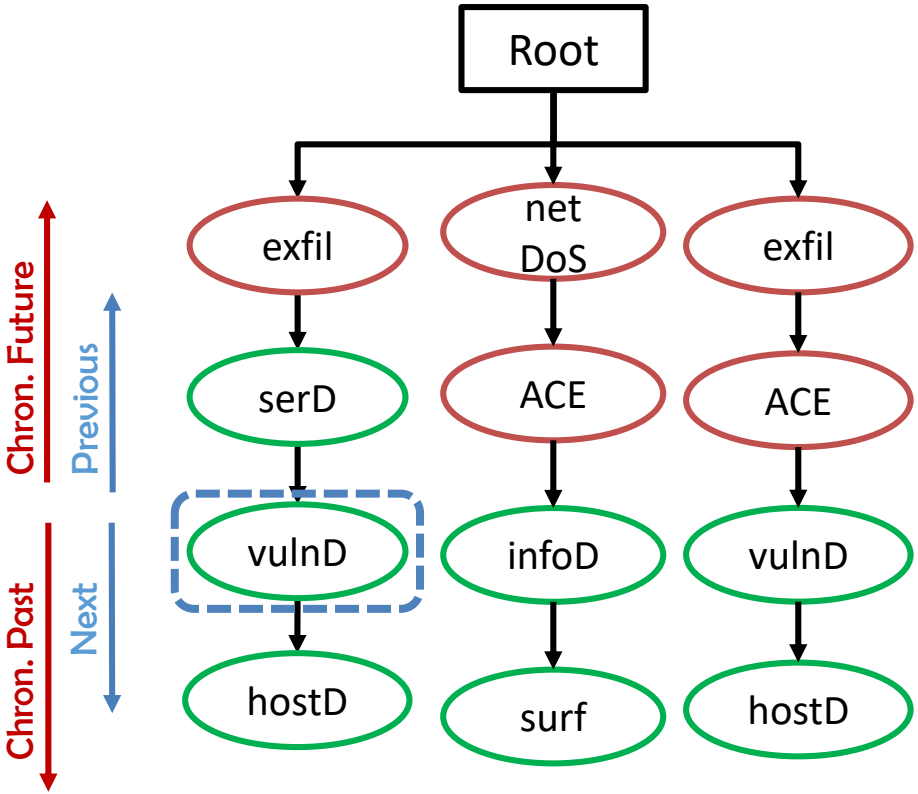

Suffix-based PDFA

- Suffix-based Probabilistic Deterministic Finite Automaton
- Summarizes attack paths
- Brings infrequent episodes to the top
  - Red → Severe | Blue → Medium severity

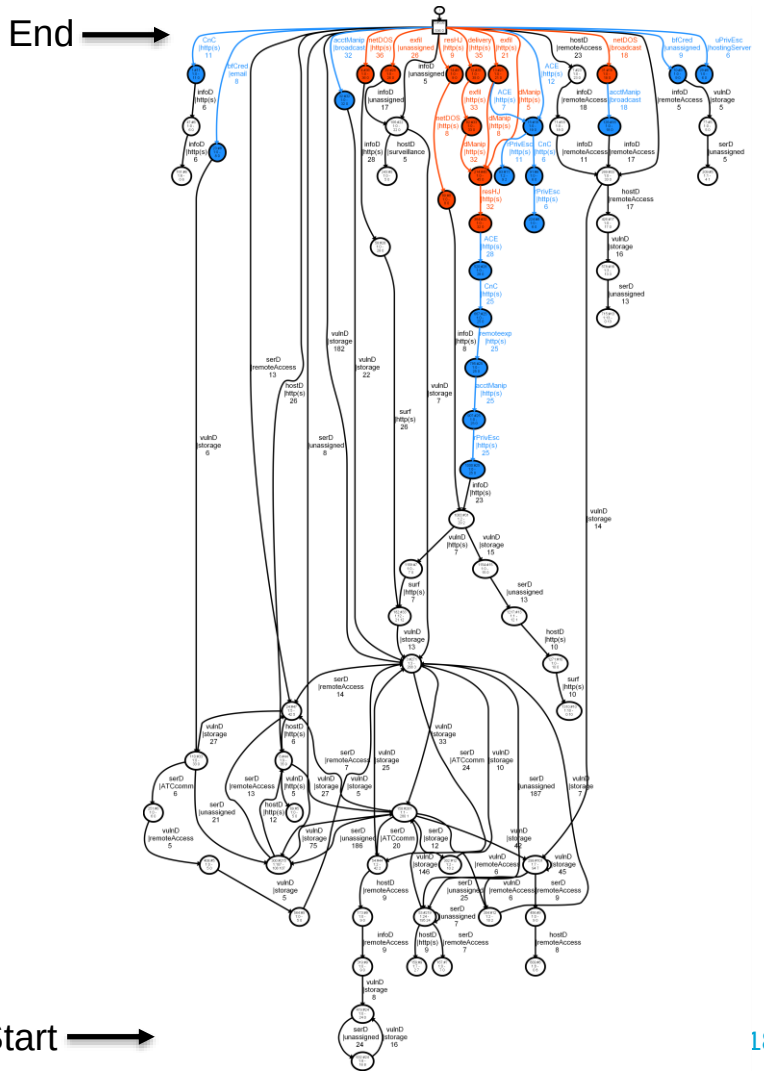

TU Delft

Suffix-based PDFA

- State identifiers model context
- Markovian properties
  - States → milestones with context
- Low-severity *Sinks* ignored
  - Cleaner model

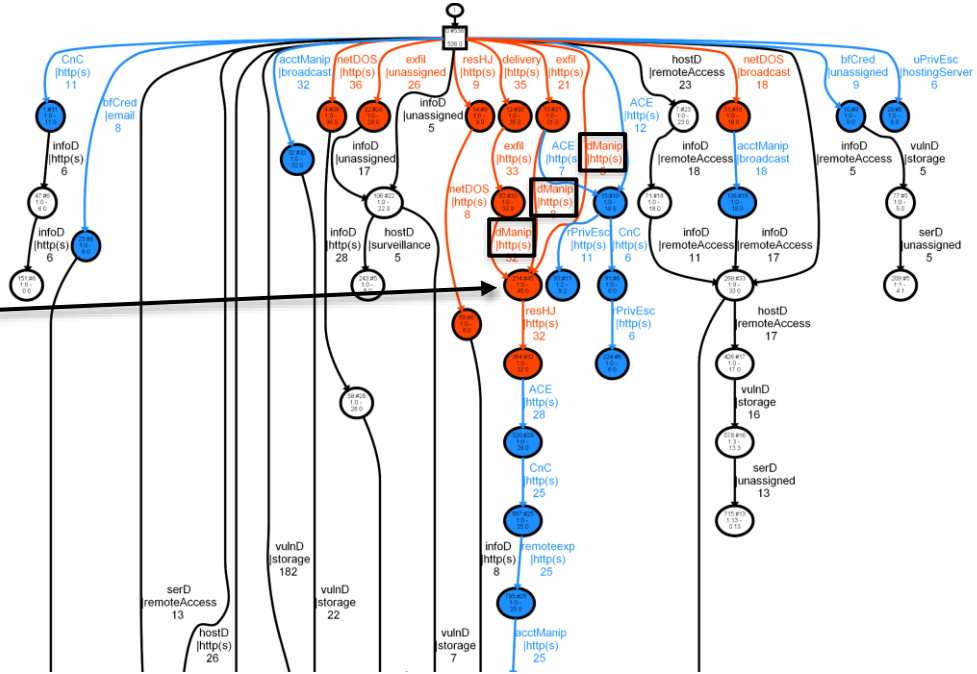

Adding context & Attack graph formation

Episode sequences

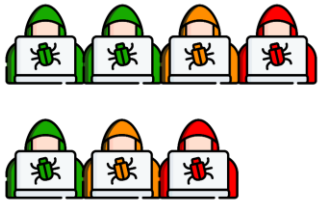

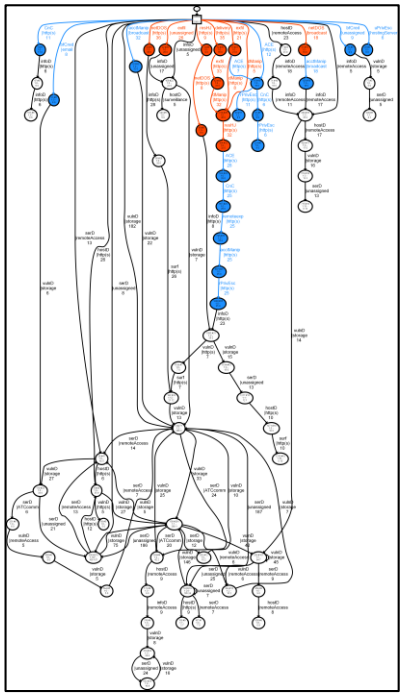

State sequences

Vertex: Group of alerts
Edge: Temporal order

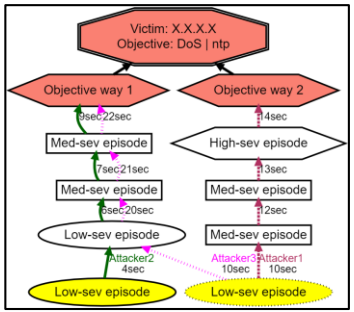

On a *per-victim, per-objective* basis

RIT
TU Delft

Experimental dataset

- Suricata alerts from Collegiate Penetration Testing Competition[1]
  - 6 multi-attacker teams
  - 1 fictitious network
  - 330,270 alerts
- Moskal's Action-Intent framework[2]
  - Alert signature → Attack stage

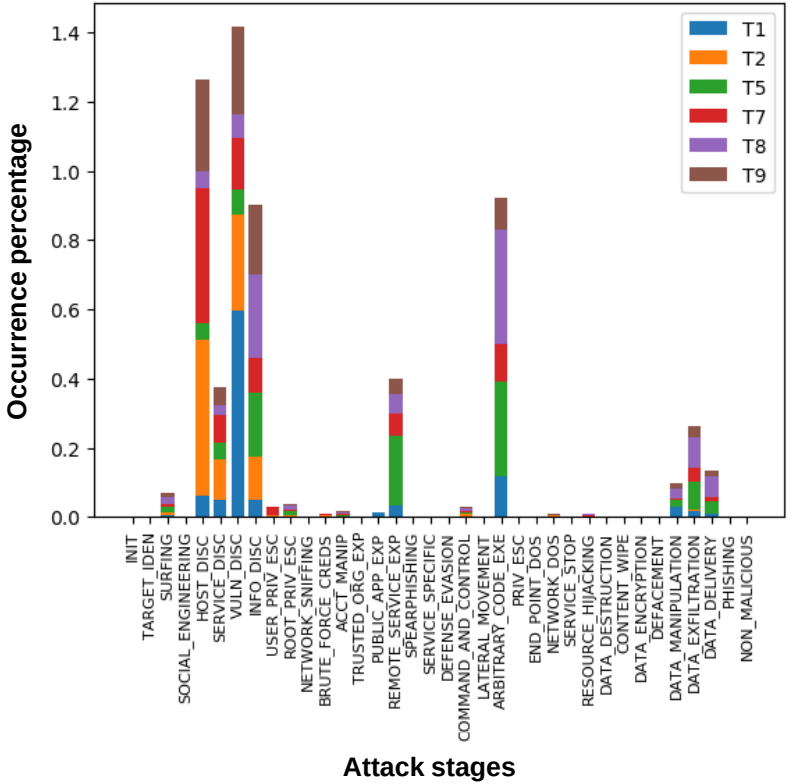

1. CPTC dataset: https://www.globalcptc.org/
2. S. Moskal and S. J. Yang, "Framework to describe intentions of a cyber attack action," arXiv preprint arXiv:2002.07838, 2020.

[1] Alert triaging

- 330,270 alerts → 93 alert-driven AGs
- Compresses ~500 alerts in < 25 vertices

	# alerts (raw)	# alerts (filtered)	#episodes	#ES/ #ESQ	#ESS	#AGs
T1	81373	26651	655	103	108	53
T2	42474	4922	609	86	92	7
T5	52550	11918	622	69	74	51
T7	47101	8517	576	63	73	23
T8	55170	9037	439	67	79	33
T9	51602	10081	1042	69	110	30

RIT
TU Delft

[2] Attacker strategy visualization

- Shows how the attack transpired
- 3 teams, 5 attempts
- 3 ways to reach objective
  - Discovered by S-PDFA

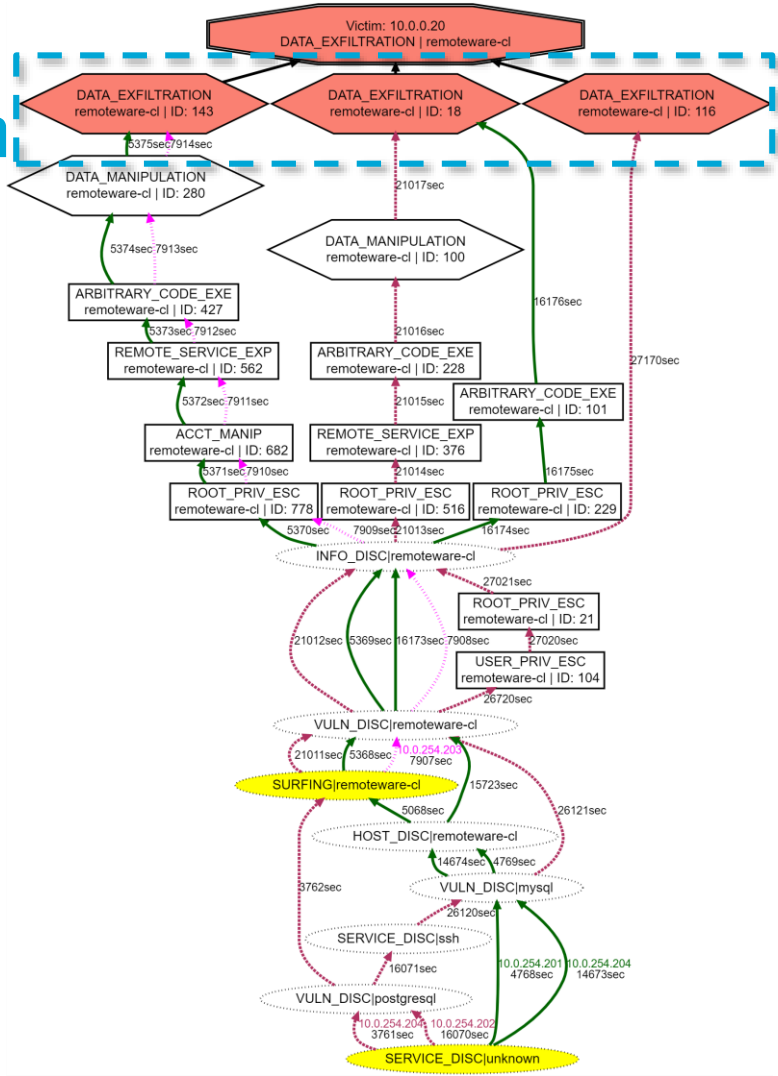

RIT
TU Delft

[3] Attacker strategy comparison

- T5 and T8 share a common strategy
- Some paths are shorter than others
- Attackers follow shorter paths to re-exploit an objective in 84.5% cases

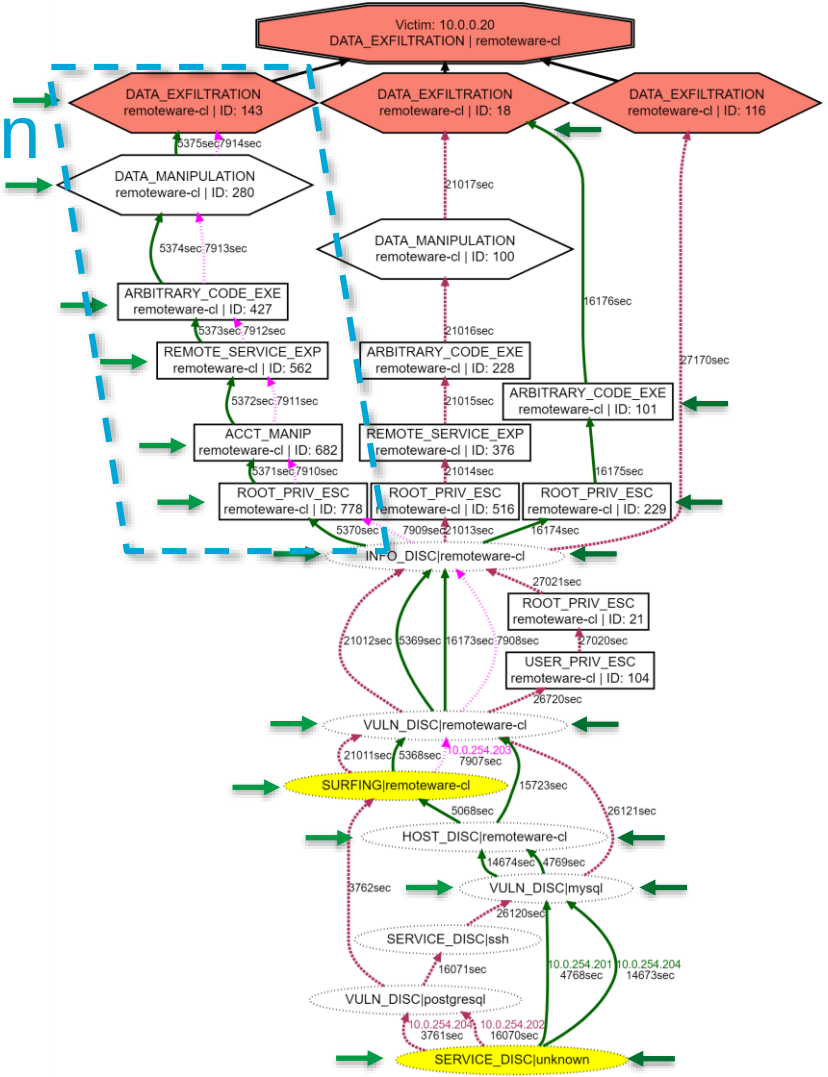

RIT
TU Delft

SAGE: Summary & Future work

- SAGE uses sequence learning for automated strategy extraction
  - S-PDFA critical for accentuating infrequent events & modeling context
- Alert-driven attack graphs
  - Compress thousands of alerts in a few AGs
  - Provide insight into attacker strategies and behavior dynamics

- Attack path prioritization
- Missing paths in AGs
- Modeling evolving strategies

TU Delft

Alert-driven attack graphs using S-PDFA. Azqa Nadeem, Sicco Verwer, Stephen Moskal, Shanchieh Jay Yang. In IEEE Transactions on Dependable and Secure Computing, 2021.

Use Case II: Malware Behavior Profiles

Inconsistent malware family labels

- Malware labels are inconsistent and black-box

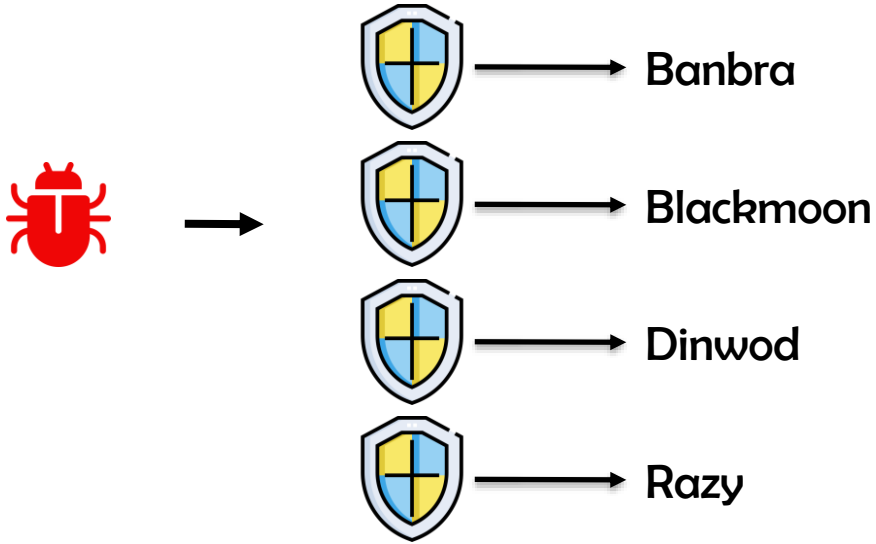

AV vendors & their naming conventions

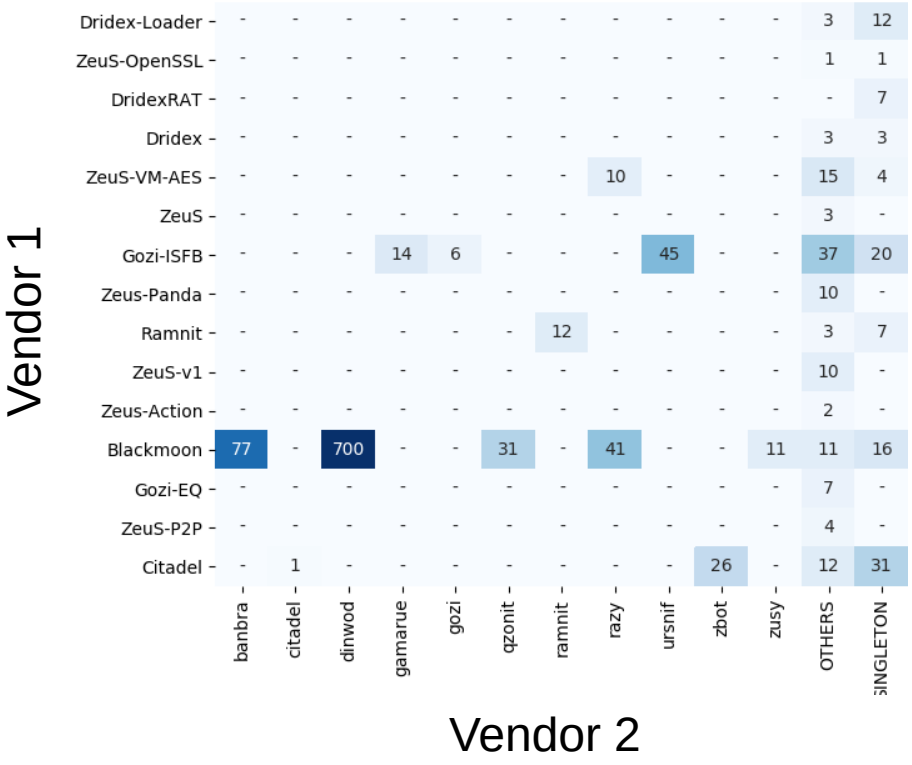

Malware behavior profiles

- Malware labels are inconsistent and black-box
- Behavior profiles are more insightful of capabilities

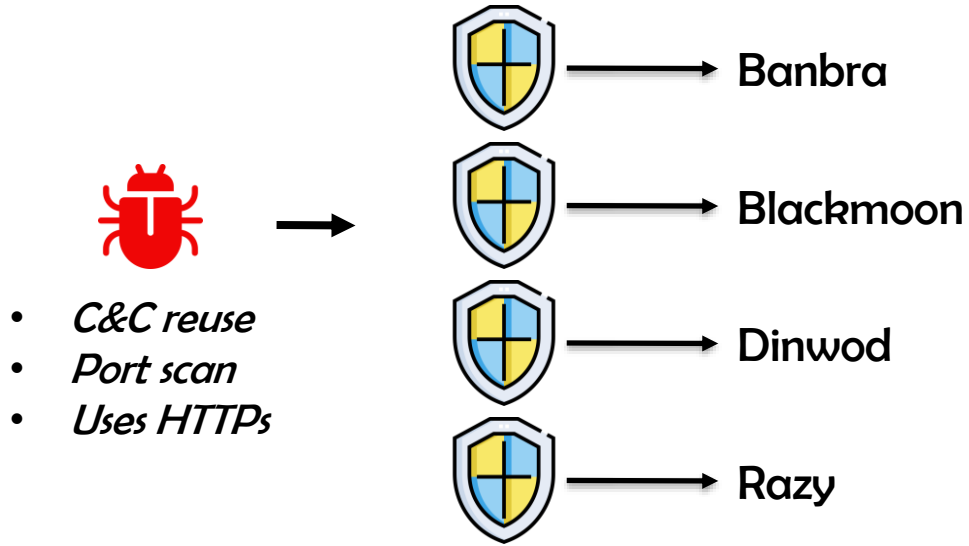

Malware capability assessment

- How to discover behaviors and build profiles?
- MalPaCA: Malware Packet Sequence Clustering and Analysis

FOX IT

TU Delft

Beyond Labeling: Using Clustering to Build Network Behavioral Profiles of Malware Families. Azqa Nadeem, Christian Hammerschmidt, Carlos H. Ganan, Sicco Verwer. In Malware Analysis using Artificial Intelligence and Deep Learning, Springer, 2021.

Network trace collection

- Malware infected machine generates network traffic

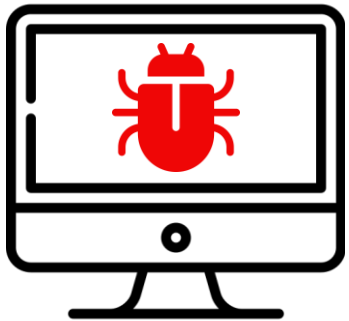

No.	Source	Destination	Protoc	Length	Info
40	192.168.1.2	192.168.1.110	ICMP	82	Redirect (Redirect for host)
41	CzNicZSP_00:0…	PcsCompu_7c:9…	ARP	60	192.168.1.1 is at d8:58:d7:00:0f:72
42	192.168.1.110	203.153.165.21	TCP	182	49191 → 8343 [PSH, ACK] Seq=1 Ack=1 Win=65700 Len=128
43	203.153.165.21	192.168.1.110	TCP	60	8343 → 49191 [ACK] Seq=1 Ack=129 Win=15744 Len=0
44	203.153.165.21	192.168.1.110	TCP	1188	8343 → 49191 [PSH, ACK] Seq=1 Ack=129 Win=15744 Len=1134
45	192.168.1.110	203.153.165.21	TCP	380	49191 → 8343 [PSH, ACK] Seq=129 Ack=1135 Win=64564 Len=326
46	192.168.1.2	192.168.1.110	ICMP	408	Redirect (Redirect for host)
47	203.153.165.21	192.168.1.110	TCP	113	8343 → 49191 [PSH, ACK] Seq=1135 Ack=455 Win=16768 Len=59
48	fd2d:ab8c:225…	fd2d:ab8c:225…	DNS	110	Standard query 0xb554 A www.download.windowsupdate.com

FOX IT

TU Delft

Beyond Labeling: Using Clustering to Build Network Behavioral Profiles of Malware Families. Azqa Nadeem, Christian Hammerschmidt, Carlos H. Ganan, Sicco Verwer. In Malware Analysis using Artificial Intelligence and Deep Learning, Springer, 2021.

Network trace collection

- Malware infected machine generates network traffic

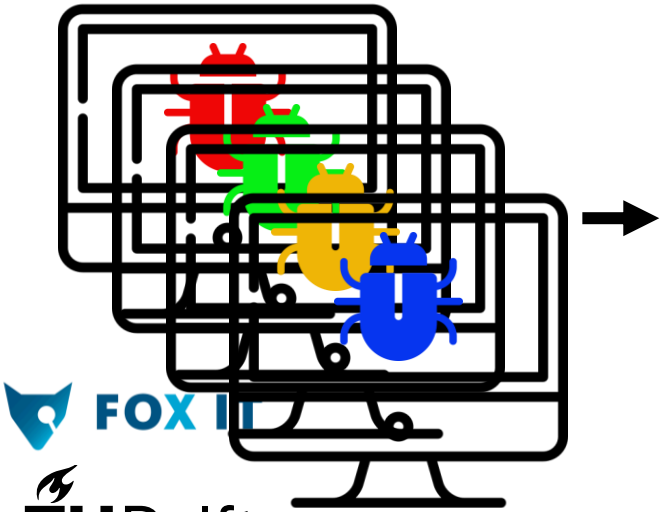

No.	Source	Destination	Protoc	Length	Info
40	192.168.1.2	192.168.1.110	ICMP	82	Redirect (Redirect for host)
41	CzNicZSP_00:0…	PcsCompu_7c:9…	ARP	60	192.168.1.1 is at d8:58:d7:00:0f:72
42	192.168.1.110	203.153.165.21	TCP	182	49191 → 8343 [PSH, ACK] Seq=1 Ack=1 Win=65700 Len=128
43	203.153.165.21	192.168.1.110	TCP	60	8343 → 49191 [ACK] Seq=1 Ack=129 Win=15744 Len=0
44	203.153.165.21	192.168.1.110	TCP	1188	8343 → 49191 [PSH, ACK] Seq=1 Ack=129 Win=15744 Len=1134
45	192.168.1.110	203.153.165.21	TCP	380	49191 → 8343 [PSH, ACK] Seq=129 Ack=1135 Win=64564 Len=326
46	192.168.1.2	192.168.1.110	ICMP	408	Redirect (Redirect for host)
47	203.153.165.21	192.168.1.110	TCP	113	8343 → 49191 [PSH, ACK] Seq=1135 Ack=455 Win=16768 Len=59
48	fd2d:ab8c:225…	fd2d:ab8c:225…	DNS	110	Standard query 0xb554 A www.download.windowsupdate.com

FOX IT

TU Delft

Beyond Labeling: Using Clustering to Build Network Behavioral Profiles of Malware Families. Azqa Nadeem, Christian Hammerschmidt, Carlos H. Ganan, Sicco Verwer. In Malware Analysis using Artificial Intelligence and Deep Learning, Springer, 2021.

Behavior catalog construction

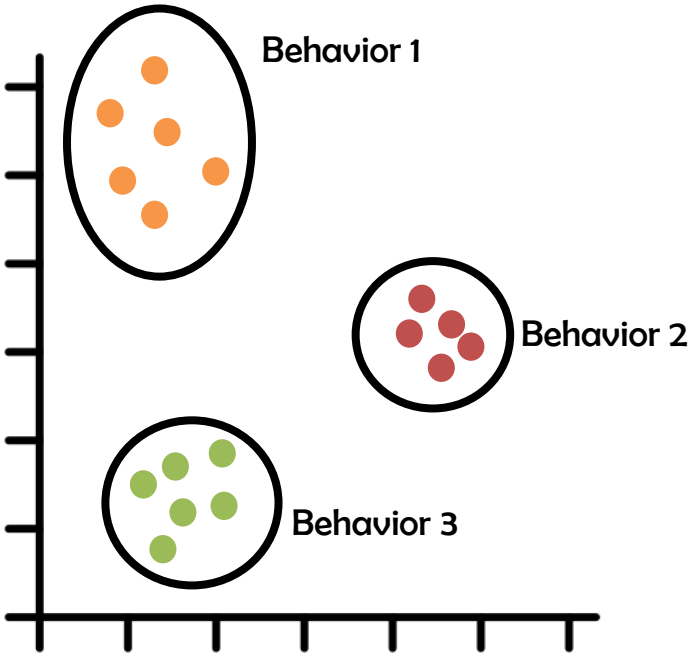

Behavior profiles

Gozi

Citadel

Blackmoon

FOX IT

TU Delft

Beyond Labeling: Using Clustering to Build Network Behavioral Profiles of Malware Families. Azqa Nadeem, Christian Hammerschmidt, Carlos H. Ganan, Sicco Verwer. In Malware Analysis using Artificial Intelligence and Deep Learning, Springer, 2021.

Similarity analysis

- Distance calculation with distortions in sequences

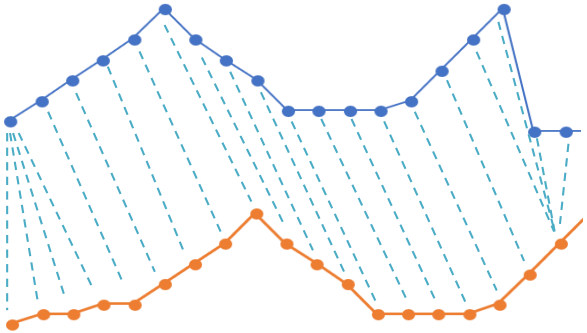

Dynamic Time Warping

+

HDBSCAN clustering

Behavior (Cluster) analysis

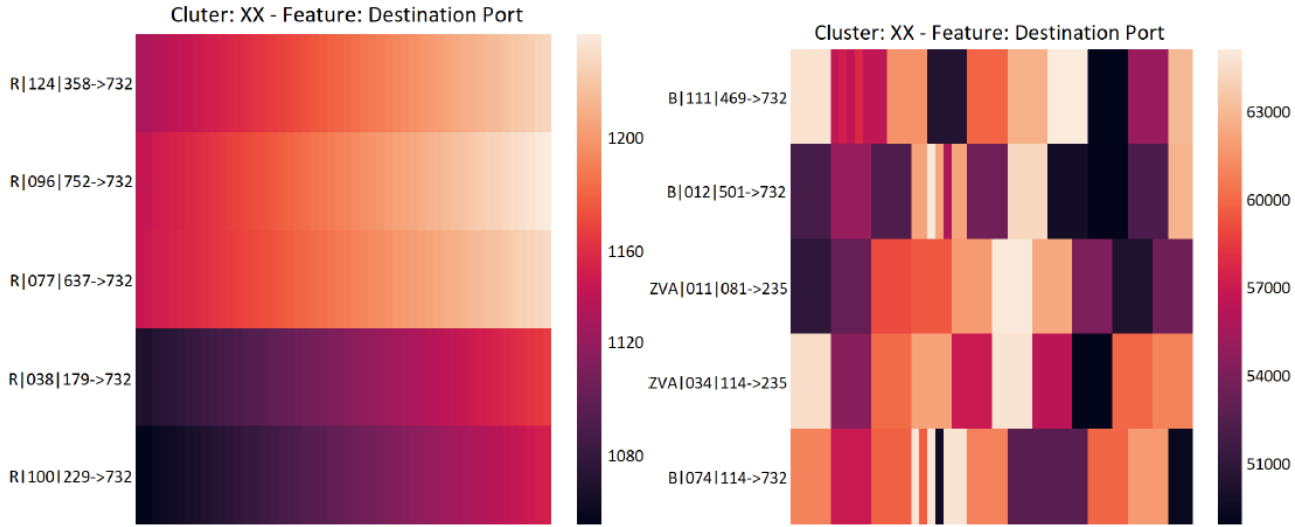

(a) Systematic port scan.

(b) Randomized port scan.

Malware Behavior Profiles

Behavior catalog

	B	C	D	DL	GE	GI	R	Z	ZP	ZPa	Zv1	ZVA
SSDP traffic	✓	✓	✓	✓	✓	✓	✓	✓	-	✓	-	✓
Broadcast traffic	✓	✓	-	✓	-	✓	✓	-	✓	-	✓	✓
LLMNR traffic	✓	✓	-	✓	-	✓	-	-	-	-	-	-
System. port scan	✓	✓	-	-	-	✓	✓	-	-	-	-	✓
Random. port scan	✓	✓	-	-	-	✓	✓	-	-	-	-	✓
In conn spam	-	-	-	-	-	✓	-	-	-	-	-	-
Out conn spam	-	-	-	-	-	✓	-	-	-	-	-	-
Malicious Subnet	-	-	-	-	-	-	-	-	-	-	-	✓
In HTTPs	-	✓	-	✓	-	✓	-	-	-	✓	-	-
Out HTTPs	-	-	-	-	-	✓	-	-	-	✓	-	-
C&C reuse	✓	-	-	-	-	-	-	-	-	✓	-	-
Misc.	✓	✓	-	✓	-	✓	-	✓	-	✓	-	✓
# Clusters	7	11	***1***	8	***1***	***16***	4	2	***1***	7	***1***	7

FOX IT

TU Delft

Beyond Labeling: Using Clustering to Build Network Behavioral Profiles of Malware Families. Azqa Nadeem, Christian Hammerschmidt, Carlos H. Ganan, Sicco Verwer. In Malware Analysis using Artificial Intelligence and Deep Learning, Springer, 2021.

Malware Behavior Profiles

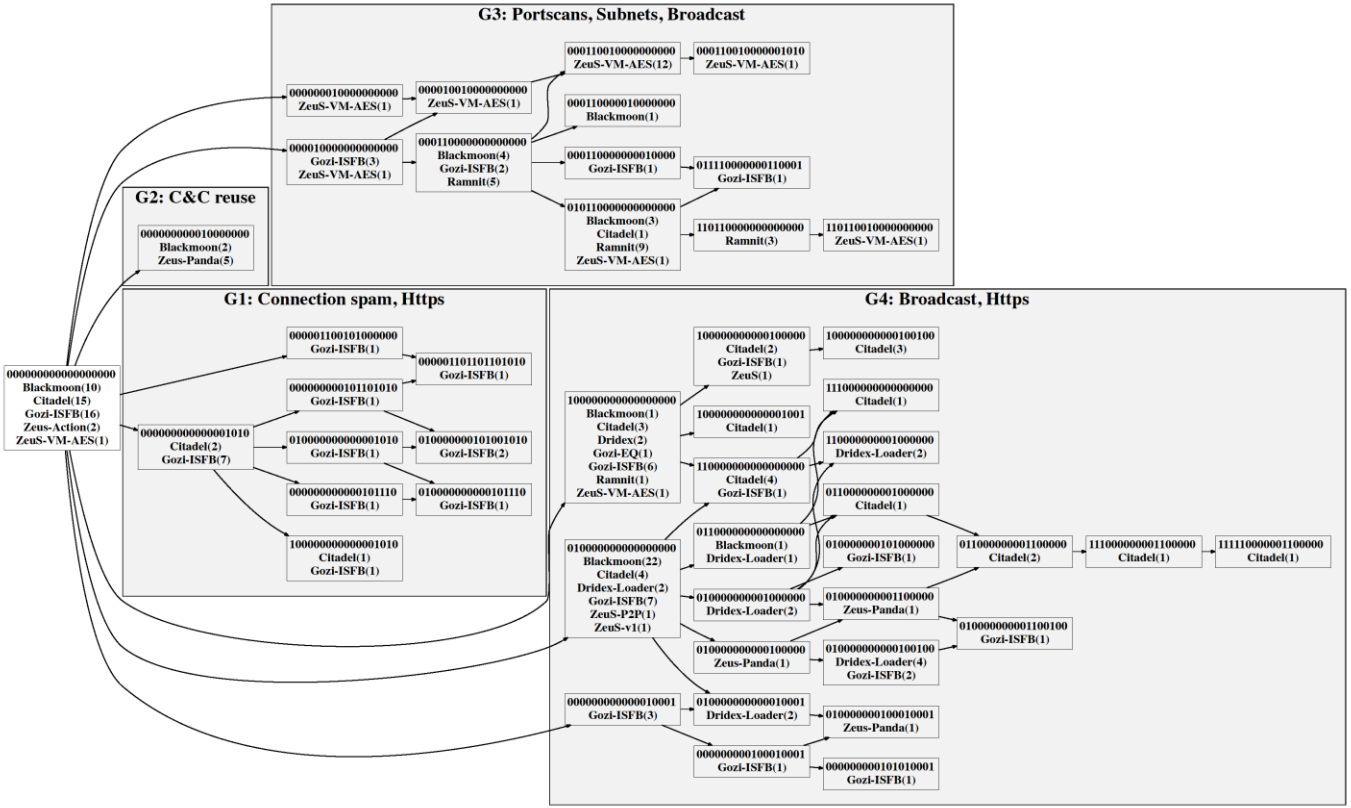

Beyond Labeling: Using Clustering to Build Network Behavioral Profiles of Malware Families. Azqa Nadeem, Christian Hammerschmidt, Carlos H. Ganan, Sicco Verwer. In Malware Analysis using Artificial Intelligence and Deep Learning, Springer, 2021.

MalPaCA: Summary & Future work

- Malware family names → noisy & inconsistent
- MalPaCA for building behavioral profiles
  - Clustering multivariate packet sequences
  - Clusters → Malware behavior catalog
  - Malware behavioral profile → cluster membership

- Network + system behavior catalog
- Continual learning
- Adversarial robustness

TU Delft

Beyond Labeling: Using Clustering to Build Network Behavioral Profiles of Malware Families. Azqa Nadeem, Christian Hammerschmidt, Carlos H. Ganan, Sicco Verwer. In Malware Analysis using Artificial Intelligence and Deep Learning, Springer, 2021.

Wrap-up

- Goal: Learning attacker behavior from temporal data
- Unsupervised setting with limited prior knowledge
- Input: Observables | Output: Intelligence

- 2 use-cases
  - Intrusion alerts → Attacker strategies via attack graphs
  - Network traffic → Malware behavior profiles

TU Delft

Questions?

Goal: Learning attacker behavior from temporal data
Unsupervised setting with limited prior knowledge
Input: Observables | Output: Intelligence

2 use-cases
- Intrusion alerts → Attacker strategies via attack graphs
- Network traffic → Malware behavior profiles

azqa.nadeem@tudelft.nl

https://cyber-analytics.nl/

TU Delft