Data Exfiltration Detection and Prevention: Virtually Distributed POMDPs for Practically Safer Networks

Sara Mc Carthy*, Arunesh Sinha,* Milind Tambe*, Pratyusa Manadhata**

*University of Southern California**
*Hewlett Packard Labs***

TOP 3 CYBER THREATS

facing organizations in 2016:

52%
Social Engineering

40%
Insider Threats

39%
Advanced Persistent Threat

SOURCE:
ISACA'S JANUARY 2016 CYBERSECURITY SNAPSHOT, GLOBAL DATA,
WWW.ISACA.ORG/2016-CYBERSECURITY-SNAPSHOT

CSX CYBERSECURITY NEXUS

ISACA®
Trust in, and value from, information systems

MOTIVATION

ADVANCED PERSISTENT THREATS

Advanced	• Attackers are **sophisticated** and **intelligent**, with large set of **resources.** • Use human ability and creativity, not just bots or worms with continuous **monitoring** and **interaction**
Persistent	• **"Low-and-slow"** approach, operates quietly over an extended period of time maintain long-term access to the target,
Threat	• Goal is often **Data Exfiltration:** obtaining and extracting financial, technological, or other information.

MOTIVATION

DNS EXFILTRATION

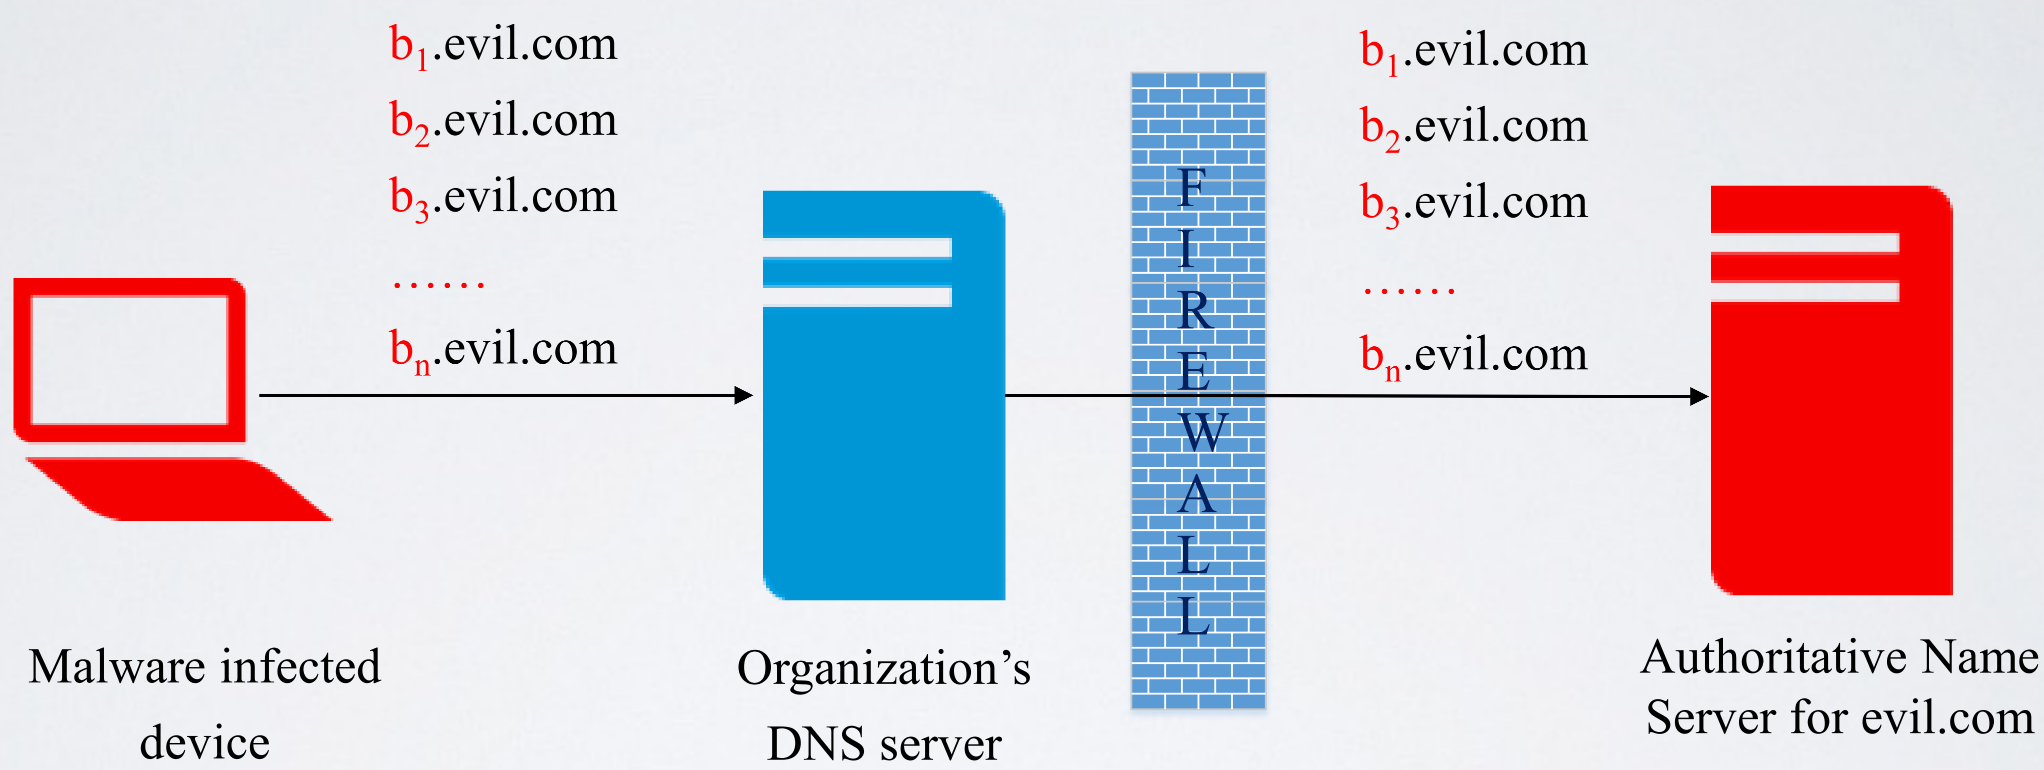

Legitimate

LONG QUERIES

5.1o19sr00ors95qo0p73415p3r8r8q777634r5o86osn295ss2rqoss3r9601ro3.1r1p7r4719o343936
48s2345nn60qnqoop45psos37n551s002n80850sr2r8n3.r1105qqq28r7pn82843rp76383qr6344qq
pq7rpnrp63o957687r980r.rrqs656p04pn614q6n76o97883op73r0p787rn92.i.02.s.sophosxl.net

g63uar2ejiq5tlrkg3zezf2fksjrxpxyvro4ce5yz65udnjn.dagbuu5pkocwcaxkntmxzwvkbulhg3qlj6ho7jw
obeddjqvv.gepxfdwfhu76on6gza2nkringxp35e6g3ftpqlpl5h6uofgo.kukjy4jvybu7jhrlhrgxe7es3lmkxd
rpmpb4lg7wmbpygjg7.gef2uoemc6pi88tz.er.spotify.com

REPEATED QUERIES

1751913.86c0ade0d13143ab83d7e4f60cbd204c.00000000.xello.xobni.com
1753942.86c0ade0d13143ab83d7e4f60cbd204c.00000000.xello.xobni.com
1756950.86c0ade0d13143ab83d7e4f60cbd204c.00000000.xello.xobni.com
1758762.86c0ade0d13143ab83d7e4f60cbd204c.00000000.xello.xobni.com

MALICIOUS

p9b-8-na-5w-2z3-djmu-7pk-qy-0-bok-re9-ym-v9h-av-njx-2es.info

Paxson, V., Christodorescu, M., Javed, M., Rao, J., Sailer, R., Schales, D., Stoecklin, M.P., Thomas, K., Venema, W., Weaver, N.: Practical comprehensive bounds on surreptitious communication over dns.

PROBLEM : CLASSICAL MACHINE LEARNING

Outlier Detection

High Cost of Error

Semantic Gap

PROBLEM : CLASSICAL MACHINE LEARNING

Outlier Detection

High Cost of Error

Semantic Gap

ML is good at identifying what is **similar** rather than discovering meaningful outliers

Lack of labelled attack data leads to too many false positives and alerts

Results in **alert fatigue**

PROBLEM : CLASSICAL MACHINE LEARNING

Outlier Detection

High Cost of Error

Semantic Gap

ML is good at identifying what is **similar** rather than discovering meaningful outliers

Lack of labelled attack data leads to too many false positives and alerts

Results in **alert fatigue**

Cost of any misclassification is extremely high compared to many other machine learning applications.

False positive requires spending time examining the reported incident.

False negatives cause serious damage

PROBLEM : CLASSICAL MACHINE LEARNING

Outlier Detection

ML is good at identifying what is **similar** rather than discovering meaningful outliers

Lack of labelled attack data leads to too many false positives and alerts

Results in **alert fatigue**

High Cost of Error

Cost of any misclassification is extremely high compared to many other machine learning applications.

False positive requires spending time examining the reported incident.

False negatives cause serious damage

Semantic Gap

How to transfer results into actionable reports for the network operator

What remedial **steps should be taken?**

PROBLEM : DECISION THEORY

MDP

POMDP

Can reason about uncertainty in environment and provide **actions to take**

Need Complex Models to Capture real world dynamics

May be infeasible to generate model from domain experts

Can become extremely difficult to solve / not scalable

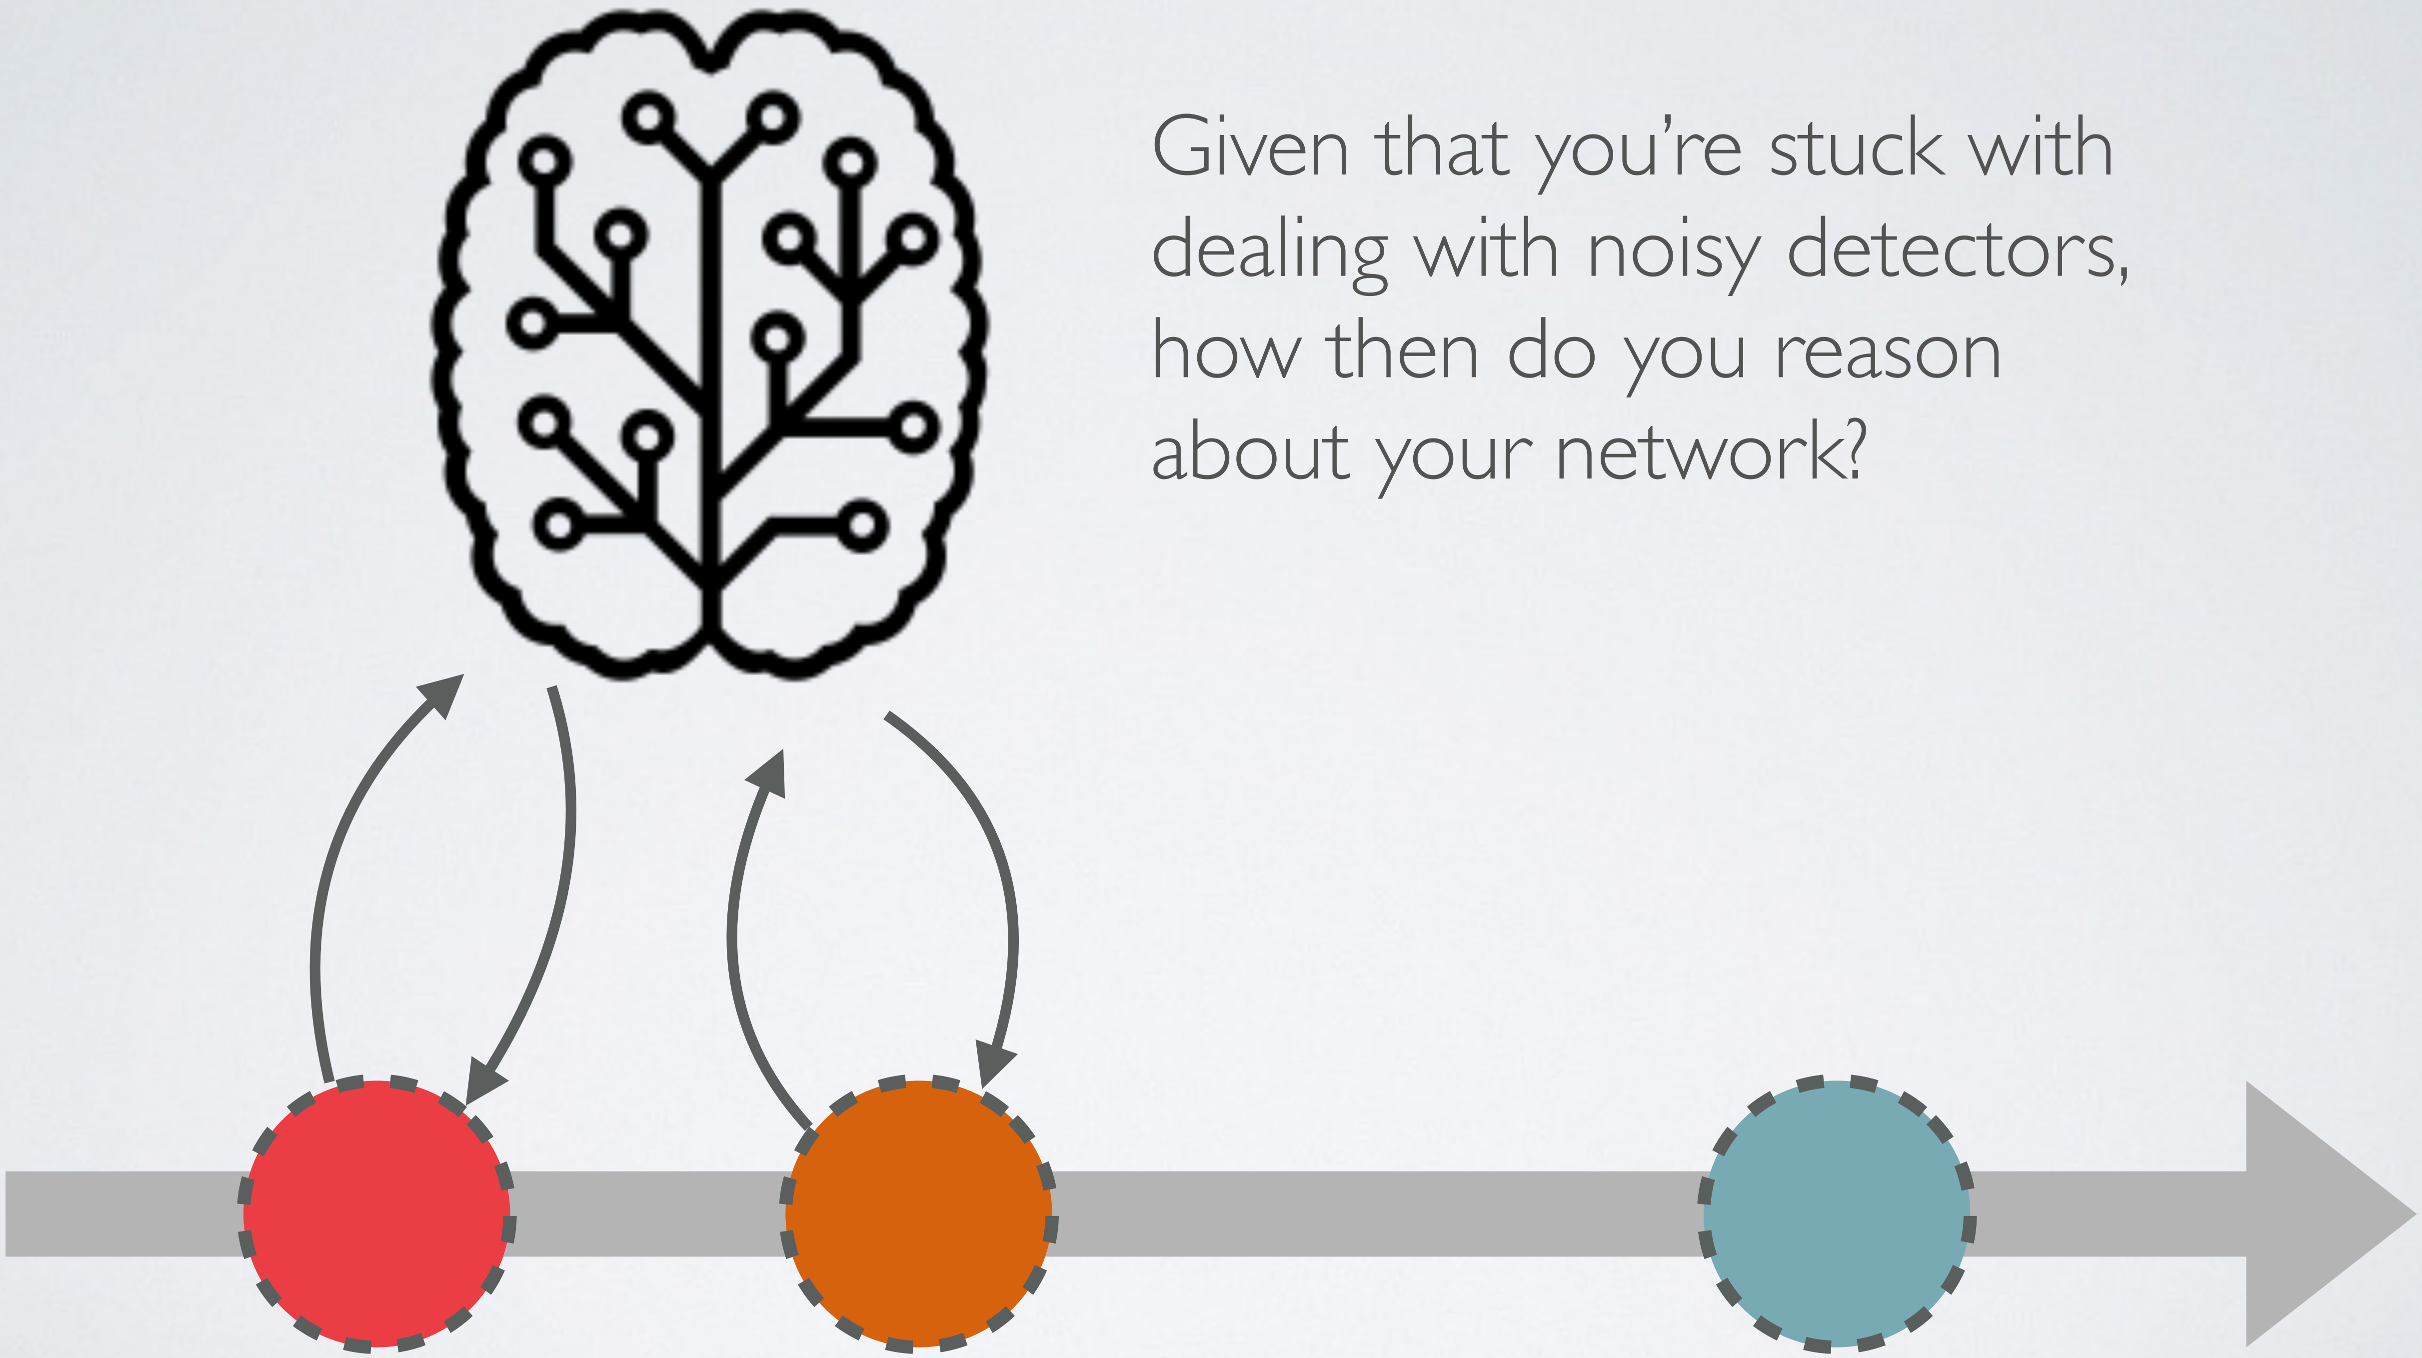

Given that you're stuck with dealing with noisy detectors, how then do you reason about your network?

CONTRIBUTIONS

- Address the Active Sensing challenge with a scalable, fast decision-theoretic model for reasoning about noisy sensors in a computer network and determine optimal sensing strategies

- Provide a novel **VD-POMDP solution method** for solving this model

- Evaluation on a **real network testbed**

PROBLEM

Data Exfiltration over DNS

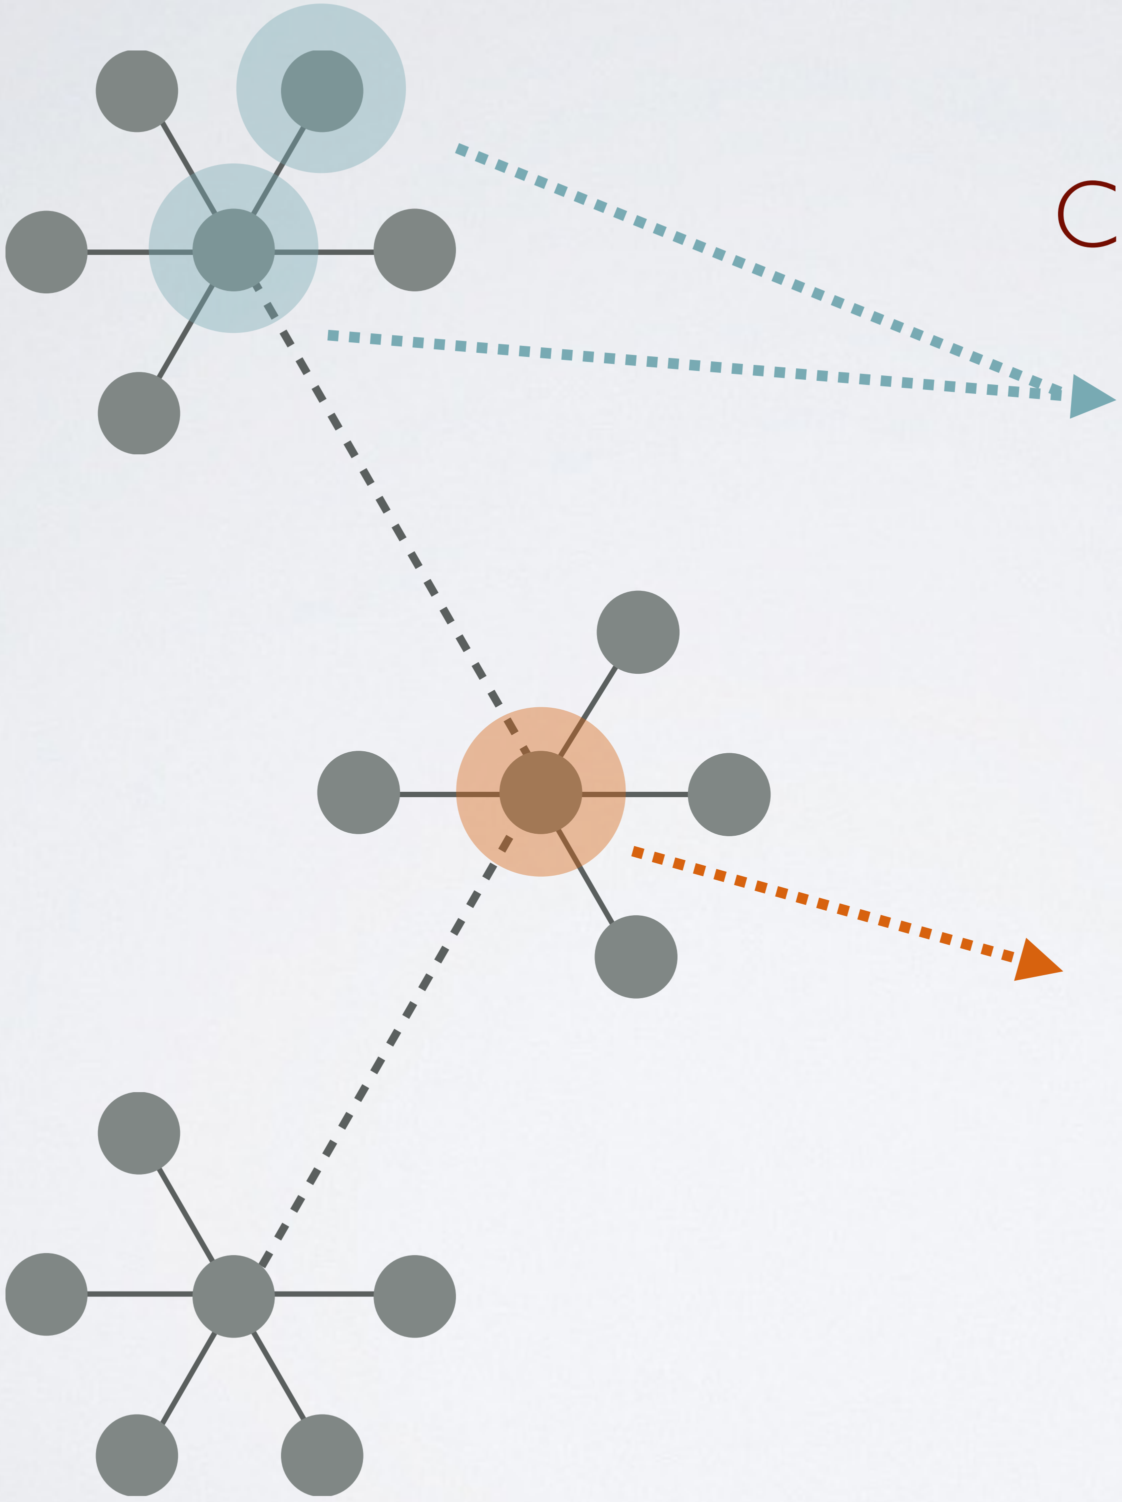

Computer Network

Nodes:
- Host Computer
- Routers, etc.

Edges:
- Channels of Communication

Acces Point:
- Nodes connected to the outside world (Internet)

PROBLEM

Data Exfiltration over DNS

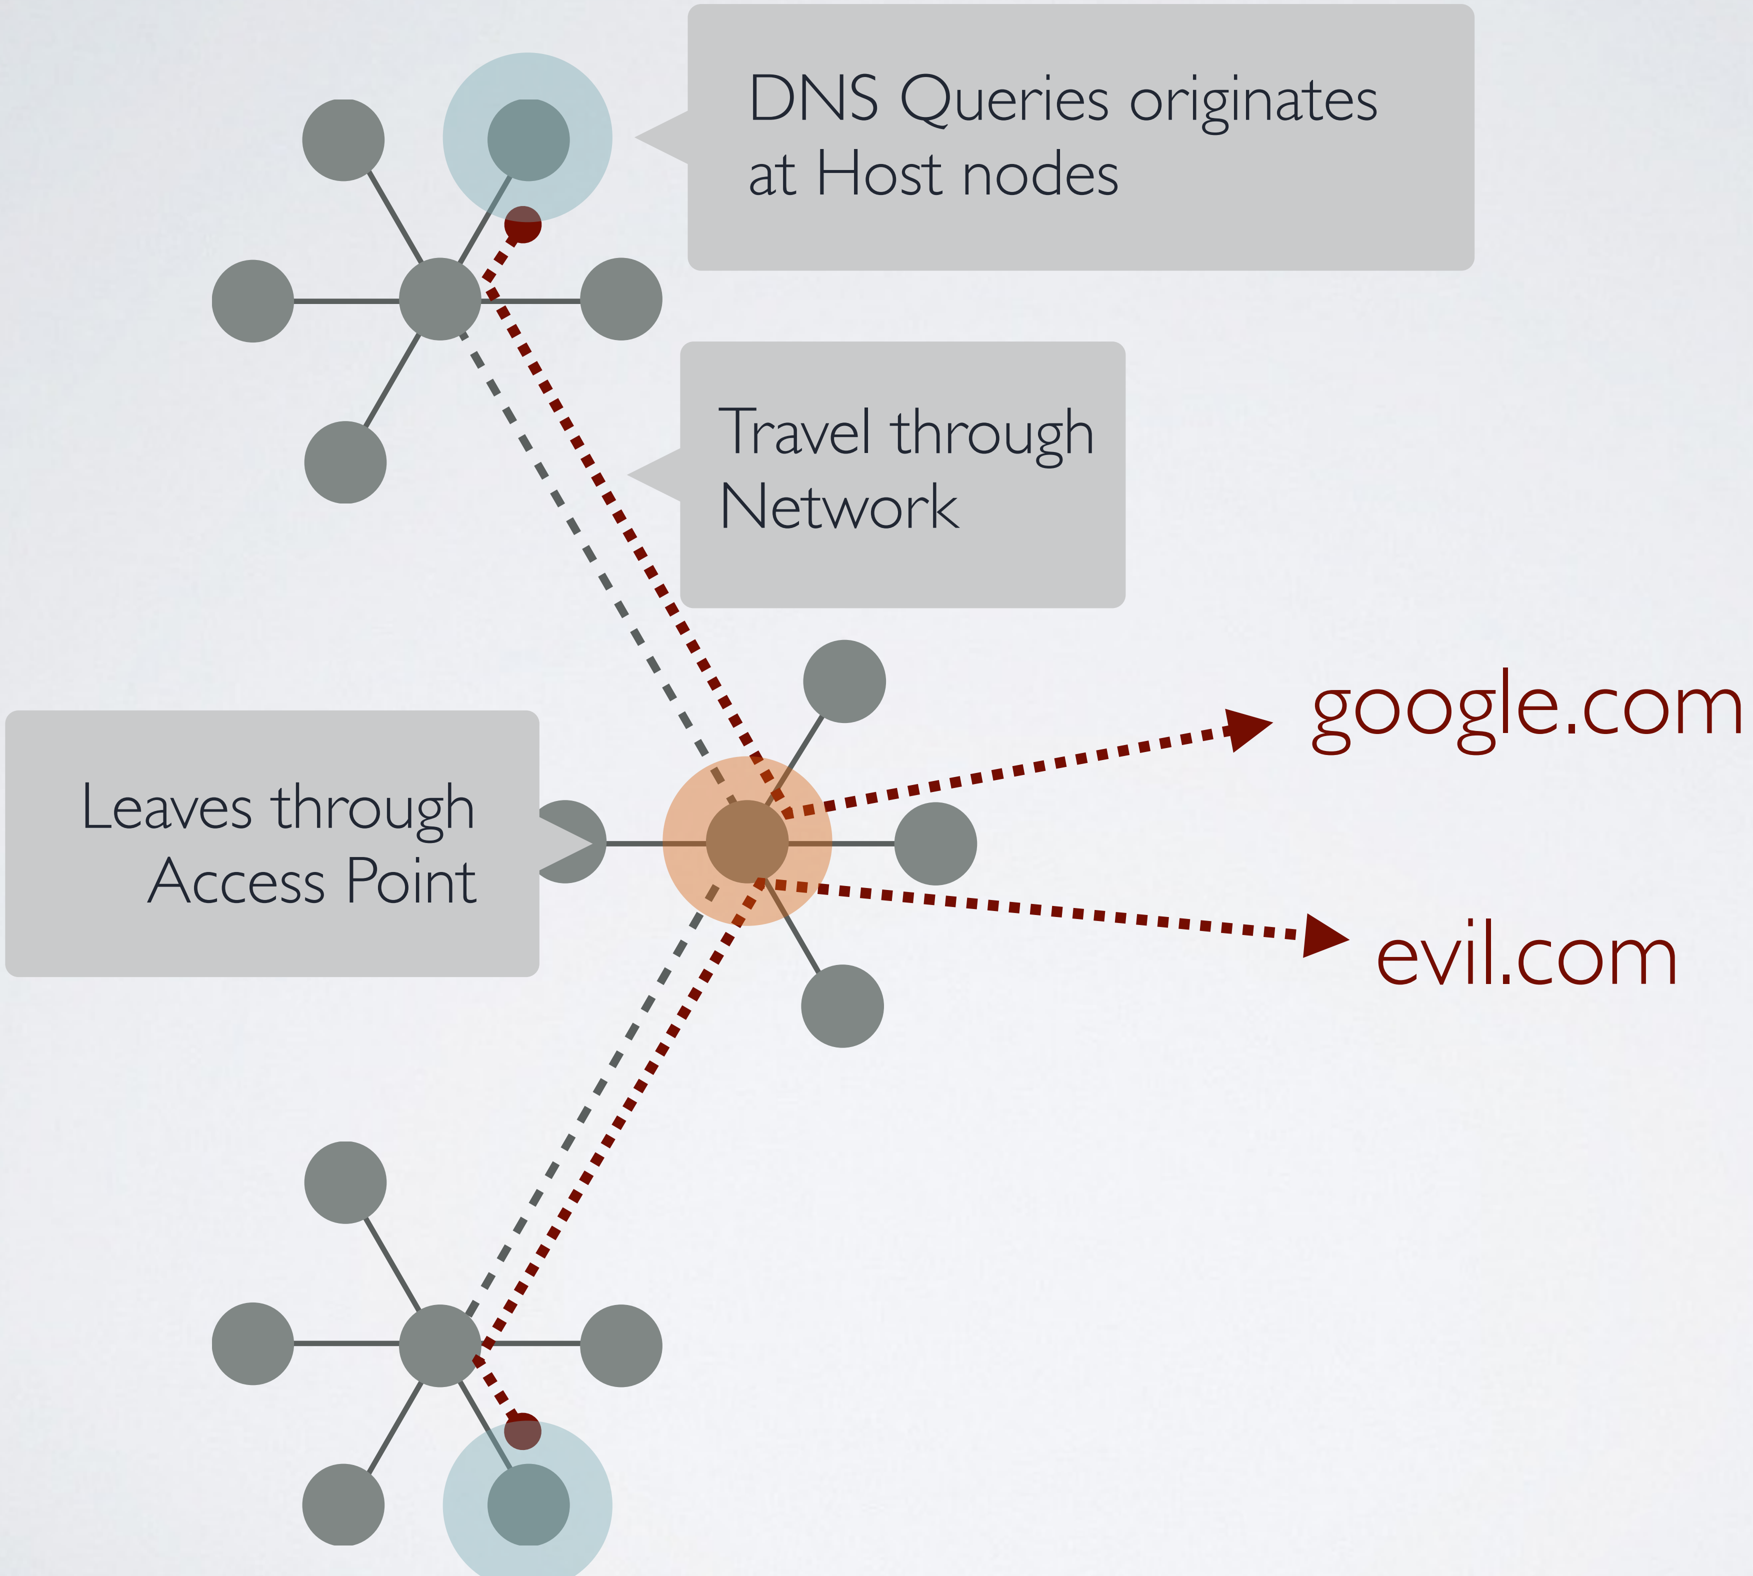

PROBLEM

Data Exfiltration over DNS

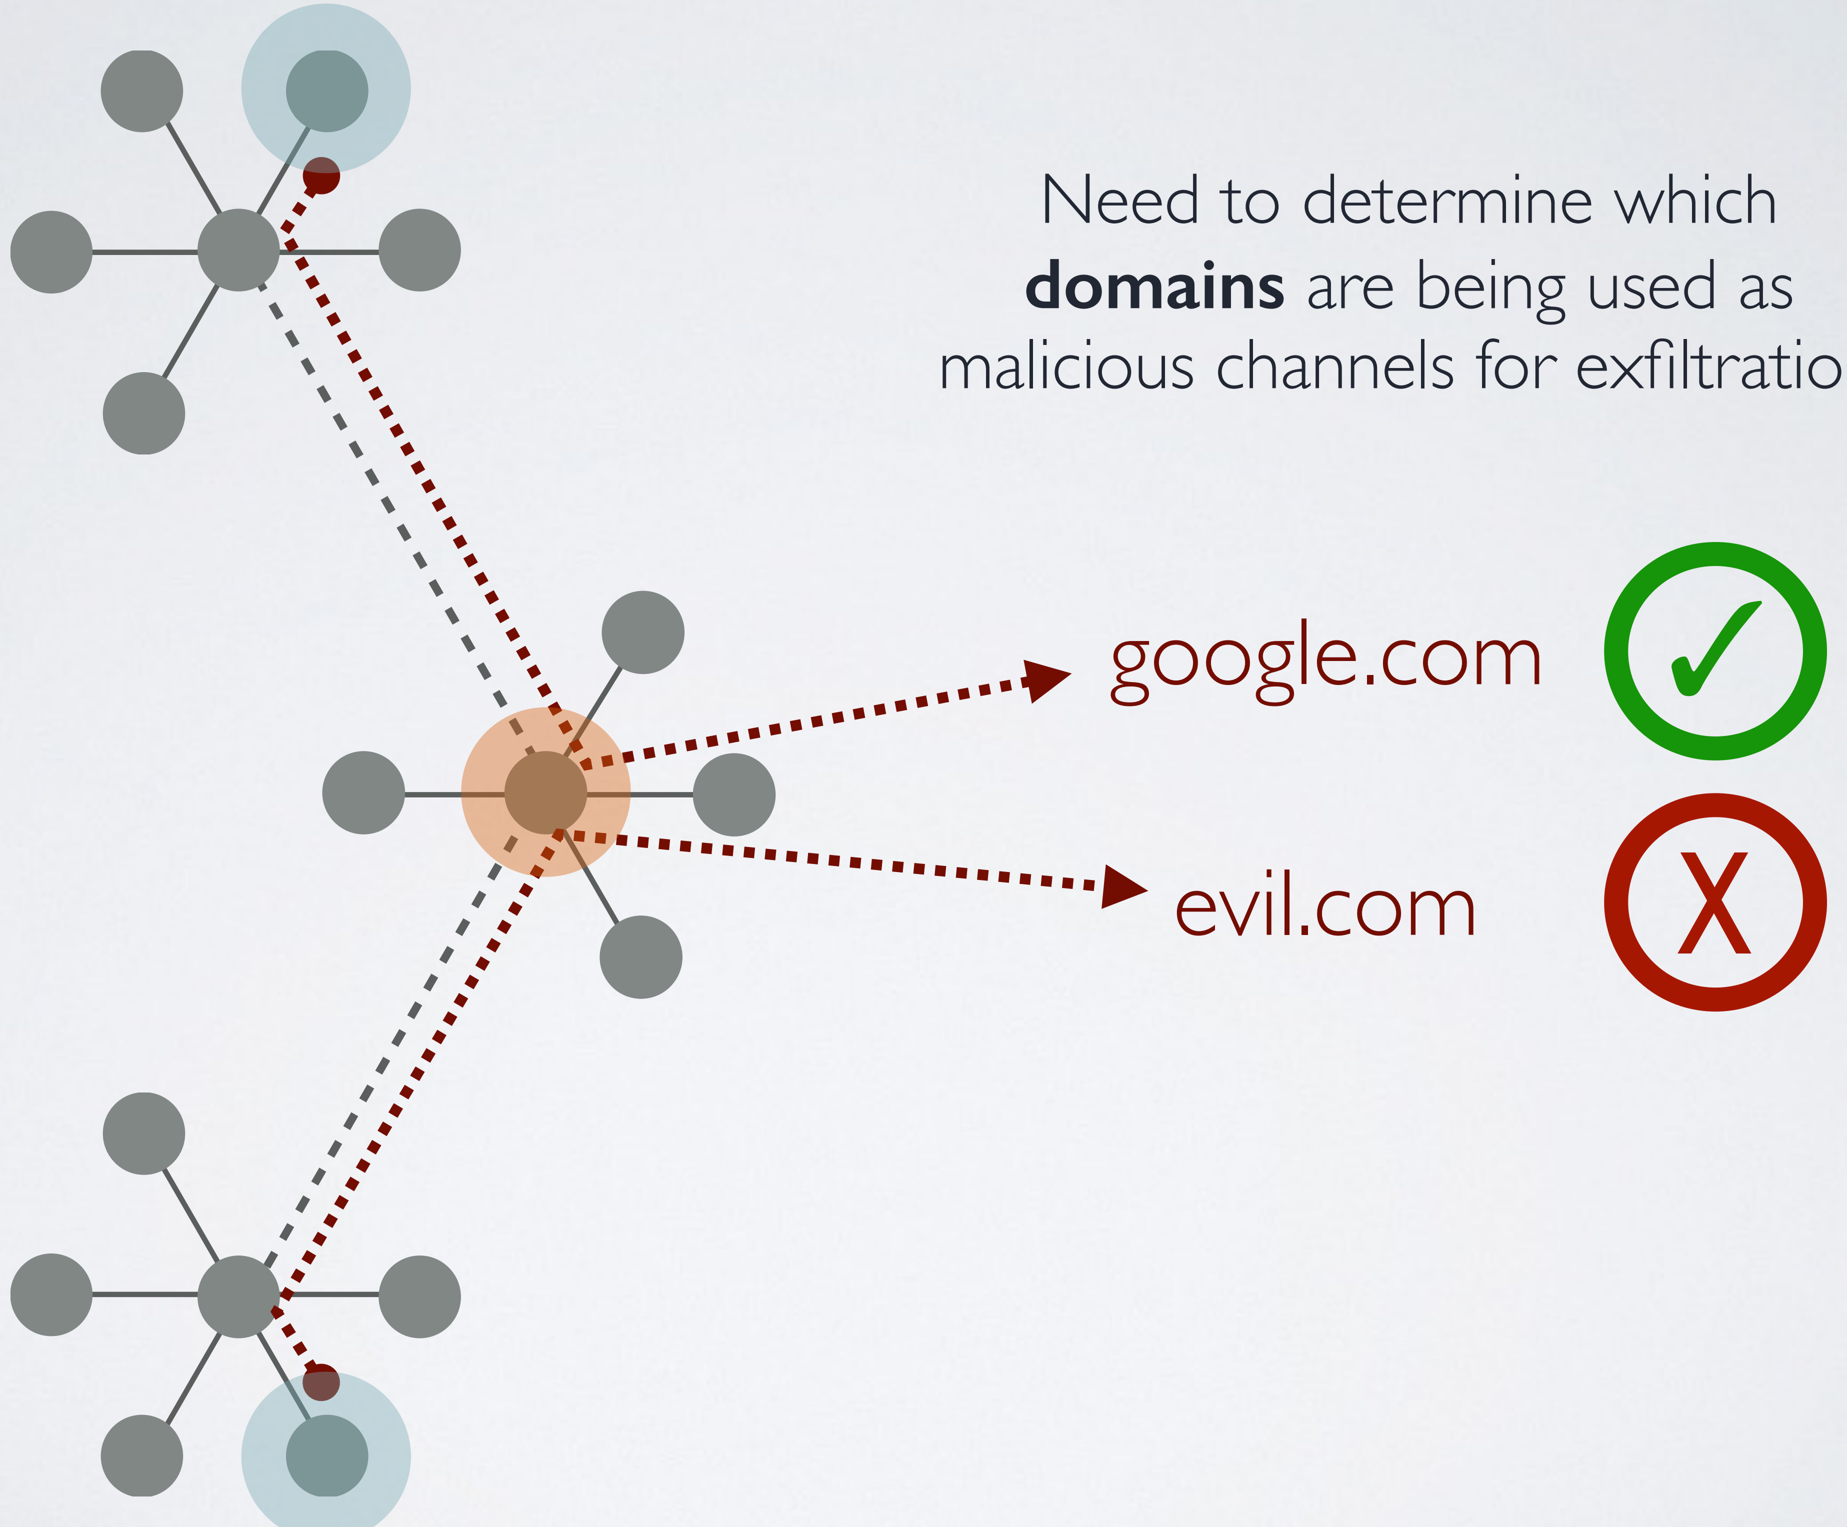

Need to determine which **domains** are being used as malicious channels for exfiltration

PROBLEM

Data Exfiltration over DNS

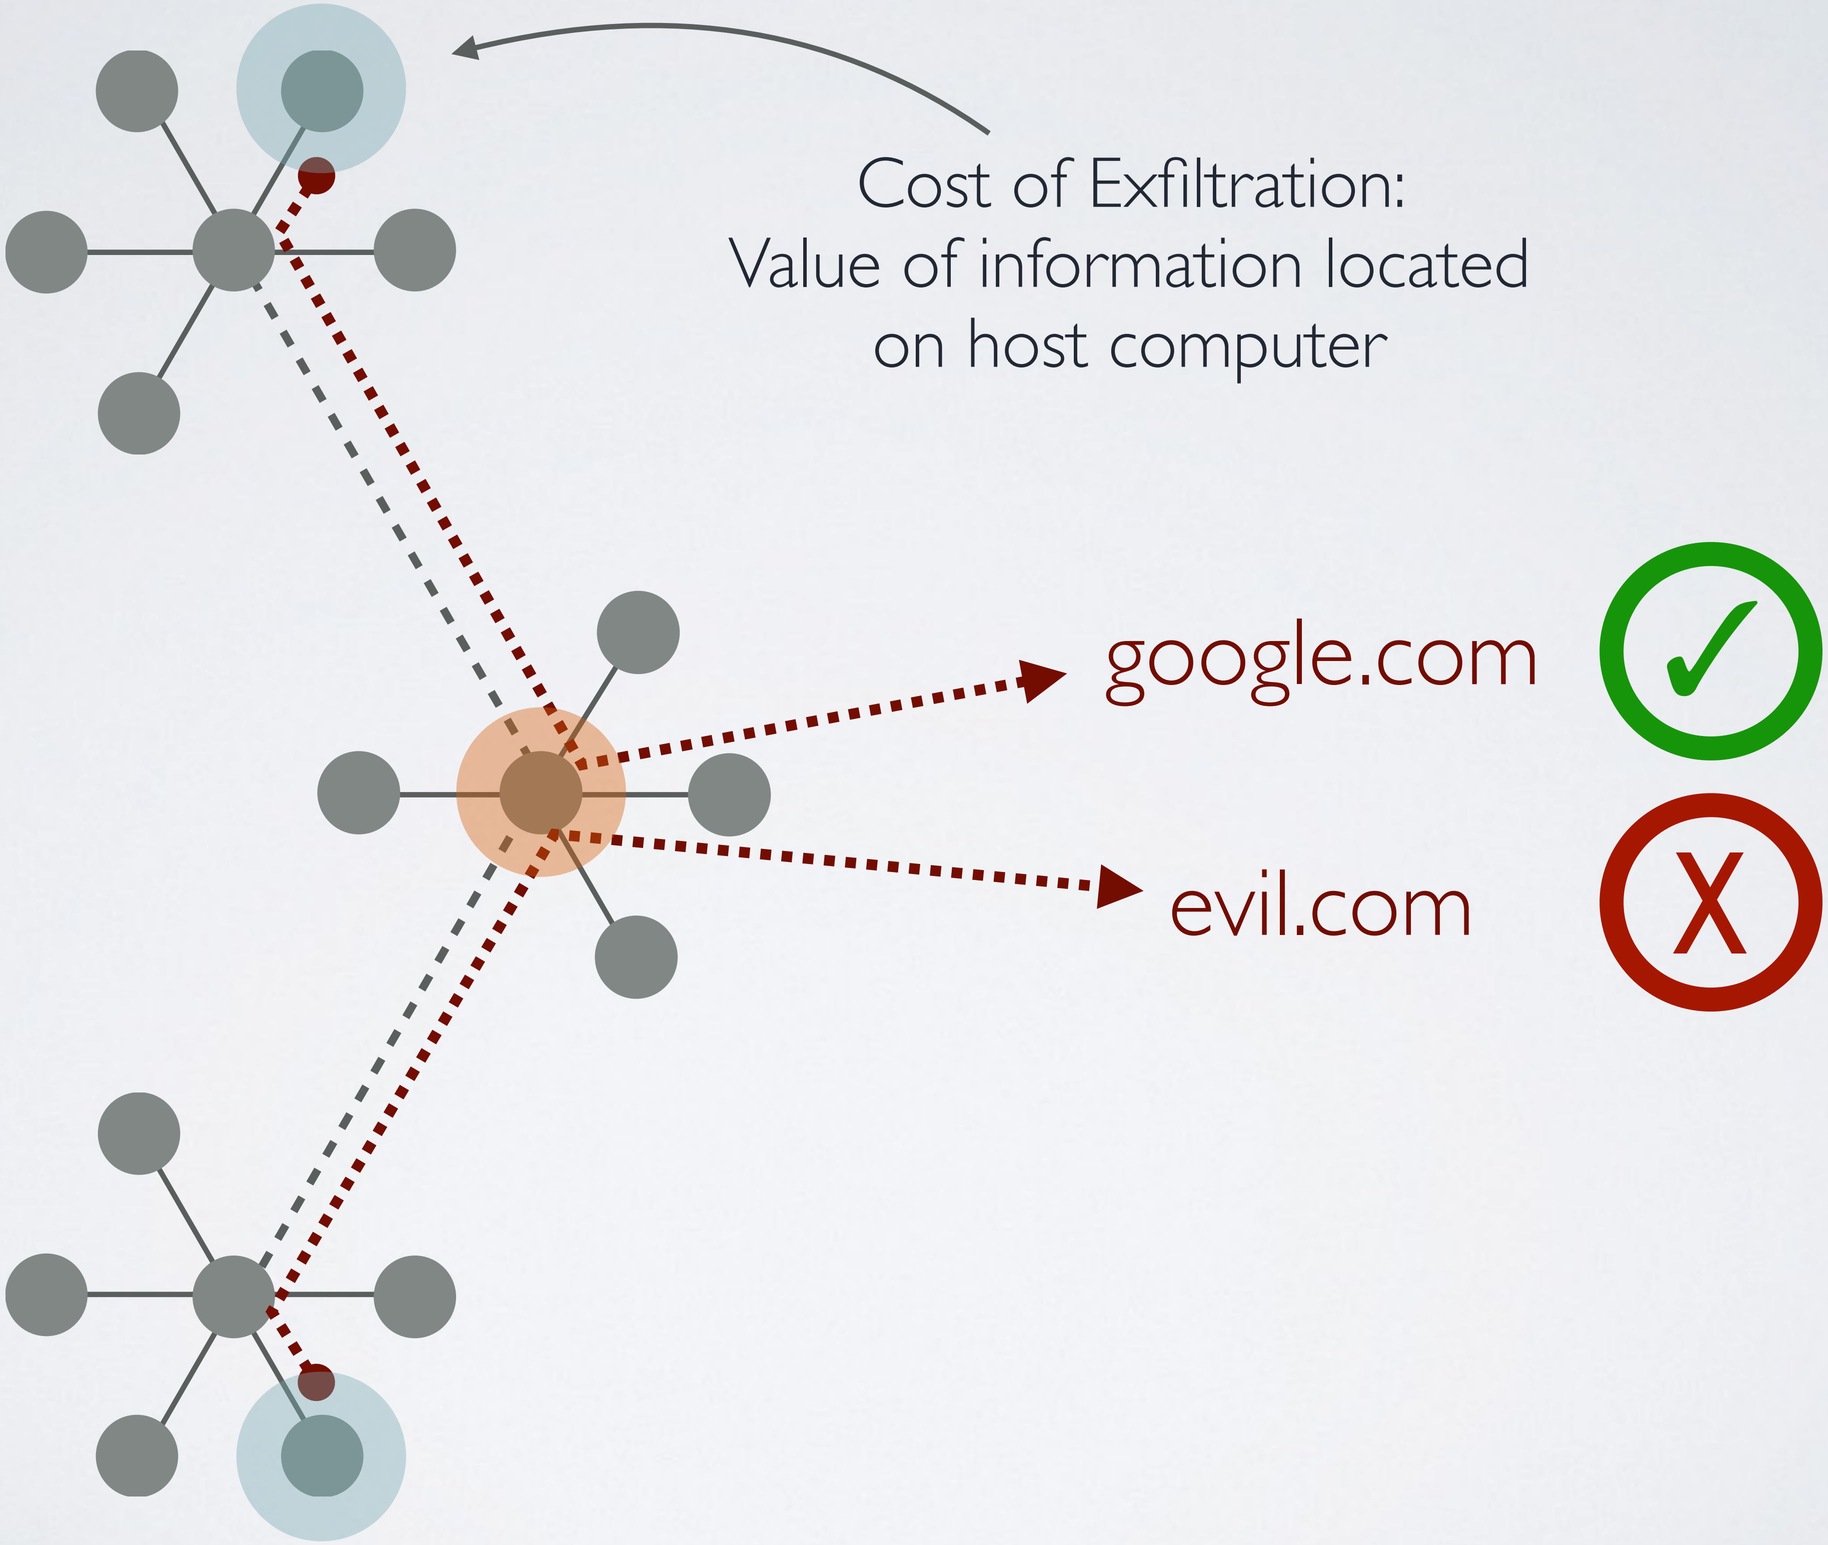

PROBLEM

Data Exfiltration over DNS

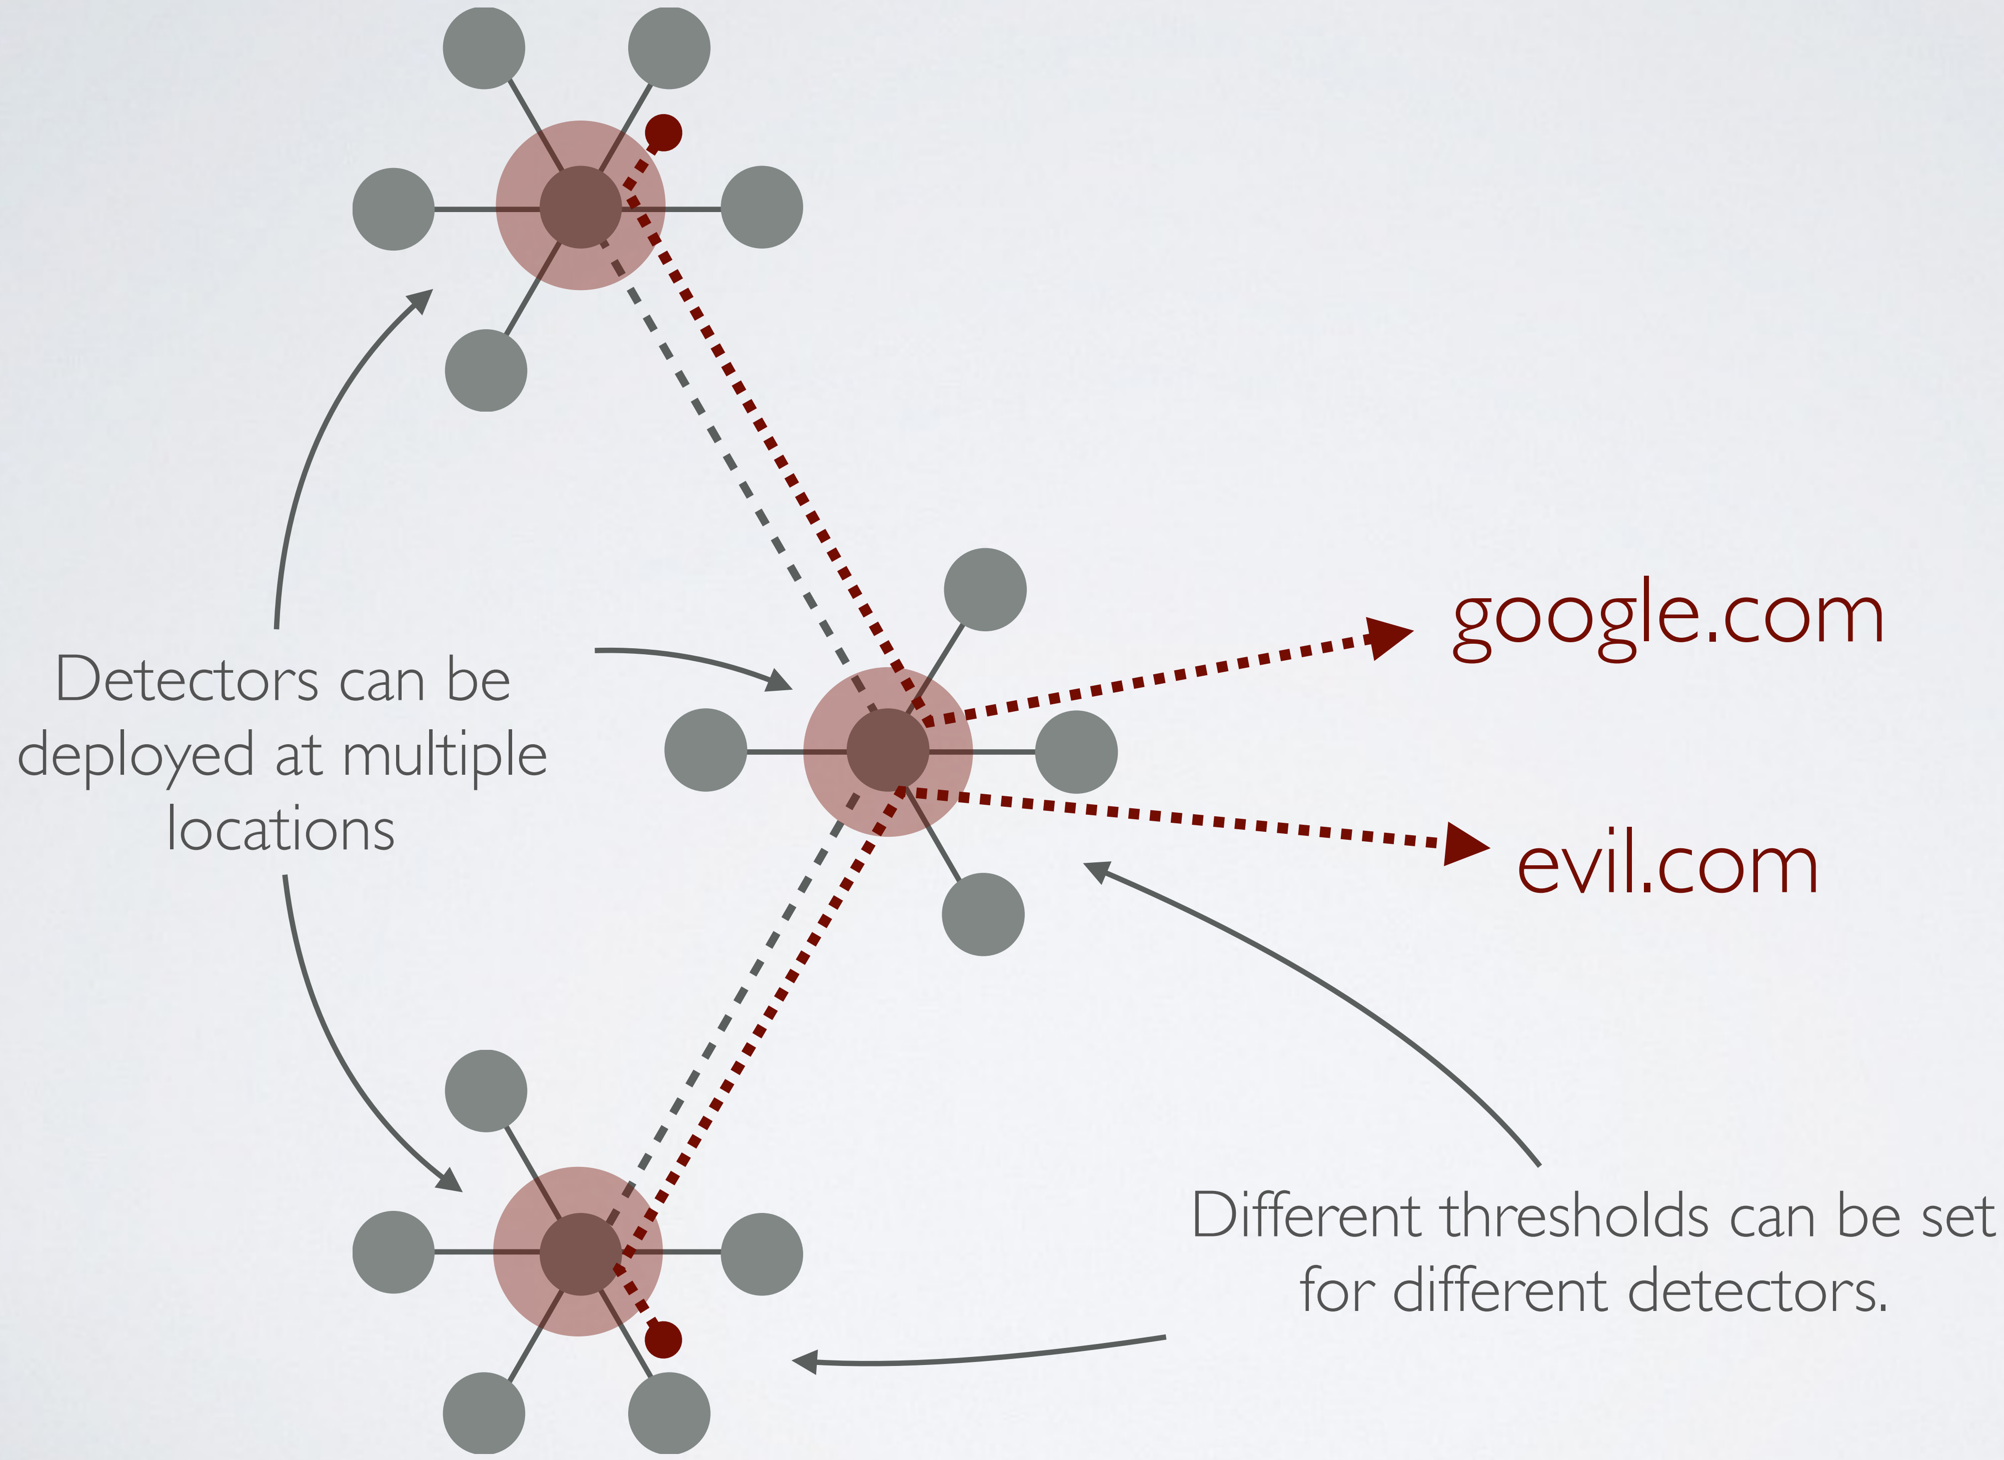

PROBLEM

Data Exfiltration over DNS

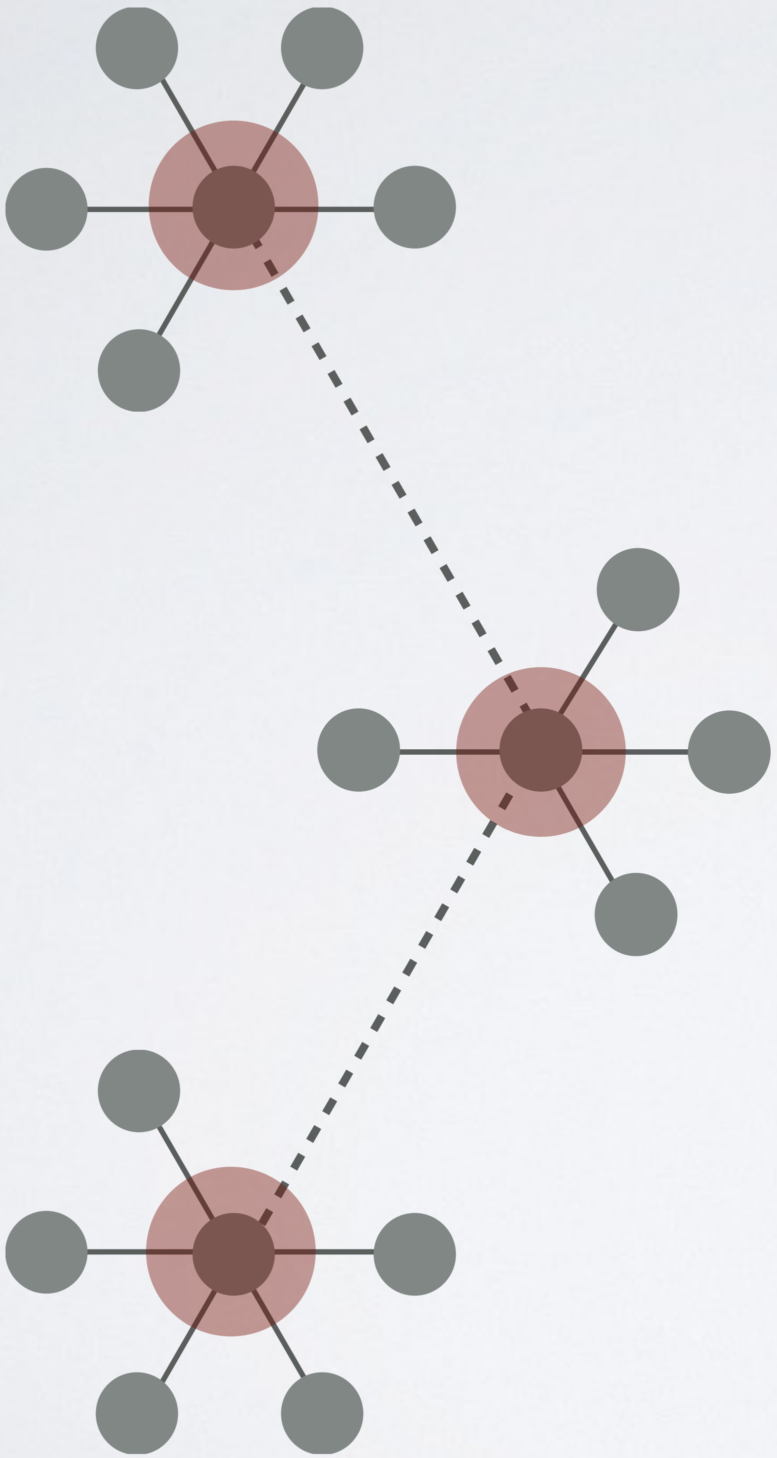

Gather Information

Detectors are imperfect:

- will often miss attacks
- have high false positive rates

Threshold determines miss rate

Need to gather information from many sources

Build up a belief over time about the network state

PROBLEM

Data Exfiltration over DNS

Ensure Network Performance

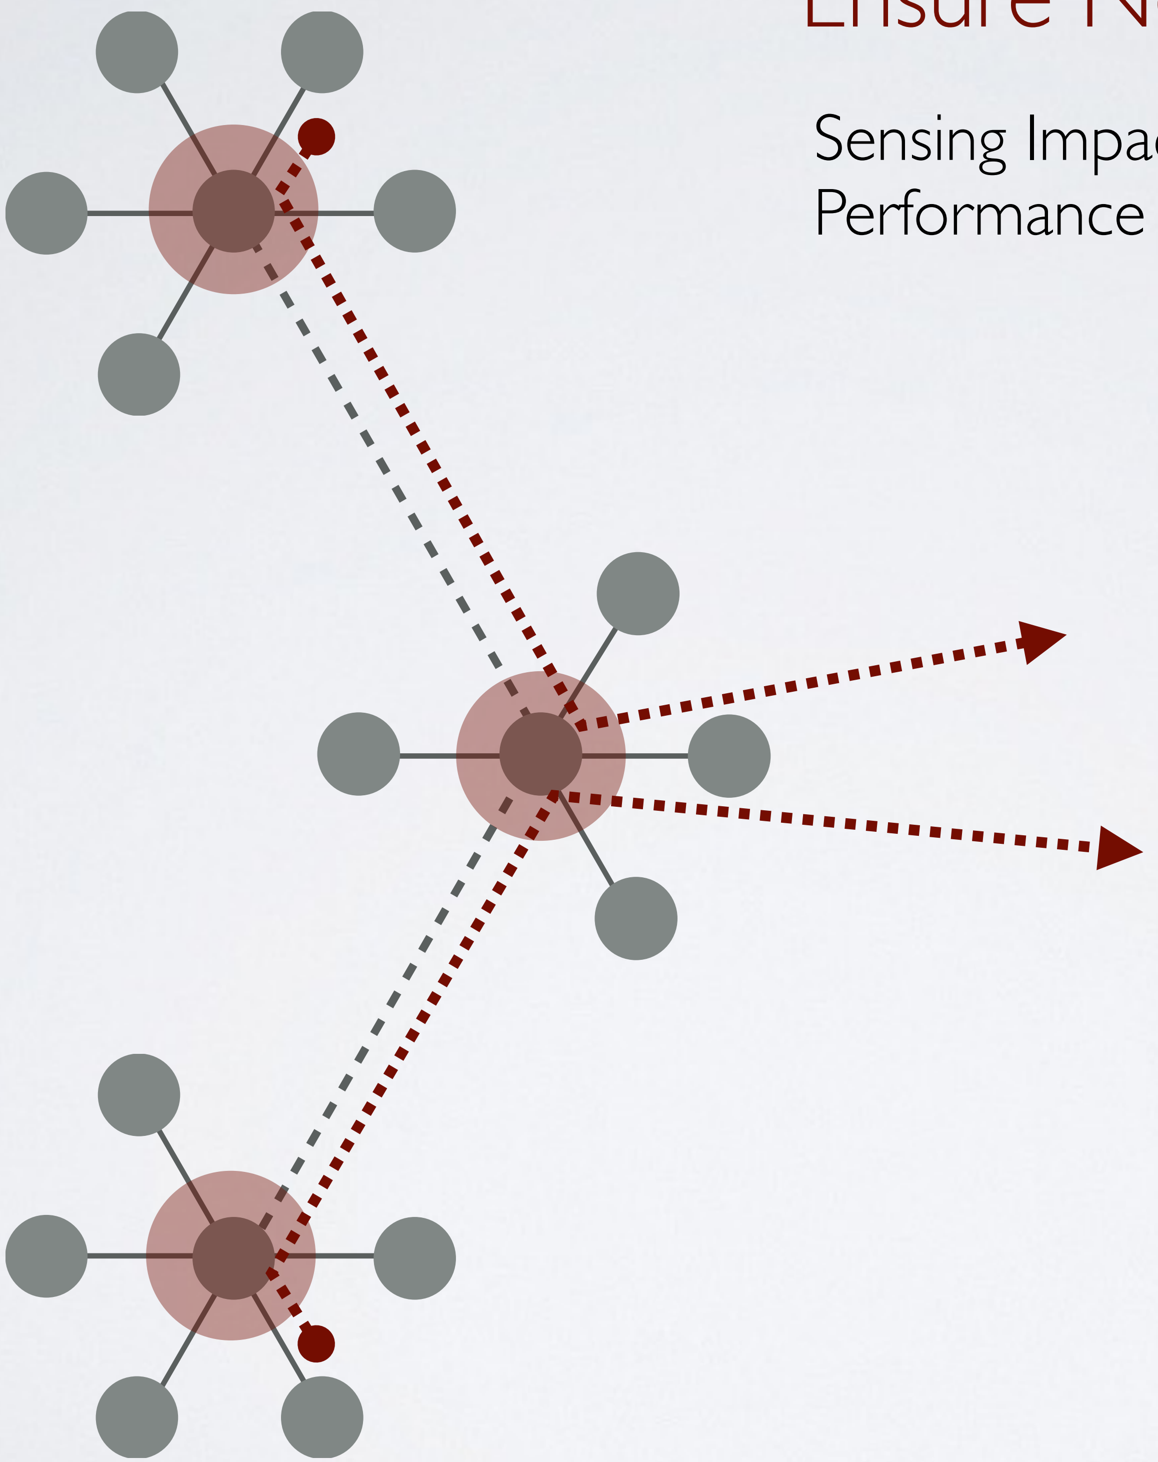

Sensing Impacts Network Performance

Cost of sensing is determined by amount of traffic through a node

VD-POMDP

Virtually distributed POMDP formulation

1. Factoring: Abstract the model to induce sparse interaction
   Divide and solve sub-POMDPs Offline

2. Policy Aggregation: resolve interactions online

3. Execute Joint Policy

1 Factor

Network representation couples channels & information about each domain

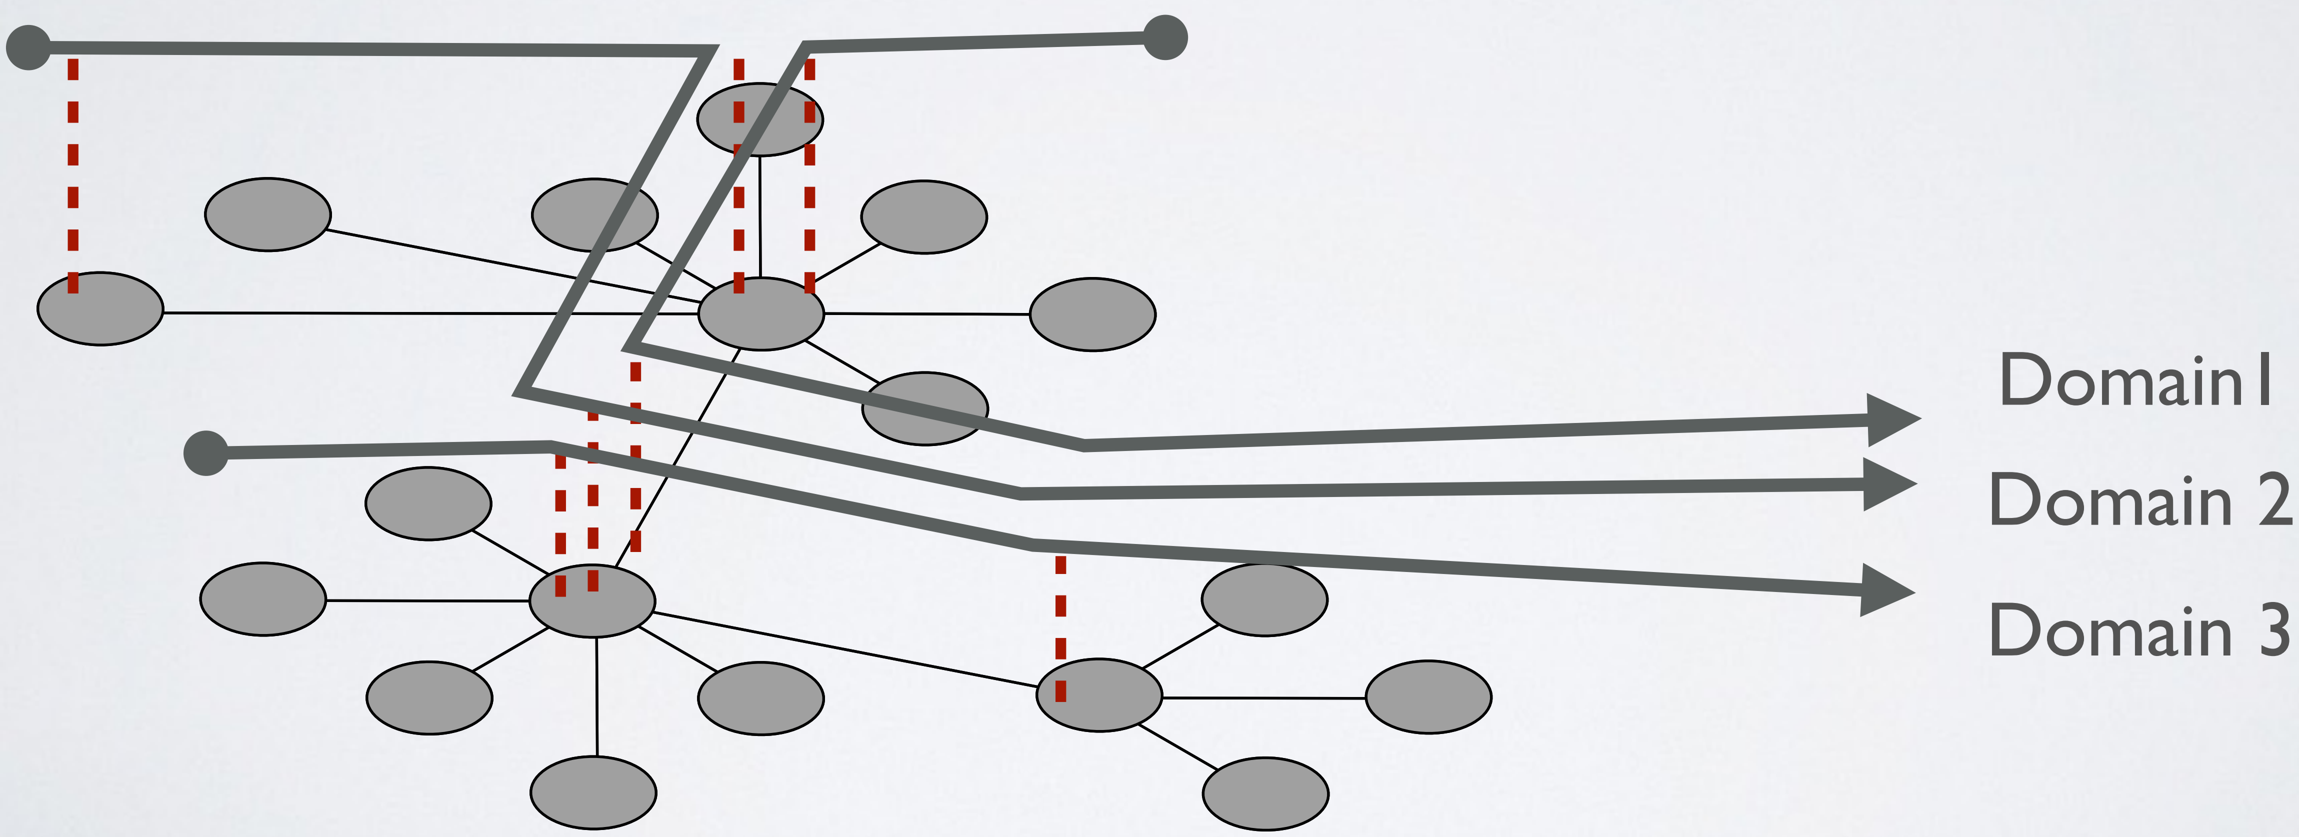

1 Factor

Network representation couples channels & information about each domain

Choice of node gives us information about many domains

Choice of channel gives information about one domain

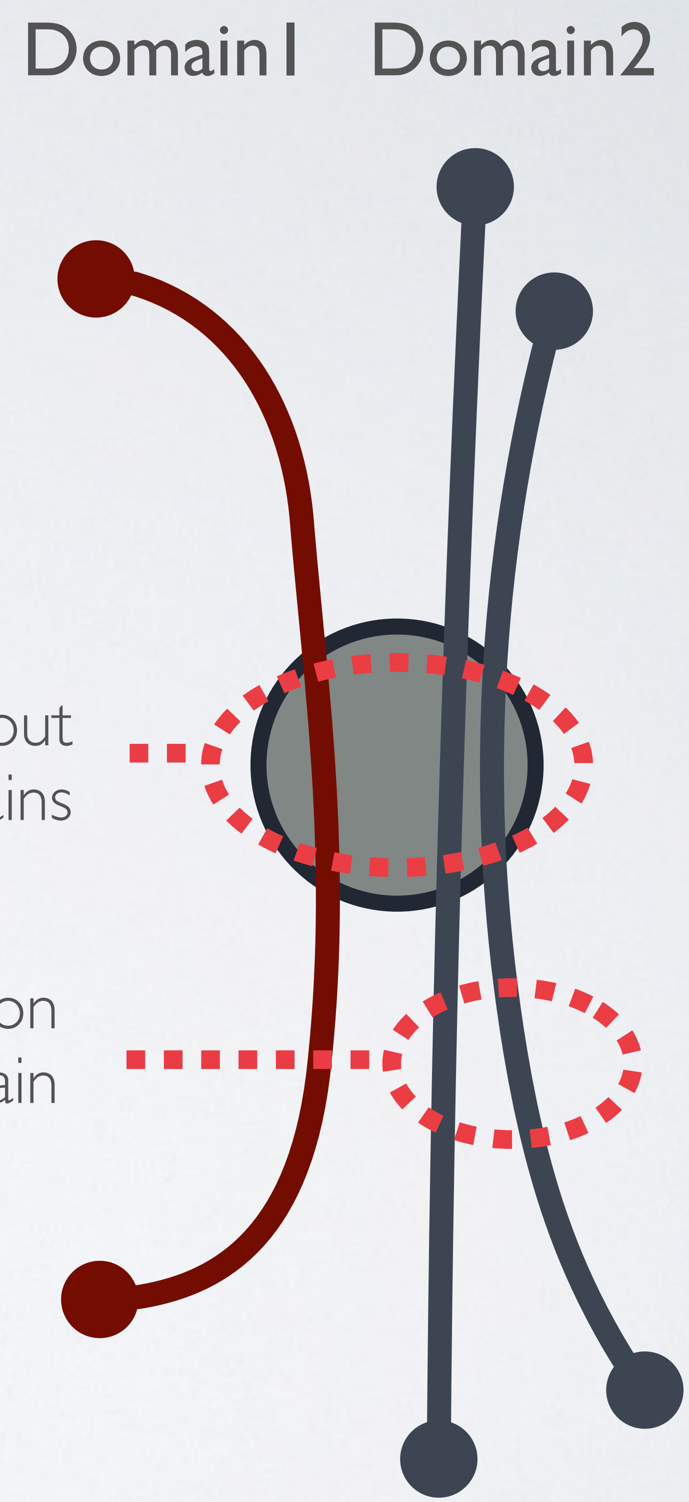

1 Factor

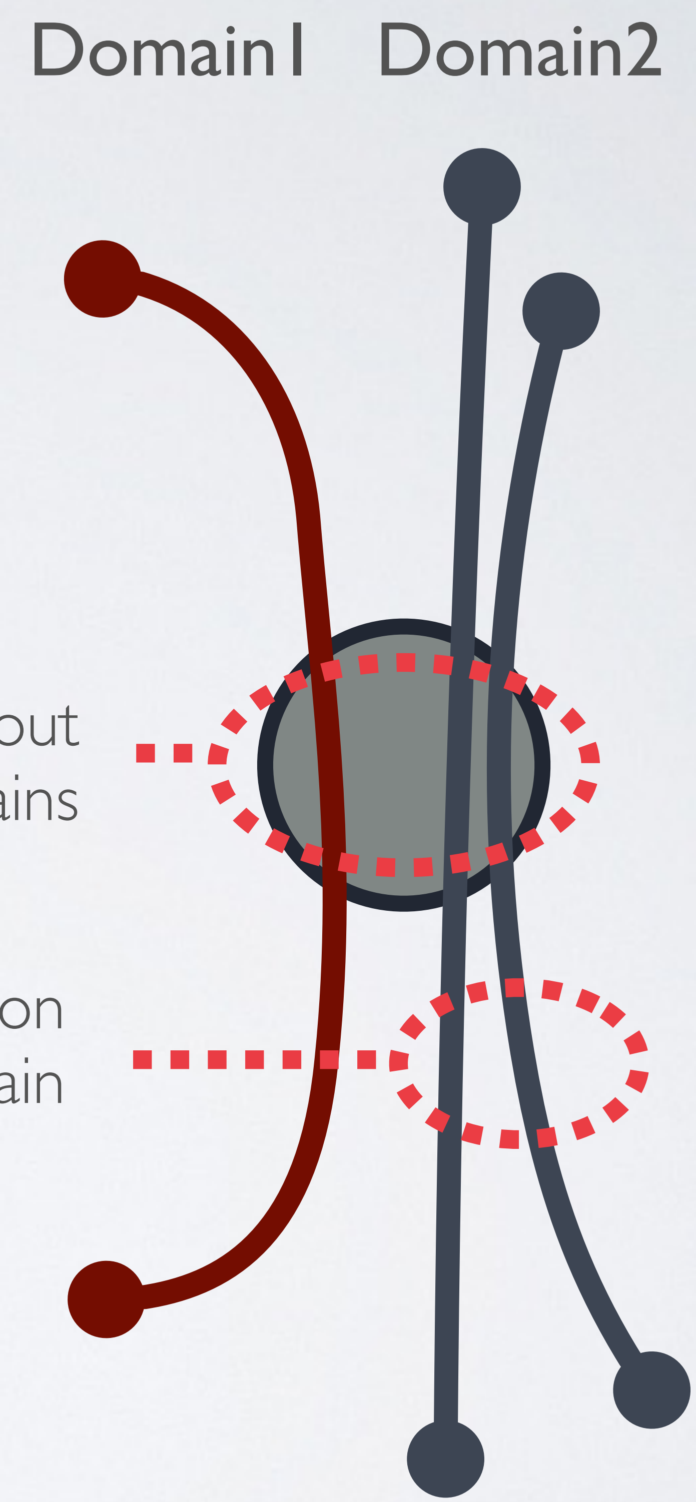

Reason about which channels to sense over instead of which nodes to sense on.

Choice of node gives us information about many domains

Choice of channel gives information about one domain

1 Factor

How does this change the cost of actions?

Cost of a node = Traffic through node
Build a lookup table mapping each action on channels to lowest cost action on nodes.
We can efficiently compute this using a linear program

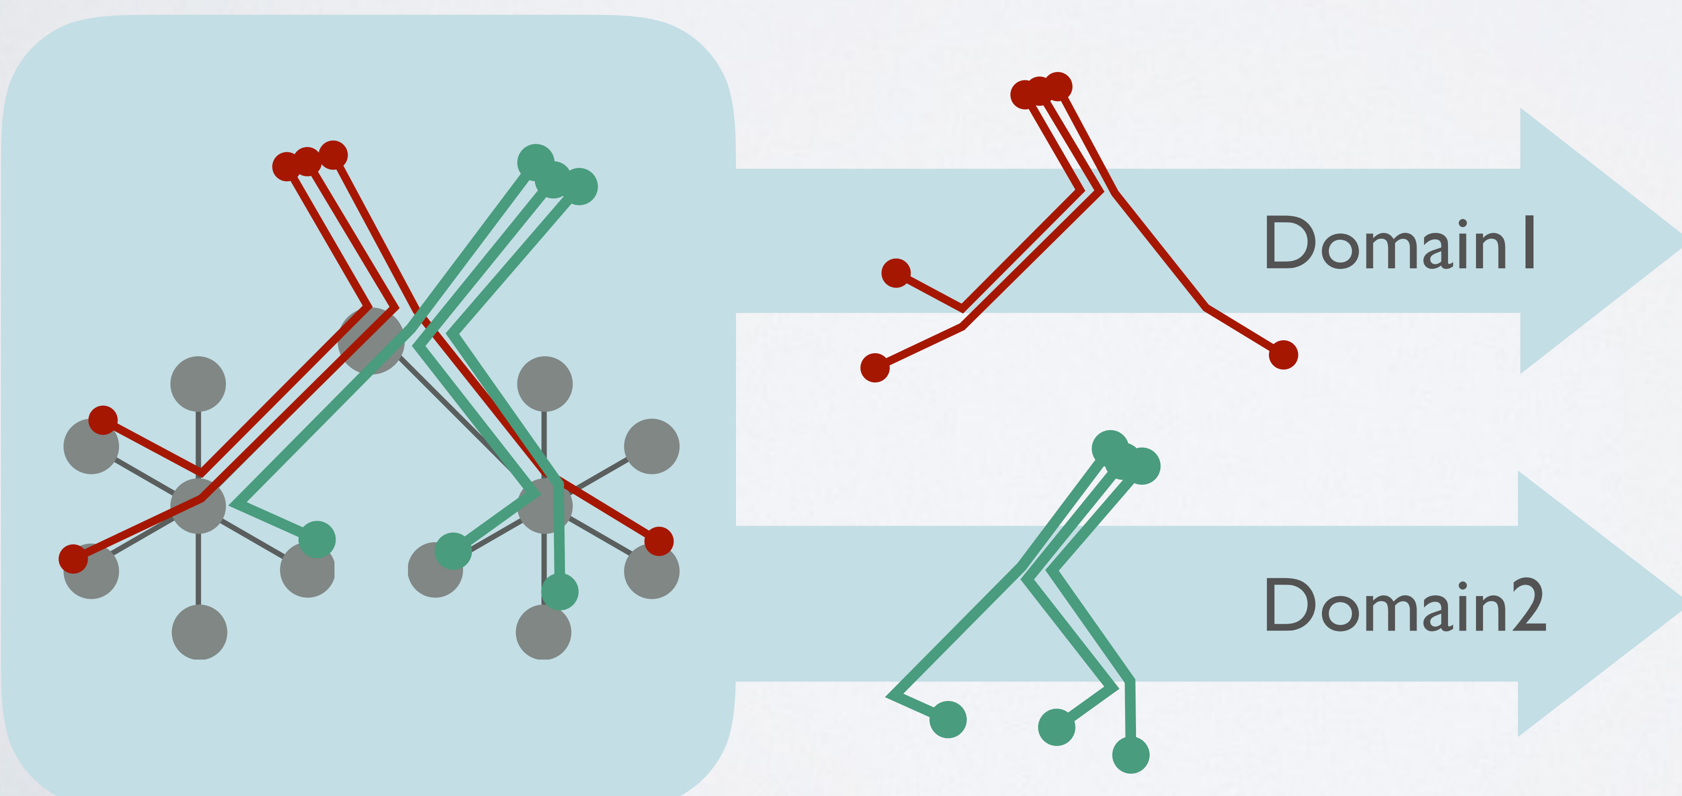

#channel	nodes	cost

#channel	nodes	cost

2 Online Policy Aggregation & Execution

Sub-Agent maintains belief for every domain

10% Bad

90% bad!

Query Sub-Agent for an action

Set of channels

Policy 1

Set of channels

Policy 2

2 Online Policy Aggregation & Execution

Aggregate actions

Turn on corresponding detectors
Get observations
Update Belief

EVALUATION

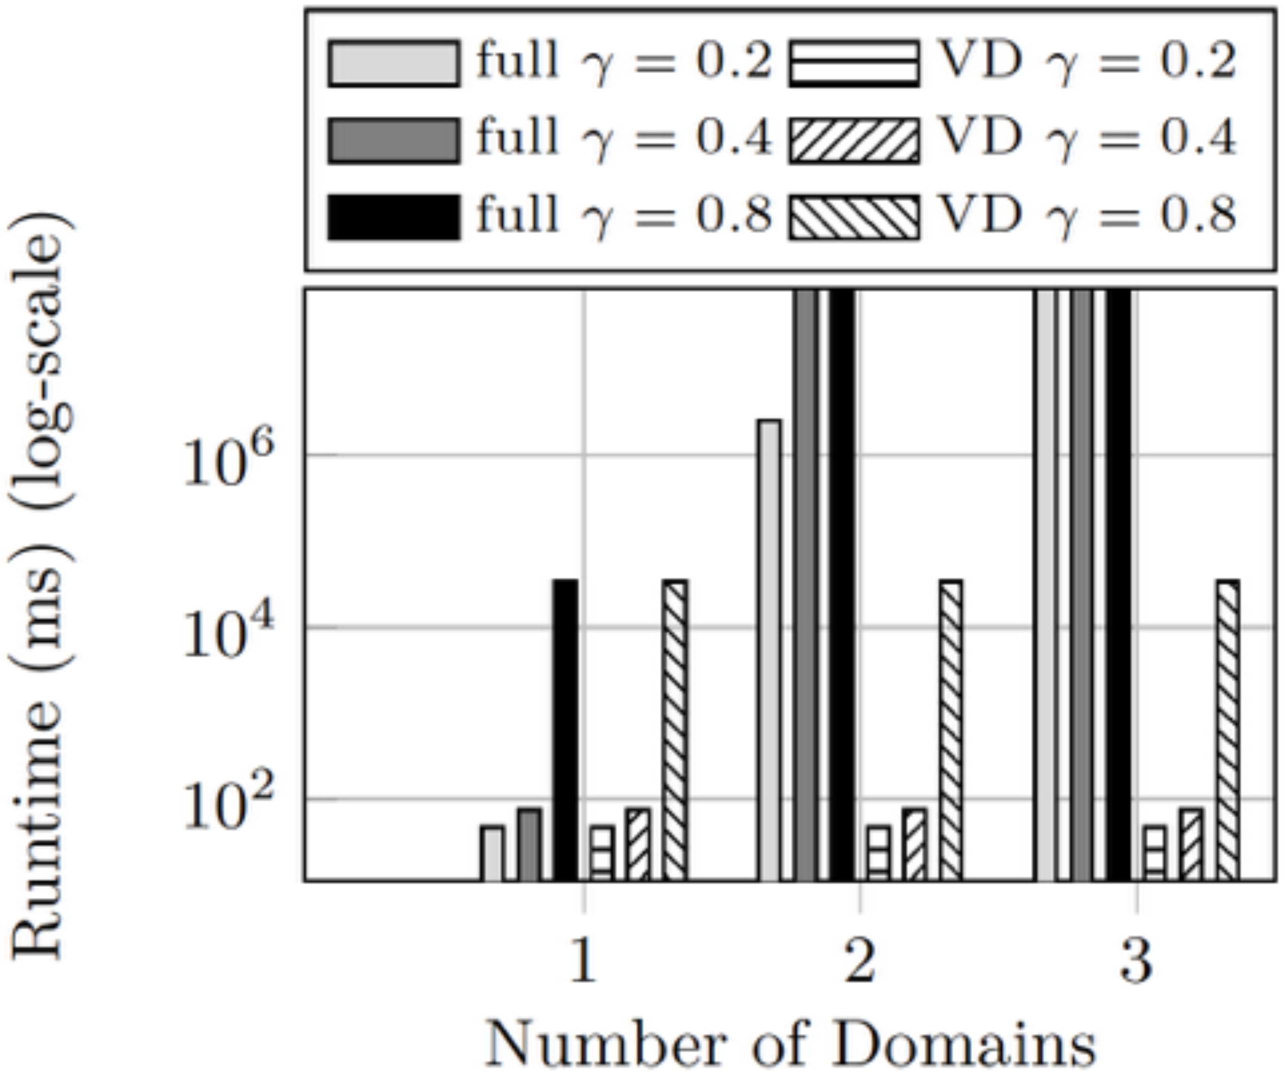

full $\gamma = 0.2$
VD $\gamma = 0.2$
full $\gamma = 0.4$
VD $\gamma = 0.4$
full $\gamma = 0.8$
VD $\gamma = 0.8$
Runtime (ms) (log-scale)
10^6
10^4
10^2
1
2
3
Number of Domains

EVALUATION

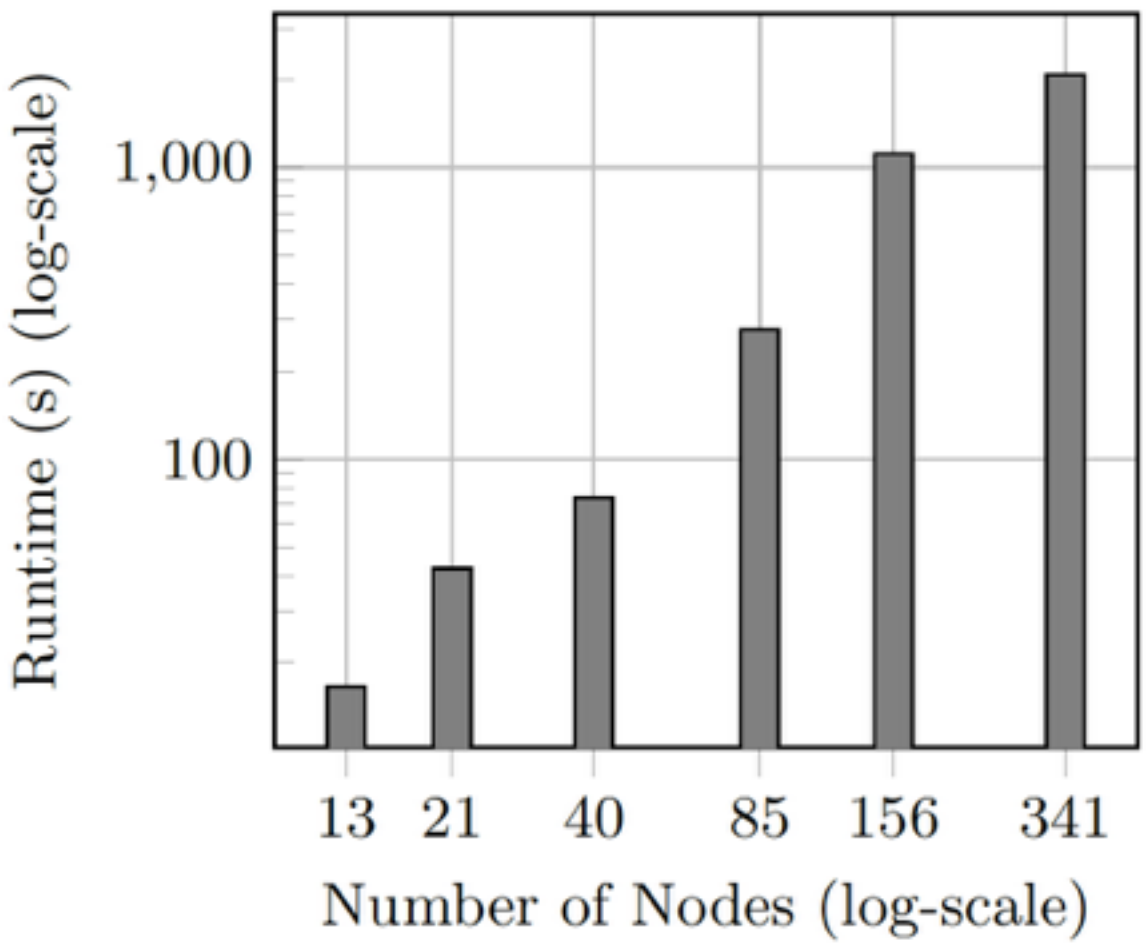

Runtime (s) (log-scale)
1,000
100
13
21
40
85
156
341
Number of Nodes (log-scale)

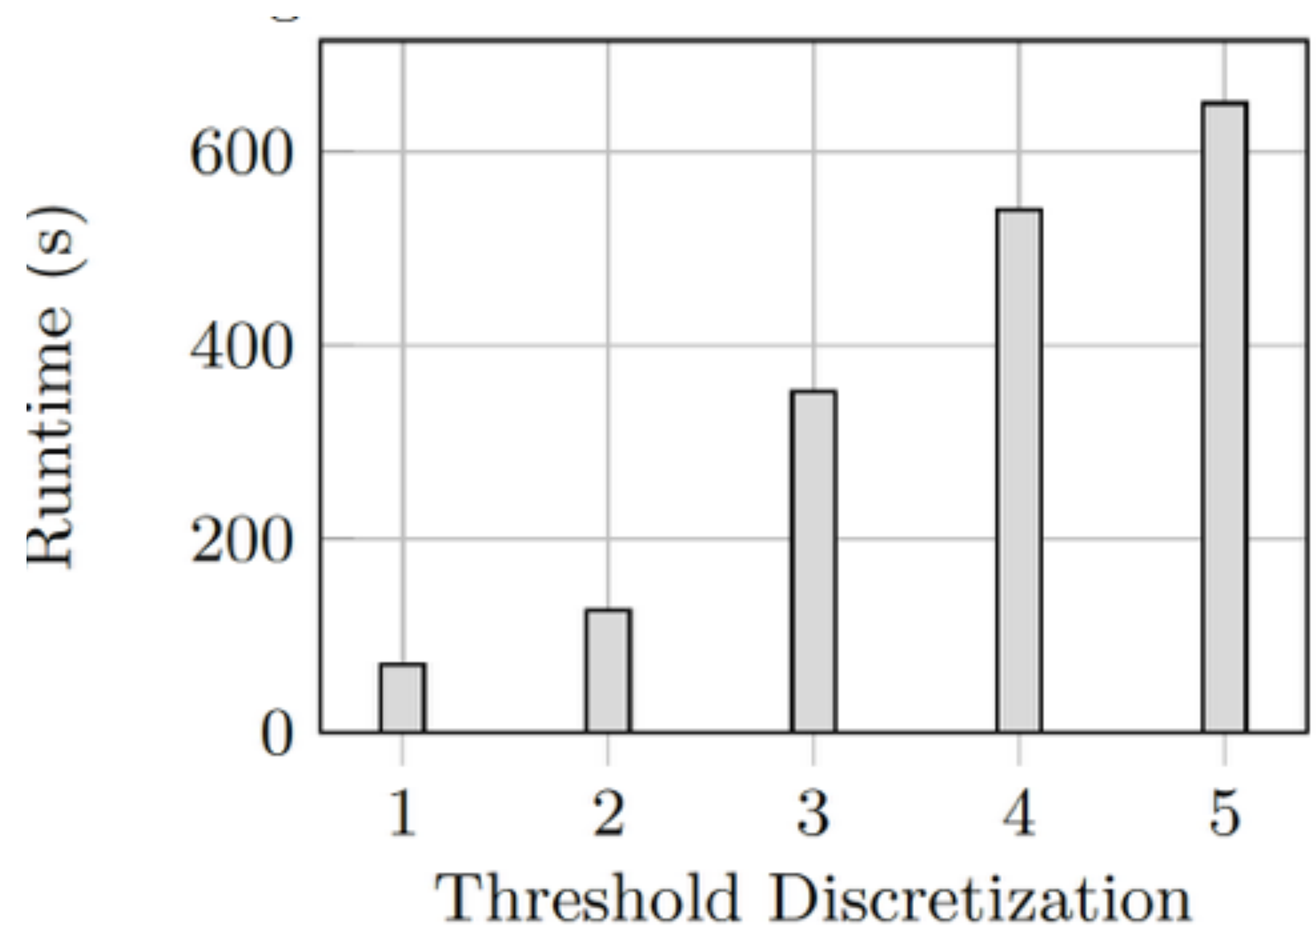

Runtime (s)
600
400
200
0
1
2
3
4
5
Threshold Discretization

DETER TESTBED

Entropy Based Detector

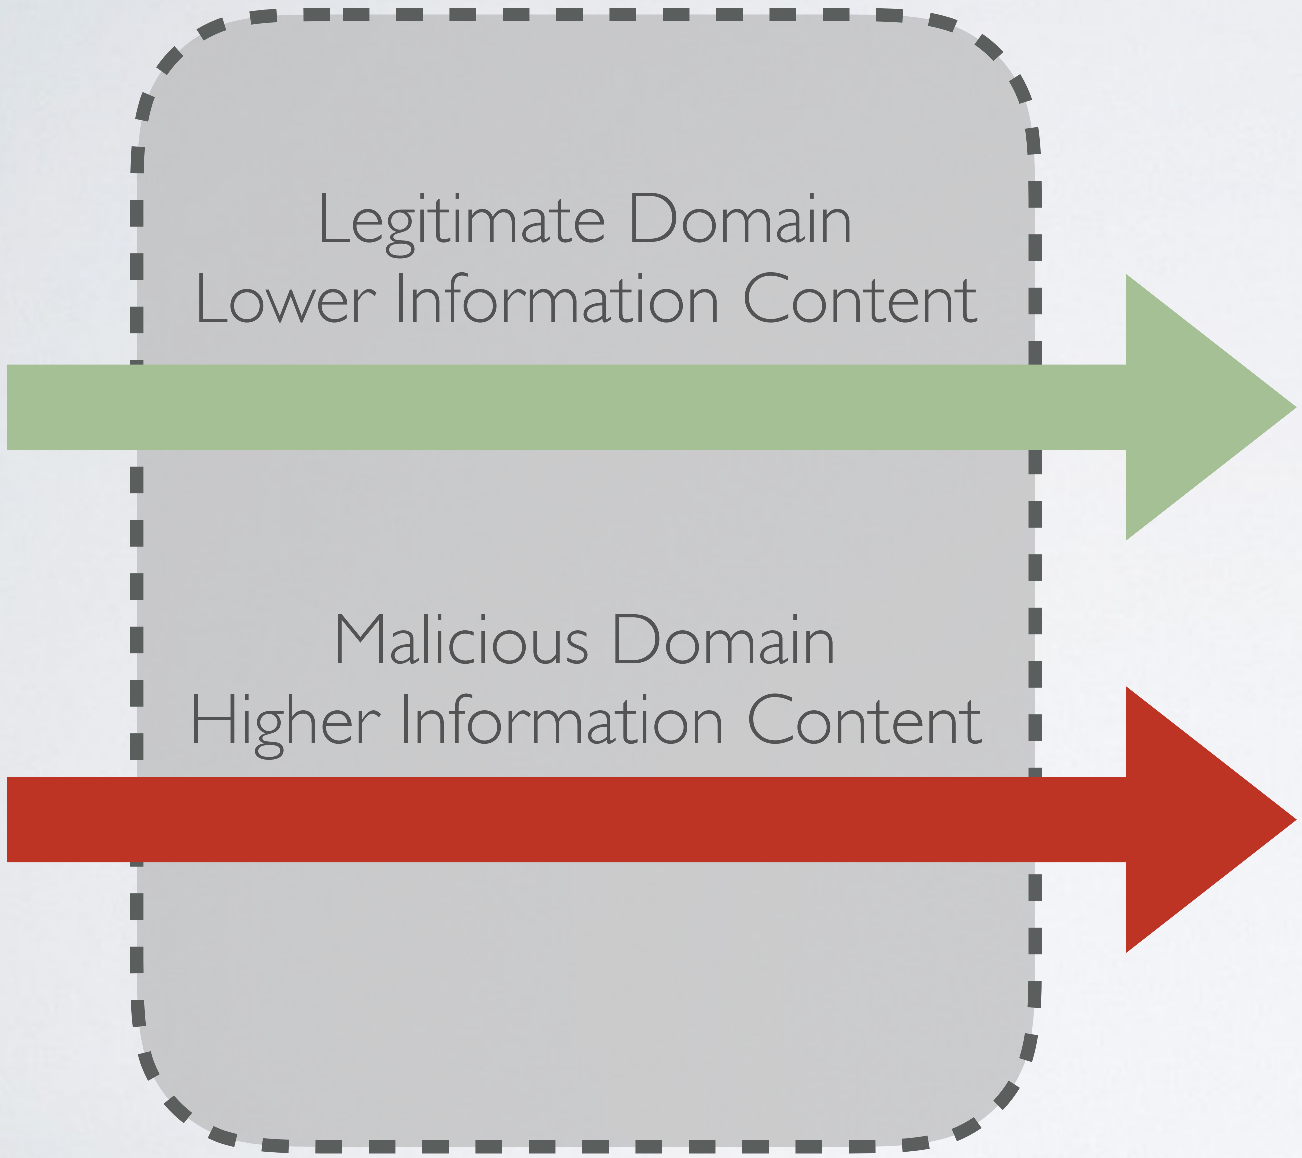

If Information Content Passes Threshold Value

Generate Alert

Paxson, V., Christodorescu, M., Javed, M., Rao, J., Sailer, R., Schales, D., Stoecklin, M.P., Thomas, K., Venema, W., Weaver, N.: Practical comprehensive bounds on surreptitious communication over dns.

DETER TESTBED

Network	Timesteps to Classify	Attack Traffic Accuracy	User Traffic Accuracy
Synthetic 40 Nodes	4.079	1.0	1.0
Synthetic 85 Nodes	3.252	1.0	1.0
Synthetic 156 Nodes	3.235	1.0	1.0
Synthetic 341 Nodes	3.162	1.0	1.0
DETER	5.3076	1.0	0.995

SUMMARY

- Decision theoretic model for reasoning about noisy sensors in a computer network and determine optimal sensing strategies
- Provide a scalable efficient solution method for solving this model
  - solving large scale pomds faster
  - introduces abstraction in planning to induce sparse interaction in factored POMDPs offline
  - interactions are resolved at execution time
- Experimental validation of our model

THANKS!

sara.m.mccarthy@gmail.com

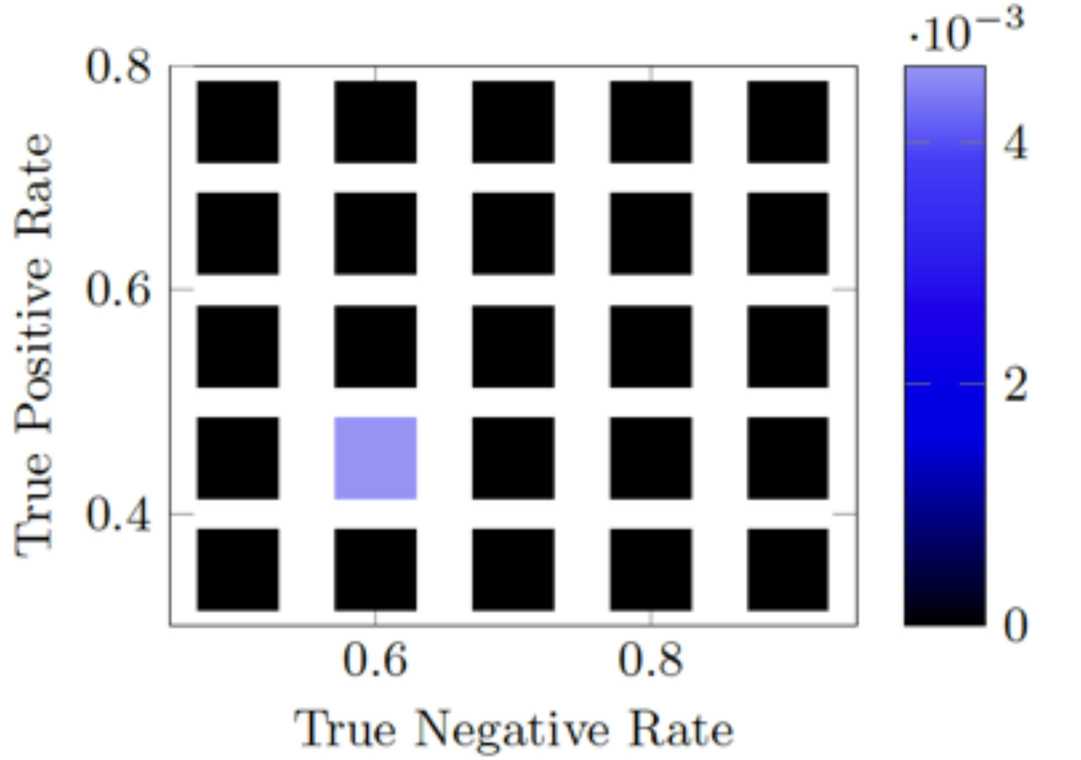

(a) Percent of Incorrect Legitimate Domain Classifications

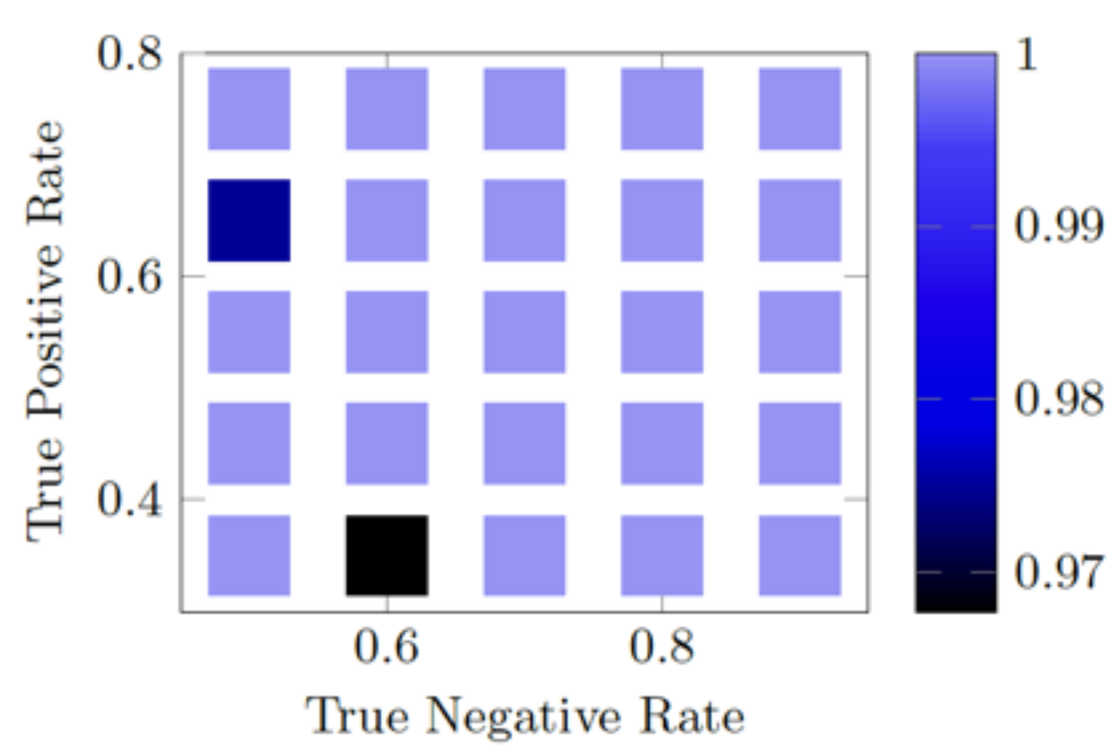

(b) Percent of Correct Malicious Domain Classifications

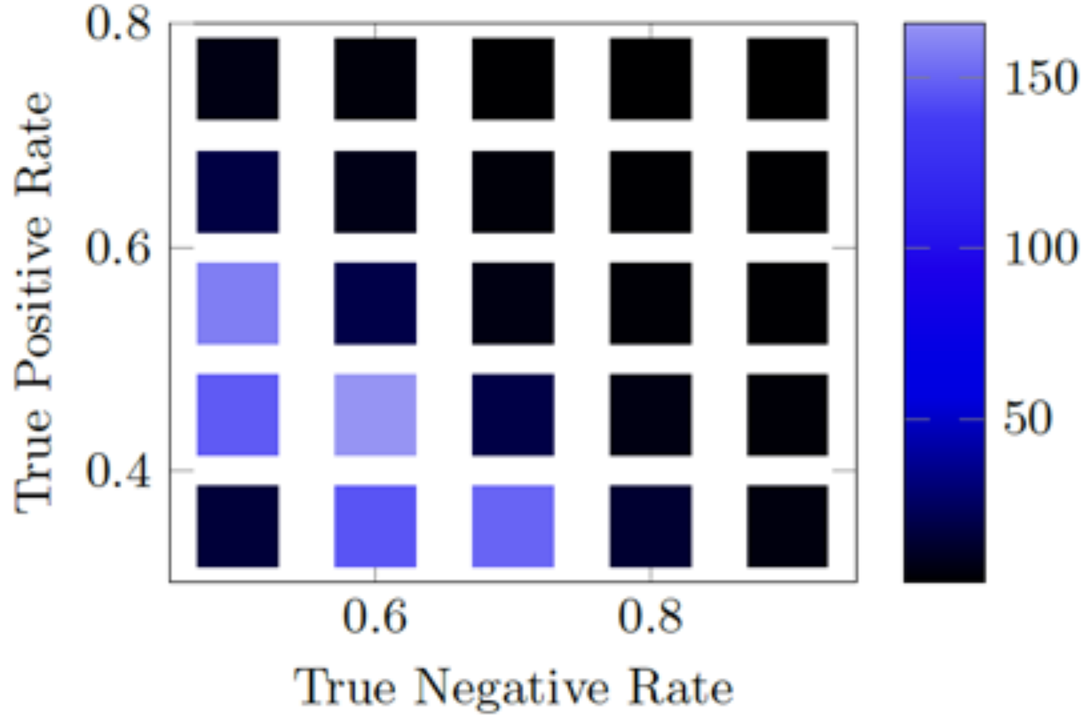

(c) Timesteps to Classify Domain

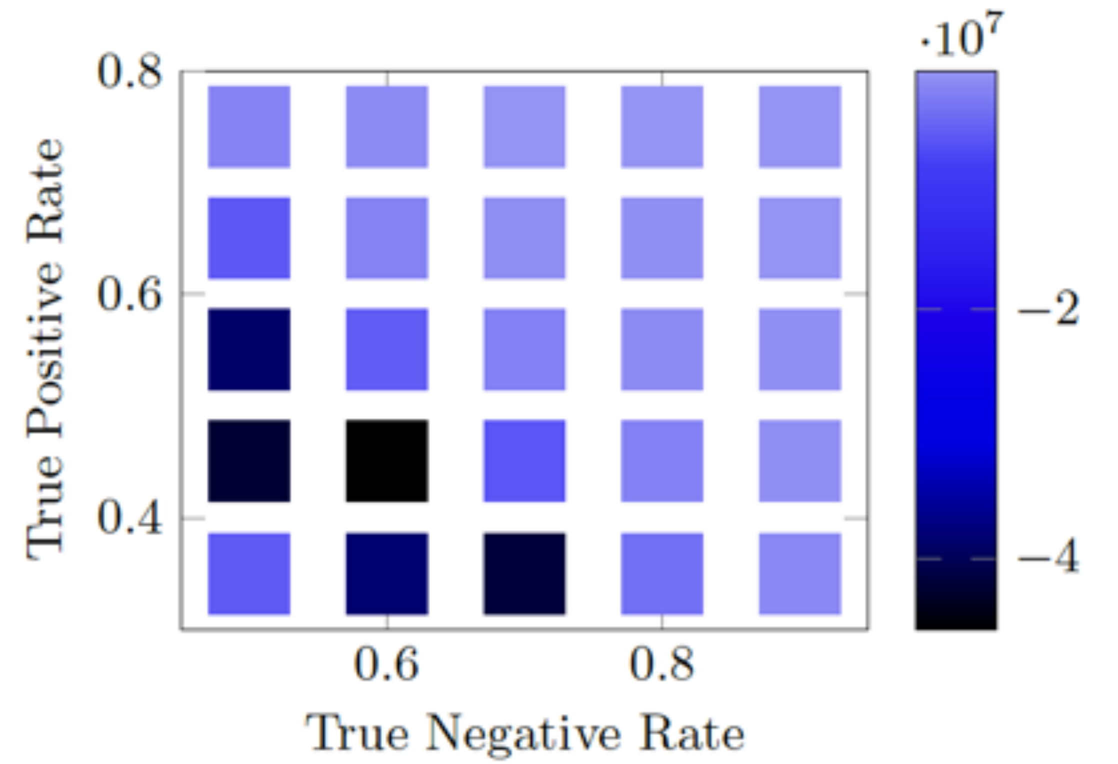

(d) Average Reward

Fig. 7: Testing the robustness with respect to error in the planned true positive and true negative rate.