Detecting Malicious Domains via Graph Inference

Pratyusa K. Manadhata, Sandeep Yadav, Prasad Rao, William Horne

HP Labs, Princeton, NJ

Enterprises collect security data

Event data is a treasure trove of information

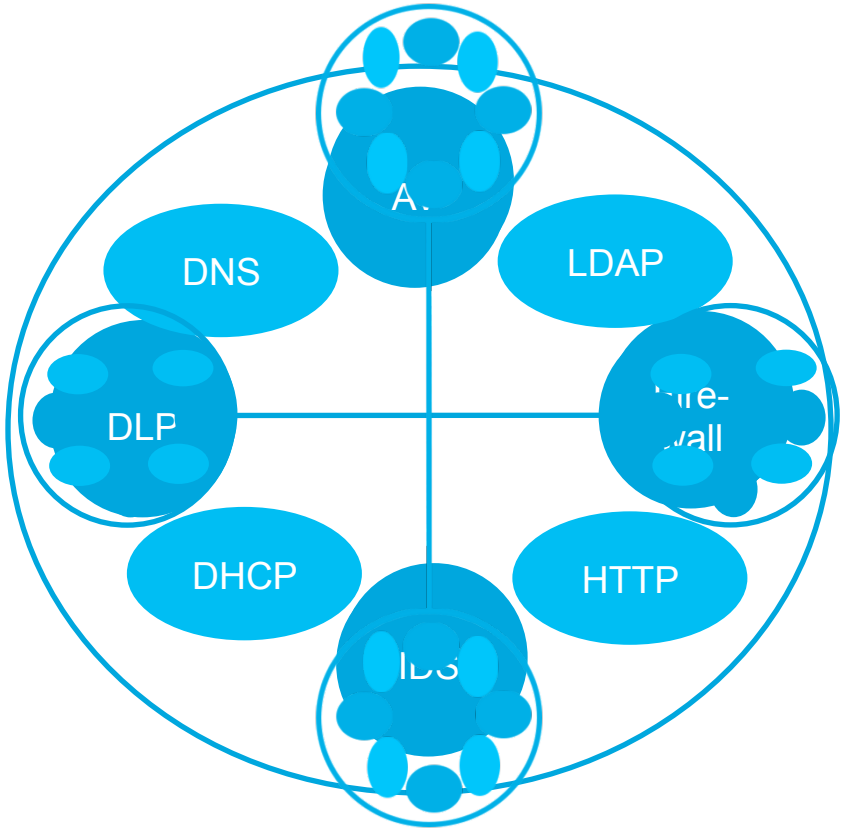

Research challenge and opportunity

Algorithms and systems to identify actionable security information from event data

Example: Malicious domain detection

Malware infection in enterprises is a big problem

Majority of the infections happen via malicious domain access [SYMC11]

State of the art

Domain blacklists

Resource intensive techniques

e.g., Learning/Statistical analysis

Malicious domain detection via graph inference

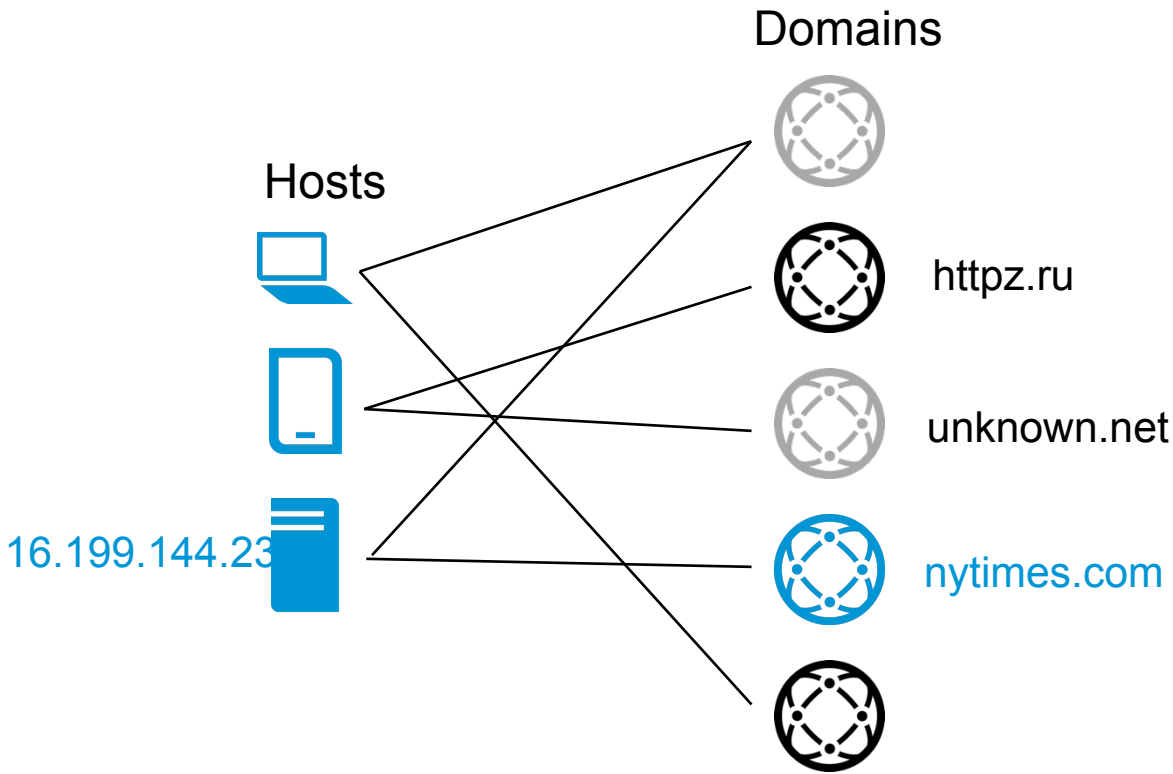

Estimating marginal probability of being malicious

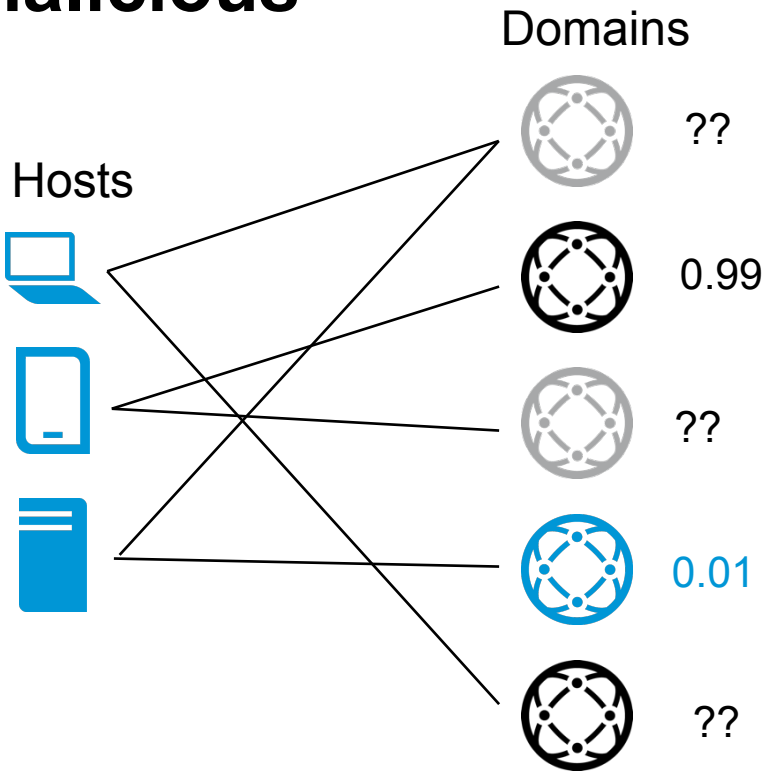

Joint Prob. Dist. $P(x_1, x_2, .., x_n)$

$$MP(x_i) = \sum_{x1} .. \sum_{x_{i-1}} \sum_{x_{i+1}} .. \sum_{x_n} P(x_1, x_2, .., x_n)$$

Belief propagation algorithm [P82, YFW01]

Marginal probability estimation in graphs

- NP-complete

Belief propagation is fast and approximate

- Iterative message passing

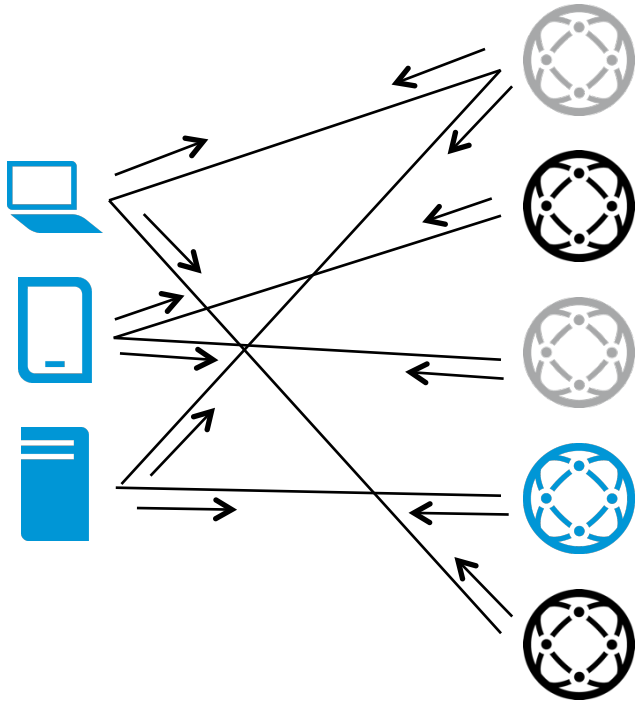

Message passing

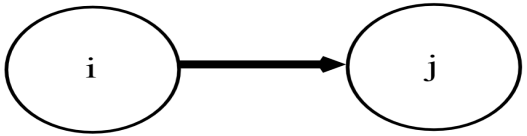

Message(i → j) ∝ (prior, edge potential, incoming messages)

$$m_{ij}(x_j) = \sum_{x_i \in S} \phi(x_i)\psi(x_i, x_j) \prod_{k \in N(i) \setminus j} m_{ki}(x_i)$$

Prior Edge potential Incoming messages

Belief computation

Belief(i) $\propto$ (prior, incoming messages)

$$b_i(x_i) = K\phi(x_i) \prod_{j \in N(i)} m_{ji}(x_i)$$

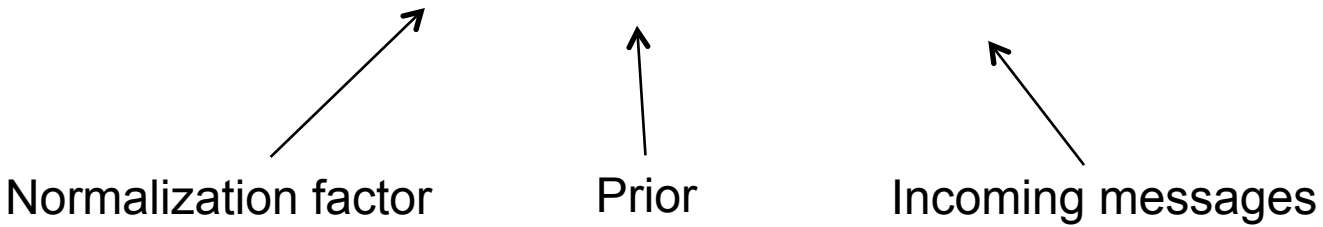

HTTP Proxy logs

Logs from a large enterprise

- 98 HTTP proxy servers, 7 months of data
- 1 day's logs : 1.29 billion events
- 2.80M nodes and 27.8M edges

Priors from ground truth (1.45% nodes) potential

- 21.6K known bad domains: 0.99
- 19.7K known good domains: 0.01
- Unknown hosts and domains: 0.5

Edge

	Benign	**Malicious**
Benign	0.51	0.49
Malicious	0.49	0.51

Scales to enterprise settings

12 core 2.67GHz desktop with 96GB of RAM

Java implementation

53GB RAM

7.8 minutes per iteration

A domain detection ROC plot

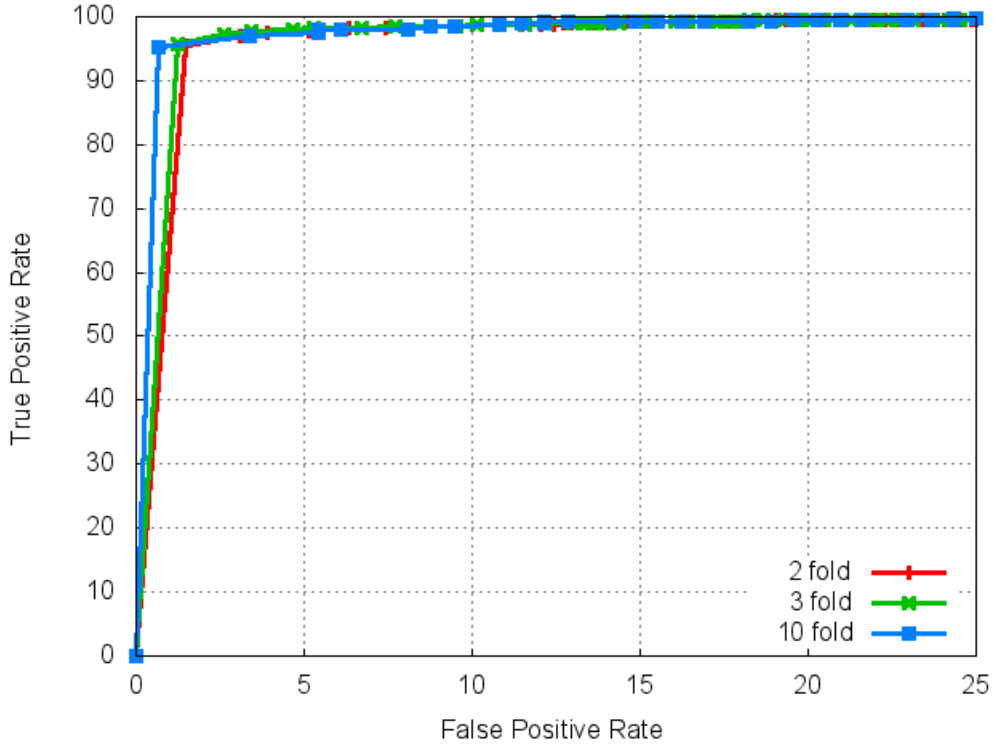

ROC plots for seven days' data

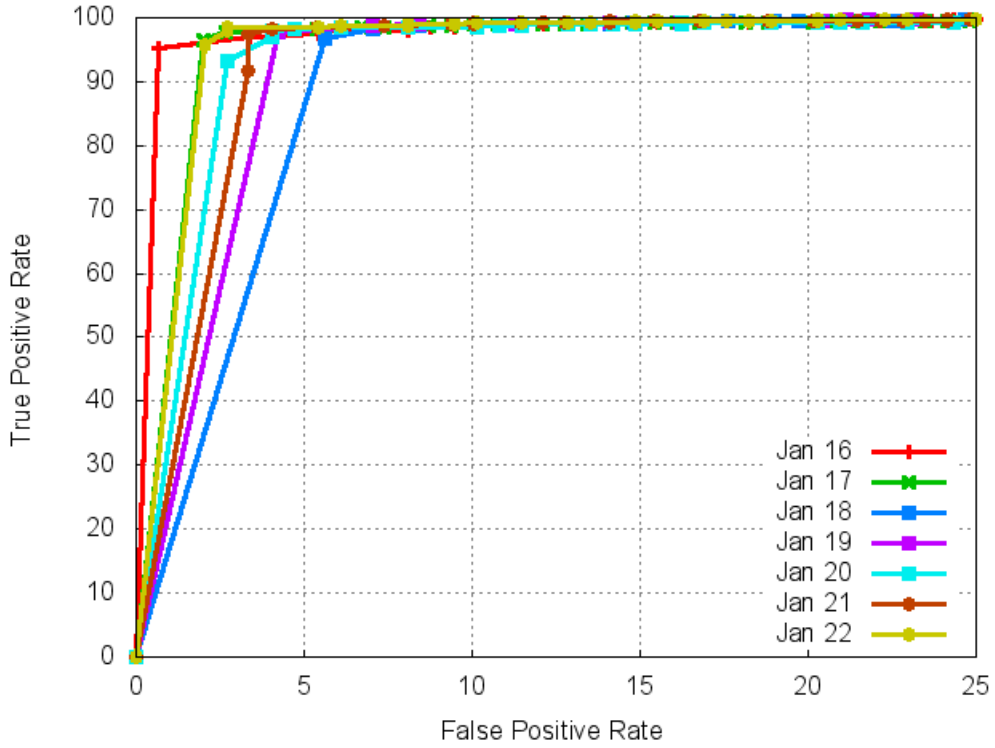

Detection details

Low degree false positives

Unknown domain detection

luo41cxjsbxfrhtbxfubxaqawhxjshsjx.info

awhvkvkzk17fxa67e51pvp42ozmyiqhvfwp12.info

Near-time detection

1 day: 115 minutes

6 hours: 38 m

3 hours: 17 m

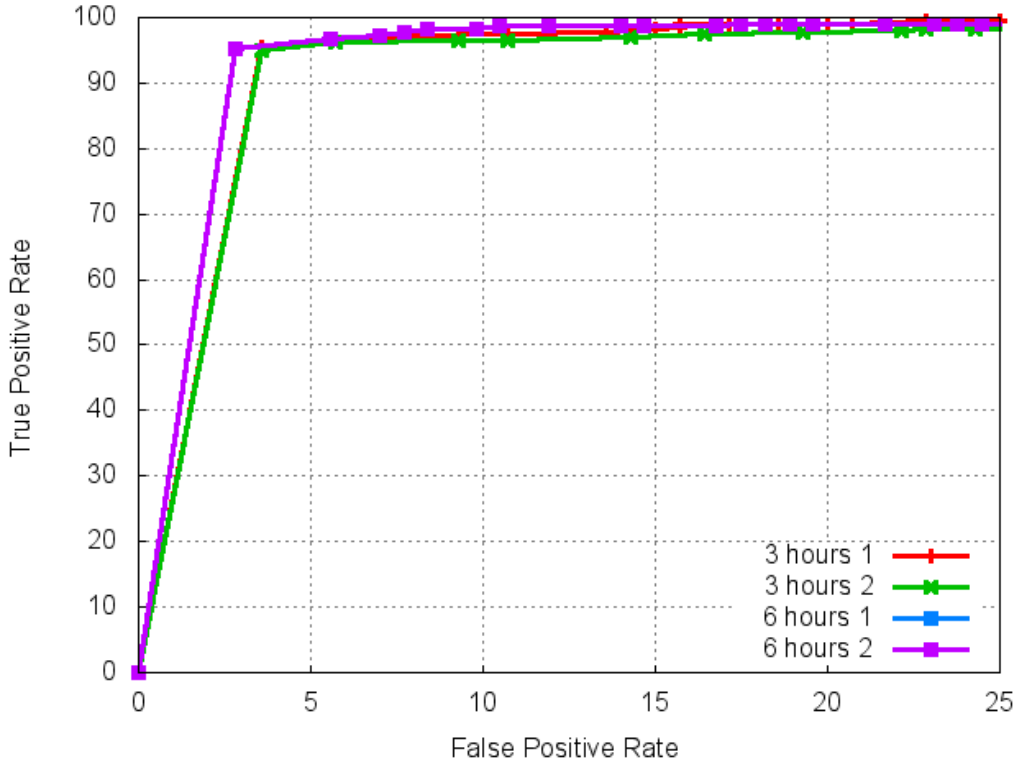

Summary

Can extract actionable security information from event data

Scales to enterprise settings, robust w.r.t. parameter choices

Works well with minimal labeled data, performs better with more

Discovers previously unknown malicious domains

Thank you

Acknowledgements: Marc Eisenbarth, Stuart Haber, and A.L. Narasimha Reddy

Pratyusa K. Manadhata
manadhata@hp.com

Seven randomly chosen days

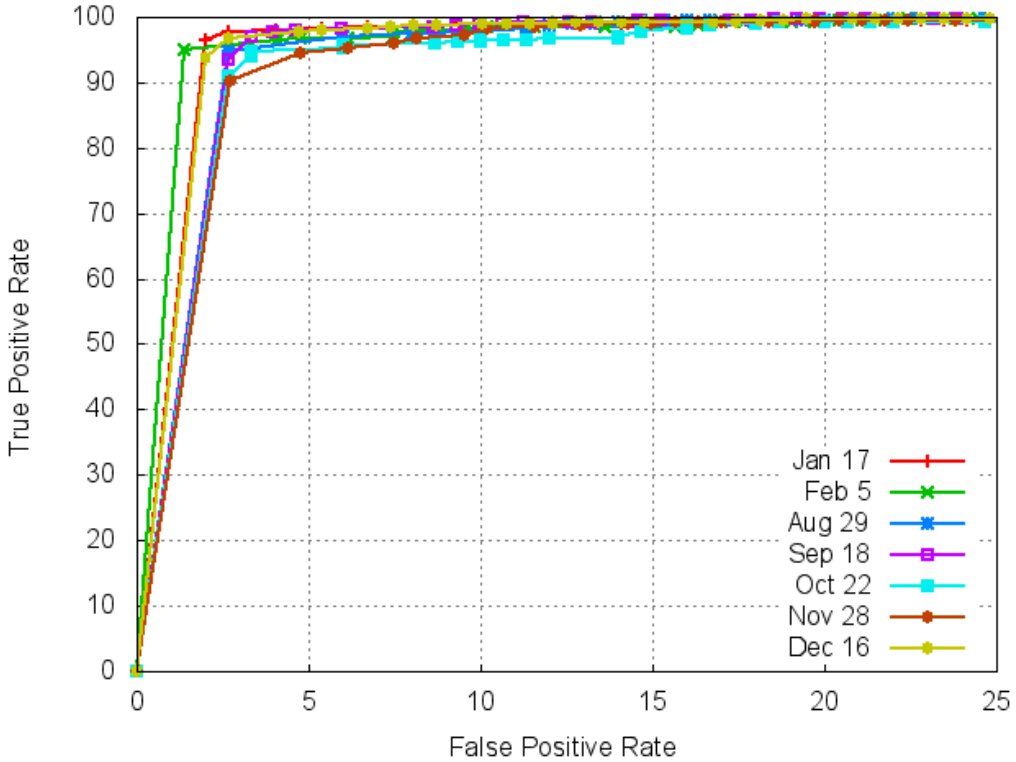

Our approach

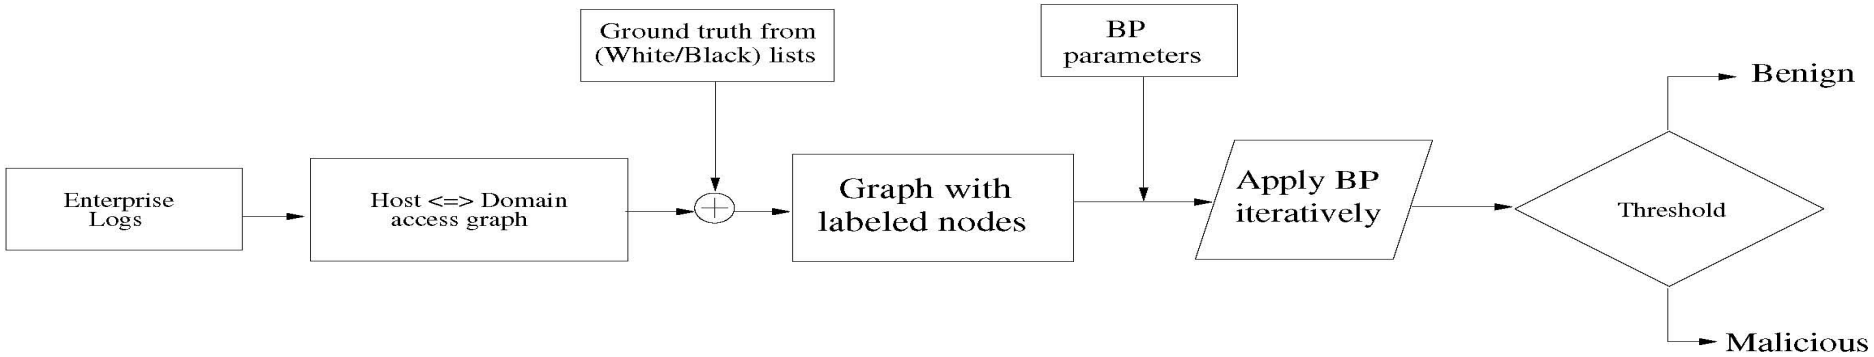

1. No additional data collection
2. No feature computation
3. Minimal labeled data