Information Decay + POMDP: Incorporating Defender's Behaviour in Autonomous Penetration Testing

Jonathon Schwartz[1], Hanna Kurniawati[1], and Edwin El-Mahassni[2]

[1] Research School of Computer Science, ANU

[2] Defence Science and Technology Group, Australian Department of Defence

Australian National University | RESEARCH SCHOOL OF COMPUTER SCIENCE

Australian Government
Department of Defence

Introduction

Autonomous penetration testing (pen-testing) aims to assess the security of a network by finding and exploiting vulnerabilities. We view pen-testing as a sequential decision problem with three sources of uncertainty (table 1). **In this work we introduce a pen-testing model that can handle all three sources of uncertainty and demonstrate its effectiveness in two benchmark scenarios (fig. 2).**

Method	Partial observability	Unreliable actions	Defender
Attack planning [1]	no	yes	no
POMDP [2]	yes	yes	no
Stochastic game [3]	no	yes	yes
This paper	**yes**	**yes**	**yes**

Table 1:Autonomous pen-testing: current state and sources of uncertainty.

Modelling the Defender

The pen-tester and defender observe each other only indirectly via changes to the network state. **We propose to model the defender as a Markovian Arrival Process (MAP) which represents the expected time the defender takes to mitigate an attack.** For this work we use the Bernoulli process with a single parameter: the *information decay factor* d (fig. 1).

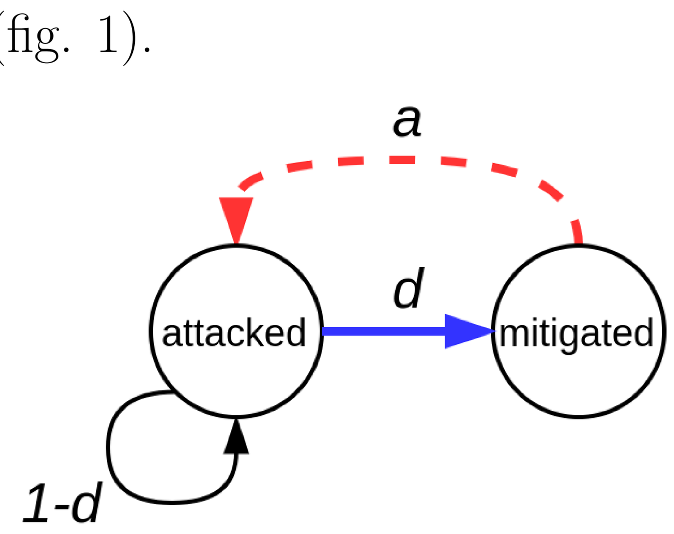

Figure 1:Bernoulli Process where d models the defender mitigating the attack.

D-PenTesting

Given a pen-testing POMDP model $\mathcal{P} = \langle S, A, T, O, Z, R, \gamma \rangle$.

Define $\mathcal{P}_d = \langle S, A, T_d, O, Z, R, \gamma \rangle$, with transition T_d for state variable s_j:

$$T_d\left(s'_j \mid s_j, a\right) = \begin{cases} T\left(s'_j \mid s_j, a\right) & \text{if } a \text{ changes or observes } s_j \\ d \cdot \frac{1}{|S'_j|-1} & \text{else if } s'_j \neq s_j \\ 1-d & \text{otherwise.} \end{cases}$$

Requires knowing d beforehand.

LD-PenTesting

Uses Bayesian Reinforcement Learning to learn the defenders model online.

Define $\mathcal{P}_{ld} = \langle S_{ld}, A, T_{ld}, O, Z_{ld}, R_{ld}, \gamma \rangle$:

$$S_{ld} = S \times D$$
$$Z_{ld}(\langle s, d \rangle, a, o) = Z(s, a, o)$$
$$R_{ld}(\langle s, d \rangle, a) = R(s, a)$$
$$T_{ld}(\langle s, d \rangle, a, \langle s', d' \rangle) = T_d(s, a, s') \cdot \Delta_{dd'}$$

Where D represents possible values of d, discretised to resolution δ and $\Delta_{dd'}$ is the Kronecker Delta (identity) function.

Results

We tested our approach on two benchmark scenarios extended to be partially observable and multi-agent [2, 3].

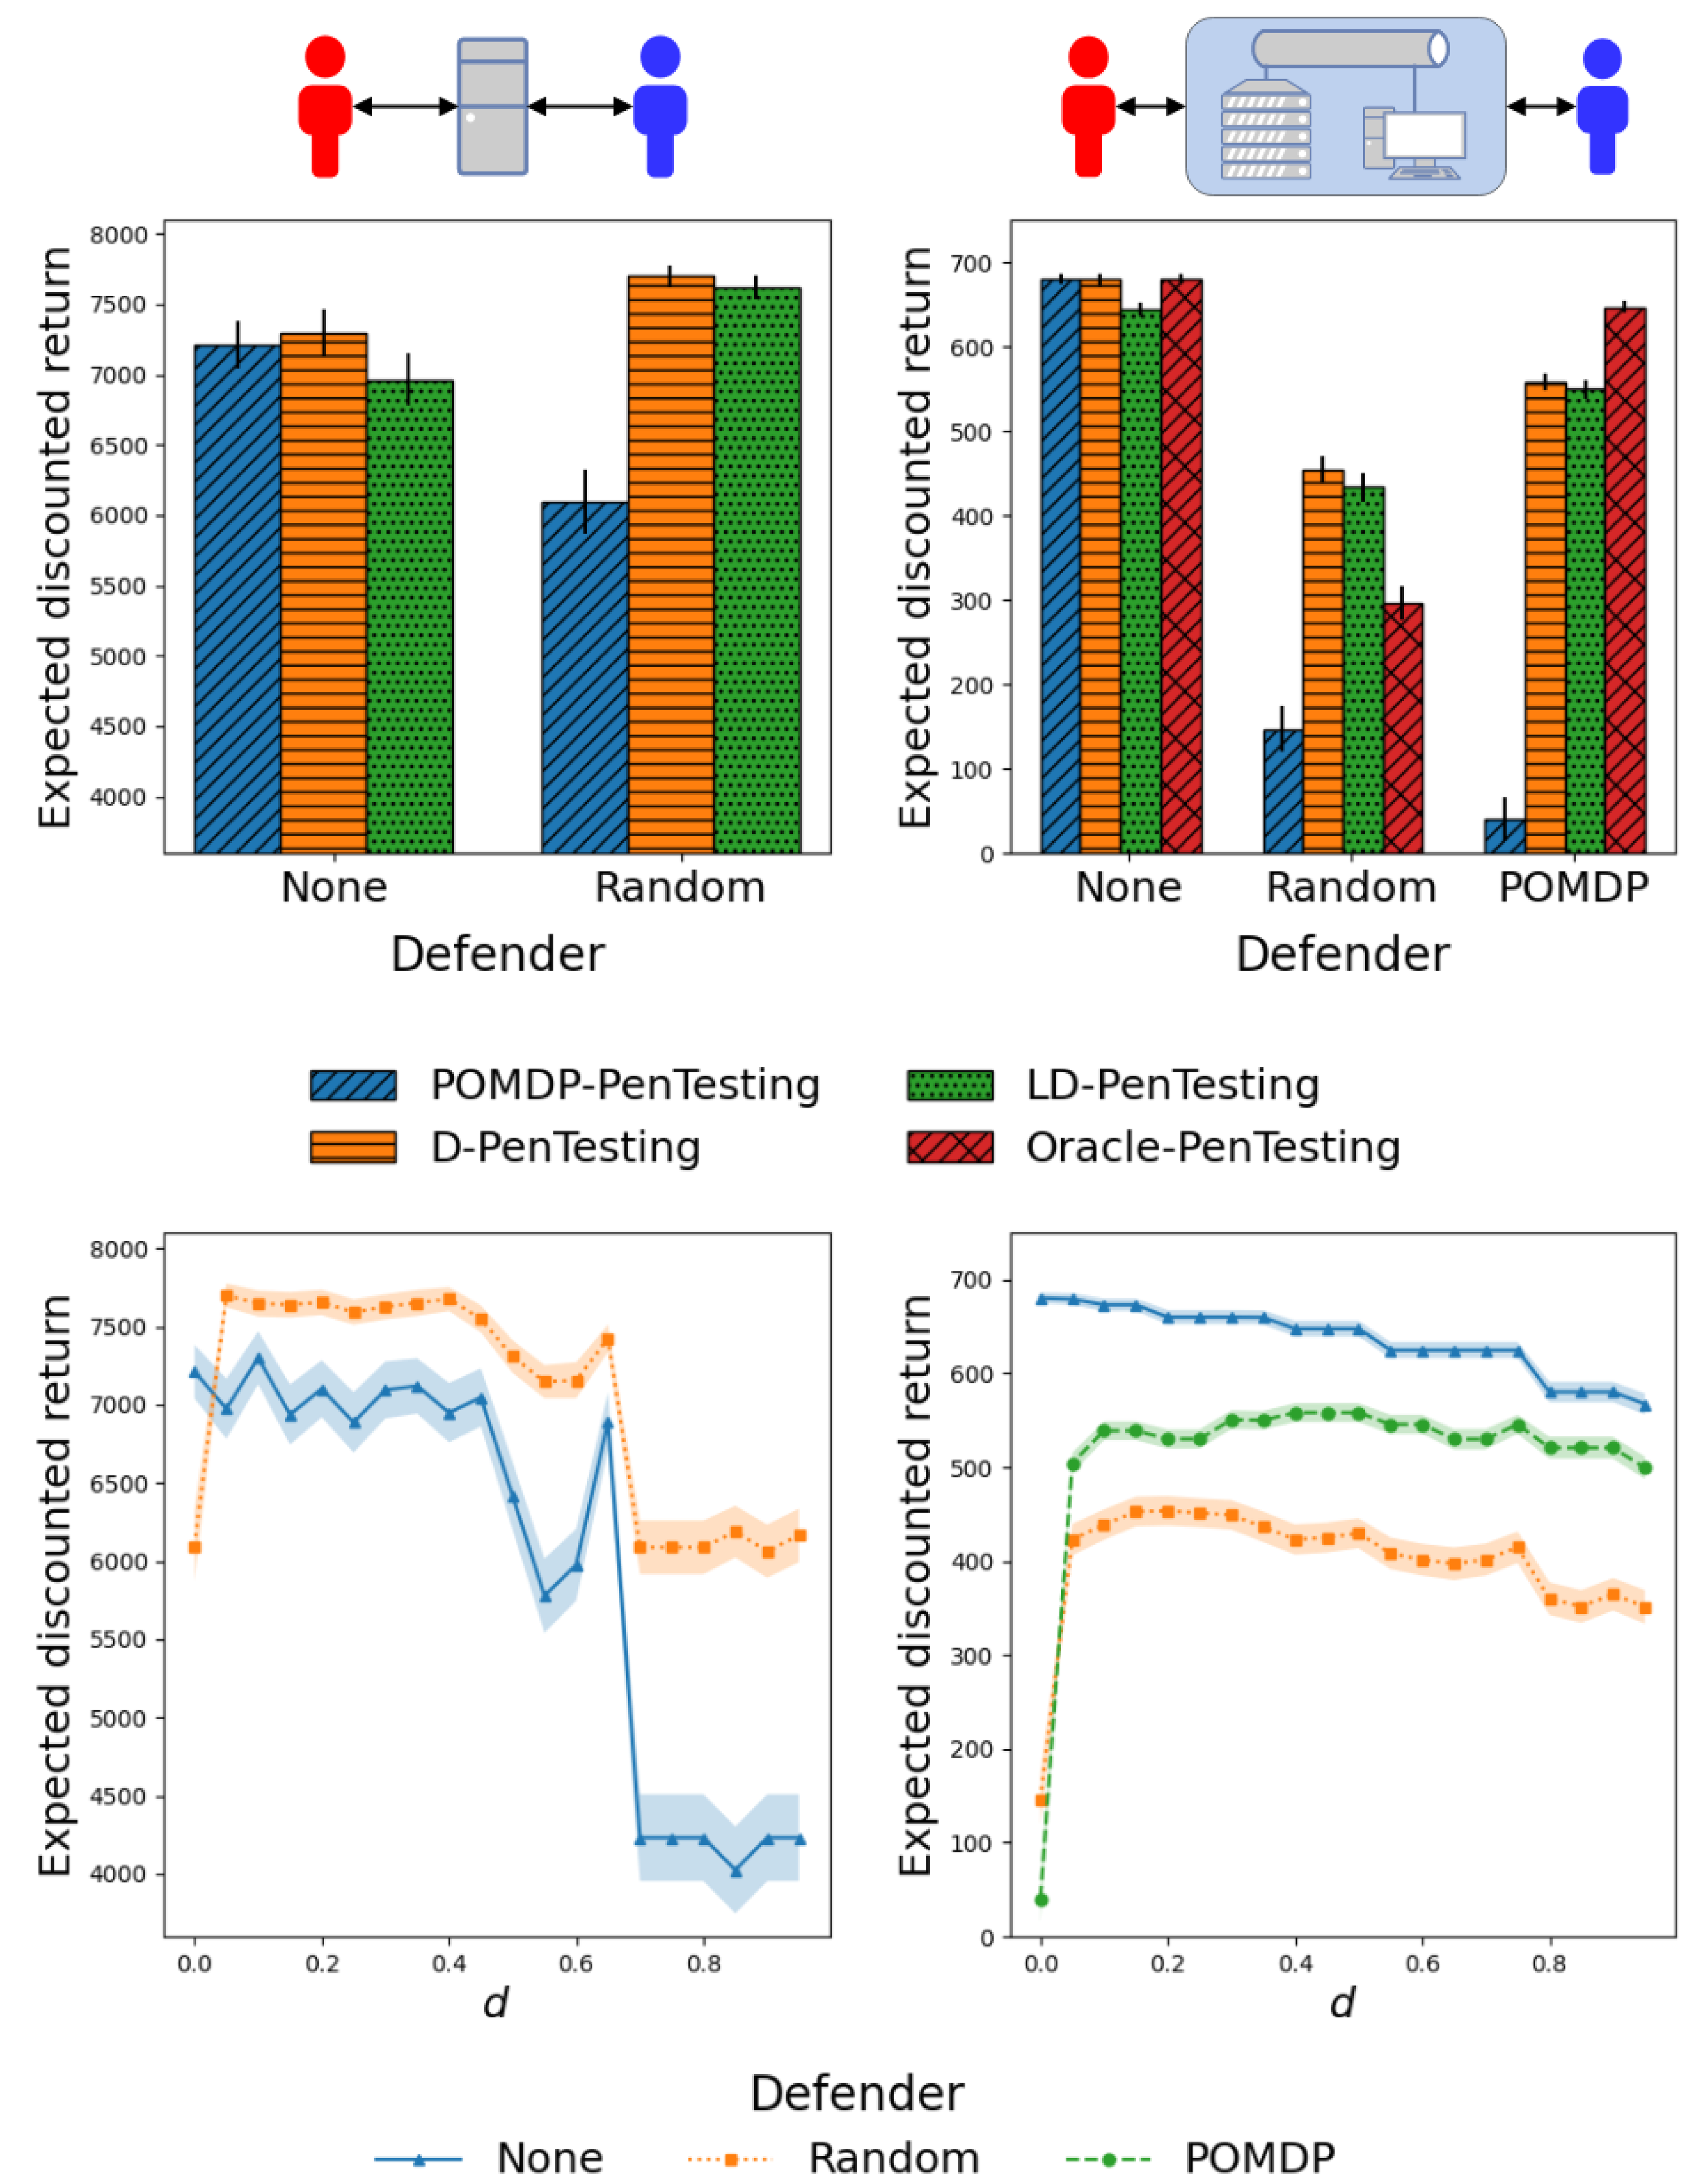

Figure 2:Performance of our approach for two benchmark scenarios (columns). The bar plot graphs compare the performance of pen-testing models for each defender. The line graphs show performance of D-PenTesting for different values of d.

Conclusion

In this work we:

1. presented an efficient abstract defender model based on a MAP,
2. used this model to create D-PenTesting and LD-PenTesting which can handle all three sources of uncertainty present in pen-testing (Table 1),
3. showed the effectiveness of our approach in two benchmark scenarios.

References

[1] J. Lucangeli, C. Sarraute, and G. Richarte, "Attack planning in the real world," in *Workshop on Intelligent Security (SecArt 2010)*, pp. 10–18, 2010.

[2] C. Sarraute, O. Buffet, and J. Hoffmann, "Pomdps make better hackers: Accounting for uncertainty in penetration testing," in *Twenty-Sixth AAAI Conference on Artificial Intelligence*, pp. 1816–1824, 2012.

[3] K.-w. Lye and J. M. Wing, "Game strategies in network security," *International Journal of Information Security*, vol. 4, no. 1-2, pp. 71–86, 2005.