Formulating Cyber-Security as Convex Optimization Problems *

Kyriakos G. Vamvoudakis, João P. Hespanha **, Richard A. Kemmerer, and Giovanni Vigna ***

University of California, Santa Barbara

Abstract. Mission-centric cyber-security analysts require a complete overview and understanding of the state of a mission and any potential threats to their completion. To facilitate this, we propose optimization-based algorithms that can be used to predict in real-time how an attacker may try to compromise a cyber-mission with a limited amount of resources, based on a model that takes into account potential damage to the mission and probabilistic uncertainty. Two different optimization schemes are considered: one where all the mission data is known a priori to the attacker and another where system identification and a moving horizon optimization is used to produce the estimates based on historical data. Our schemes are compared with real attacks carried our by human players in the 2011 international Capture The Flag (iCTF) hacking competition.

Keywords: Cyber-Security, Convex Optimization, System Identification, iCTF

1 Introduction

Guaranteeing the security of cyber-missions is a complex, multi-dimensional challenge that demands a multi-faceted, strategic solution. The terminology *cyber-mission* refers to a set of computer transactions aimed at accomplishing a specific purpose or task, such as placing an online shopping order, submitting a paper to a conference through an online submission system, or printing a bank statement at an ATM machine. Cyber-missions typically require a large number of computer services, including encryption services, authentication servers, database engines, web servers. We are especially interested in cyber-missions

* This material is based upon work supported by ARO MURI Grant number W911NF0910553.

** K. G. Vamvoudakis, and J. P. Hespanha are with the Center for Control, Dynamical-systems and Computation (CCDC), University of California, Santa Barbara, CA 93106-9560 USA e-mail: kyriakos@ece.ucsb.edu, hespanha@ece.ucsb.edu

*** R. A. Kemmerer, and G. Vigna are with the Computer Security Lab, University of California, Santa Barbara, CA 93106-9560 USA e-mail: kemm@cs.ucsb.edu, vigna@cs.ucsb.edu

that go through several states, each of which may require one or more computer services. Cyber-missions are especially vulnerable to attacks because it may be possible to prevent the mission's completion by compromising just one of the multiple services required by the mission, provided that the right service is compromised at the right time.

Cyber-missions are pervasive and can be found in trading, banking, power systems management, road traffic managements, healthcare, online shopping, business-to-business transactions, etc. The disruption to cyber-missions can thus result in cyber or physical consequences that threaten National and economic security, critical infrastructure, public health, and welfare. Moreover, stealthy cyber-attackers can lay a hidden foundation for future exploitation or attack, which they can later execute at a time of greatest advantage. Securing cyberspace requires a layered security approach across the public and private sectors.

In the cyber-mission security domain, the security analyst is interested in making decisions based on the potential damage that attacks can inflict to the mission and also on the probability that the potential damage is realized. To focus their attention and coordinate defensive actions, security professionals must be able to determine which attacks presents the biggest threat and prioritize which services to defend, a problem often referred to as *cyber situation awareness*. Situation awareness [3] is a common feature of many cyber-security solutions but most of them are fragmented. In this paper, we present a model that can be used to predict how an attacker may try to compromise a cyber-mission with a limited amount of resources, based on a model that takes into account potential damage to the mission and probabilistic uncertainty.

This approach followed here motivated by the need to avoid flooding the security analyst with raw data about complex missions and detailed logs from intrusion detection systems (IDSs). Instead, an automated or semi-automated system should process this data and present the analyst with high-level information about the computer services that are currently most crucial for mission completion and thus most likely to be the target of attacks, based on the current state of the mission and its future expected evolution. To achieve this we propose a relatively general model to describe the damage to a cyber-mission caused by potential attacks. This model can be utilized in optimization schemes to discover optimal policies to distribute attack resources over time and over the different computer services relevant to the mission so as to maximize damage to the cyber mission. The models proposed, need mission parameters that typically vary with time according to complex dynamics, which are difficult to determine in an analytic fashion. To avoid this difficulty, we learn such parameters using system identification of low-order state-space models that are used to make predictions of the parameter evolution for a reasonable future time horizon.

Security competitions are exceptional venues for researchers to discover and validate novel security solutions. The international Capture The Flag (iCTF) [5] is a distributed wide-area security exercise whose goal is to test the security skills of the participants. The iCTF contest is organized by the Security Lab of the Department of Computer Science at UCSB and is held once a year. The Capture

the Flag contest is a multi-site, multi-team hacking contest in which a number of teams compete independently against each other. The 2011 edition of iCTF was aimed at Cyber-Situation Awareness and, to our knowledge, produced the first experimental dataset that includes mission descriptions as well as attack logs and the statuses of computer services required by missions [2, 5]. We have used this data to validate the algorithms presented in this paper and show their efficacy in predicting attacks to cyber missions by the human participants in the exercise.

The results presented in this paper were also used in the design of a high-level visualization tool to help security analysts to protect the computer systems under attack in the 2011 iCTF competition [4]. We are in the process of developing human subject experiments to demonstrate the benefits of using the predictions generated by the methodology proposed in this paper, instead of searching through mission traces and security logs.

The remainder of the paper is structured as follows. Section 2 develops the general mathematical framework for cyber-security and then describes how one can use classical system identification techniques to identify the completely unknown or partially known time-varying processes. Section 3 describes an optimization problem to discover how an attacker would optimally allocate her resources through all the services as time evolves for two different scenarios. The first assumes that the all the mission data is known to the attacker, whereas the second one uses a moving horizon optimization scheme that estimates this data online to predict when and where to attack. In Section 4, the algorithms proposed are applied to data from the 2011 iCTF competition. Comparison results between how the teams in the competition attacked and the results obtained by the optimization schemes are presented in Section 5. Finally, Section refse:conclusions concludes and discusses about future work.

2 General Framework for Cyber-Security

This section presents a general framework to model mission-critical cyber-security scenarios.

2.1 Cyber-Mission Damage Model

Suppose that the *(potential) damage* that an attacker can inflict to a cyber mission is quantified by a scalar $x_{\mathrm{PD}} \geqslant 0$ that is a function of the level of *attack resources* $u_{\mathrm{AR}} \geqslant 0$ devoted to the attack. The mapping from attack resources to potential damage is expressed by the so called *potential damage equation* that we approximate by a linear map:

$$x_{\mathrm{PD}} = f(u_{\mathrm{AR}}) \coloneqq a + b\, u_{\mathrm{AR}}, \tag{1}$$

where $a \in \mathbb{R}^+$ can be viewed as the zero-resource damage level, and $b \in \mathbb{R}^+$ the marginal damage per unit of attack resources.

Whether or not the potential damage to the mission x_{PD} is realized is assumed to be a stochastic event that occurs with a given probability $\rho \in [0,1]$ that also depends on the attack resources $u_{\mathrm{AR}} \in \mathbb{R}^+$, according to the so-called *uncertainty equation* that we approximate by a linear map projected to the interval $[0,1]$:

$$\rho = g(u_{\mathrm{AR}}) \coloneqq \Pi_{[0,1]}(c - d\,u_{\mathrm{AR}}) \tag{2}$$

where $\Pi_{[0,1]} : \mathbb{R} \to \mathbb{R}$ denotes the projection function

$$\Pi_{[0,1]}(x) = \begin{cases} 0 & x < 0 \\ x & x \in [0,1] \\ 1 & x > 1, \end{cases}$$

the scalar $c \geqslant 0$ corresponds to a zero-resource probability of damage, and the scalar $d \geqslant 0$ to the marginal *decrease* in the probability of damage per unit of attack resources. We note that an increase in attack resources u_{AR} leads to an increase in the potential damage x_{PD} [expressed by the + sign before the b term in (1)], but may actually decrease the probability that the potential damage will actually be realized [expressed by the − sign before the d term in (2)], which is motivated by the fact that a large-scale attack is more likely to trigger defense mechanisms that can prevent the potential damage from being realized.

The total expected damage y_{TD} to the mission can be found by multiplying equations (1) and (2), leading to the *expected damage equation*

$$y_{\mathrm{TD}} = f(u_{\mathrm{AR}})g(u_{\mathrm{AR}}). \tag{3}$$

In the context of cyber-missions that evolve over time and require multiple computer services, the potential damage equation (1) and the uncertainty equation (2) need to be augmented with an index $t \in \{1, 2, \ldots, T\}$ that parameterizes mission time and an index $s \in \{1, 2, \ldots, S\}$ that parameterizes the required computer services, as in

$$x_{\mathrm{PD}}{}_t^s = f_t^s(u_{\mathrm{AR}}{}_t^s) = a_t^s + b_t^s u_{\mathrm{AR}}{}_t^s, \tag{4}$$

$$\rho_t^s = g_t^s(u_{\mathrm{AR}}{}_t^s) = \Pi_{[0,1]}(c_t^s - d_t^s u_{\mathrm{AR}}{}_t^s) \tag{5}$$

where $u_{\mathrm{AR}}{}_t^s$ denotes the attack resources committed to attack service s at time t, $x_{\mathrm{PD}}{}_t^s$ the potential damage at time t due to an attack to the service s, and P_t^s the probability of realizing this damage. The corresponding *expected damage equation* then becomes:

$$y_{\mathrm{TD}} = \sum_{t=1}^{T} \sum_{s=1}^{S} f_t^s(u_{\mathrm{AR}}{}_t^s) g_t^s(u_{\mathrm{AR}}{}_t^s). \tag{6}$$

3 Optimization

An intelligent attacker would seek to optimally allocate her available resources to maximize the total expected missing damage. We shall consider here several

options for this optimization that differ on the information that is available to the attacker.

3.1 Optimization Scheme with Known Mission Damage Data

When all the data $\{a_t^s, b_t^s, c_t^s, d_t^s : \forall s, t\}$ that define the potential damage and uncertainty equations is known a-priori, optimal attack resource allocation can be determined by solving the following optimization.

$$\begin{aligned} \text{maximize} \quad & \sum_{t=1}^{T}\sum_{s=1}^{S} f_t^s(u_{\mathrm{AR}_t}{}^s) g_t^s(u_{\mathrm{AR}_t}{}^s) \\ \text{subject to} \quad & \sum_{t=1}^{T}\sum_{s=1}^{S} u_{\mathrm{AR}_t}{}^s \leqslant U_{\mathrm{TR}} \\ \text{w.r.t.} \quad & u_{\mathrm{AR}_t}{}^s \in [0,\infty),\ \forall t,\ \forall s, \end{aligned} \tag{7}$$

where U_{TR} denotes the total budget of attack resources available to the attacker. As stated in the following proposition, this optimization can be converted into the following concave maximization.

Proposition 1. *When the functions f_t^s, g_t^s are of the form (4)–(5) with $a_t^s, b_t^s, c_t^s, d_t^s \geqslant 0$, $\forall t, s$. The value and optimum of (7) can be obtained through the following concave maximization problem:*

$$\begin{aligned} \textit{maximize} \quad & \sum_{t=1}^{T}\sum_{s=1}^{S} (a_t^s + b_t^s u_{\mathrm{AR}_t}{}^s)(c_t^s - d_t^s u_{\mathrm{AR}_t}{}^s - \sigma_t^s) \\ \textit{subject to} \quad & \sum_{t=1}^{T}\sum_{s=1}^{S} u_{\mathrm{AR}_t}{}^s \leqslant U_{\mathrm{TR}}, \quad c_t^s - d_t^s u_{\mathrm{AR}_t}{}^s - \sigma_t^s \leqslant 1,\ \forall t,\ \forall s \\ \textit{w.r.t.} \quad & u_{\mathrm{AR}_t}{}^s \in \left[0, \frac{c_t^s}{d_t^s}\right],\ \sigma_t^s \geqslant 0,\ \forall t,\ \forall s. \end{aligned} \tag{8}$$

When $c_t^s \in [0,1]$, one can set the corresponding $\sigma_t^s = 0$ in (8). Moreover, when $c_t^s \in [0,1]$, $\forall t, s$ and all the constraints on the $u_{\mathrm{AR}_t}{}^s$ are inactive, the solution to this optimization can be found in closed form and is equal to

$$u_{\mathrm{AR}_t}{}^s = \bar{u}_t^s - \bar{\mu}_t^s \max\left\{0, \sum_{\bar{t}}\sum_{s=1}^{S} \bar{u}_{\bar{t}}^s - U_{\mathrm{TR}}\right\}, \quad \bar{u}_t^s \coloneqq \frac{b_t^s c_t^s - a_t^s d_t^s}{2 b_t^s d_t^s}, \quad \bar{\mu}_t^s \coloneqq \frac{\frac{1}{2 b_t^s d_t^s}}{\sum_{\bar{t}}\sum_{s=1}^{S} \frac{1}{2 b_{\bar{t}}^s d_{\bar{t}}^s}}.$$

Note that, if any of the constraints on the attack resources are active, a closed-form solution may not be easy and one has to solve the optimization problem (8) instead.

Proof. 1 To prove that (7) and (8) are equivalent, we start by noting that

$$g_t^s(u_{\mathrm{AR}_t}{}^s) = \begin{cases} 0 & c_t^s - d_t^s u_{\mathrm{AR}_t}{}^s < 0 \quad \Leftrightarrow \quad u_{\mathrm{AR}_t}{}^s > \frac{c_t^s}{d_t^s} \\ 1 & c_t^s - d_t^s u_{\mathrm{AR}_t}{}^s > 1 \quad \Leftrightarrow \quad u_{\mathrm{AR}_t}{}^s < \frac{c_t^s - 1}{d_t^s} \\ c_t^s - d_t^s u_{\mathrm{AR}_t}{}^s & u_{\mathrm{AR}_t}{}^s \in \left[\frac{c_t^s - 1}{d_t^s}, \frac{c_t^s}{d_t^s}\right]. \end{cases}$$

Suppose, by contradiction, that (8) could lead to a larger maximum than (7). The condition $u_{\mathrm{AR}t}{}^s \in \left[0, \frac{c_t^s}{d_t^s}\right]$ guarantees that the same set of $u_{\mathrm{AR}t}{}^s$ satisfy the constraints of (7) and that

$$g_t^s(u_{\mathrm{AR}t}{}^s) = \begin{cases} 1 & c_t^s - d_t^s u_{\mathrm{AR}t}{}^s > 1 \quad \Leftrightarrow \quad u_{\mathrm{AR}t}{}^s < \frac{c_t^s - 1}{d_t^s} \\ c_t^s - d_t^s u_{\mathrm{AR}t}{}^s & c_t^s - d_t^s u_{\mathrm{AR}t}{}^s \leqslant 1 \quad \Leftrightarrow \quad u_{\mathrm{AR}t}{}^s \geqslant \frac{c_t^s - 1}{d_t^s} \end{cases}$$

and the condition $c_t^s - d_t^s u_{\mathrm{AR}t}{}^s - \sigma_t^s \leqslant 1$ guarantees that

$$\begin{cases} c_t^s - d_t^s u_{\mathrm{AR}t}{}^s - \sigma_t^s \leqslant 1 = g_t^s(u_{\mathrm{AR}t}{}^s) & u_{\mathrm{AR}t}{}^s < \frac{c_t^s - 1}{d_t^s} \\ c_t^s - d_t^s u_{\mathrm{AR}t}{}^s - \sigma_t^s = g_t^s(u_{\mathrm{AR}t}{}^s) - \sigma_t^s \leqslant g_t^s(u_{\mathrm{AR}t}{}^s) & u_{\mathrm{AR}t}{}^s \geqslant \frac{c_t^s - 1}{d_t^s}, \end{cases}$$

which shows that $c_t^s - d_t^s u_{\mathrm{AR}t}{}^s - \sigma_t^s \leqslant g_t^s(u_{\mathrm{AR}t}{}^s)$ and therefore (8) cannot lead to a larger maximum than (7).

Suppose now, also by contradiction, that (7) could lead to a larger maximum than (8). First note that if a few of the $u_{\mathrm{AR}t}{}^s$ that maximize (7) were larger than $\frac{c_t^s}{d_t^s}$, for those $u_{\mathrm{AR}t}{}^s$ we would have $g_t^s(u_{\mathrm{AR}t}{}^s) = 0$ and the same exact cost could be obtained for (7) by replacing each of these $u_{\mathrm{AR}t}{}^s$ with $\frac{c_t^s}{d_t^s}$. So we may assume, without loss of generality, that all the $u_{\mathrm{AR}t}{}^s$ are smaller than or equal to $\frac{c_t^s}{d_t^s}$.

In this case, we could use the same $u_{\mathrm{AR}t}{}^s$ in (8) and set

$$\sigma_t^s = \begin{cases} 0 & c_t^s - d_t^s u_{\mathrm{AR}t}{}^s \leqslant 1 \\ c_t^s - d_t^s u_{\mathrm{AR}t}{}^s - 1 & c_t^s - d_t^s u_{\mathrm{AR}t}{}^s > 1. \end{cases}$$

This selection of σ_t^s would satisfy the constraints of (8) and guarantee that

$$g_t^s(u_{\mathrm{AR}t}{}^s) = c_t^s - d_t^s u_{\mathrm{AR}t}{}^s - \sigma_t^s,$$

and therefore (7) and (8) would lead to the same maximum. This completes the proof that (7) and (8) are equivalent.

The optimization scheme just defined is a concave maximization problem (convex minimization) with linear constraints. The dual problem is given by,

$$
\begin{aligned}
J^{\perp} \coloneqq & \max_{\lambda_1 \geqslant 0, \eta_t^s \geqslant 0, \zeta_t^s \geqslant 0} \max_{u_{\mathrm{AR}_t^s} \in \mathbb{R}} \sum_{t=1}^{T} \sum_{s=1}^{S} (a_t^s + b_t^s u_{\mathrm{AR}_t^s})(c_t^s - d_t^s u_{\mathrm{AR}_t^s}) \\
& - \lambda_1 \Big(\sum_{t=1}^{T} \sum_{s=1}^{S} u_{\mathrm{AR}_t^s} - U_{\mathrm{TR}} \Big) - \sum_{t=1}^{T} \sum_{s=1}^{S} \eta_t^s \Big(u_{\mathrm{AR}_t^s} - \frac{c_t^s}{d_t^s} \Big) \\
& + \sum_{t=1}^{T} \sum_{s=1}^{S} \zeta_t^s u_{\mathrm{AR}_t^s} \\
= & \max_{\lambda_1 \geqslant 0, \eta_t^s \geqslant 0, \zeta_t^s \geqslant 0} \max_{u_{\mathrm{AR}_t^s} \in \mathbb{R}} \sum_{t=1}^{T} \sum_{s=1}^{S} \Big(a_t^s c_t^s - a_t^s d_t^s u_{\mathrm{AR}_t^s} + b_t^s c_t^s u_{\mathrm{AR}_t^s} - b_t^s d_t^s u_{\mathrm{AR}_t^s}{}^2 \\
& - \lambda_1 u_{\mathrm{AR}_t^s} - \eta_t^s u_{\mathrm{AR}_t^s} + \zeta_t^s u_{\mathrm{AR}_t^s} \Big) + \lambda_1 U_{\mathrm{TR}} + \sum_{t=1}^{T} \sum_{s=1}^{S} \eta_t^s \frac{c_t^s}{d_t^s} \\
= & \max_{\lambda_1 \geqslant 0, \eta_t^s \geqslant 0, \zeta_t^s \geqslant 0} \max_{u_{\mathrm{AR}_t^s} \in \mathbb{R}} \sum_{t=1}^{T} \sum_{s=1}^{S} \Big(a_t^s c_t^s - b_t^s d_t^s u_{\mathrm{AR}_t^s}{}^2 + (b_t^s c_t^s - a_t^s d_t^s + \zeta_t^s - \eta_t^s - \lambda_1) u_{\mathrm{AR}_t^s} \Big) \\
& + \lambda_1 U_{\mathrm{TR}} + \sum_{t=1}^{T} \sum_{s=1}^{S} \eta_t^s \frac{c_t^s}{d_t^s}.
\end{aligned}
$$

The inner maximization can be solved using standard calculus and is achieved for

$$
u_{\mathrm{AR}_t^s} = \frac{b_t^s c_t^s - a_t^s d_t^s + \zeta_t^s - \eta_t^s - \lambda_1}{2 b_t^s d_t^s},
$$

yielding

$$
J^{\perp} \coloneqq \max_{\lambda_1 \geqslant 0, \eta_t^s \geqslant 0, \zeta_t^s \geqslant 0} \sum_{t=1}^{T} \sum_{s=1}^{S} \Big(\frac{(b_t^s c_t^s - a_t^s d_t^s + \zeta_t^s - \eta_t^s - \lambda_1)^2}{4 b_t^s d_t^s} + \eta_t^s \frac{c_t^s}{d_t^s} \Big) + \lambda_1 U_{\mathrm{TR}}.
$$

For this problem the Karush-Kuhn-Tucker (KKT) conditions [1] lead to

$$
\begin{aligned}
\frac{\partial J^{\perp}}{\partial \lambda_1} = 0 \quad & \Leftrightarrow \quad \lambda_1 = \frac{\sum_{t=1}^{T} \sum_{s=1}^{S} \frac{b_t^s c_t^s - a_t^s d_t^s + \zeta_t^s - \eta_t^s}{2 b_t^s d_t^s} - U_{\mathrm{TR}}}{\sum_{t=1}^{T} \sum_{s=1}^{S} \frac{1}{2 b_t^s d_t^s}} \quad \text{or} \quad \lambda_1 = 0 \\
\frac{\partial J^{\perp}}{\partial \eta_t^s} = 0 \quad & \Leftrightarrow \quad \eta_t^s = \zeta_t^s - a_t^s d_t^s - b_t^s c_t^s - \lambda_1 \quad \text{or} \quad \eta_t^s = 0 \\
\frac{\partial J^{\perp}}{\partial \zeta_t^s} = 0 \quad & \Leftrightarrow \quad \zeta_t^s = -b_t^s c_t^s + a_t^s d_t^s + \eta_t^s + \lambda_1 \quad \text{or} \quad \zeta_t^s = 0.
\end{aligned}
$$

Let us assume that ${u_{\mathrm{AR}}}_t^s$ is inside the interval $\left[0, \frac{c_t^s}{d_t^s}\right]$, which would lead to all the η_t^s and ζ_t^s equal to zero (inactive constraints) and therefore we would need

$$\sum_{t=1}^{T}\sum_{s=1}^{S}\frac{b_t^s c_t^s - a_t^s d_t^s - \lambda_1}{2b_t^s d_t^s} = U_{\mathrm{TR}} \quad \Leftrightarrow \quad \lambda_1 = \frac{\sum_{t=1}^{T}\sum_{s=1}^{S}\frac{b_t^s c_t^s - a_t^s d_t^s}{2b_t^s d_t^s} - U_{\mathrm{TR}}}{\sum_{t=1}^{T}\sum_{s=1}^{S}\frac{1}{2b_t^s d_t^s}} \geqslant 0 \tag{9}$$

$$\text{or} \qquad \lambda_1 = 0$$

and

$${u_{\mathrm{AR}}}_t^s = \bar{u}_t^s - \bar{\mu}_t^s \max\left\{0, \sum_{\bar{t}}\sum_{s=1}^{S}\bar{u}_{\bar{t}}^s - U_{\mathrm{TR}}\right\}, \quad \bar{u}_t^s \coloneqq \frac{b_t^s c_t^s - a_t^s d_t^s}{2b_t^s d_t^s}, \quad \bar{\mu}_t^s \coloneqq \frac{\frac{1}{2b_t^s d_t^s}}{\sum_{\bar{t}}\sum_{s=1}^{S}\frac{1}{2b_{\bar{t}}^s d_{\bar{t}}^s}}.$$

We can view the term being subtracted from ${u_{\mathrm{AR}}}_t^s$ as a normalizing term that makes sure that the ${u_{\mathrm{AR}}}_t^s$ add up to the constraint U_{TR}.

Note that if the closed-form formula shown above for ${u_{\mathrm{AR}}}_t^s$ ever becomes negative, then the corresponding ζ_t^s will become active and we must have

$$\frac{\partial J^{\perp}}{\partial \zeta_t^s} = 0 \quad \Rightarrow \quad \zeta_t^s = \lambda_1 + a_t^s d_t^s - b_t^s c_t^s \quad \Rightarrow \quad {u_{\mathrm{AR}}}_t^s = 0.$$

Similarly if the formula for ${u_{\mathrm{AR}}}_t^s$ ever becomes larger than $\frac{c_t^s}{d_t^s}$, then the corresponding η_t^s will become active and we must have

$$\frac{\partial J^{\perp}}{\partial \eta_t^s} = 0 \quad \Rightarrow \quad \eta_t^s = -b_t^s c_t^s - \lambda_1 - a_t^s d_t^s \quad \Rightarrow \quad {u_{\mathrm{AR}}}_t^s = \frac{c_t^s}{d_t^s}. \qquad \blacksquare$$

Remark 1. Note that, if any of the constraints on the attack resources are active, a closed-form solution is not possible and one has to solve the optimization problem instead. □

3.2 Unknown Mission Damage Data

Often the mission-specific parameters $\{a_t^s, b_t^s, c_t^s, d_t^s : \forall s, t\}$ that define the potential damage and uncertainty equations are not known a-priori and, instead, need to be estimated online.

One approach that can be used to address this scenario is to assume that these parameters are generated by linear dynamics of the form

$$x_{at+1}^s = A_a^s x_{at}^s + B_a^s w_t^s, \qquad a_t^s = C_a^s x_{at}^s, \tag{10}$$
$$x_{bt+1}^s = A_b^s x_{bt}^s + B_b^s w_t^s, \qquad b_t^s = C_b^s x_{bt}^s, \tag{11}$$
$$x_{ct+1}^s = A_c^s x_{ct}^s + B_c^s w_t^s, \qquad c_t^s = C_c^s x_{ct}^s, \tag{12}$$
$$x_{dt+1}^s = A_d^s x_{dt}^s + B_d^s w_t^s, \qquad d_t^s = C_d^s x_{dt}^s, \tag{13}$$

where the $\{w_t^s, \forall s, t\}$ are sequences of zero-mean random processes with variances σ_w^s. One can then use historical data to estimate these dynamics using black-box identification techniques. Once estimates for the dynamics are available, one can use online data to predict future values for the mission-specific parameters $\{a_t^s, b_t^s, c_t^s, d_t^s : \forall s, t\}$, based on past observations.

Suppose that at some time $k < T$ the attacker has observed the values of the past mission-specific parameters $\{a_t^s, b_t^s, c_t^s, d_t^s : \forall s, t \leqslant k\}$ and needs to make decisions on the future attack resources $u_{\mathrm{AR}_t^s}$, $t \geqslant k$. One can use (10)–(13) to construct estimates $\{\hat{a}_t^s, \hat{b}_t^s, \hat{c}_t^s, \hat{d}_t^s : \forall s, t > k\}$ for the future mission-specific parameters and obtain the future $u_{\mathrm{AR}_t^s}$, $t \geqslant k$ using the following optimization:

$$\text{maximize} \quad \sum_{t=1}^{k}\sum_{s=1}^{S} f_t^s(u_{\mathrm{AR}_t^s}) g_t^s(u_{\mathrm{AR}_t^s}) + \sum_{t=k+1}^{T}\sum_{s=1}^{S} \hat{f}_t^s(u_{\mathrm{AR}_t^s}) \hat{g}_t^s(u_{\mathrm{AR}_t^s}) \tag{14}$$

$$\text{subject to} \quad \sum_{t=1}^{T}\sum_{s=1}^{S} u_{\mathrm{AR}_t^s} \leqslant U_{\mathrm{TR}} \tag{15}$$

$$\text{w.r.t.} \quad u_{\mathrm{AR}_t^s} \in [0, \infty), \ \forall t \in \{k, ..., T\}, \ \forall s, \tag{16}$$

where f_t^s and g_t^s denote the functions defined in (4) and (5), respectively, whereas $\hat{f}_t^s$ and $\hat{g}_t^s$ are estimates of these functions computed using the estimated mission-specific parameters $\{\hat{a}_t^s, \hat{b}_t^s, \hat{c}_t^s, \hat{d}_t^s : \forall s, t > k\}$.

The optimization (14) can be solved at each time step $k \in \{1, 2, \ldots, T-1\}$, allowing the attacker to improve her allocation of attack resources as new information about the missing parameters becomes available. Note that one could remove from the (double) summations in (14) any terms that do not depend on the optimization variables.

4 iCTF Competition

The international Capture The Flag (iCTF) is a distributed wide-area security exercise to test the security skills of the participants. This contest is organized by the Security Lab of the Department of Computer Science at UCSB and it has been held yearly since 2003. In traditional editions of the iCTF (2003-2007), the goal of each team was to maintain a set of services such that they remain available and uncompromised throughout the contest. Each team also had to attempt to compromise the other teams' services. Since all the teams received an identical copy of the virtual host containing the vulnerable services, each team had to find the vulnerabilities in their copy of the hosts and possibly fix the vulnerabilities without disrupting the services. At the same time, the teams had to leverage their knowledge about the vulnerabilities they found to compromise the servers run by other teams. Compromising a service allowed a team to bypass the service's security mechanisms and to "capture the flag" associated with the service. During the 2008-2010 iCTFs, new competition designs have been introduced. More precisely, in 2008 a separate virtual network was created

for each team. The goal was to attack a terrorist network and defuse a bomb after compromising a number of hosts. In 2009, the participants had to compromise the browsers of a large group of simulated users, steal their money, and create a botnet. In 2010, the participants had to attack the rogue nation Litya, ruled by the evil Lisvoy Bironulesk. The teams' goal was to attack the services supporting Litya's infrastructure only at specific times, when certain activities were in progress. In addition, an intrusion detection system would temporarily firewall out the teams whose attacks were detected. The 2011 iCTF competition is briefly summarized below from the perspective of one team playing against the rest of the world. The 2010 [2] and 2011 [5] iCTF competitions were designed closely match practical cyber-security mission scenarios.

4.1 2011 iCTF

The 2011 iCTF was centered around the theme of illegal money laundering. This activity is modeled after cyber-criminal money laundering operations and provided a perfect setting for risk-reward analysis, as the trade-offs are very intuitively understood.

The general idea behind the competition was the conversion ("laundering") of money into points. The money was obtained by the teams by solving security-related challenges (e.g., decrypting an encrypted message, find hidden information in a document, etc.) The conversion of money into points was performed by utilizing data captured from an exploited service. Therefore, first a team had to obtain money by solving challenges, and then the money had to be translated into points by exploiting the vulnerability in a service of another team. Successful conversion of money to points depended on a number of factors, calculated together as the "risk function", which is described in detail below. Note that, at the end of the game, the money had no contribution to the final stand of a team: only points mattered.

One challenge with the formulation "one-against-world" is that in the 2011 iCTF game, winning was not just about maximizing points. Winning was about getting more points than each of the opponents (individually).

The game was played in rounds 255 (each takes about 2min), but we only have data for 248 rounds since the logging server was temporarily down. Each team hosts a server that runs 10 services each with its own (unknown) vulnerabilities. Each service $s \in \{1, 2, \ldots, 10\}$ of each hosting team is characterized by three time-varying quantities $\forall t \in \{1, 2, \ldots, 248\}$:

- the *cut* C_t^s, which is the percentage of money that goes to the team when money is laundered through service s (same values for every team),

- the *payoff* P_t^s, which is the percentage of money that will be transformed into points for the team that launders the money (same value for every team);

$$P_t^s = 0.9e^{-\frac{TicksActive}{10}}$$

– the *risk* R_t^s, which is the probability of losing all the money (instead of getting a conversion to points).

The generation of the time series for the cuts, payoffs, and risks for the different services was based on an underlying set of cyber missions that were running while the game was played. Essentially, when the states of the cyber missions required a particular service, the cut, payoff, and risk would make that service attractive for attackers from the perspective of converting money to points. However, the players were not informed about the state of the cyber-missions and, instead, at the beginning of each round t, the team is informed of the values of C_t^s, P_t^s, R_t^s for every s, and t.

4.2 Actions Available to Every Team

A team (we) has the following key actions in the actual competition:

1. *Defensive actions*: Activate/deactivate one of its own services.
   In the iCTF competition a team could also correct any vulnerability that it discovered in its services. We assumed here that all known vulnerabilities had been corrected.
2. *Money laundering*: Select
   (a) team to attack (mute decision within the "one-against-world" formulation);
   (b) service s to compromise, which implicitly determines the payoff P_t^s, the risk R_t^s, and the cut C_t^s;
   (c) amount of money to launder $u_{\mathrm{AR}t}^s$ at time t through the service s.
   This action results in a number of points given by

$$X_t^s = \begin{cases} P_t^s(1-C_t^s)D_t u_{\mathrm{AR}t}^s & \text{w.p. } 1-\min\{\rho_t^s,1\} \\ 0 & \text{w.p. } \min\{\rho_t^s,1\} \end{cases} \tag{17}$$

   where D_t is the team's defense level and ρ_t^s is the probability that the conversion of money to points will succeed, as given by the formula

$$\rho_t^s \coloneqq \frac{R_t^s u_{\mathrm{AR}t}^s}{30} + \frac{1}{6}\Big(\frac{N_t^j - 700}{300 + |N_t^j - 700|} + 1\Big) + \frac{1}{6}\Big(\frac{Q_t^s - 1500}{300 + |Q_t^s - 1500|} + 1\Big)$$

   where N_t^j is the overall amount of money that has been laundered by the team j through the particular team being exploited and Q_t^s is the overall amount of money that has been laundered by the team through the particular service being exploited. Because we do not model each team individually we will consider the "worst" case scenario for the following quantities, $N = 492$, $Q = 2257$ (according to data from the competition), and defense level of the team as $D = 1$.

To map this game with the general framework described in Section 2, we associate the money to launder $u_{\mathrm{AR}_t^s}$ at time t through service s with the resources $u_{\mathrm{AR}_t^s}$ devoted to attack service s at time t, and associate the points X_t^s in (17) with damage to the mission.

The total attack resources U_{TR} available to each team in the general framework described in Section 2, now corresponds to the money available to each team. While we could model more accurately the process by which teams get money, for simplicity we assumed that each team had available a fixed amount of money ($5060) that could be spend throughout the duration of the game which is given by the average money of all the teams during the competition. The results regarding which services where attacked and when proved to be relatively insensitive to this parameter.

4.3 Optimization Schemes and iCTF

In this section we apply the optimization schemes defined in Sections 3.1 and 3.2 to the iCTF game. We are seeking to optimally allocate our available resources in the competition such that the total number of points is maximized while meeting the specified constraints. The maximization of the expected reward by a team can be formulated as follows

$$\begin{aligned}
\text{maximize} \quad & \sum_{t=1}^{248}\sum_{s=1}^{10} \rho_t^s P_t^s (1-C_t^s) D_t^s u_{\mathrm{AR}_t}{}^s \\
\text{subject to} \quad & \sum_{t=1}^{248}\sum_{s=1}^{10} u_{\mathrm{AR}_t}^s \leqslant U_{\mathrm{TR}} \coloneqq 5060 \\
\text{w.r.t.} \quad & u_{\mathrm{AR}_t}^s \in [0,\infty),\ \forall s \in \{1,2,\ldots,10\},\ t \in \{1,2,\ldots,248\},
\end{aligned}$$

where,

$$\begin{aligned}
&\rho_t^s = \min\{(\beta_t^s + \frac{R_t^s}{30} u_{\mathrm{AR}_t}{}^s), 1\}, \\
&\beta_t^s \coloneqq \frac{1}{6}\Big(\frac{N_t - 700}{300 + |N_t - 700|} + 1\Big) + \frac{1}{6}\Big(\frac{Q_t - 1500}{300 + |Q_t - 1500|} + 1\Big) = 0.4
\end{aligned}$$

and the parameters $P_t^s, C_t^s, D_t^s, \beta_t^s$ can either be considered known or unknown.

By using Proposition 1, and setting the constraint $\sigma_t^s = 0$ in (8) (since $(1-\beta_t^s) \in [0,1]$), we can write the equivalent optimization problem as,

$$\begin{aligned}
\text{maximize} \quad & \sum_{t=1}^{248}\sum_{s=1}^{10} (1 - \beta_t^s - \frac{R_t^s}{30} u_{\mathrm{AR}_t}{}^s) P_t^s (1-C_t^s) u_{\mathrm{AR}_t}{}^s \\
\text{subject to} \quad & \sum_{t=1}^{248}\sum_{s=1}^{10} u_{\mathrm{AR}_t}^s \leqslant U_{\mathrm{TR}} \\
\text{w.r.t.} \quad & u_{\mathrm{AR}_t}^s \in \Big[0, \frac{1-\beta_t^s}{\frac{R_t^s}{30}}\Big],\ \forall s \in \{1,2,\ldots,10\},\ t \in \{1,2,\ldots,248\},
\end{aligned}$$

which is a concave maximization problem with linear constraints that is easy to solve numerically as described in Section 3.1.

The above optimization depends on the following assignments, $a_t^s \coloneqq 0$, $b_t^s \coloneqq P_t^s(1-C_t^s)$, $c_t^s \coloneqq 1-\beta_t^s$, $d_t^s \coloneqq \frac{R_t^s}{30}$. When these are not known, one can estimate $b_t^s \coloneqq P_t^s(1-C_t^s)$, $c_t^s \coloneqq 1-\beta_t^s$, $d_t^s \coloneqq \frac{R_t^s}{30}$ using a low order state space models given by (11)-(13). By then applying the optimization scheme described in Section 3.2, with a horizon of $N = 5$, one can still make accurate predictions of when and how to distribute the available attack resources. The optimization model just described, results in an optimization to obtain the future $u_{\mathrm{AR}_t}^s, \forall t \geqslant k$ and performed under a moving horizon of 5 ticks,

$$
\begin{aligned}
\text{maximize} \quad & \sum_{t=1}^{k}\sum_{s=1}^{10} b_t^s(c_t^s - d_t^s {u_{\mathrm{AR}_t}}^s){u_{\mathrm{AR}_t}}^s + \sum_{t=k+1}^{248}\sum_{s=1}^{10} \hat{b}_t^s(\hat{c}_t^s - \hat{d}_t^s {u_{\mathrm{AR}_t}}^s){u_{\mathrm{AR}_t}}^s \\
\text{subject to} \quad & \sum_{t=1}^{248}\sum_{s=1}^{10} {u_{\mathrm{AR}_t}}^s \leqslant U_{\mathrm{TR}k} \\
& x_{bt+1}^s = A_b^s x_{at}^s + B_b^s w_t^s \\
& \hat{b}_t^s = C_b^s x_{bt}^s \\
& x_{ct+1}^s = A_c^s x_{ct}^s + B_c^s w_t^s \\
& \hat{c}_t^s = C_c^s x_{ct}^s \\
& x_{dt+1}^s = A_d^s x_{dt}^s + B_d^s w_t^s \\
& \hat{d}_t^s = C_d^s x_{dt}^s \\
\text{w.r.t.} \quad & {u_{\mathrm{AR}_t}}^s \in \left[0, \frac{\hat{c}_t^s}{\hat{d}_t^s}\right], \ \forall t \in \{k, \ldots, 248\}, \forall s \in \{1, 2, \ldots, 10\}.
\end{aligned}
$$

5 iCTF Results

This section presents numerical results obtained from the optimizations described above to data from the attack logs of the 2011 iCTF competition. All the optimizations have been implemented through a Matlab-based convex optimization solver such as CVX [1]. The optimization scheme described in Section 3.2 yielded very close results to the scheme described in Section 3.1 for a predicting horizon of $N = 5$.

Initially we will assume that a "sophisticated" attacker would be able to compromise any one of the 10 services. Figure 1 show the points and the money collected by such an optimal attacker, whereas Figure 2 shows the same (aggregate) data for the teams that participated in the competition.

One can also consider attackers with different level of sophistication, e.g., attackers that are only able to find vulnerabilities in a subset of the 10 services that the "sophisticated" was able to attack. By observing the data of the top 20 teams in the competition we were able to partition the sophistication in two levels. For comparison, we show the behavior of an attacker A that was

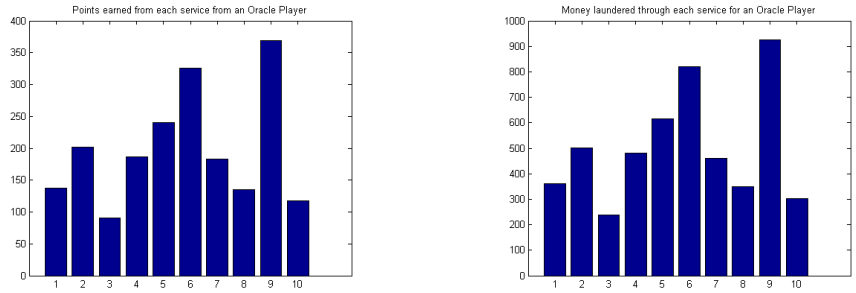

(a) points earned (b) money laundered

Fig. 1. Behavior of an optimal "sophisticated" attacker able to attack all 10 services

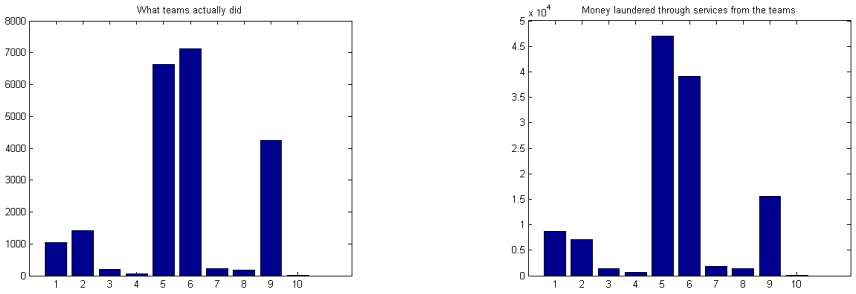

(a) points earned (b) money laundered

Fig. 2. Aggregate behavior of all teams that participated in the competition

only able to attack the services $1, 2, 4, 5, 6, 9$ (similar to the first 10 teams in the competition); and another attacker B that was only able to attack services $1, 2, 5, 6, 7, 8$ (similar to the teams from place 11 to 20 in the competition). The "sophisticated" attacker was able to gather with 1987 points, whereas the two other attackers were able to get 1821 and 1721 points, respectively.

The results in Figure 1(a) show that the most profitable services to attack were $5, 6$ and 9. The top 10 teams in the competition attacked mostly 5 and 6 because 9 was a hard service to get into. Only the top 3 teams discovered how to attack service 9 and only at the end of the game so they had relatively little time to explore that vulnerability. Aside from this, the prediction based on the optimization framework developed here qualitatively reflect the actions of the good teams. In fact, the top two team in the competition followed attack strategies qualitatively close to that of attacker A in Figure 3 as seen in Figure 5.

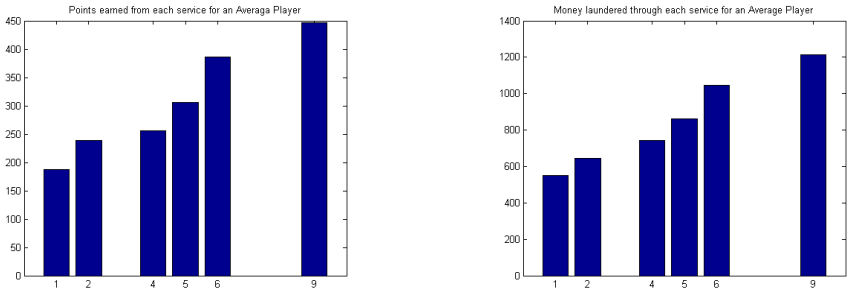

(a) points earned (b) money laundered

Fig. 3. Behavior of an optimal attacker A able to attack services 1,2,4,5,6,9

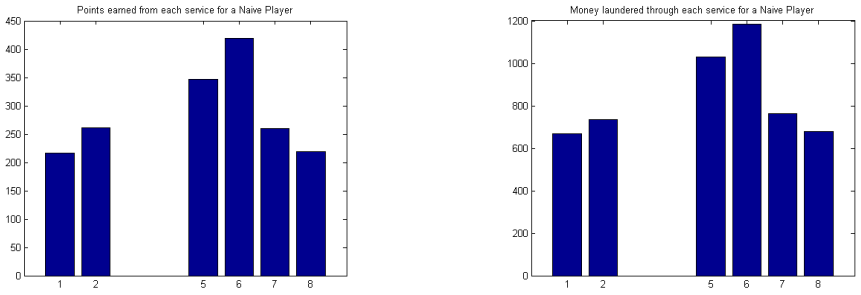

(a) points earned (b) money laundered

Fig. 4. Behavior of an optimal attacker B able to attack services 1,2,5,6,7,8

6 Future Work

Our future work in this area is focused on developing analysis tools to explore what-if scenarios based on past data and the structure of the cyber-mission. To this end, we are developing optimization schemes for the defender's possible actions, such as taking a service off-line when the service is not needed or extending the duration of a state that would be unable to progress if a certain service is compromised. We are also developing human-computer interfaces to demonstrate the useful of this type of analysis for security analysts.

References

1. S. Boyd and Lieven Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004
2. A. Doupe, M. Egele, B. Caillat, G. Stringhini,G. Yakin, A. Zand, L. Cavedon, and G. Vigna, "Hit'em where it hurts: A live security exercise on cyber situational

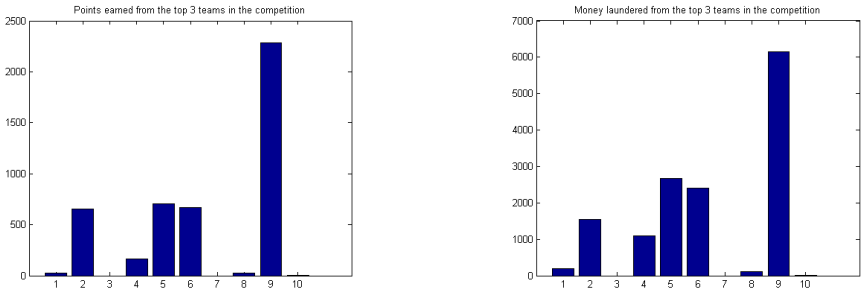

(a) points earned (b) money laundered

Fig. 5. Behavior of the top 3 teams during the competition

awareness," *Proceedings of the Annual Computer Security Applications Conference (ACSAC 2011)*, Orlando, FL, December 2011
3. M. Endsley, *Theoretical Underpinnings of Situation Awareness: A Critical Review*, chapter 1, pages 3-32, L. Erlbaum Assoc., 2000
4. N. Stockman, K. G. Vamvoudakis, L. Devendorf, T. Hllerer, R. Kemmerer, J. P. Hespanha, "A Mission-Centric Visualization Tool for Cybersecurity Situation Awareness," Technical Report, University of California, Santa Barbara, August 2012
5. G. Vigna, *The 2011 UCSB iCTF: Description of the game, http://ictf.cs.ucsb.edu/*, 2011