European Journal of Operational Research 252 (2016) 888–899

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Decision Support

Multithreat multisite protection: A security case study

David Ríos Insua[a], Javier Cano[b,*], Michael Pellot[c], Ricardo Ortega[c]

[a] *Instituto de Ciencias Matemáticas ICMAT-CSIC, Spain*
[b] *Department of Computer Science and Statistics, Rey Juan Carlos University, Spain*
[c] *TMB, Spain*

ARTICLE INFO

Article history:
Received 28 April 2015
Accepted 20 January 2016
Available online 30 January 2016

Keywords:
Risk analysis
Multiple intelligent attackers
Adversarial risk analysis
Multisite protection
Transportation

ABSTRACT

We provide a novel adversarial risk analysis approach to security resource allocation decision processes for an organization which faces multiple threats over multiple sites. We deploy a Sequential Defend-Attack model for each type of threat and site, under the assumption that different attackers are uncoordinated, although cascading effects are contemplated. The models are related by resource constraints and results are aggregated over the sites for each participant and, for the Defender, by value aggregation across threats. We illustrate the model with a case study in which we support a railway operator in allocating resources to protect from two threats: fare evasion and pickpocketing. Results suggest considerable expected savings due to the proposed investments.

1. Introduction

Crime and terrorism constitute major global issues. As an example, among the threats considered in the World Economic Forum (2015) Global Risks report, there are several related with security, including large-scale terrorist attacks or a major escalation in organized crime. Similarly, we may find security among the seven thematic H2020 priorities for European research (ec.europa.eu/programmes/horizon2020). Governments and organizations worldwide are indeed increasingly committed to protecting themselves against various security threats. Recent large-scale terrorist events like 9/11 or the Madrid train bombings have led to significant national investments in protective responses, see (Haberfeld & von Hassell, 2009). However, public opinion has not always seen such expenditures as prudent or effective, see (Parnell et al., 2008) or (Sunstein, 2007).

In turn, this has motivated great interest in modeling issues in relation with security, with varied tools from areas such as reliability analysis, data mining, game theory or complex dynamic systems. Recent accounts of various techniques and applications in the field of counterterrorism may be seen in e.g. Ezell, Bennett, von Winterfeldt, Sokolowski, and Collins (2010) or Wein (2009). Parnell et al. (2008) and Enders and Sandler (2011) provide overviews on strategies, models, and research issues in security risk analysis. Other relevant work include e.g. Zhuang and Bier (2007), who discuss resource allocation for countering terrorism and natural disasters; (Brown, Carlyle, Salmerón, & Wood, 2006), where the protection of critical infrastructures is addressed; or (Yang, Kiekintveld, Ordóñez, Tambe, & John, 2013), who present mathematical models of adversarial behavior to support security forces in their fight against different adversaries.

We consider problems in which an organization needs to protect multiple sites from multiple threats. Our case study refers to deciding the security resource allocation for a railway system whose operator faces threats from fare evaders and pickpockets. The figures presented in the paper have been modified from actual figures to protect the confidentiality of the case study provider. Therefore, *the data is realistic data but not actual data*. We assume that the relevant multiple threats are uncoordinated, in that different attackers do not make common cause, although the outcome of different types of attacks might affect each other. In our case study, fare evaders and pickpockets will not be coordinated, although pickpockets alone and a group of fare evaders will be organized. Hausken and Levitin (2012) provide a classification of systems defense and attack models. Within such classification, we shall be facing a case of protection from attacks over multiple elements with incomplete information. For earlier work on protecting from multiple attackers, see (Hausken & Bier, 2011) and references therein. Haphuriwat and Bier (2011); Hausken (2014b) and Levitin, Hausken, and Dai (2014) provide ideas in relation with multiple site protection. Bier, Oliveros, and Samuelson (2007) and Hausken (2014a) refer to uncertainty in attacker resources and asset valuation. All of them perform game theoretical analyses under convenient common knowledge assumptions.

[*] Corresponding author. Tel.: +34 914 888 411.
E-mail addresses: david.rios@icmat.es (D. Ríos Insua), javier.cano@urjc.es (J. Cano), mpellot@tmb.cat (M. Pellot), rortegap@tmb.cat (R. Ortega).

http://dx.doi.org/10.1016/j.ejor.2016.01.041
0377-2217/© 2016 Elsevier B.V. All rights reserved.

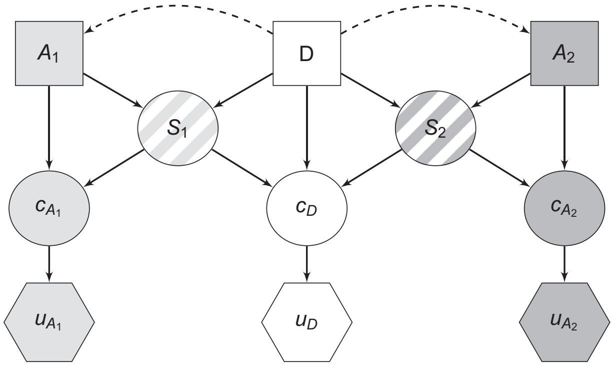

Fig. 1. Multiagent influence diagram for a bi-threat problem.

In contrast, we provide an adversarial risk analysis (ARA) approach, see (Ríos Insua, Ríos, & Banks, 2009), for such problems, combining multiple sites, multiple attackers and taking into account all relevant uncertainty sources. ARA builds a decision analysis model for one of the agents (she, the Defender), who forecasts the actions of her adversaries. Then, she will be able to decide her optimal defensive actions. Our approach will be based on the Sequential Defend-Attack model, see e.g. Brown et al. (2006) or Ríos and Ríos Insua (2012). In it, the Defender first chooses a defense and, then, after observing it, the Attacker decides his attack. We deploy one of such models for each type of threat and site, which we relate through resource constraints and aggregation of results over various sites for each participant and, for the case of the Defender, also by value aggregation over the various threats. We assume no particular spatial structure relating the sites, e.g. through proximity or a neighboring structure, see (Gil, Ríos, & Ríos Insua, 2016).

In Section 2, we provide a general framework for the basic problem of protecting a single site from multiple threats, illustrating it with our case in Section 3. Section 4 extends the previous model to the protection of multiple sites, applying it to an expanded version of the case study in Section 5. As described below, all the involved parameters have been assessed with the aid of transportation experts, using expert judgment elicitation techniques, see (O'Hagan et al., 2006) or Farquhar (1984), then validated at a security transportation workshop, and finally checked for robustness through sensitivity analysis.

2. Single site multithreat protection

We start with the basic multithreat protection problem over a single site. We consider a Defender, D, who needs to deploy defensive resources $d \in \mathcal{D}$ to protect the site from m uncoordinated attackers $A_1, \ldots, A_m$. These observe her decision d and, then, respectively, launch attacks $a_i \in \mathcal{A}_i$, $i = 1, \ldots, m$. The interaction between D and A_i through their corresponding decisions d and a_i, leads to a random result $S_i \in \mathcal{S}_i$. The Defender faces multiattribute consequences c_D which depend on her defense effort d and the results $s_1, \ldots, s_m$. She then gets her utility u_D. Each attacker will get his multiattribute consequences c_{A_i}, which depend on his attack effort a_i and his result s_i, and then gets his utility u_{A_i}.

The problem is illustrated in the multiagent influence diagram in Fig. 1, see (Koller & Milch, 2003). For simplicity, we only display two attackers, that is $m = 2$. White nodes correspond to the Defender, solid (light and dark) gray nodes to attackers A_1 and A_2, respectively. Striped nodes refer to interactions between the Defender and the attackers. Dashed arrows between node D and nodes A_1 and A_2 indicate that the attackers decide their alternatives after having observed the decision by D.

As an example, a port authority (D) is trying to protect a port against actions from drug smugglers (A_1) and terrorists (A_2) ready to introduce nuclear weapons. The Defender decisions d are portfolios which could include sniffer dogs, metal detectors, inspectors and others. Drug smuggler decisions a_1 typically would refer to drug smuggling (timing, placing, quantities) strategies. Terrorist decisions a_2 could refer to weapon smuggling strategies, like whether or not to infiltrate a nuclear weapon. S_1 could refer to the amount of drugs actually smuggled and S_2 could refer to the number of weapons smuggled.

The Defender aims at finding her optimal defense strategy d^*. She evaluates her consequences through her utility $u_D(d, s_1, \ldots, s_m)$. Assuming conditional independence between the outcomes S_i of different attacks, given the defensive resources d and attacks a_i, she needs to assess the probability models $p_D(s_i|d, a_i)$, $i = 1, \ldots, m$, reflecting which outcomes she finds more likely when attacker A_i launches attack a_i and she has deployed defensive resources d. She gets her expected utility, given the attacks, integrating out the uncertainty over the outcomes of the attacks:

$$\psi_D(d|a_1, \ldots, a_m) = \int \cdots \int u_D(d, s_1, \ldots, s_m) \times p_D(s_1|d, a_1) \cdots p_D(s_m|d, a_m) \, \mathrm{d}s_1 \ldots \mathrm{d}s_m. \quad (1)$$

Suppose that the Defender is able to build the models $p_D(a_i|d)$, $i = 1, \ldots, m$, expressing her beliefs about which attack a_i will be chosen by the ith attacker after having observed d. Since attacks are uncoordinated, we assume conditional independence of $a_1, \ldots, a_m$ given d. Then, D may compute

$$\psi_D(d) = \int \cdots \int \psi_D(d|a_1, \ldots, a_m) \, p_D(a_1|d) \cdots p_D(a_m|d) \times \mathrm{d}a_1 \ldots \mathrm{d}a_m,$$

and solve $\max_d \psi_D(d)$ to find her optimal defense resource allocation d^*.

The only nonstandard assessments in this formulation are those of $p_D(a_i|d)$. To obtain them, the Defender may put herself into the shoes of each attacker, and solve their corresponding problem separately, since they are uncoordinated. For instance, for the problem faced by attacker A_1, assuming that he is an expected utility maximizer, see (French & Ríos Insua, 2000), the Defender would need his utility $u_{A_1}(a_1, s_1)$ and probabilities $p_{A_1}(s_1|d, a_1)$. Then, she would solve

$$a_1^*(d) = \operatorname*{argmax}_{a_1 \in \mathcal{A}_1} \int u_{A_1}(a_1, s_1) \, p_{A_1}(s_1|d, a_1) \, \mathrm{d}s_1, \quad (2)$$

to find his optimal attack given that she has implemented d. However, she lacks knowledge about u_{A_1} and p_{A_1}. Suppose we may model her uncertainty about them through random utilities and probabilities $\left(U_{A_1}, P_{A_1}\right)$, and propagate that uncertainty to obtain the random optimal attack, given her defense d,

$$A_1^*(d) = \operatorname*{argmax}_{a_1 \in \mathcal{A}_1} \int U_{A_1}(a_1, s_1) \, P_{A_1}(s_1|d, a_1) \, \mathrm{d}s_1. \quad (3)$$

Then, we would get $p_D(a_1|d) = \Pr(A_1^*(d) \le a_1)$, which may be approximated by Monte Carlo through Algorithm 1.

A similar scheme could be implemented in parallel for the other attackers, $A_2, \ldots, A_m$, leading to estimates $\widehat{p_D}(a_i|d)$ of the required probabilities $p_D(a_i|d), i = 2, \ldots, m$.

The approach may be generalized in several ways. For example, the simultaneous, but uncoordinated, implementation of attacks $a_1, \ldots, a_m$ could be jointly detrimental in face of defensive resources d, which could be shared against various types of attacks, see Fig. 2a. Then, we could rewrite the probability model in (1) as

$$p_D(s_1|d, a_1, \ldots, a_m) \cdots p_D(s_m|d, a_1, \ldots, a_m),$$

and proceed in a similar fashion.