Security Games with Probabilistic Constraints on the Agent's Strategy

Corine M. Laan[1,2,3](✉), Ana Isabel Barros[2,4], Richard J. Boucherie[1], and Herman Monsuur[3]

[1] Stochastic Operations Research, University of Twente, Enschede, Netherlands
c.m.laan@utwente.nl
[2] TNO - Defense, Safety and Security, Den Hague, Netherlands
[3] Netherlands Defense Academy, Den Helder, Netherlands
[4] Institute for Advanced Study, Amsterdam, Netherlands

Abstract. This paper considers a special case of security games dealing with the protection of a large area divided in multiple cells for a given planning period. An intruder decides on which cell to attack and an agent selects a patrol route visiting multiple cells from a finite set of patrol routes such that some given operational conditions on the agent's mobility are met. For example, the agent might be required to patrol some cells more often than others. In order to determine strategies for the agent that deal with these conditions and remain unpredictable for the intruder, this problem is modeled as a two-player zero-sum game with probabilistic constraints such that the operational conditions are met with high probability. We also introduce a variant of the basic constrained security game in which the payoff matrices change over time, to allow for the payoff that may change during the planning period.

Keywords: Game theory · Probability constraints · Defense applications

1 Introduction

This paper considers a special case of a security game dealing with the protection of a large area for a given time period where the agent's strategy set is restricted. The area consists of several cells containing assets to be protected. An intruder decides on which cell to attack, while the agent needs to select a patrol route that visits multiple cells. The agent's strategy is constrained by existing governmental guidelines that require that some cells should be patrolled more often than others. This problem can be modeled as a two-player zero-sum game with probabilistic constraints.

In the literature there are several models considering patrolling games (e.g., [1,5,8]). Also, many models consider constraints on the agent's or intruder's strategy set. For example in [2,6,15], the authors require constraints on the agent's strategy because only a limited number of resources is available, and

S. Rass et al. (Eds.): GameSec 2017, LNCS 10575, pp. 481–493, 2017.
DOI: 10.1007/978-3-319-68711-7_25

in [17] the authors consider constraints on both the agent's and the intruder's strategy set.

Often, linear constraints are considered in constrained games. For instance, in [3] a two-person zero-sum game with linear constraints is introduced. More recently, [10] described a bimatrix game with linear constraints on the strategy of both players. In [14], the author considers nonlinear ratio type constraints. Our security game models situations where operational conditions have to be met with high probability, which results in nonlinear probabilistic constraints.

An example application of our model lies in countering illegal or unreported and unregulated fishing. These illicit activities endanger the economy of the fishery sector, fish stocks and the marine environment and require the monitoring of large areas with scarce resources subject to national regulations. To support the development of patrols against illegal fishing, in [7] a decision support system is developed. This system models the interaction between different types of illegal fishers and the patrolling forces as a repeated game. More recently, [4] introduced a game theoretical approach wherein a generalization of Stackelberg games is used to derive sequential agent strategies that learn from adversary behavior. However, these papers do not consider constraints to the patroller's strategy.

The main contribution of this paper is that we introduce a new model to cope with the conditions on the agent's random strategy that have to be met with high probability. Because of the random nature of the strategies, it cannot be guaranteed that the conditions are always met. By introducing probabilistic constraints, we assure that the conditions are met with high probability. In practice the payoff matrices may change over time, in the fishery case, due to weather conditions, seasonal fluctuations or other circumstances. Therefore, we introduce an extension of the model to deal with multiple payoff matrices.

This paper is organized as follows. In Sect. 2, we introduce the new security game model with constraints on the agent's strategy. In Sect. 3, we present an extension of the model in which multiple payoff matrices are considered. Finally, in Sect. 4 we give examples of the model and present computational results.

2 Model with Constant Payoff

This section describes the model assuming that the gain an intruder obtains by successfully visiting a cell is constant over the planning period. We first provide a general description of a constrained security game over multiple cells in Sect. 2.1. For each cell, there is a condition on the minimal number of visits per time period for that cell. We discuss the probability that these conditions are met for each cell separately in Sect. 2.2, which gives a lower bound for the game value. In the application of countering illegal fishing, governmental guidelines require that some cells should be patrolled more than others because some regions are more vulnerable. The conditions on the number of visits have to be met for all cells simultaneously. These simultaneous conditions are discussed in Sect. 2.3.

2.1 Constrained Game

We consider a security game with constraints on the strategy sets (see [11], Chap. 3.7). Let $C = \{1, ..., N_C\}$ be the set of cells that can be attacked by an intruder and let $R = \{1, ..., N_R\}$ be the set of routes that can be chosen by the agent. The matrix A indicates which cells are visited by each route, such that a_{ij} equals 1 if route i includes cell j and 0 otherwise. Let M be the payoff matrix, such that m_{ij} is the payoff for the intruder if the agent chooses route i and the intruder attacks cell j, $i = 1, ..., N_R$, $j = 1, ..., N_C$:

$$m_{ij} = ((1 - d_j)a_{ij} + (1 - a_{ij}))\, g_j, \quad i = 1, ..., N_R,\ j = 1, ..., N_C, \tag{1}$$

where g_j is the intruder's gain if the intruder successfully attacks cell j and d_j is the probability that the intruder is caught if the agent's chosen route i includes cell j. The game is repeated N_D times (e.g. days), our planning period. We assume that only one intruder is present in the area. If that intruder is caught, then another will replace him. The overall aim from an intruders perspective is to maximize the total payoff over the time period.

Remark 1. Note that the model described in this section assumes that each intruder attacks one cell each day. By changing the payoff matrix and the actions of the agent and the intruder, the model can be extended to other situations. □

The intruder attempts to maximize the payoff by choosing which cell to attack, so the action set of the intruder is given by C. The agent tries to catch the intruder by selecting a route, so the action set of the agent is given by R. The agent minimizes the payoff by deciding on the probability p_i, $i = 1, ..., N_R$, that route i is chosen, while the intruder maximizes the payoff by selecting the probability q_j, $j = 1, ..., N_C$, that cell j is attacked. The strategy of the agent is constrained by the conditions $f(p) \geq 0$, determined by the minimum number of times each cell is visited by the agent. In Sects. 2.2 and 2.3, we will elaborate on these conditions. The value of the game, V, equals the expected payoff per day. Optimal strategies can be found by solving the following mathematical program:

$$\begin{aligned} V = \min_p \max_q \quad & p^T M q \\ \text{s.t.} \quad & f(p) \geq 0, \\ & \sum_{i=1}^{N_R} p_i = 1,\ \sum_{j=1}^{N_C} q_j = 1, \\ & p, q \geq 0. \end{aligned} \tag{2}$$

Taking the dual of the inner linear program $\max_q\{p^T Mq | \sum_{j=1}^{N_C} q_j = 1, q \geq 0\}$, the minmax formulation (2) can be rewritten to obtain the value of the game and optimal strategies for the agent:

$$\begin{aligned} V = \min_{p,z} \quad & z \\ \text{s.t.} \quad & e^T z \geq p^T M, \\ & f(p) \geq 0, \\ & \sum_{i=1}^{N_R} p_i = 1,\ p \geq 0, \end{aligned} \tag{3}$$

where e is the row vector with only ones. Note that there only exists a value for this game if the set $\{p|f(p) \geq 0, \sum_{i=1}^{N_R} p_i = 1, p \geq 0\}$ is not empty.

Remark 2. For clearness of presentation, we model the game as a zero-sum game. Note that a similar model applies if we consider a bimatrix game in which the agent and the intruder have different payoff matrices. In bimatrix games, the game value is calculated using quadratic programming (see for example [12], Chap. 13.2) instead of linear programming, but the probabilistic constraints can be implemented similarly. In addition, in the same manner, conditions on the intruder's strategy set can be added. □

2.2 Conditions on the Number of Visits to a Cell

In this subsection, we consider conditions on the number of visits for each cell separately to obtain a lower bound for V. Let N_D be the number of days in the planning period. The strategy of the agent is constrained by the minimum number of visits v_j to each cell j, $j = 1, ..., N_C$, over the entire period N_D, that must be realized with at least probability $1 - \epsilon$. Given any strategy p, the probability that cell j is visited by the agent is $a_j p$, where a_j is the row vector of the j-th column of A.

Let X_j, $j = 1, ..., N_C$, be the random variable that records the number of visits to cell j during the planning period. The probability that cell j is visited equals $a_j p$. As there are N_D successive days, X_j is binomially distributed with parameters N_D and $a_j p$. The constraint on the number of visits then reads $P(X_j \geq v_j) \geq (1 - \epsilon)$, i.e.,

$$\sum_{k=v_j}^{N_D} \frac{N_D!}{k!(N_D - k)!}(a_j p)^k (1 - a_j p)^{N_D - k} \geq 1 - \epsilon,$$

which can be implemented in (3) by choosing $f(p) = (f_1(p), f_2(p), ...f_{N_C}(p))$ with $f_j(p) = P(X_j \geq v_j) - (1 - \epsilon)$.

For large N_D, the binomial distribution becomes intractable for implementation. Therefore, we use the following approximation. For large N_D, the binomially distributed X_j can be approximated by the normally distributed $\tilde{X}_j$ with mean $N_D a_j p$ and variance $N_D a_j p(1 - a_j p)$ (see [13], Chap. 1.8):

$$P(X_j \geq v_j) = 1 - P(X_j < v_j) \approx 1 - P(\tilde{X}_j \leq v_j),$$

yielding

$$f_j(p) = \epsilon - \Phi\left(\frac{v_j - N_D a_j p}{\sqrt{N_D a_j p(1 - a_j p)}}\right), \tag{4}$$

where $\Phi(x)$ is the cumulative distribution function for the standard normal distribution.

Considering the conditions for each cell separately gives a relaxation of the original conditions, where the minimum number of visits has to be obtained for all cells simultaneously. If we replace $f(p)$ in (3) by the constraints in (4), we obtain the following lower bound for the game value V:

$$\begin{aligned} V_L = \min_{p,z} \quad & z \\ \text{s.t.} \quad & e^T z \geq p^T M, \\ & \Phi\left(\frac{v_j - N_D a_j p}{\sqrt{N_D a_j p(1 - a_j p)}}\right) \leq \epsilon, \quad j = 1, ..., N_C, \\ & \sum_{i=1}^{N_R} p_i = 1, \ p \geq 0. \end{aligned} \tag{5}$$

In order to linearize these constraints, we can determine for each cell j all possible values of $a_j p$ such that $\epsilon - P(\tilde{X}_j \leq v_j) \geq 0$ using the table of the standard normal distribution. The constraints in (5) can be replaced by the linear constraint $p^T A \geq \tilde{b}$, where $\tilde{b}_j$ is determined by the minimum probability for each cell such that the conditions are met with probability $1 - \epsilon$.

Visits to cells are correlated via the routes. Therefore, we are interested in the joint probability:

$$P(X_1 \geq v_1, X_2 \geq v_2, ..., X_{N_C} \geq v_{N_C}),$$

that we will discuss in the next section.

2.3 Conditions on All Cells Simultaneously

In this section, we discuss the condition on the minimum number of visits for all cells simultaneously. Let Y_i, $i = 1, ..., N_R$, be the random variable that specifies the number of times that route i is selected. $Y = (Y_1, Y_2, ..., Y_{N_R})$ is multinomially distributed with parameters N_D and p:

$$P(Y_1 = v_1, Y_2 = v_2, ..., Y_{N_R} = v_{N_R}) = N_D! \prod_{i=1}^{N_R} \frac{p_i^{v_i}}{v_i!}.$$

For large N_D, Y_i, $i = 1, ..., N_R$ can be approximated by the multivariate normally distributed $\tilde{Y}_i$ with expectation $N_D p_i$, variance $N_D p_i(1 - p_i)$ and covariance $Cov(\tilde{Y}_i, \tilde{Y}_{i'}) = -N_D p_i p_{i'}$, $i' = 1, ..., N_R$ (see [13], Chap. 1.8).

The number of times cell j is visited, X_j, can then be expressed as $X_j = \sum_{i=1}^{N_R} a_{ij} Y_i$ and using the approximation $\tilde{Y}$ for Y, X_j can be approximated by a normally distributed $\tilde{X}_j$ with expectation, variance and covariance (see [13], Chap. 1.4), $j = 1, .., N_C$:

$$E(\tilde{X}_j) = N_D a_j p, \qquad Var(\tilde{X}_j) = N_D a_j p(1 - a_j p),$$
$$Cov(\tilde{X}_j, \tilde{X}_{j'}) = \sum_{i=1}^{N_R} \sum_{i'=1}^{N_R} a_{ij} a_{i'j'} Cov(\tilde{Y}_i, \tilde{Y}_{i'}).$$

The probability that the conditions are met for all cells is:

$$\begin{aligned} P(X_1 \geq v_1, X_2 \geq v_2, ..., X_{N_C} \geq v_{N_C}) &\approx P(\tilde{X}_1 \geq v_1, \tilde{X}_2 \geq v_2, ..., \tilde{X}_{N_C} \geq v_{N_C}) \\ &= \frac{1}{\sqrt{|\Sigma|(2\pi)^{N_C}}} \int_{v_1}^{\infty} \int_{v_2}^{\infty} ... \int_{v_{N_C}}^{\infty} e^{-\frac{1}{2}(v-\mu)'\Sigma^{-1}(v-\mu)} dv_{N_C} ... dv_1, \end{aligned} \tag{6}$$

where Σ is the covariance matrix and μ is a vector with all expected values. This can be implemented in (3) by choosing $f(p)$ as

$$f(p) = P(\tilde{X}_1 \geq v_1, \tilde{X}_2 \geq v_2, ..., \tilde{X}_{N_C} \geq v_{N_C}) - (1 - \epsilon). \tag{7}$$

The constraint described above is not linear and cumbersome to implement in a mathematical program. To simplify the model, we use a lower bound for the probability that the conditions are met and implement this lower bound.

A lower bound for the probability that the conditions for all cells are met is:

$$P(\tilde{X}_1 \geq v_1, ..., \tilde{X}_{N_C} \geq v_{N_C}) \geq 1 - \sum_{j=1}^{N_C} P(\tilde{X}_j < v_j). \tag{8}$$

This lower bound can be used to simplify the mathematical program as follows:

$$f(p) = \epsilon - \sum_{j=1}^{N_C} \Phi\left(\frac{v_j - N_D a_j p}{\sqrt{N_D a_j p(1 - a_j p)}}\right).$$

Replacing $f(p)$ in (3) by a lower bound in the condition, results in an upper bound for the game value V:

$$\begin{aligned} V_U = \min_{p,z} \quad & z \\ \text{s.t.} \quad & e^T z \geq p^T M, \\ & \sum_{j=1}^{N_C} \Phi\left(\frac{v_j - N_D a_j p}{\sqrt{N_D a_j p(1 - a_j p)}}\right) \geq \epsilon, \\ & \sum_{i=1}^{N_R} p_i = 1, \ p \geq 0, \end{aligned} \tag{9}$$

Combining this upper bound and the lower bound obtained in Sect. 2.2, we readily obtain the following result:

Lemma 1. *For* V_L *given in* (5) *and* V_U *given in* (9) *we have* $V_L \leq V \leq V_U$ □

In Sect. 4, we investigate the impact of this approximation modeling approach on the game value.

Remark 3. We may linearize this program by approximating the normal distribution for each cell j by a piecewise linear function as described in [16], Chap. 9.2. However, we use in the result section the mathematical program stated in (9) since this model is still solvable for realistic instances. □

3 Generalization: Multiple Payoff Matrices

The previous section considers games with constant payoff. This section considers a generalization to situations where payoff can change over time due to, e.g., weather conditions or seasonal fluctuations resulting in multiple payoff matrices.

3.1 Constrained Game

Consider the game with multiple payoff matrices $M^{(k)}$, $k = 1, ..., N_M$, of size $N_R \times N_C$. Let $\mu^{(k)}$ be the probability that the payoff matrix is $M^{(k)}$, with $\sum_{k=1}^{N_M} \mu^{(k)} = 1$. Moreover let $q^{(k)}$ and $p^{(k)}$ be strategies of the agent and the intruder when the payoff matrix is $M^{(k)}$. The value of the game is the expected payoff per day and can be found by solving the following optimization problem:

$$
\begin{aligned}
V = \min_p \max_q \quad & \sum_{k=1}^{N_M} \mu^{(k)} (p^{(k)})^T M^{(k)} q^{(k)} \\
\text{s.t.} \quad & f(p) \geq 0, \\
& \sum_{i=1}^{N_R} p_i^{(k)} = 1, \ \sum_{i=1}^{N_C} q_i^{(k)} = 1, \quad k = 1, ..., N_M, \\
& p, q \geq 0,
\end{aligned}
\tag{10}
$$

where $p^T = (p^{(1)}, ..., p^{(N_M)})$ and $q^T = (q^{(1)}, ..., q^{(N_M)})$. In the next section, we discuss the constraint $f(p) \geq 0$ if multiple payoff matrices are considered.

3.2 Conditions for Games with Multiple Payoff Matrices

The conditions on the minimal number of visits for all cells during the planning period can be constructed following the same reasoning as in Sect. 2. Now, the number of visits for cell j is the sum of the number of visits for cell j for each payoff matrix. Let $X_j^{(k)}$, $j = 1, ..., N_C$, $k = 1, ..N_M$, be the random variable describing the number of visits to cell j when the payoff matrix is M_k and let

$\tilde{X}_j^{(k)}$ be the approximation of $X_j^{(k)}$. $N_D^{(k)}$ is the number of periods that the payoff matrix is $M^{(k)}$. We are interested in the following probability:

$$P(\tilde{X}_1^{(1)} + ... + \tilde{X}_1^{(N_M)} \geq v_1, ..., \tilde{X}_{N_C}^{(1)} + ... + \tilde{X}_{N_C}^{(N_M)} \geq v_{N_C}),$$

with $E(\tilde{X}_j^{(k)})$, $Var(\tilde{X}_j^{(k)})$, and $Cov(\tilde{X}_j^{(k)})$ calculated as in Sect. 2.3 with $N_D^{(k)}$ and $p^{(k)}$. Since $\tilde{X}_j^{(k)}$ and $\tilde{X}_{j'}^{(k)}$ are independent if $j \neq j'$, we have:

$$E(\tilde{X}_j) = \sum_{k=1}^{N_M} N_D^{(k)} a_j p^{(k)}, \qquad Var(\tilde{X}_j) = \sum_{k=1}^{N_M} N_D^{(k)} a_j p^{(k)} (1 - a_j p^{(k)}),$$

$$Cov(\tilde{X}_j, \tilde{X}_{j'}) = \sum_{k=1}^{N_M} \sum_{k'=1}^{N_M} Cov(X_j^{(k)}, X_{j'}^{(k')}).$$

To make sure that the conditions are met with high probability we define,

$$f(p) = P(\tilde{X}_1 \geq v_1, ..., \tilde{X}_{N_C} \geq v_{N_C}) - (1 - \epsilon),$$

where $P(\tilde{X}_1 \geq v_1, ..., \tilde{X}_{N_C} \geq v_{N_C})$ equals (6). Similarly as in Sect. 2.3, a lower bound for this probability is given in (8). Taking the dual of the inner LP of (10) and using this lower bound, optimal strategies for the agent and the intruder can be found by solving:

$$\begin{aligned} V_U = \min_{p,z} \quad & \sum_{k=1}^{N_M} z^{(k)} \\ \text{s.t.} \quad & e^T z^{(k)} \geq \mu^{(k)} (p^{(k)})^T M^{(k)}, \quad k = 1, ..., N_M, \\ & \sum_{j=1}^{N_C} \Phi \left(\frac{v_j - \sum_{k=1}^{N_M} N_D^{(k)} a_j p^{(k)}}{\sqrt{\sum_{k=1}^{N_M} N_D^{(k)} a_j p^{(k)} (1 - a_j p^{(k)})}} \right) \geq \epsilon, \\ & \sum_{i=1}^{N_R} p_i^{(k)} = 1, \quad k = 1, ..., N_M, \\ & p \geq 0, \end{aligned} \tag{11}$$

where $z = (z^{(1)}, ..., z^{(N_M)})$. In the next section, we will illustrate this model.

4 Results

In this section, we give computational results and examples to illustrate our models. In Sect. 4.1, we investigate the approximation error introduced in Sect. 2.3. Thereafter, we give two examples to illustrate our model in Sect. 4.2.

4.1 Computational Results

This section investigates the error introduced by using the lower bound in (8). Solving (3) with $f(p)$ given in (7) numerically is computationally intractable for networks with more than two or three routes and cells. Therefore, we have compared the relative difference between the lower and upper bounds of V, see Lemma 1. We have randomly generated 100 payoff matrices, conditions and routes for different network sizes. Table 1 shows the average relative difference between the upper and lower bound with 95%-confidence interval between brackets. The last columns gives the average running time in seconds for (9). The results are implemented in Matlab version R2016b [9] on an Intel(R) Core(TM) i7 CPU, 2.4 GHz, 8 GB of RAM. As the results in Table 1 show, (9) gives a good approximation of the game value V and can be solved in reasonable time. The size of more realistic examples, as encountered in the patrolling against illegal fishing context, is comparable to the size of these randomly generated instances.

Table 1. Average relative difference of upper bound V_U and lower bound V_L ($\epsilon = 0.05$).

# Cells	# Routes	Error	Running time
10	5	0.8% (± 1.0%)	0.217 s
20	15	1.9% (± 1.8%)	0.347 s
30	25	2.2% (± 1.4%)	0.819 s

4.2 Illustrative Examples

This section presents some examples to illustrate the models described in this paper. The results in this section are obtained by implementation of (9) and (11). Consider an area with 12 cells and 9 routes. The routes are chosen such that the cells are evenly spread over all routes, see Table 2. Suppose $N_M = 2$ and the payoff matrices are constructed using (1), where $d_j = 0.9$, $j = 1, ..., N_C$ and $g^{(k)}$ is the intruder's gain. Figure 1 depicts payoff matrices $M^{(1)}$, $M^{(2)}$ and two example routes, Routes 1 and 8. The white cells have a gain of 1, the light gray cells have a gain of 2 and the dark gray cell have a gain of 3.

Constant Payoff Matrix. Consider the games with payoff matrices $M^{(1)}$ and $M^{(2)}$ separately. Suppose that the planning period for both payoff matrices is $N_D = 100$. Table 3 shows the game values for different conditions. For example, a condition of 0.1 means that the minimum number of visits equals 10. The second and the third column give the game value of both games for the conditions specified in the first column. The first row shows the value of the game without conditions on the number of visits to the cells, the second row considers the game in which all nodes must be visited at least 10 times, and the third row

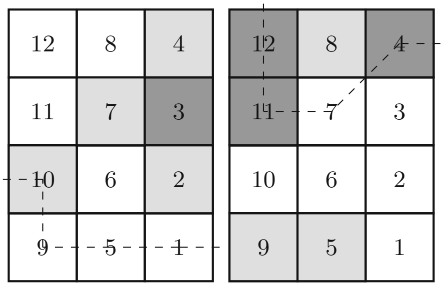

Fig. 1. Payoff matrices $M^{(1)}$, Route 1 (left) and $M^{(2)}$, Route 8 (right).

Table 2. Possible routes.

Routes	Cells visited by route
1	1, 5, 9, 10
2	2, 3, 8, 12
3	3, 7, 6, 10
4	4, 7, 6, 9
5	1, 2, 3, 4
6	3, 4, 7, 12
7	2, 5, 6, 9
8	4, 7, 11, 12
9	2, 5, 10, 11

considers the game in which Nodes 1-4 must be visited at least 30 times and the other nodes at least 10 times.

Table 3 indicates that the game value increases if more conditions are imposed on the agent's strategy. However, the increase of the game value depends on the payoff matrix. For example, the extra condition on Nodes 1–4 does not increase the game value for payoff matrix $M^{(1)}$, since the intruder's gain for these nodes is high and the agent is already patrolling these cells more often, as the results below indicate.

Table 3. Expected payoff per day for different conditions ($\epsilon = 0.05$).

Conditions (fraction)	Payoff $M^{(1)}$	Payoff $M^{(2)}$	Average	Combined
None	1.10	1.58	1.34	1.34
All nodes: 0.1	1.23	1.64	1.44	1.34
Nodes 1–4: 0.3, Nodes 5–12: 0.1	1.23	2.14	1.69	1.35

Figure 2 displays the agent's strategy for the different payoff matrices without conditions. The color of each cell is determined by the gain of the intruder and the number within each cell shows the fraction of the time period that the cell

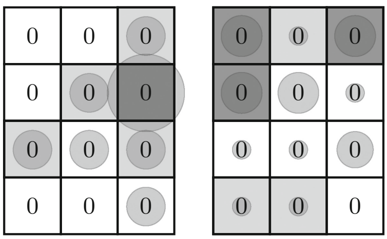

Fig. 2. Agent's strategy for the game without conditions.

should be visited. The agent's strategy is shown by the circles in each cell. The probability that a cell is visited is proportional to the radius of the circle in that specific cell. For example in Fig. 2, the probability that cell 3 is visited equals 1 for $M^{(1)}$ and 0.24 for $M^{(2)}$. Figure 3 displays the agent's strategy when conditions as given in Table 3 are considered. For all cases, it is clear that cells with a high gain for the intruder are visited more often.

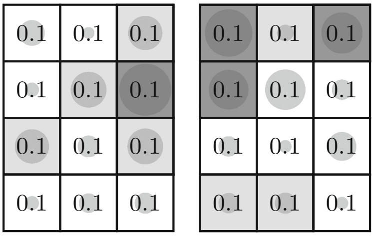

(a) All nodes: 0.1

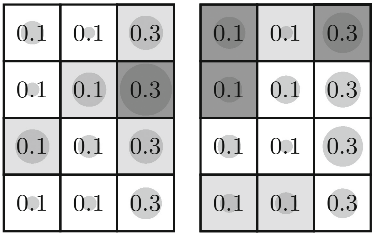

(b) Nodes 1-4: 0.3, Nodes 5-12: 0.1

Fig. 3. Agent's strategy for different conditions.

Multiple Payoff Matrices. The previous example considers the game with a constant payoff matrix such that for each game the conditions on the minimum number of visits have to be met. Now, we consider multiple payoff matrices simultaneously. Suppose that the total planning period $N_D = 200$ and both payoff matrices $M^{(1)}$ and $M^{(2)}$ have equal probability, so $\mu^{(1)} = \mu^{(2)} = 0.5$. Again, routes and conditions are given in Tables 2 and 3. A condition of 0.1 means that the total number of visits is 20, but it is, for example, allowed that there are only 5 visits when the payoff matrix is $M^{(1)}$ and 15 when the payoff matrix is $M^{(2)}$. This is the benefit of playing the game repeatedly and considering multiple payoff matrices simultaneously. In the last column of Table 3 the value of the game in which the conditions are combined for multiple payoff matrices is shown. If there are no conditions on the number of visits to the cells, the game value is just the average of both games with constant payoff, which is shown

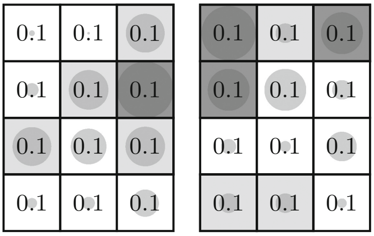

(a) All nodes: 0.1

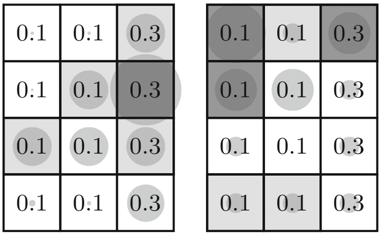

(b) Nodes 1-4: 0.3, Nodes 5-12: 0.1

Fig. 4. Agent's strategy if multiple payoff matrices are considered simultaneously.

in the second last column of Table 3. However, when conditions are considered, the value of the combined game is lower than the average of both games with constant payoff, because the agent has more flexibility in meeting the conditions.

Figure 4 shows the agent's strategy for the combined game with conditions given in Table 3. Comparing the results with those in Fig. 3 reveals that the agent has more flexibility in meeting the constraints when multiple payoff matrices are considered. Indeed the agent visits a cell less often when the gain is low and compensates this lack of visits when the gain of that cell is high.

5 Concluding Remarks

Patrolling a region with conditions on the frequency of visits to specific parts of that area while taking into account the optimal payoff of the intruder or agent can be modeled as a zero-sum security game with probabilistic constraints on the agent's strategy. These constraints prohibit exact solutions for large (realistic) instances. Therefore, we have developed a model yielding an upper bound and a lower bound for the game value. Computational results reveal that the relative difference between the upper and lower bound for the instances considered is less than 2.5% and that instances of realistic size can be solved within seconds.

In practice, the agent's strategy is constrained by existing guidelines. Numerical examples show that as the number of conditions increases, the agent's loss will increase. However, if multiple payoff matrices are considered, the agent has more flexibility in meeting the conditions and the loss of the agent is reduced.

In this paper, we have assumed that only one intruder is present in the area, that the payoff of intruders is known and that the agent decides on a strategy in advance. For future research, it would be interesting to investigate the case where not all payoff matrices are known in advance and multiple intruders attack simultaneously. Also, considering a more dynamic strategy of the agent, for example by taking into account extra information about the payoff and cells that already have been visited, should be pursued.

References

1. Alpern, S., Morton, A., Papadaki, K.: Patrolling games. Operat. Res. **59**(5), 1246–1257 (2011)
2. Brown, G., Carlyle, M., Salmeron, J., Wood, K.: Defending critical infrastructure. Interfaces **36**(6), 530–544 (2006)
3. Charnes, A.: Constrained games and linear programming. Proc. Nat. Acad. Sci. **39**(7), 639–641 (1953)
4. Fang, F., Stone, P., Tambe, M.: When security games go green: Designing defender strategies to prevent poaching and illegal fishing. In: IJCAI, pp. 2589–2595 (2015)
5. Gatti, N.: Game theoretical insights in strategic patrolling: model and algorithm in normal-form. In: ECAI, pp. 403–407 (2008)
6. Golany, B., Goldberg, N., Rothblum, U.G.: A two-resource allocation algorithm with an application to large-scale zero-sum defensive games. Comput. Oper. Res. **78**, 218–229 (2017)

7. Haskell, W., Kar, D., Fang, F., Tambe, M., Cheung, S., Denicola, E.: Robust protection of fisheries with compass. In: Twenty-Sixth IAAI Conference (2014)
8. Lin, K.Y., Atkinson, M.P., Chung, T.H., Glazebrook, K.D.: A graph patrol problem with random attack times. Oper. Res. **61**(3), 694–710 (2013)
9. MATLAB. version 9.1 (R2016b). The MathWorks Inc., Natick, Massachusetts (2016)
10. Meng, F., Zhan, J.: Two methods for solving constrained bi-matrix games. Open Cybern. Syst. J. **8**, 1038–1041 (2014)
11. Owen, G.: Game Theory, 3rd edn. Academic Press, London (1995)
12. Peters, H.: Game Theory: A Multi-leveled Approach, 1st edn. Springer, Heidelberg (2008)
13. Ross, S.: Stochastic Processes. John Wiley & Sons Inc, New York (1996)
14. Semple, J.: Constrained games for evaluating organizational performance. Eur. J. Oper. Res. **96**(1), 103–112 (1997)
15. Washburn, A., Lee, E.L.T.C.: Allocation of clearance assets in IED warfare. Nav. Res. Logistics (NRL) **58**(3), 180–187 (2011)
16. Winston, W.L.: Operation Research, Applications and Algorithms. Brooks/Cole, Belmont (2004)
17. Wood, K.R.: Bilevel network interdiction models: formulations and solutions. In: Wiley Encyclopedia of Operations Research and Management Science (2011)