Decision Support to Facilitate Cost-Optimal Response in Time- and Safety-Critical Situations *

Jonathan Ott †
Universität Karlsruhe (TH)

Jutta Hild ‡
Fraunhofer IITB

Alexander Bauer §
Fraunhofer IITB

Abstract

Coping with threats such as terrorism, crime, natural hazards or industrial accidents requires early detection and correct assessment of critical situations, followed by appropriate actions to avoid potential damage to person or property. We here focus on threat prevention in closed properties of traffic infrastructures (railway stations, harbours, airports) or logistics centres. To early detect and prevent damage, the property is continuously monitored by sensors and security staff. We propose to model the threat situation of the property as a continuous-time Markov decision process (CMDP). The CMDP-model is used as core of a decision support system (DSS), which assists the security staff in two aspects: it facilitates to identify the threat situation of the property and suggests the cost-optimal action by which the property will be kept in or returned into a safe state. The concept is also shown in a numerical example.

1 Introduction

As politically or criminally motivated acts of violence have increased in the last years, there is increasing demand for technical support in order to early detect arising threats. The scenario in the REX (Risk-controlled area EXploration) project is threat prevention for a closed property which is continuously monitored.

A common method to perform this surveillance task is to use sensor-collected data to continuously keep track of the whole property. Sensor technology is well developed and easily available, however, in addition intelligent methods to control the sensor data are needed. This is necessary, as an overall surveillance generates high data volumes, which eventually have to be exploited by a human decision maker. However, the security staff are not able to evaluate all sensor signals such as cameras, microphones, fire detectors, light barriers, etc. permanently and simultaneously. Therefore, it is necessary to focus on the evaluation of sensor signals which are relevant to clarify the current situation. An intelligent surveillance system should draw the attention of the security staff to the most threatened locations inside of the property. Thus, the decision maker is relieved from keeping track of unthreatened locations, which allows to focus completely on efficient actions to deal with the main threat.

Besides sensor selection, it is also necessary to decide on appropriate actions. As security staff perform their task under mental and time pressure, there is general agreement on the usefulness of supporting the decision making process by a software system. Simple decision support systems offer predefined recommended actions on predefined situations, allowing an operator to react quickly and appropriately, once a critical situation has been correctly assessed. To provide an optimal coordination between sensor evaluation for situation assessment and the corresponding actions, the here presented approach unifies those two different decision processes into a single concept for decision support. The cost-optimal action is determined using a probabilistic model on threat events, sensors, actions and corresponding costs, which represents the potential threat situation of the monitored property. The algorithm, based upon a continuous-time Markov decision process (CMDP), is cast into a decision support system (DSS), being developed to facilitate efficient risk assessment in large closed properties. Once the model of the property is defined, the system continuously visualizes the current threat situation as a risk map and gives recommendations on the cost-optimal action, either the employment of sensors or the initiation of countermeasures, for each subarea of the property in order to return and keep it in a safe state.

To present the surveillance task at a glance, Figure 1 shows its components. On the right hand side the property is sketched together with the surveillance means, on the left hand side the decision maker, and in between the interface consisting of tools which help the decision maker to monitor and act on the current situation: surveillance monitor, decision support system and phone. The surveillance

*The underlying projects to this article are funded by the Bundesministerium für Bildung und Forschung under promotional references 03BAPAC1 and 03GEPAC2. The authors are responsible for the content of this article.

†Universität Karlsruhe (TH), Institut für Stochastik,
e-mail: Jonathan.Ott@stoch.uni-karlsruhe.de

‡Fraunhofer-Institut für Informations- und Datenverarbeitung (IITB),
e-mail: Jutta.Hild@iitb.fraunhofer.de

§Fraunhofer-Institut für Informations- und Datenverarbeitung (IITB),
e-mail: Alexander.Bauer@iitb.fraunhofer.de

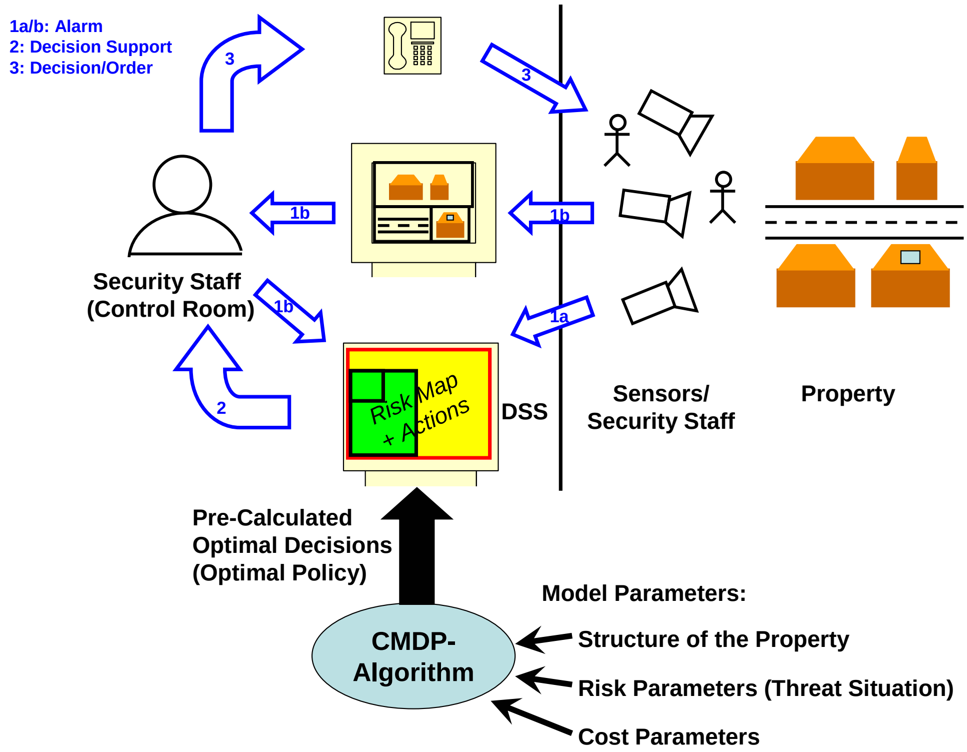

Figure 1: Components of the surveillance task.

circle works as follows: alarms are set off either automatically by sensors (1a) or by the security staff (1b: report from security guards on-scene or detection of the control-room security staff, who for example detect a suspicious object on a camera picture). Then the DSS processes these data to determine the status quo of the property's threat situation. The DSS displays this information as a risk map of the property, where the property is divided in meaningful subareas. Furthermore, the DSS identifies the corresponding cost-optimal action for all subareas on the basis of the pre-calculated optimal decisions (optimal policy). Thus, the core of the DSS is the CMDP-algorithm. As mentioned before, to calculate the optimal policy for a particular property, the CMDP has to model the property's threat situation threat situation. Therefore, the components of the CMDP need several input parameters. These define on the one hand the structural characteristics of the property including available surveillance means, on the other hand the threat potential which could occur inside of the property; additionally, the corresponding costs of both have to be considered. The DSS displays the cost-optimal action together with the risk map to the coordinating security staff in the control-room (2). Finally, the security staff give the order to perform the appropriate actions (3).

In the following sections, the core of the DSS, the CMDP-model of a threat situation for a closed property. In section 2, the model is developed generically, whereas section 3 applies the model assumptions to a numerical example. Section 4 summarizes our results in a brief conclusion and sketches future work.

2 Mathematical Model

In the first subsection, we describe a mathematical model, designed to deal with the tasks described in the introduction. In the following, some variants on the generic model are introduced, which refine the model for more realistic applications. Finally, some further comments on the model are made.

2.1 Model Description

The model we describe here is rather general, such that it can be applied to various surveillance tasks. The property needs to fulfill some requirements. It has to be divisible in certain sectors. Each sector has to be equipped with sensors which record the situation of the sector, such that the decision maker has access to the information of the past. Our model is based upon sets of threat events, from which we get the parameters for the CMDP.

As already mentioned, the mathematical model for our problem is a CMDP as it is described for example in

[Ber01] or [Put05]. Formally, a CMDP is a tuple of the form $(S, A, D, \lambda, p, K, R, \alpha)$, where the components are the following: S is a finite state space and A is a finite action space. $D := \{(s, a) \in S \times A : a \in D(s)\}$ is the restriction set, in which $D(s) \neq \emptyset$ is the set of all admissible actions in state $s \in S$. So, $(s, a) \in D$ means that action a is admissible in state s. The components p and λ give the stochastic dynamics as follows. Suppose, the decision maker chooses in state s action $a \in D(s)$. With this choice, the system remains an exponentially distributed random time in state s and then jumps to a state s' with transition probability $p^a_{ss'}$. $\lambda(s, a)$ is the parameter for the exponentially distributed sojourn time. The remaining components are used to determine the costs. $K(s, a)$ are the set-up costs, which have to be paid right after action $a \in D(s)$ is chosen. $R(s, a)$ is the cost rate, which has to be paid as long as the system remains in s and action a is carried out. $\alpha > 0$ is a discount factor.

We now present a parameterization of the CMDP, which has an adequate interpretation for our surveillance task and from which we easily get the components of the CMDP. We define the state space to be $S := G^\Sigma$. Thereby, let Σ be the finite set of all sectors of the property. Further, we assume that threat is measured as a threat level contained in the finite set $G := \{0, 1, \ldots, g_{\max}\}$, where $g_{\max} \in \mathbb{N}$ and $g_{\max} \geq 1$. Since Σ and G are finite, S is finite, too. By definition, a state $s \in S$ assigns a threat level to each sector of the property. Here, $s(\sigma)$ denotes the threat level of sector σ. In this manner, a state is a risk map, where the current risk of a sector is just the corresponding component of the state.

To model the state transitions of S close to a realistic surveillance scenario, we do not only model the set of sectors but also the threat dependencies between sectors, like they appear for example when sectors are close-by to each other or because of other structural circumstances. In this case, 'dependency' means that if threat increases in σ, then threat will increase in σ^*, too. We therefore introduce a dependency matrix N, where N is a $\Sigma \times \Sigma$-matrix with entry $N(\sigma, \sigma^*) = 1$ if σ^* is dependent on σ, and $N(\sigma, \sigma^*) = 0$ if σ^* is not dependent on σ. Note, that the dependencies given in N do not need to be symmetric, for example a sector which is essential for electric power supply can affect far-off sectors without being dependent vice versa. For definiteness, define $N(\sigma, \sigma) = 0$ ($\sigma \in \Sigma$). Thus, the objects Σ and N are the structural components of the model.

To give a meaningful interpretation of the threat levels in G we introduce sets of events symbolized by $\mathcal{E}(\sigma)$ that model external threat for all sectors $\sigma \in \Sigma$. They increase or decrease threat in a certain sector at the time they occur and thus influence the subsequent state. For each $\sigma \in \Sigma$ let $\mathcal{E}(\sigma)$ be a non-empty finite set of threat events in σ. For consistency let $\mathcal{E}(\sigma) \cap \mathcal{E}(\sigma^*) = \emptyset$, $\sigma \neq \sigma^*$, such that from $e \in \bigcup_{\sigma\in\Sigma} \mathcal{E}(\sigma) =: \mathcal{E}$ the related sector, in which e can occur, is known. We assume that all threat events are independent.

Let $g \in G$ be the current threat level in sector σ. Threat event $e \in \mathcal{E}(\sigma)$ occurs after an exponentially distributed time with rate $\lambda_e(g)$ which affects the components λ and p of the CMDP definition. Furthermore, its occurrence causes the following changes of the sectors' threat levels: threat level change in σ is given by a deterministic function $\psi_e : G \to G$ and changes to $\psi_e(g)$. For any dependent sector σ^* with $N(\sigma, \sigma^*) = 1$ and threat level g^*, the threat level changes to $\varphi_e(g^*)$, where $\varphi_e : G \to G$ is another deterministic function. Otherwise, i. e. $N(\sigma, \sigma^*) = 0$, the threat level remains unchanged. These functions influence the component p of the CMDP definition as it will be shown below. Additionally, when e occurs, the decision maker has to pay a cost $C_e \geq 0$.

To perform his task of returning the property into a safe state, a decision maker can choose between certain actions. Let $A_0 := \{0, 1, 2\}$ be the set of elementary actions the decision maker can assign to each of the sectors. These elementary actions have the following interpretations: elementary action 0 means 'Do nothing', elementary action 1 means 'Analyse sensor' and elementary action 2 means 'Send guard'. Thus, the action space is $A := A_0^\Sigma$. This means, that the decision maker has to assign one elementary action to each sector. Note, that the elementary action 'Do nothing' is an action only in a formal sense. It can be assigned to any sector, meaning that the decision maker does not have necessarily to take physical action in the sectors.

Let the restriction set be $D := S \times A$. This means that there is a supply of guards and analysers that covers all demands at any time. The elementary actions 1 and 2 are assigned with specific durations $1/\lambda_1(\sigma) > 0$ and $1/\lambda_2(\sigma) > 0$ respectively, depending on the sector $\sigma \in \Sigma$ where the elementary action is taken. To fit into the CMDP frame, these times have to be exponentially distributed random times with rates $\lambda_1(\sigma)$ and $\lambda_2(\sigma)$, since the expectation of an exponentially distributed random variable with rate $\nu > 0$ is just $1/\nu$. It is assumed that elementary action 0 requires infinite time such that elementary 0 will never be accomplished. According to $\lambda_e, e \in \mathcal{E}$, the parameters $\lambda_1(\sigma), \lambda_2(\sigma), \sigma \in \Sigma$, are used to determine λ and p of the CMDP.

To entirely determine the transition probabilities p we introduce further parameters. For elementary action 1, i. e. 'Analyse sensor', $\sigma \in \Sigma$ and $g \in G$ let $\psi_1^{\sigma g}$ be a probability mass function on G. So, for $g, g' \in G$ the number $\psi_1^{\sigma g}(g')$ is the probability, that after analysis of the sensor data of sector σ with threat level g the analyser asserts that the sector's threat level has to be g'. In the following, the $\psi_1^{\sigma g}(g')$ are called analysis probabilities. After completing elementary action 1 in sector σ and identifying $s'(\sigma)$ as subsequent threat level, the subsequent threat levels of the other sectors $\sigma^* \neq \sigma$ are given by

$$s'(\sigma^*) = \min\left\{g_{\max}, \max\left\{0,\ s(\sigma^*) + N(\sigma, \sigma^*)\left\lfloor \frac{s'(\sigma) - s(\sigma)}{2} \right\rfloor\right\}\right\}, \tag{1}$$

where for $x \in \mathbb{R}$ $[x]$ denotes the greatest integer not greater than x. In equation (1), min and max composition assures, that the new threat level of σ^* lies in G. Equation (1) takes into account, that after investigation of a sector there is some information about its dependent sectors, too. Their threat levels increase or decrease approximately half of the threat level change of σ itself. Since $N(\sigma,\sigma^*) = 0$ for any sector σ^* which is not dependent on σ, the threat level does not change in σ^*.

When elementary action 2 is completed in a certain sector, the threat level of this sector will be set to 0, since the guards are assumed to work perfectly. This assumption is formally expressed as

$$\psi_2(g) = 0,\ g \in G.$$

Similarly to elementary action 1, the guard has some influence on the dependent sectors. This is expressed by a function $\varphi_2 : G \to G$ which gives the subsequent threat levels of the depending sectors after elementary action 2 has ended. If the current threat level of a dependent sector is g, the subsequent threat level will be $\varphi_2(g)$, as soon as the guard finishes his job.

The cost rates of elementary actions 1 and 2 are $c_1, c_2 \geq 0$ respectively. So, for example if sensor analysis takes duration T, the operator will have to pay $c_1\,T$. In addition to the cost rates for the elementary actions, there is another parameter set, that influences the action costs, such that they depend on the current state, too. So far, the costs of an action just depend on the action itself. Since we need some more notation for the remaining cost parameters, we describe the dynamics of the model at first.

For an arbitrary state $s \in S$ and an arbitrary action $a \in A$, the mechanism of the model works as follows. Suppose, state $s = (s(\sigma))_{\sigma\in\Sigma} \in S$ is given and the decision maker chooses action $a = (a(\sigma))_{\sigma\in\Sigma} \in A$. For each $\sigma \in \Sigma$, define the independent random variables $X_e^{(s,a)} \sim \mathrm{Exp}(\lambda_e(s(\sigma)))$, $e \in \mathcal{E}(\sigma)$, and $X_\sigma^{(s,a)} \sim \mathrm{Exp}(\lambda_{a(\sigma)})$ with $\lambda_0 := 0$ and $\mathrm{Exp}(0) := \delta_\infty$, with δ_∞ the Dirac measure in ∞. These random variables give the duration, until a threat event occurs or until an elementary action is accomplished respectively. Then, the time until the first of these events occurs, which is given by $X_{s,a} := \min\left\{\min_{e\in\mathcal{E}}\left\{X_e^{(s,a)}\right\}, \min_{\sigma\in\Sigma}\left\{X_\sigma^{(s,a)}\right\}\right\}$, is exponentially distributed with rate $\sum_{\sigma\in\Sigma}\left(\sum_{e\in\mathcal{E}(\sigma)}\lambda_e(s(\sigma)) + \lambda_{a(\sigma)}(\sigma)\right) =: \lambda(s,a)$. After time $X_{s,a}$, the system jumps to a new state according to the causing event, i. e. the one which attains the minimum. The decision maker now has to choose a new action. The probability that a threat event $e \in \mathcal{E}(\sigma)$ occurs before anything else happens is $\lambda_e(s(\sigma))/\lambda(s,a)$. The probability that the elementary action $a(\sigma)$ in σ is completed before anything else happens is $\lambda_{a(\sigma)}(\sigma)/\lambda(s,a)$.

In a certain threat situation $s \in S$ and a chosen action $a \in A$, the mentioned mechanism can be seen as a race between the threat events and the elementary actions. In an informal way, the actions have to be chosen appropriately, such that the property is very likely switched over to a safe state.

Considering the CMDP components once more, we established the rates $\lambda(s,a)$ as the rates of the random variables $X_{s,a}$. To establish the transition probabilities, we have to take into account that the only way for a state change is the occurrence of a threat event or the finishing of an elementary action. A subsequent state is reached with the transition probabilities $p_{ss'}^a$, $(s,a) \in D$, $s' \in S$:

$$\begin{aligned} p_{ss'}^a = \frac{1}{\lambda(s,a)} \sum_{\sigma\in\Sigma} \Bigg[\sum_{e\in\mathcal{E}(\sigma)} \lambda_e(\sigma) \cdot \mathbb{1}\Big(& s'(\sigma) = \psi_e(s(\sigma)) \\ & \text{and } s'(\sigma^*) = N(\sigma,\sigma^*)\,\varphi_e(s(\sigma^*)) \\ & +(1 - N(\sigma,\sigma^*))\,s(\sigma^*),\ \sigma^* \neq \sigma \Big) \\ + \lambda_1(\sigma)\,\psi_1^{\sigma s(\sigma)}(s'(\sigma)) \cdot \mathbb{1}\Big(& a(\sigma) = 1 \text{ and } s'(\sigma^*) = \\ & \min\Big\{g_{\max}, \max\Big\{0, \\ & s(\sigma^*) + N(\sigma,\sigma^*)\left[\frac{s'(\sigma) - s(\sigma)}{2}\right]\Big\}\Big\},\ \sigma^* \neq \sigma \Big) \\ + \lambda_2(\sigma) \cdot \mathbb{1}\Big(& a(\sigma) = 2,\ s'(\sigma) = \psi_2(s(\sigma)) \text{ and } s'(\sigma^*) \\ & = N(\sigma,\sigma^*)\,\varphi_2(s(\sigma^*)) \\ & +(1 - N(\sigma,\sigma^*))\,s(\sigma^*),\ \sigma^* \neq \sigma \Big) \Bigg]. \end{aligned} \tag{2}$$

The function $\mathbb{1}(\mathcal{S})$ equals 1 if the statement $\mathcal{S}$ is true, otherwise, it equals 0. From equation (2) can be seen how the model parameters are used. Thus, $p_{ss'}^a$ is the sum of the probabilities of threat events, which lead from s to s', and of the probabilities of completing the elementary actions out of a, which lead to s'. The indicators $\mathbb{1}$ in (2) check, if s' is reachable from s with the appropriate threat events or elementary actions.

Up to this point, the cost components of the CMDP have still to be declared. With the previous specifications, now the state-dependent part of the action costs can be described. For this purpose, let $\gamma(\sigma,g) \in [0,1]$, $\sigma \in \Sigma, g \in G$, be the probability of finding a dangerous object in sector σ when taking elementary actions 1 or 2, while σ has threat level g. Further, let C_σ be the expected costs for removing such an object from sector σ. So, the costs for completing elementary action $a(\sigma) \in \{1,2\}$ in sector σ are given by $\gamma(\sigma,g)\,C_\sigma\,\lambda_{a(\sigma)}(\sigma)/\lambda(s,a)$, because $\lambda_{a(\sigma)}(\sigma)/\lambda(s,a)$ is just the probability, that elementary action $a(\sigma)$ is responsible for the state change.

Now, we are able to compute the expected discounted costs from 0 until $X_{s,a}$ with a chosen discount factor $\alpha > 0$. These so called one-step costs are given by

$$r(s,a) = E\left[\sum_{\sigma\in\Sigma}\int_0^{X_{s,a}} (\delta_{1a(\sigma)}\,c_1 + \delta_{2a(\sigma)}\,c_2)\,e^{-\alpha t}\,dt\right]$$

$$+ E\left[\sum_{\sigma\in\Sigma} e^{-\alpha X_{s,a}} \eta_{0a(\sigma)} \frac{\lambda_{a(\sigma)}(\sigma)}{\lambda(s,a)} \gamma(\sigma, s(\sigma))\, C_\sigma\right]$$

$$+ E\left[\sum_{\sigma\in\Sigma} \sum_{e\in\mathcal{E}(\sigma)} e^{-\alpha X_{s,a}} \frac{\lambda_e(s(\sigma))}{\lambda(s,a)} C_e\right]$$

$$= \frac{1}{\lambda(s,a)+\alpha} \sum_{\sigma\in\Sigma} \Bigg[\delta_{1a(\sigma)}\, c_1 + \delta_{2a(\sigma)}\, c_2$$

$$+ \eta_{0a(\sigma)}\, \lambda_{a(\sigma)}(\sigma)\, \gamma(\sigma, s(\sigma))\, C_\sigma + \sum_{e\in\mathcal{E}(\sigma)} \lambda_e(s(\sigma))\, C_e \Bigg]. \tag{3}$$

In formula (3), δ denotes the Kronecker delta, which is defined as $\delta_{xy} = 1$, if $x = y$, and $\delta_{xy} = 0$, otherwise, $x, y \in M$, where M is an arbitrary set, and $\eta_{xy} := 1 - \delta_{xy}$.

Each time a threat event has occurred or an elementary action is finished, all elementary actions, which are not yet completed, are stopped, and the decision maker has to choose a new action. Because of the memory-lack property of the exponential distribution setting up a just stopped elementary action can be interpreted as to keep it going on. So, set-up costs for elementary actions make no sense, since the decision maker would have to pay it each time he would start this action, even if he just continued this elementary action. Thus, define $K(s,a) := 0$ ($(s,a) \in D$).

Since there are no set-up costs, the cost rates are given by $R(s,a) := (\lambda(s,a) + \alpha)\, r(s,a)$ for $(s,a) \in D$, which is the sum in equation (3). The cost rates are the costs given by the cost rates of the elementary actions, the costs of the discovery of dangerous objects, and by the rates of the occurrence of a threat event or a completed elementary action together with their corresponding costs.

Now having specified all components of a CMDP, we can use the known theory of CMDP. We can look for a policy, such that the expected discounted costs are minimal. There are several methods to solve this problem, like value iteration and linear programming. It is known that for the expected discounted cost criterion, there exists a decision rule $f^* : S \to A$, which gives the optimal action for a given state, such that the expected discounted costs using f^* at each decision point are optimal, see [Ber01], pp. 258–261, or [Put05], Theorem 11.3.2. Thus, an action plan can be computed from the model which contains the optimal actions for all states.

2.2 Variants of the Model

Since in many surveillance tasks the number of the personnel is fixed, there are constraints to the actions. Let $n_P \leq |\Sigma|$ be the number of the personnel which is trained in sensor analysis and inspection. Here, $|\Sigma|$ denotes the number of sectors of the property. So, the restriction set is chosen to be $D = \{(s,a) \in S \times A : \sum_{\sigma\in\Sigma} (\delta_{1a(\sigma)} + \delta_{2a(\sigma)}) \leq n_P\}$. In this case, only the salaried personnel is available to survey the property. The costs c_1 and c_2 should be set to 0, since the costs for the actions are fixed by the salary per time unit, which shall be denoted by $c_P \geq 0$. When comparing the optimal expected discounted costs of the original model and the variant, the known discounted costs of the personnel $n_P\, c_P/\alpha$ must be added to the costs of the variant model.

The variant could be used, for instance, to optimize the number of the personnel. This can be done by solving problems corresponding to $1, 2, \ldots, |\Sigma|$ employees. Then, choose the number of employees resulting in minimal costs.

Alternative variants could be thought of, for example, restrictions on the number of analysis actions due to a lack of surveillance monitors. In this case, the CMDP components are determined as described in subsection 2.1, despite restriction set D which has to be declared appropriately.

2.3 Remarks on the Model

Some problems arise from the model which make the computation of an optimal policy a hard task. This is due to the so called 'curse of dimensionality'. As the property is getting bigger, the number of states and actions in our model is getting bigger in an exponential relationship. More precisely, it holds $|S| = |G|^{|\Sigma|}$ and $|A| = 3^{|\Sigma|}$. For example, if $g_{\max} = 4$ and $|\Sigma| = 4$, the number of states of the model is $5^4 = 625$ and the number of actions is $3^4 = 81$. If we increase $|\Sigma|$ to more realistic 10, then the number of states will be $5^{10} \approx 10^7$ and the number of actions $3^{10} \approx 6 \cdot 10^4$. With growing number of sectors, there are three problems. First, the number of model parameters, which have to be determined, increases. Second, the number of parameters of the CMDP, which have to be computed, increases. Third, the solving methods, such as linear programming, need more memory and more time.

It is crucial that in a CMDP there is an optimal policy, which is a decision rule, such that the optimal action depends on the current state only. This is because all components of the CMDP just depend on the current state and the current action. In our model, optimal decisions are made only because of the current threat levels in the sectors. No other information influences the decisions. Even if the decision maker knows more about the specific threat than what is modelled, the precise information will be irrelevant for the decision making process of the system. With some modification of a given model this problem can be circumvented which will be presented in the next paragraph.

Another point is, how to interpret a threat level. In the model, a threat level is defined by the rates of certain threat events in a certain sector. In addition, a threat level also contains information about the existence of a dangerous object. So, the model is working on an aggregation of different threat types and it does not include the exact situation. One possibility to deal with the aggregation of several threat types is to define a new sector for each threat type. These sectors are threatened by one specific threat only. In this way, physical sectors of the property are split up into several sectors, which represent specific threats. Splitting up sectors, again, highly increases the dimensionality of the model. Of course, it is possible to define threat event sets and the rates of the threat events for each sector in another way. In so doing, each sector has its own interpretation of the threat levels. Further, the model includes the uncertainty of the decision maker which is modelled by the analysis probabilities $\psi_1^{\sigma g}$.

The discount factor α determines how strong the future affects the optimal decision. For $\alpha \gg 0$, the future does not play an important role and the optimal action will be chosen very myopic. Whereas in the case of for example $\alpha = 0.001$, the future states of the system become very important, and the optimal action is chosen with a long-run view on the system.

3 Numerical Example

In the first subsection, a numerical example is introduced, which models the threat situation of a simplified airport. The example is set up along the lines of subsection 2.1. In the second subsection, there will be a discussion of the resulting expected discounted cost-optimal policy.

3.1 Definition of the Numerical Example

Since we do not have real data, we estimated the data somewhat arbitrary. Assume, our task is to survey an airport, which has the four sectors 'Technical Sector' (TS), 'Terminal' (T), 'Outskirt Area' (OA) and 'Fence' (F), such that $\Sigma = \{\mathrm{TS}, \mathrm{T}, \mathrm{OA}, \mathrm{F}\}$. Their dependencies are given by

$$N = \begin{array}{c} \\ \\ \end{array} \begin{array}{cccc} \mathrm{TS} & \mathrm{T} & \mathrm{OA} & \mathrm{F} \end{array} \\ N = \begin{pmatrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \begin{array}{c} \mathrm{TS} \\ \mathrm{T} \\ \mathrm{OA} \\ \mathrm{F} \end{array},$$

as shown in Figure 2. We define five threat levels by $G = \{0, 1, \ldots, 4\}$. With this definition, it is possible to distinguish between several threat levels leading from 'Everything is OK' to 'Urgent threat'. The state space is $S = \{(g_{\mathrm{TS}}, g_{\mathrm{T}}, g_{\mathrm{OA}}, g_{\mathrm{F}}) : g_\sigma \in G\ (\sigma \in \Sigma)\}$.

We assume that for each sector, there is exactly one camera, which has full insight to the sector and which records the situation of the sector continuously. The recordings are stored long enough to always allow the analyser to determine the current threat level. Additionally, we assume that in each sector, there is exactly one sensor, which sets off a signal automatically when detecting a threat situation. The types of the sensors are not specified further.

In this example, there are three threat events for each sector. The first one is the destruction of a sector, denoted by e_1^σ, $\sigma \in \Sigma$. The second one is, that an alarm is set off, denoted by e_2^σ, $\sigma \in \Sigma$. The third one is, that after a certain random time nothing has happened at all, denoted by e_3^σ, $\sigma \in \Sigma$. The last threat event is a virtual event, which is given by a random generator, that is linked to the surveillance system. The virtual event represents the risk assessment of a sector when there is no new information about the sector during a period of time. So, the sets of threat events are

$$\mathcal{E}(\mathrm{TS}) = \left\{e_1^{\mathrm{TS}}, e_2^{\mathrm{TS}}, e_3^{\mathrm{TS}}\right\},$$
$$\mathcal{E}(\mathrm{T}) = \left\{e_1^{\mathrm{T}}, e_2^{\mathrm{T}}, e_3^{\mathrm{T}}\right\},$$
$$\mathcal{E}(\mathrm{OA}) = \left\{e_1^{\mathrm{OA}}, e_2^{\mathrm{OA}}, e_3^{\mathrm{OA}}\right\},$$
$$\mathcal{E}(\mathrm{F}) = \left\{e_1^{\mathrm{F}}, e_2^{\mathrm{F}}, e_3^{\mathrm{F}}\right\}.$$

The rates of the threat events are shown in Table 1. In the lowest threat level 0, destruction of a sector is assumed to happen in the mean once a year, and in the highest level 4, destruction is assumed to happen in the mean twice per hour. In every threat level, an alarm occurs in the mean twice as much as a destruction. So, an alarm occurs sig-

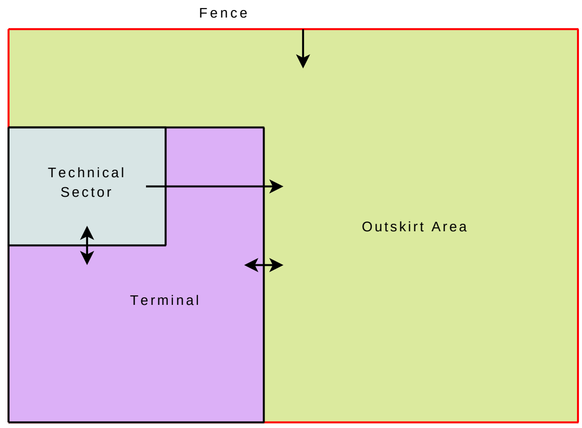

Figure 2: The airport with four sectors from the numerical example, with an arrow leading from one sector to another indicating that the latter is dependent from the first.

	threat level g				
	0	1	2	3	4
$\lambda_{e_1^{\mathrm{TB}}}(g)$	1/8760	1/168	1/24	1	2
$\lambda_{e_2^{\mathrm{TB}}}(g)$	1/4380	1/84	1/12	2	4
$\lambda_{e_3^{\mathrm{TB}}}(g)$	1	1	1	1	1
$\lambda_{e_1^{\mathrm{T}}}(g)$	1/8760	1/168	1/24	1	2
$\lambda_{e_2^{\mathrm{T}}}(g)$	1/4380	1/84	1/12	2	4
$\lambda_{e_3^{\mathrm{T}}}(g)$	1	1	1	1	1
$\lambda_{e_1^{\mathrm{A}}}(g)$	1/8760	1/168	1/24	1	2
$\lambda_{e_2^{\mathrm{A}}}(g)$	1/4380	1/84	1/12	2	4
$\lambda_{e_3^{\mathrm{A}}}(g)$	1	1	1	1	1
$\lambda_{e_1^{\mathrm{Z}}}(g)$	1/8760	1/168	1/24	1	2
$\lambda_{e_2^{\mathrm{Z}}}(g)$	1/4380	1/84	1/12	2	4
$\lambda_{e_3^{\mathrm{Z}}}(g)$	1	1	1	1	1

Table 1: The rates of the threat events for the numerical example.

nificantly more frequently than an urgent threat causing a destruction, since there are probably false alarms. The rate of the virtual event is 1 per hour for all threat levels.

When a sector is destructed, it is assumed to be rebuilt in no time at the costs of this sector. After that, the sector is left with threat level 0. When an alarm is set off, the threat level increases by 2. When nothing has happened, in current threat levels 0 and 1 the threat level is set to 1. This stands for a relatively conservative risk assessment. If the threat level is 2, 3 or 4, it will decrease by 1. This definition represents the assumption that a threat, for example a terroristic attack, would occur shortly after an alarm. So, repeated occurrence of the virtual threat event in all sectors will put the system to the threshold state $(1, 1, 1, 1)$. Formally, this means for $g \in G$

$$\psi_{e_1^{\text{TS}}}(g) = \psi_{e_1^{\text{T}}}(g) = \psi_{e_1^{\text{OA}}}(g) = \psi_{e_1^{\text{F}}}(g) = 0,$$
$$\psi_{e_2^{\text{TS}}}(g) = \psi_{e_2^{\text{T}}}(g) = \psi_{e_2^{\text{OA}}}(g) = \psi_{e_2^{\text{F}}}(g) = \min\{g+2, 4\},$$
$$\psi_{e_3^{\text{TS}}}(g) = \psi_{e_3^{\text{T}}}(g) = \psi_{e_3^{\text{OA}}}(g) = \psi_{e_3^{\text{F}}}(g) = \begin{cases} 1, & \text{if } g = 0, 1 \\ g-1, & \text{else} \end{cases}.$$

For the dependent sectors, the transition functions of the threat events are defined as follows. If a sector is destroyed, all dependent sectors will be left with the highest threat level 4. When an alarm is set off, the threat level of dependent sectors increases by 1. So, an alarm in a sector will cause a 'little' alarm in the dependent sectors. If there is a threat level change because of the virtual event, there will be no change of the threat level in dependent sectors. Formally, this means for $g \in G$

$$\varphi_{e_1^{\text{TS}}}(g) = \varphi_{e_1^{\text{T}}}(g) = \varphi_{e_1^{\text{OA}}}(g) = \varphi_{e_1^{\text{F}}}(g) = 4,$$
$$\varphi_{e_2^{\text{TS}}}(g) = \varphi_{e_2^{\text{T}}}(g) = \varphi_{e_2^{\text{OA}}}(g) = \varphi_{e_2^{\text{F}}}(g) = \min\{4, g+1\},$$
$$\varphi_{e_3^{\text{TS}}}(g) = \varphi_{e_3^{\text{T}}}(g) = \varphi_{e_3^{\text{OA}}}(g) = \varphi_{e_3^{\text{F}}}(g) = g.$$

The elementary actions A_0 remain as defined in the model description in subsection 2.1, with the exception that elementary action 1 now means 'Analyse camera recording'. The action space is given by $A = A_0^{\Sigma}$. The restriction set is $D = S \times A$, meaning that the decision maker can choose any elementary action for any sector. The action rates are chosen as $\lambda_1(\sigma) = 60$ $(\sigma \in \Sigma)$ and $\lambda_2(\sigma) = 1$ $(\sigma \in \Sigma)$, is interpreted as follows. Let the time unit be 1 hour. Then, analysis of one camera recording requires 1 minute in the mean and inspection of a sector requires 1 hour in the mean.

For elementary action 1, the analysis probabilities are assumed to be the same for all sectors, which means $\psi_1^{\sigma g} = \psi_1^g$ $(\sigma \in \Sigma, g \in G)$, see Table 2. It was taken care, that it is very likely to end up in neighbouring threat levels. The table reads as follows: the probability, that in current threat level g the threat level will change to g', is given in the cell of the g-th row and the g'-th column. For instance, starting in threat level 4, the analyser will find the subsequent threat level to be 2 with probability 0.3. Recall, that if there is an

		g'				
		0	1	2	3	4
g	0	0.5	0.3	0.1	0.1	0.0
	1	0.4	0.1	0.3	0.1	0.1
	2	0.3	0.2	0.2	0.2	0.1
	3	0.1	0.3	0.2	0.1	0.3
	4	0.0	0.1	0.3	0.2	0.4

Table 2: The analysis probabilities ψ_1^g, $g \in G$, for the numerical example.

analysis done in a certain sector, the change of the threat levels of the other sectors is given by equation (1), when elementary action 1 is completed.

After elementary action 2 is accomplished in sector σ, the new threat level of σ is 0, which is a general assumption of the model. Further, we assume that the threat level of dependent sectors decreases by 1. So, the guard does not only clear the situation in σ, but his presence also influences dependent sectors by the function φ_2 given for $g \in G$ by

$$\varphi_2(g) = \max\{0, g-1\}.$$

To finish the description of the example, it remains to define the costs and associated parameters. In detail, the threat event costs are

$$C_{e_1^{\text{TS}}} = 1{,}000{,}000,\ C_{e_1^{\text{T}}} = 10{,}000{,}000,$$
$$C_{e_1^{\text{OA}}} = 5{,}000{,}000,\ C_{e_1^{\text{F}}} = 100{,}000,$$
$$C_{e_2^{\text{TS}}} = C_{e_2^{\text{T}}} = C_{e_2^{\text{OA}}} = C_{e_2^{\text{F}}} = C_{e_3^{\text{TS}}} = C_{e_3^{\text{T}}} = C_{e_3^{\text{OA}}} = C_{e_3^{\text{F}}} = 0.$$

The only threat events, which lead to costs are those of type 'destruction'. There are no costs neither for the alarm nor for the virtual alarm. The cost rates of the elementary actions are

$$c_1 = 100,\ c_2 = 100.$$

The costs of removing dangerous objects are

$$C_{\text{TS}} = 50{,}000,\ C_{\text{T}} = 100{,}000,\ C_{\text{OA}} = 50{,}000,\ C_{\text{F}} = 10{,}000,$$

while the probabilities of discovering dangerous objects are

$$\gamma(\text{TS}, 0) = \gamma(\text{T}, 0) = \gamma(\text{OA}, 0) = \gamma(\text{F}, 0) = 0.001,$$
$$\gamma(\text{TS}, 1) = \gamma(\text{T}, 1) = \gamma(\text{OA}, 1) = \gamma(\text{F}, 1) = 0.005,$$
$$\gamma(\text{TS}, 2) = \gamma(\text{T}, 2) = \gamma(\text{OA}, 2) = \gamma(\text{F}, 2) = 0.01,$$
$$\gamma(\text{TS}, 3) = \gamma(\text{T}, 3) = \gamma(\text{OA}, 3) = \gamma(\text{F}, 3) = 0.1,$$
$$\gamma(\text{TS}, 4) = \gamma(\text{T}, 4) = \gamma(\text{OA}, 4) = \gamma(\text{F}, 4) = 0.5.$$

The discount factor is set to

$$\alpha = 0.001.$$

Following the remarks of the model in subsection 2.3, the interpretation for example of threat level 2 is the following: with rate 1/24 per hour the sector will be destroyed, with rate 1/12 per hour the sensor sets off an alarm, and with rate 1 per hour the system's virtual event will occur, see Table 1; the probability of finding a dangerous object is 0.01. Additionally, a state gets some uncertainty in the sense, that analysing camera recordings would yield a different threat level according to the analysis probabilities given in Table 2.

3.2 Results and Discussion

The expected discounted cost-optimal policy for the numerical example is given in Figure 3. There are four subfigures, one for each sector, indicating its optimal elementary actions. Each dot stands for a state of the system. Assume, we want to know the optimal action for the state $s = (1, 3, 4, 0)$. First, we have to find the corresponding dot in every subfigure. Since $s(\mathrm{TS}) = 1$, we have to look at the second grouped column from the left. Since $s(\mathrm{T}) = 3$, the wanted state can be found in the fourth grouped row from below. Since $s(\mathrm{OA}) = 4$, it must be in the fifth row of that row group. And finally, since $s(\mathrm{F}) = 0$, the state is found in the first column of the first grouped row. So, state $(1, 3, 4, 0)$ matches the sixth dot from the left in the 20th row from below. In all subfigures, the corresponding dot is yellow.

The elementary actions are encoded by the colours of the dots and have the following meanings. A green dot corresponds to elementary action 1 and thus denotes, that it is optimal to do nothing in the specific sector. A yellow dot corresponds to elementary action 2 and denotes, that it is optimal to analyse the camera recordings of the specific sector. Finally, a red dot corresponds to elementary action 2 and denotes, that it is optimal to send a guard to the sector. Thus, for the state $(1, 3, 4, 0)$ the action $(1, 1, 1, 1)$ is optimal, which corresponds to the analysis of the camera recordings of all sectors.

Even in this little example, the optimal policy is not obvious and has little structure. For example, watch the optimal elementary actions for sector F, which look rather arbitrary, which they are not, since they minimize the costs of the model. Especially, there is no scheme like 'monotonicity' in the optimal policy. In state $(2, 3, 0, 0)$, it is optimal to analyse the camera recordings of F. In state $(2, 3, 0, 1)$, it is optimal to send the guards to F. Whereas in state $(2, 3, 0, 2)$, it is optimal to analyse the recordings of F, again. The optimal elementary actions of the other sectors do not change along the considered states. Although in the considered states, only the threat level of F increases and the elementary actions of the other sectors remain the same: at first, it is optimal to analyse, then send guards, and then analyse again, while one could have expected, that it could not be optimal to analyse the camera recordings again in state $(2, 3, 0, 2)$. This is due to the definitions of the example and the complex behaviour of the resulting CMDP. For sector F, it can be optimal to send a guard even in threat level 4. Probably, this is because F is rather cheap in contrast to the other sectors, and that there is a chance to finish first, which would decrease the threat levels of OA, since OA is dependent of F.

It is remarkable, that it will be optimal to analyse the camera recordings in the sectors TS, T and OA, if the threat level is 3 or 4. This may be caused by the rate of elementary action 1 which is of higher magnitude than the rates of the threat events. Therefore, it is very likely, that the sensor analysis will cause the next state change, which will lead to a lower subsequent threat level with high probability. By contrast, the rate of elementary action 2 has the same magnitude. So, it is much more likely, that the threat event 'destruction' occurs at high costs.

4 Conclusion and Outlook

The aim of our investigation is to contribute to threat prevention in large, closed properties by assisting the security staff in their surveillance task. A decision support system should help to identify the current threat situation and suggest how to best deal with the current threats. For that purpose, the decision support system requires a description of the property's threat situation. As a solution, a CMDP-model for surveillance applications for closed properties has been introduced. It is based upon a concept of threat events, which can occur randomly in the sectors of the property. The model defines a risk assessment for the property as well as admissible actions to deal with the threat. From the model, the optimal policy, the core of the decision support system, can be derived. Additionally, a numerical example has been given for a simplified airport. Since the expected discounted cost-optimal policy has no simple structure, it could hardly be predicted by the decision maker himself. Thus, it is beneficial to base the decision support system upon the introduced model.

The model described in this article suffers from the curse of dimensionality. With an increasing number of sectors, it will be more challenging to determine the optimal policy for realistic-sized properties exactly. So, heuristics or approximations must be developed to find some sub-optimal policies as well as estimations for the quality of such sub-optimal policies.

References

[Ber01] Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control*, volume II. Athena Scientific, Belmont, Massachusetts, 2nd edition, 2001.

[Put05] Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 2005.

Figure 3: The expected discounted cost-optimal policy for the numerical example. The squares mark the state $(1, 3, 4, 0)$, used in subsection 3.2 to explain how to read the representation of the optimal policy.