

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

6-2016

Protecting the NECTAR of the Ganga River through game-theoretic factory inspections

Benjamin FORD

Matthew BROWN

Amulya YADAV

Amandeep SINGH

Arunesh SINHA

Singapore Management University, aruneshs@smu.edu.sg

See next page for additional authors

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#)

Citation

FORD, Benjamin; BROWN, Matthew; YADAV, Amulya; SINGH, Amandeep; SINHA, Arunesh; SRIVASTAVA, Biplav; KIEKINTVELD, Christopher; and MILLIND, TAMBE. Protecting the NECTAR of the Ganga River through game-theoretic factory inspections. (2016). *Proceedings of International Conference on Practical Applications of Agents and Multi-Agent Systems (PAAMS)*. 97-108.

Available at: https://ink.library.smu.edu.sg/sis_research/4662

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

Author

Benjamin FORD, Matthew BROWN, Amulya YADAV, Amandeep SINGH, Arunesh SINHA, Biplav SRIVASTAVA, Christopher KIEKINTVELD, and TAMBE MILLIND

Protecting the NECTAR of the Ganga River through Game-Theoretic Factory Inspections

Benjamin Ford¹, Matthew Brown¹, Amulya Yadav¹, Amandeep Singh², Arunesh Sinha¹, Biplav Srivastava³, Christopher Kiekintveld⁴, and Milind Tambe¹

¹ University of Southern California, Los Angeles, CA, US

{benjamif,mattheab,amulyaya,aruneshs,tambe}@usc.edu

² Columbia University, New York, NY, US as4330@columbia.edu

³ IBM Research, New Delhi, Delhi, India sbiplav@in.ibm.com

⁴ University of Texas at El Paso, El Paso, TX, US cdkiekintveld@utep.edu

Abstract. Leather is an integral part of the world economy and a substantial income source for developing countries. Despite government regulations on leather tannery waste emissions, inspection agencies lack adequate enforcement resources, and tanneries' toxic wastewaters wreak havoc on surrounding ecosystems and communities. Previous works in this domain stop short of generating executable solutions for inspection agencies. We introduce NECTAR - the first security game application to generate environmental compliance inspection schedules. NECTAR's game model addresses many important real-world constraints: a lack of defender resources is alleviated via a secondary inspection type; imperfect inspections are modeled via a heterogeneous failure rate; and uncertainty, in traveling through a road network and in conducting inspections, is addressed via a Markov Decision Process. To evaluate our model, we conduct a series of simulations and analyze their policy implications.

Keywords: Game Theory, Inspection, Security Games, Human-robot/agent interaction, Multiagent Systems

1 Introduction

The leather industry is a multi-billion dollar industry [14], and in many developing countries such as India and Bangladesh, the tanning industry is a large source of revenue. Unfortunately, the chemical byproducts of the tanning process are highly toxic, and the wastewater produced by tanneries is sent to nearby rivers and waterways. As a result, the Ganga River (along with many others) has become extremely contaminated, leading to substantial health problems for the large populations that rely on its water for basic needs (e.g., drinking, bathing, crops, livestock) [11]. Tanneries are required by law to run wastewater through sewage treatment plants (STPs) prior to discharge into the Ganga. In many cases, however, the tanneries either do not own or run this equipment, and it is

up to regulatory bodies to enforce compliance. However, inspection agencies have a severe lack of resources; the combination of tanneries’ unchecked pollution and inspection agencies’ failure to conduct inspections forced India’s national environment monitoring agency to ban the operation of 98 tanneries near Kanpur, India while threatening the closure of approximately 600 tanneries [13]. It is our goal to provide agencies with randomized inspection plans so tanneries reduce harmful effluents and an important facet of India’s economy can operate.

In this paper, we introduce a new game-theoretic application, NECTAR (Nirikshana for **E**nforcing **C**ompliance for **T**oxic wastewater **A**batement and **R**eduction)⁵, that incorporates new models and algorithms to support India’s inspection agencies by intelligently randomizing inspection schedules. We build on previous deployed solutions based on Stackelberg Security Games (SSG) for counter-terrorism [17] and traffic enforcement [6]. Our SSG models are also the first to focus on the problem of pollution prevention by modeling the interaction between an inspection agency (the leader) and leather tanneries (many followers) - an interaction which poses a unique set of challenges. (i) Because there is a large disparity between the number of inspection teams and the number of tanneries, inspection plans must be efficient. (ii) We cannot assume that inspectors can catch 100% of violations. (iii) Inspectors must travel to the tanneries via a road network so solutions must be robust to delays (e.g., traffic). Finally, current fine policies may not be sufficient to induce compliance, and (iv) it is important to investigate alternative fine structures.

NECTAR addresses these new challenges of tannery inspections. (i) Our SSG model captures the inspection process and accounts for two types of inspections: thorough inspections and simple (i.e., quick) inspections. While thorough inspections take longer to conduct (and thus less of them can be conducted), they are more likely to detect violations than simple, surface-level inspections which may only be able to check for obvious violations. To model the imperfect nature of these inspections, we (ii) introduce two failure rates: one for thorough inspections and one for simple inspections, with simple inspections failing at a higher rate. (iii) We also address the uncertainty involved with road networks by using a Markov Decision Process (MDP) that will represent and ultimately generate the game solution. Finally, (iv) we also investigate how tannery compliance is affected by two fine structures: fixed fines and variable fines, where the latter will result in larger tanneries receiving larger fines. For the evaluation of our model, we apply NECTAR to a real-world network of tanneries in Kanpur, India, we evaluate the quality of NECTAR’s generated solutions, and we demonstrate how NECTAR’s solutions can be visualized via a Google Earth overlay.

⁵ Nirikshana, the Hindi word for inspect. As many mythological stories and even popular Bollywood songs attest, Ganga water is supposed to be NECTAR (or Amrit, the Hindi antonym of poison) which has inspired our project. The project name is intentionally chosen to fit this international and inter-cultural theme.

2 Related Work

Several theoretical papers have used game theory to model the impact of environmental policies. Environmental games [18] use Stackelberg Games to model interactions between a regulator and a polluting firm, while [7] used game theory to study the effect of environmental policies in the Chinese electroplating industry. *Inspection games* consider the general problem of scheduling inspections, and have been extensively studied in the literature. For example, [8] models cases where an inspector must travel to multiple sites and determine violations as a stochastic game. A general theory of inspection games for problems such as arms control and environmental policy enforcement has been studied in [2], including analysis of whether inspectors can benefit from acting first. [16] also considered inspection games with sequential inspections, including compact recursive descriptions of these games. However, most of these works do not focus on concrete applications and thus, unlike our work, do not provide executable inspection schedules to inspectors.

Other areas of research have considered various models of patrolling strategies and scheduling constraints. These include patrolling games [1, 5, 3] and security games with varying forms of scheduling constraints on resources [19, 12, 6]. There has also been recent work on utilizing MDPs to represent strategies in security games [15, 4]. However, none of these efforts have focused on environmental inspections and have not investigated topics important in this domain, such as the impact of fine structures on adversary behavior (i.e., compliance).

3 Motivating Domain

The pollution of India’s rivers is a major concern. The waters of India’s largest river, the Ganga (or Ganges) River, are used by over 400 million people – roughly one-third of India’s population and more than the entire population of the United States. Unfortunately, the Ganga is ranked the fifth dirtiest river in the world. Generated from various sources such as sewage and industrial effluents, the pollution inflicts serious health conditions on all life that depends on the river. In Kanpur, villagers suffer from health conditions (e.g., cholera, miscarriages), and livestock yield less milk and die suddenly [9].

Situated around the city of Kanpur, the various leather tanneries are a major source of pollution in the Ganga River [9]. While there are a few sewage treatment plants (STPs) in Kanpur, they can neither treat the full volume nor the full range of produced pollutants [10]. In particular, treating heavy metals like chromium, mercury, arsenic, and nickel is costly and needs specialized personnel (in addition to the personnel required to operate the STPs). The government has put in regulations requiring the tanneries to own and operate effluent plants to remove the pollutants before they discharge their sewage. However, the tanneries have not been willing to undertake the additional cost of installing and operating the treatment units. Even when tanneries have installed the units, they avoid operating them whenever possible.

To address non-compliance issues, the government sends inspection teams to visit the tanneries. Inspecting the tanneries is a time-consuming, quasi-legal activity where the “as-is” situation is carefully recorded and samples are collected that can later be subjected to judicial scrutiny. It is also costly because, apart from the inspectors themselves, help from local police is requisitioned for safety, lab work is done for sample testing, and movement logistics are carefully planned; a full inspection is costly to conduct. Due to these costs, the number of inspectors that can be sent out on a patrol is very limited. Our application seeks to help by (1) generating randomized inspection patrols that maximize the effectiveness of available inspectors, and (2) introducing limited inspection teams which conduct simple inspections - a low-cost alternative to full inspection teams which conduct thorough inspections. While limited inspection teams cannot replace the needed capabilities of a full inspection team, they can still inspect tanneries and issue a fine for obvious violations (e.g., the site not owning an STP). We will refer to full inspection teams and limited inspection teams as thorough inspection resources and simple inspection resources, respectively.

4 Model

In this section, we model this pollution prevention problem as a defender-attacker Stackelberg Security Game (SSG). The task of the defender is to send resources to different tannery sites (i.e., the multiple adversaries) on a road network. The defender must devise a patrol strategy to maximize compliance among a number of sites (each site denoted by l), where each site has a number of factories f_l and each site’s compliance cost increases with the number of factories. In addition, the defender must take into account the time it takes to travel to and inspect each site. We model the road network as a graph where the nodes represent sites and the edges represent the roads connecting each site. Each edge also has a cost, e_{ab} , associated with it that represents the travel time from a site a to another site b . Using publicly available data regarding tannery locations in Kanpur, we constructed a graph consisting of 50 sites.

The defender has two types of resources: r_1 number of thorough inspection resources and r_2 simple inspection resources. For thorough inspection resources, the inspector conducts a detailed inspection that takes i time units. We model imperfect inspections such that even if a violation exists, the inspectors will fail to detect it with a low probability γ_1 . For simple inspection resources, the inspector will conduct a superficial inspection that takes d time units. Since the inspection is not detailed, simple inspection resources will not detect anything but obvious violations. Thus, such resources have a higher probability of failure γ_2 . Each of the defender’s resources (thorough and simple) have a maximum time budget, t_1 and t_2 respectively, to conduct inspections and travel to sites.

In the SSG framework, the defender will commit to a randomized patrol strategy (a mixed strategy) which is a probability distribution over the executable daily inspection patrols (the pure strategies for all resources). The adversaries (the sites) can fully observe the defender’s mixed strategy and know the proba-

bility of being inspected by a thorough inspection team or a simple inspection team on a given day. Formulating the mixed strategy requires enumerating all feasible pure strategies for the defender. However, this approach is impractical for two main reasons: (1) for any realistically-sized patrolling problem, the defender pure strategy space is so large that it cannot fit into memory. For example, with our Kanpur graph of 50 tanneries, only one defender resource, and a time horizon of 10 hours, the pure strategy space size would be too large to enumerate (approximately $50 \text{ choose } 10$). Therefore, we adopt a compact representation (a transition graph) that will allow our approach to scale to large problem sizes. (2) Inspectors must travel to sites via a road network (with potential delays), and the corresponding uncertainty cannot be handled by a standard SSG formulation. Rather than reasoning about mixed strategies, we instead use the compact representation to reason about spatio-temporal flow through a transition graph. To account for stochasticity and uncertainty in the outcome of actions, we use a Markov Decision Process (MDP) to represent the defender’s inspection patrolling problem. We can solve the corresponding linear program (LP) to compute the optimal inspection strategy, i.e., the optimal MDP policy.

4.1 Compact Game Representation: Transition Graph

Brown et al. also faced the challenge of large state spaces for a traffic enforcement domain [6]. Since their game also takes place on a road network, there are sufficient similarities between our approach and theirs to apply their techniques, based on transition graphs, to improve the scalability of our model.

Instead of enumerating an exponential number of pure strategies, we need only enumerate a polynomial number of states and edges in the transition graph. We then compute the optimal probability flow (as seen in the next section), also called a marginal coverage vector, and sample from the vector to create inspection schedules. As the defender resource types (thorough and simple) have different time constraints, each has its own transition graph.

We discretize time into a granularity of h hours. In the thorough inspection resource transition graph, a vertex is added for each site l every h hours until the resource time budget t_1 has been expended. Similarly for the simple resource’s transition graph, vertices are added until the time budget t_2 has been expended.

4.2 MDP Formulation

We present an MDP $\langle S, A, T, R \rangle$ to incorporate uncertainty into the transition graph. An example MDP is shown in Figure 1 to illustrate these definitions.

- S : Finite set of states. Each state $s \in S$ is a tuple (l, τ) , where l is the site that the resource is located, and τ is the current time step. For example, an inspector at site A at hour 1 is represented as $s_{A,1}$. Each vertex in the transition graph corresponds to a state s .




Fig. 1. Illustrative MDP Example

- A : Finite set of actions. $A(s)$ corresponds to the set of actions available from state s (i.e., the set of sites reachable from l) that the resource can travel to and inspect. For example, at site A at hour 1, the only available action is to move to site B (i.e., the solid arrow from A to B in Figure 1).
- $T_1(s, a, s')$: Probability of an inspector ending up in state s' after performing action a while in state s . Travel time and inspection time are both represented here. As a simple example, there could be probability 0.7 for transition $T_1(s_{A,1}, a_B, s_{B,2})$: a transition from site A at hour 1 to move to and inspect site B will, with a probability of 0.7, finish at hour 2 (a travel + inspection time of 1 hour). The dashed lines in Figure 1 represent the remaining probability (0.3) that the same action will instead finish at hour 3 (due to a delay). Note that the two resource types have separate transition functions due to the difference in action times (i for thorough inspection resources and d for simple inspection resources).
- $R(s, a, s')$: The reward function for ending in state s' after performing action a while in state s . As we are interested in the game-theoretic reward, we define the reward in the LP and define $R = 0 \forall s, a, s'$.

5 Inspection Patrol Generation

We provide a linear program (LP) to compute the optimal flow through the MDP (i.e., the transition graph with uncertainty). By normalizing the outgoing flow from each state in the MDP, we obtain the optimal MDP policy from which we can sample to generate dynamic patrol schedules. In the following LP formulation, we make use of the following notation. A site l has a number of factories f_l , and if a site is caught violating during an inspection, they receive a fine, α_l . On the other hand, if a site wants to remain in compliance, they will need to pay a compliance cost β for each factory (total cost = βf_l). We represent the expected cost for each site l as v_l . As defined in the following LP, the expected cost corresponds to the lowest of either the site's expected fine or the site's full cost of compliance; we assume that these adversaries are rational and that they will choose to pay the lowest of those two values (expected fine or cost of compliance). Finally, we denote as S_l the set of all states that correspond to site l (i.e., all time steps associated with site l).

As discussed in the transition graph definition, the optimal flow through the graph corresponds to the optimal defender strategy, and that flow is represented

by a marginal coverage vector. We denote the marginal probability of a resource type i (either thorough or simple inspection team) reaching state s and executing action a as $w_i(s, a)$. We also denote, as $x_i(s, a, s')$, the marginal probability of a resource type i reaching state s , executing action a , and ending in state s' .

$$\max_{w, x} \sum_l v_l \quad (1)$$

$$s.t. x_i(s, a, s') = w_i(s, a)T_i(s, a, s'), \forall s, a, s', i \quad (2)$$

$$\sum_{s', a', i} x_i(s', a', s) = \sum_{a, i} w_i(s, a), \forall s, i \quad (3)$$

$$\sum_{a, i} w_i(s_i^+, a) = r_i \quad (4)$$

$$\sum_{s, a, i} x_i(s, a, s_i^-) = r_i \quad (5)$$

$$w_i(s, a) \geq 0 \quad (6)$$

$$v_l \leq \alpha_l(p_{l1} + p_{l2}) \quad (7)$$

$$p_{l1} = (1 - \gamma_1) \sum_{s \in S_{l,a}} w_1(s, a) \quad (8)$$

$$p_{l2} = (1 - \gamma_2) \sum_{s \in S_{l,a}} w_2(s, a) \quad (9)$$

$$p_{l1} + p_{l2} \leq 1 \quad (10)$$

$$0 \leq v_l \leq \beta f_l \quad (11)$$

The objective function in Equation 1 maximizes the total expected cost over all sites. Constraints 2-5 detail the transition graph flow constraints (for thorough inspections and simple inspections). Constraint 2 defines that x is equal to the probability of reaching a state s and performing action a multiplied by the probability of successfully transitioning to state s' . Constraint 3 ensures that the flow into a state s is equal to the flow out of the state. Constraints 4-5 enforce that the total flow in the transition graph, corresponding to the number of defender resources r_i , is held constant for both the flow out of the dummy source nodes s_i^+ and into the dummy sink nodes s_i^- .

Constraint 7 constrains the expected cost for site l . Constraints 8-9 define the probability of successfully inspecting a given site l and is the summation of probabilities of reaching any of l 's corresponding states (thus triggering an inspection) and taking any action a . Note that the failure probability γ means that even if a violating site is inspected, there may not be a fine issued. Constraint 10 limits the overall probability of a site being inspected. If a site is visited by both thorough and simple inspection resources, the site will only have to pay a fine, at most, once. Constraint 11 defines the bounds for the adversary's expected cost; if the adversary's expected cost is at the upper bound ($v_l = \beta f_l$), we assume

that the adversary would prefer to have a positive public perception and choose to comply rather than pay an equivalent amount in expected fines.

6 Evaluation

In order to explore the strategic tradeoffs that exist in our model of the tannery domain, we ran a series of experiments on our Kanpur tannery graph. For each experiment, we generated 3 distinct patrolling strategy types. 1. NECTAR’s strategy, 2. the Uniform Random (UR) strategy: at each time step, every site has an equal probability of being chosen, and 3. an Ad-Hoc (AH) strategy: a deterministic strategy where sites are visited in numerical order (by ID number).

In order to analyze how different resource types affect performance, for each experiment we generated six defender strategies: the first three (NECTAR, UR, AH) correspond to when the defender had twice as many simple inspection resources as thorough inspection resources, and the last three (again NECTAR, UR, AH) correspond to when the defender had no simple inspection resources.

In addition to running experiments where each site l has the same fine (α), we ran a set of experiments where each site’s fine α_l was: $\alpha_l = \alpha f_l$ or, in other words, the fine amount is a constant α multiplied by the number of factories f_l at that site – sites with more factories will be penalized for violations more harshly than sites with fewer factories. As this type of analysis requires heterogeneous sites, we randomize the number of factories at each site.

Ultimately, we are interested in inducing compliance in sites, and for our performance metric, we compute the number of sites that would be in full compliance given the defender strategy (i.e., how many sites’ cost $v_l = \beta f_l$). The maximum number of sites in compliance for each experiment is 50 (i.e., the number of sites on our graph). The default parameter values for each experiment (unless otherwise specified) are listed in Table 1.

Table 1. Default Experiment Values

	Variable Value
Compliance Cost β	10
Fixed Fine Amount α	100
Number of Factories at Each Site f_i	2-5
Number of Simple Inspections r_2	2
Number of Sites	50
Number of Thorough Inspections r_1	1
Patrol duration (hours) t_1, t_2	6
Simple Inspection Failure Rate γ_2	0.6
Thorough Inspection Failure Rate γ_1	0.1
Time granularity (hours) h	1
Time steps to complete simple inspection	1
Time steps to complete thorough inspection	2
Variable Fine Amount α_l	30

Fixed Fine Amount In Figure 2, we analyze the effects of the fixed fine amount α on the number of complying sites. The x-axis shows the fixed fine




Fig. 2. Fixed Fine: Number of Sites in Compliance




Fig. 3. Variable Fine: Number of Sites in Compliance

amount, and the y-axis shows the number of sites that are complying (i.e., $v_l = \beta f_l$).

From the figure, we observe the following trends: (1) the NECTAR strategy does not achieve any compliance until the fine amount is 350, with all sites in compliance at 400. This is due to the objective function attempting to maximize expected cost over all sites simultaneously with a homogeneous fine. (2) While the UR and AH strategies achieve compliance from some of the sites for smaller fine amounts, they do not achieve compliance for all of the sites as quickly as the NECTAR strategy. (3) The inclusion of simple inspection resources improve performance for every strategy as expected.

Variable Fine Amount In Figure 3, we analyze the effects of the variable fine amount α_l on the number of complying sites. The x-axis shows the variable fine amount, and the y-axis shows the number of sites in compliance (i.e., $v_l = \beta f_l$).

From the figure, we observe the following trends: (1) both the NECTAR and UR strategies achieve compliance from all sites for the same variable fine amount; (2) as the fines are not homogeneous for all sites, it is beneficial for NECTAR to try to maximize expected cost in sites with many factories first (unlike with the fixed fine, there is no “water filling” effect); the NECTAR approach achieves faster compliance from larger sites, and (3) the NECTAR achieves compliance from the most sites at every point.

Number of Resources: Variable Fine In Figure 4, we analyze the effect of the number of resources when there is a variable fine amount α_l on the number of complying sites. The x-axis shows the number of thorough inspection resources,




Fig. 4. Number of Resources: Variable Fine: Number of Sites in Compliance




Fig. 5. Patrol Duration: Variable Fine: Number of Sites in Compliance

r_1 (for the strategies with simple inspection resources, the number of simple inspection resources is $r_2 = 2 \times r_1$), and the y-axis shows the number of sites that are complying (i.e., $v_l = \beta f_l$).

From the figure, we observe the following trends: (1) the NECTAR and AH strategies achieve compliance from some sites even with few thorough inspection resources, but NECTAR achieves compliance from the most sites at every point, (2) both the NECTAR and UR strategies achieve compliance from all sites for the same number of thorough inspection resources, and (3) even when there are many resources, the AH strategy does not achieve compliance from all sites.

Patrol Duration: Variable Fine In Figure 5, we analyze the effects of the patrol duration when there is a variable fine amount α_l on the number of complying sites. The x-axis shows the patrol duration, and the y-axis shows the number of sites that are complying (i.e., $v_l = \beta f_l$).

From the figure, we observe the following trends: (1) while the NECTAR strategy performs the best for lower values of patrol duration, it is eventually outpaced by the AH strategy, (2) regardless of the strategy, there is not much change in the number of sites in compliance as a function of patrol duration. For this experiment, the default values for the other parameters result in a low compliance rate regardless of the value of the variable of interest, and (3) having simple inspection resources is helpful for the NECTAR and AH strategies, but it is not very helpful for the UR strategy.

7 Discussion and Results Visualization

Based on these simulations, we make the following conclusions: (1) when the number of resources or variable fine amount is the experiment variable, NECTAR makes the most efficient use of its resources, regardless of whether it is using only thorough inspections or a combination of simple and thorough inspections; (2) having more resources (more manpower) is more useful than increasing the duration of patrols (longer work hours). This is intuitive when considering that each resource must spend time traveling to each site; two resources can each cover a separate sub-section of the graph whereas one resource will be forced to spend more time traveling. Finally, (3) using a variable fine (in which sites are fined according to their number of factories) leads to better compliance rates. This observation makes sense when put in the context of our LP’s objective function: maximize the sum of the expected costs v_l over all sites.

Since our goal is to assist inspection agencies with patrol planning, it is useful to visualize the proposed inspection patrols. In Figure 6, we show a simple graph and strategy visualization in Google Earth (a visualization for the Kanpur area is shown in Figure 7). The lines represent edges on the graph (i.e., straight line connections between sites). Each line also has a time step and a coverage probability associated with it, where the probability represents the value of the MDP’s transition function, $T(s, a, s')$. In other words, this information answers the question: “If the defender resource starts at site l at the beginning of this edge at time step t (i.e., state s), what is the probability that the defender resource will take action a and arrive at site l' , at the end of this edge, in a following time step t' (i.e., state s')?” By clicking on an edge, the user can call up the aforementioned defender strategy information (shown in Figure 6).



Fig. 6. Visualization example




Fig. 7. A Kanpur inspection patrol plan

NECTAR has been proposed to decision makers in governments, pollution control boards, and funding agencies that cover cleaning of large river basins. While field inspectors have not used randomized inspection schemes in the past, they have given positive feedback on this approach. These proposals are still in a preliminary state, and experience from literature suggests that the success of

such initiatives, potentially lasting years, will greatly depend on the collaboration of multiple stakeholders so that the tannery industry and economy can continue to grow while the urgent need to protect the environment is also satisfied.

Acknowledgments: This research was supported by MURI Grant W911NF-11-1-0332.

References

1. Alpern, S., Morton, A., Papadaki, K.: Patrolling games. *Operations research* 59(5), 1246–1257 (2011)
2. Avenhaus, R., von Stengel, B., Zamir, S.: Inspection games. In: *Handbook of Game Theory with Economic Applications*, vol. 3, chap. 51, pp. 1947–1987 (2002)
3. Basilio, N., Gatti, N., Amigoni, F.: Patrolling security games: Definition and algorithms for solving large instances with single patroller and single intruder. *Artificial Intelligence Journal* 184–185, 78–123 (2012)
4. Bosansky, B., Jiang, A., Tambe, M., Kiekintveld, C.: Combining compact representation and incremental generation in large games with sequential strategies. In: *AAAI* (2015)
5. Bošanský, B., Lisý, V., Jakob, M., Pěchouček, M.: Computing time-dependent policies for patrolling games with mobile targets. In: *AAMAS*. pp. 989–996 (2011)
6. Brown, M., Saisubramanian, S., Varakantham, P.R., Tambe, M.: Streets: game-theoretic traffic patrolling with exploration and exploitation. In: *IAAI* (2014)
7. Dong, X., Li, C., Li, J., Wang, J., Huang, W.: A game-theoretic analysis of implementation of cleaner production policies in the chinese electroplating industry. *Resources, Conservation and Recycling* 54(12), 1442–1448 (2010)
8. Filar, J., et al.: Player aggregation in the traveling inspector model. *IEEE Transactions on Automatic Control* 30(8), 723–729 (1985)
9. Gupta, S., Gupta, R., Tamra, R.: Challenges faced by leather industry in kanpur (2007), http://home.iitk.ac.in/~sgupta/tannery_report.pdf
10. India, P.C.o.: Evaluation study on the function of state pollution control boards (2013)
11. Institute, B.: Top ten toxic pollution problems: Tannery operations. Report (2011), http://www.worstpolluted.org/projects_reports/display/88
12. Jain, M., Conitzer, V., Tambe, M.: Security scheduling for real-world networks. In: *AAMAS* (2013)
13. Jainani, D.: Kanpur leather industry in danger as ngt cracks whip on pollution, In: *Financial Express* (2015)
14. Mwinyihija, M.: Emerging world leather trends and continental shifts on leather and leathers goods production. In: *World leather congress proceedings* (2011)
15. Shieh, E., Jiang, A.X., Yadav, A., Varakantham, P., Tambe, M.: Unleashing decmdps in security games: Enabling effective defender teamwork. In: *ECAI* (2014)
16. von Stengel, B.: Recursive inspection games. *arXiv preprint arXiv:1412.0129* (2014)
17. Tambe, M.: *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, New York, NY (2011)
18. Tapiero, C.S.: Environmental quality control and environmental games. *Environmental Modeling & Assessment* 9(4), 201–206 (2005)
19. Yin, Z., Jiang, A., Johnson, M., Tambe, M., Kiekintveld, C., Leyton-Brown, K., Sandholm, T., Sullivan, J.: Trusts: Scheduling randomized patrols for fare inspection in transit systems. In: *IAAI* (2012)