

TRUSTS: Scheduling Randomized Patrols for Fare Inspection in Transit Systems Using Game Theory

*Zhengyu Yin, Albert Xin Jiang, Milind Tambe, Christopher Kiekintveld,
Kevin Leyton-Brown, Tuomas Sandholm, John P. Sullivan*

■ *In proof-of-payment transit systems, passengers are legally required to purchase tickets before entering but are not physically forced to do so. Instead, patrol units move about the transit system, inspecting the tickets of passengers, who face fines if caught fare evading. The deterrence of fare evasion depends on the unpredictability and effectiveness of the patrols. In this article, we present TRUSTS, an application for scheduling randomized patrols for fare inspection in transit systems. TRUSTS models the problem of computing patrol strategies as a leader-follower Stackelberg game where the objective is to deter fare evasion and hence maximize revenue. This problem differs from previously studied Stackelberg settings in that the leader strategies must satisfy massive temporal and spatial constraints; moreover, unlike in these counterterrorism-motivated Stackelberg applications, a large fraction of the ridership might realistically consider fare evasion, and so the number of followers is potentially huge. A third key novelty in our work is deliberate simplification of leader strategies to make patrols easier to execute. We present an efficient algorithm for computing such patrol strategies and present experimental results using real-world ridership data from the Los Angeles Metro Rail system. The Los Angeles County Sheriff's Department is currently carrying out trials of TRUSTS.*

In the Los Angeles (LA) Metro Rail system and other proof-of-payment transit systems worldwide, passengers are legally required to buy tickets before boarding, but there are no gates or turnstiles. There are, quite literally, no barriers to entry, as illustrated in figure 1. Instead, security personnel are dynamically deployed throughout the transit system, randomly inspecting passenger tickets. This proof-of-payment fare collection method is typically chosen as a more cost-effective alternative to direct fare collection, that is, when the revenue lost to fare evasion is believed to be less than what it would cost to make fare evasion impossible.

For the LA Metro, with approximately 300,000 riders daily, this revenue loss can be significant; the annual cost has been estimated at \$5.6 million.¹ The Los Angeles Sheriff's Department (LASD) deploys uniformed patrols onboard trains and at stations for fare checking (and for other purposes such as crime prevention), in order to discourage fare evasion. With limited resources to devote to patrols, it is impossible to cover all loca-



Figure 1. Entrance of an LA Metro Rail Station.

tions at all times. The LASD thus requires some mechanism for choosing times and locations for inspections. Any predictable patterns in such a patrol schedule are likely to be observed and exploited by potential fare evaders. The LASD's current approach relies on humans for scheduling the patrols. However, human schedulers are poor at generating unpredictable schedules (Wagenaar 1972, Tambe 2011); furthermore such scheduling for LASD is a tremendous cognitive burden on the human schedulers who must take into account all of the scheduling complexities (for example, train timings, switching time between trains, and schedule lengths). Indeed, the sheer difficulty of even enumerating the trillions of potential patrols makes any simple automated approach — such as a simple dice roll — inapplicable.

The result of our investigation is a novel application called TRUSTS (tactical randomization for urban security in transit systems) for fare-evasion deterrence in urban transit systems, carried out in collaboration with the LASD. We take a game-theoretic approach, which studies systems with multiple self-interested parties and aims to predict the likely outcomes of the system under rational behavior of the players. In particular, we model this problem as a Stackelberg game with one leader (the LASD) and many followers, in which each metro rider (a follower) takes a fixed route at a fixed time. The leader precommits to a mixed

patrol strategy (a probability distribution over all pure strategies), and riders observe this mixed strategy before deciding whether to buy the ticket or not (the decision to ride having already been made) in order to minimize their expected total cost, following for simplicity the classic economic analysis of rational crime (Becker and Landes 1974). Both ticket sales and fines issued for fare evasion translate into revenue to the government. Therefore the optimization objective we choose for the leader is to maximize total revenue (total ticket sales plus penalties).

There are exponentially many possible pure patrol strategies, each subject to both the spatial and temporal constraints of travel within the transit network. Explicitly representing a mixed strategy would be impractical. To remedy this difficulty, TRUSTS uses the transition graph, which captures the spatial as well as temporal structure of the domain, and solves for the optimal (fractional) flow through this graph, using linear programming (LP). Such a flow can be interpreted as a marginal coverage vector. Additionally, we show that a straightforward approach to extracting patrol strategies from the marginals faces important challenges: it can create infeasible patrols that violate the constraint on patrol length, and it can generate patrols that switch too frequently between trains, which can be difficult for patrol personnel to carry out. Thus, we present a novel technique to

overcome these difficulties using an extended formulation on a history-duplicate transition graph that (1) forbids patrols that are too long and (2) penalizes patrols with too many switches.

Finally, we perform simulations based on actual ridership data provided by the LASD for four LA Metro train lines (Blue, Gold, Green, and Red). Our results suggest the possibility of significant fare-evasion deterrence and hence prevention of revenue loss with very few resources. The LASD is currently testing TRUSTS in the LA Metro system by deploying patrols according to our schedules and measuring the revenue recovered.

Related Work

There has been research on a wide range of problems related to game-theoretic patrolling on graphs. One line of work considers games in which one player, the patroller, patrols the graph to detect and catch the other player, the evader, who tries to minimize the detection probability. This includes work on hide-seeker games (Halvorson, Conitzer, and Parr 2009) for the case of mobile evaders, and search games (Gal 1979) for the case of immobile evaders.

Another line of research considers games in which the patroller deploys resources (static or mobile) on the graph to prevent the other player, the attacker, from reaching certain target vertices. There are a few variations depending on the set of possible sources and targets of the attacker. Infiltration games (Alpern 1992) considered one source and target. Asset protection problems (Dickerson et al. 2010) consider multiple sources and multiple equally weighted targets. Networked security games (Tsai et al. 2010) consider targets with different weights.

The leader-follower Stackelberg game model has been the topic of much recent research (Tambe 2011) and has been applied to a number of real-world security domains, including the Los Angeles International Airport (Jain et al. 2010), the Federal Air Marshals Service (Jain et al. 2010), and the Transportation Security Administration (Pita et al. 2011).

Urban transit systems, however, present unique computational challenges. First, unlike in existing work on graph patrolling games, and unlike in previous deployed applications on counterterrorism, here the followers we seek to influence are potentially very many: large numbers of train riders might plausibly consider fare evasion. Booz Allen Hamilton (see note 1) estimates that 6 percent of riders are ticketless in the metro system overall; anecdotal reports suggest that on some lines this percentage could be far greater, even a majority. Second, the leader has exponentially many possible patrol strategies, corresponding to all the feasi-

ble trips within the transit network subject to certain restrictions and preferences. Similar to FAMS (Jain et al. 2010), we represent patrol strategies compactly as a marginal coverage vector. But unlike the FAMS problem in which a patrol consists of a very limited number of flights (often a pair of flights), TRUSTS allows much more complex patrols and thus uses a novel compact representation based on history-duplicate transition graphs.

Game Theory Background

We begin with a brief overview of the game-theoretic concepts that we use in this article. At a high level, game theory studies games, which are mathematical models of interactions among multiple players, each trying to advance their self-interest by choosing among a set of strategies. Formally, a game has a set N of players; each player $i \in N$ has a set S_i of pure strategies; the outcome of the game is determined by the strategy profile $s \in S$, where S is the Cartesian product of the player's sets of pure strategies. Each player i 's preference over the outcomes of the game is specified by her utility function $u_i : S \rightarrow \mathbb{R}$, and her objective is to maximize her expected utility. Each player can choose a pure strategy, or a mixed strategy that is a probability distribution over pure strategies. Player i 's best response strategy, given strategies of the other players, is a strategy of i that maximizes i 's expected utility. It is always possible to pick a best response that is a pure strategy.

In a Stackelberg game between a leader and a follower, the leader commits to a mixed strategy, and the follower observes the leader's mixed strategy and plays a best response (Conitzer and Sandholm 2006). This setting has been applied to security domains where the defender's daily patrol strategies are observed by attackers. The optimal mixed strategy for the leader to commit to, together with the follower's best response, forms a Stackelberg equilibrium of the game.

Problem Setting

TRUSTS addresses the challenge of generating randomized schedules for LASD patrols for four LA Metro lines (Blue, Gold, Green, and Red). For simplicity, TRUSTS treats the LA Metro system's multiple lines as independent. Indeed, currently, the LA Metro lines have just a few transfer points. Dealing with the impact of transfers will be a topic for future work. We model the patrol scheduling problem on each line as a leader-follower Stackelberg game with one leader (the LASD) and multiple followers (riders). In this game, a pure leader strategy is a patrol, that is, a sequence of patrol actions (defined below), of constant bounded duration.

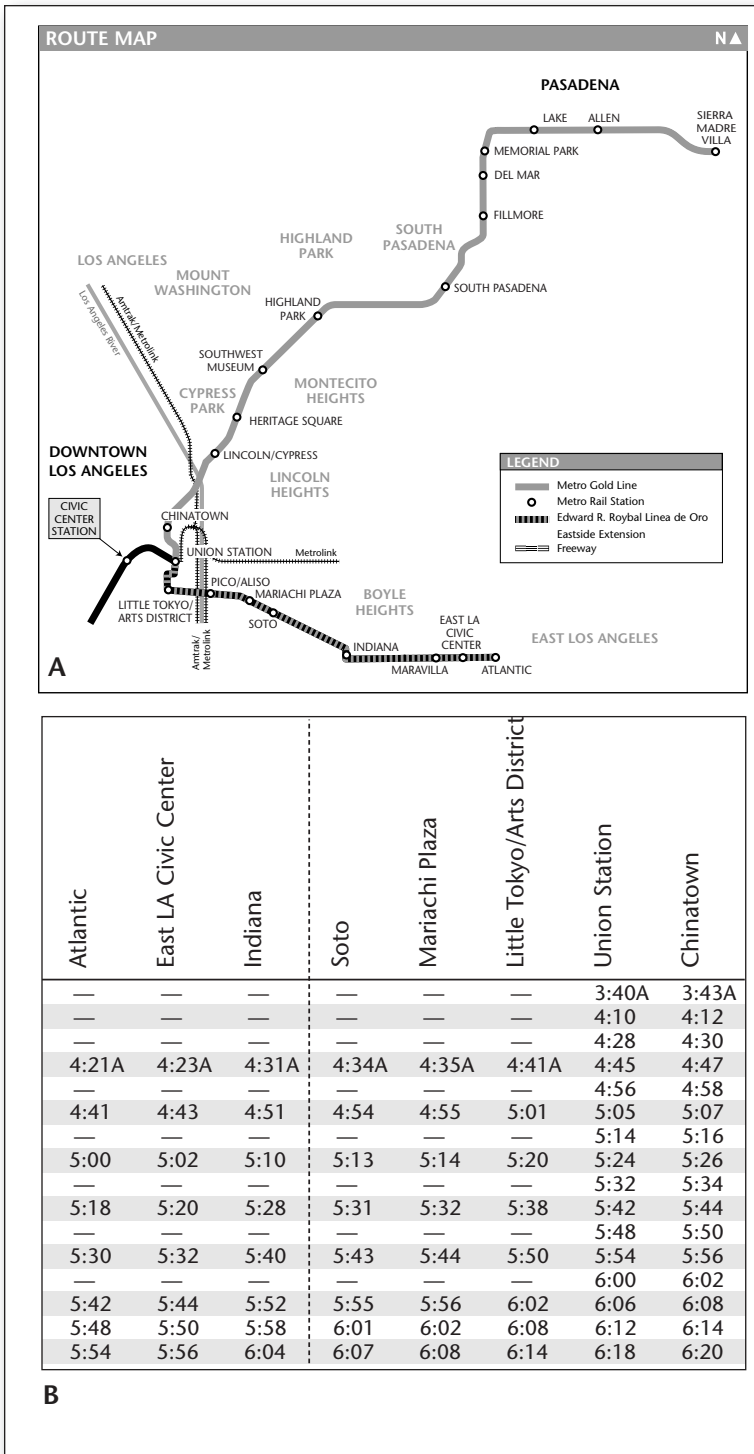


Figure 2. The Los Angeles Metro's Gold Line.

The two possible pure follower strategies are buying and not buying. There are many types of followers, one for each source, destination, and departure time triple (corresponding to the set of all riders who take such a trip).

Train System

The train system consists of a single line (for example, the Gold line as shown in figure 2a on which trains travel back and forth, in general with multiple trains traveling simultaneously. The system operates according to a fixed daily schedule (a snippet is shown in figure 2b), with trains arriving at stations at (finitely many) designated times throughout the day. Therefore we can model time as discrete, focusing only on the time steps at which some train arrival/departure event occurs. We use the (directed) transition graph $G = \langle V, E \rangle$ to encode the daily timetable of the metro line, where a vertex $v = \langle s, t \rangle$ corresponds to some pair of station s and time point t . An edge in G represents a possible (minimal) action. In particular, there is an edge from $\langle s, t \rangle$ to $\langle s', t' \rangle$ if: s' is either the predecessor or successor of s in the station sequence and $\langle s, t \rangle$ and $\langle s', t' \rangle$ are two consecutive stops for some train in the train schedule (traveling action), or $s' = s$, $t < t'$, and there is no vertex $\langle s, t'' \rangle$ with $t < t'' < t'$ (staying action). We refer to the entire path that a given train takes through G , from the start station to the terminal station, as a train path.

Patrols

There are a fixed number γ of deployable patrol units, each of which may be scheduled on a patrol of duration at most κ hours (with, for example, $\kappa = 7$). There are two sorts of patrol actions, which a given patrol unit can alternate between on its shift: on-train inspections (in which patrollers ride the train, inspecting passengers), and in-station inspections (in which they inspect passengers as they exit the station). A pure patrol strategy is represented mathematically as a path in G for each patrol unit, in which an edge e represents an atomic patrol action, that is, inspecting in-station from the time of one train event at that station to the next (at that station) or inspecting on-train as it travels from one station to the next. Each edge e has a length l_e equal to the corresponding patrol action duration and an effectiveness value f_e , which represents the percentage of the relevant ridership inspected by this action. For both in-station and on-train inspections, f_e depends on the ridership volume at that location, the time of day, and the duration. A valid pure patrol strategy is a set of paths P_1, \dots, P_γ each of size at most κ .

Example

A simple scenario with three stations (A, B, C) and four discrete time points ($6pm, 7pm, 8pm, 9pm$) is given in figure 3. The dashed lines represent staying actions; the solid lines represent traveling actions. There are four trains in the system; all edge durations are 1 hour. A sample train path here is $(A, 6pm) \rightarrow (B, 7pm) \rightarrow (C, 8pm)$. In this example, if $\kappa = 2$ and $\gamma = 1$, then the valid pure leader strate-

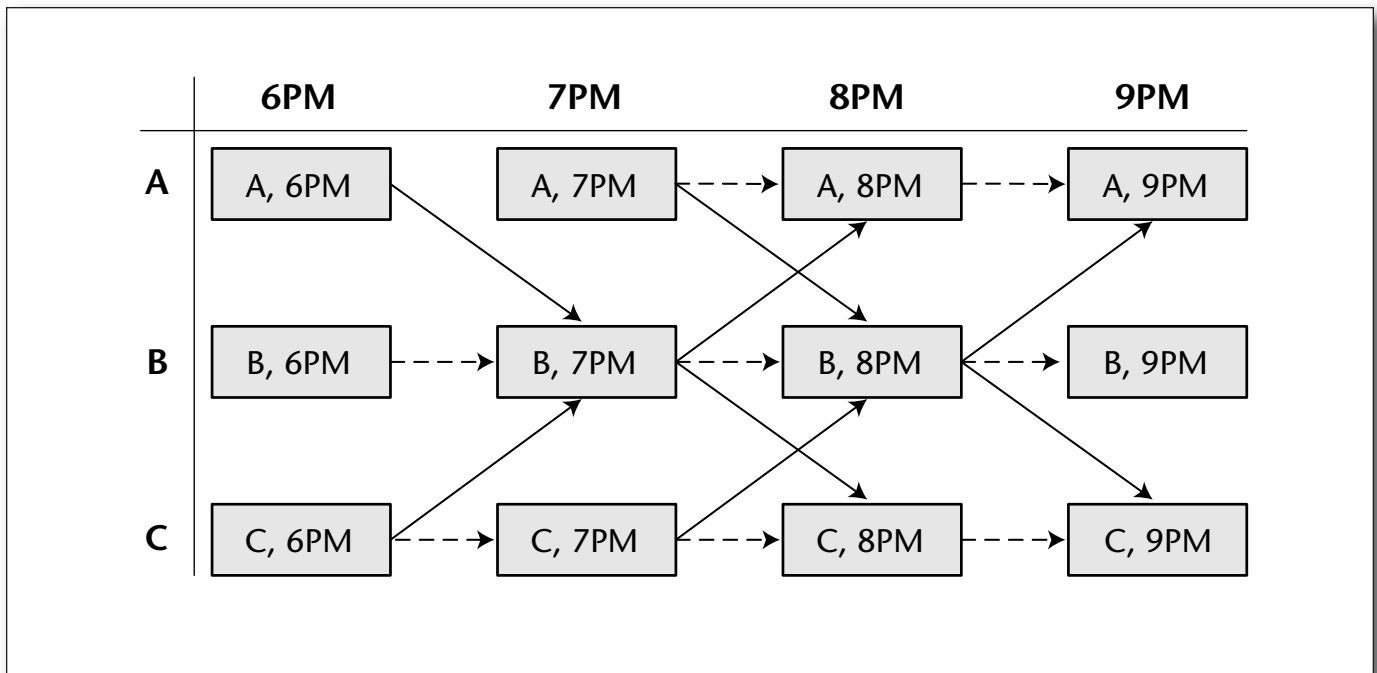


Figure 3. The Transition Graph of a Toy Problem Instance.

gies (pure patrol strategies) consist of all paths of length 2.

Riders

The riders are assumed to be daily commuters who take a fixed route at a fixed time. Horizon Research Corporation² estimates that more than 82 percent of riders use the system at least three days a week. The ticket price (for any trip within the transit system) is a nominal fee ρ , with the fine for fare evasion τ much greater. As the riders follow the same route every day, they could estimate the likelihood of being inspected, based on which they make a decision as to whether to buy a ticket. We assume the riders know the inspection probability perfectly and are rational, risk-neutral economic actors (Becker and Landes 1974) who make this choice in order to minimize expected cost.

A rider's type λ is defined by the path he or she takes in the graph. Because there is a single train line, riders never linger in stations, that is, do not follow any "stay" edges (staying at a station) mid-journey; the last edge of every follower type is a (short) stay edge, representing the action of "exiting" the destination station, during which the rider may be subject to in-station inspection. Therefore the space Λ of rider types corresponds to the set of all subpaths of train paths. (When G is drawn as in figure 3, all rider paths are "diagonal" except for the last edge.) A metro line with N stops and M scheduled trains will have $O(MN^2)$ rider types.

Given a pure patrol strategy of the γ units, (P_1, \dots, P_γ) , the inspection probability for a rider of type $\lambda \in \Lambda$ is:

$$\min\{1, \sum_{i=1}^{\gamma} \sum_{e \in P_i \cap \lambda} f_e\}, \quad (1)$$

and therefore his or her expected utility is the negative of the expected amount he or she pays: $-\rho$ if she or he buys the ticket and $-\tau$ multiplied by the expression in equation (1) otherwise. The inspection probability for a mixed strategy is then the expectation of equation (1), taken over the distribution of pure strategies.

We justify the inspection probability in equation (1) as follows. First, consider on-train inspections. The fraction of the train that is inspected in a given inspection action is determined by f_e (which depends on ridership volume). The key is that in the next inspection action, a patrol will not reinspect the fraction of the train that is already inspected in a previous inspection action. Therefore, unlike in settings where patrollers may repeatedly draw a random sample from the same set of train passengers to inspect, in our setting, the probabilities f_e are added rather than multiplied. Now also consider in-station inspections. Since a rider taking a journey only exits a single station, a rider will encounter at most one in-station inspection. Finally, when multiple patrol units cover the same edge e , the inspection probability given by

(1) is the sum of the contributions from each patrol unit, capped at 1. This is a reasonable assumption when the number of patrol units on each edge e is small, as multiple patrol units on the same train could check different cars or different portions of the same car, and multiple patrol units inspecting at the same station could check different exits.

Objective

The leader's utility, equal to total expected revenue, can be decomposed into utilities from bilateral interactions with each individual follower; furthermore the followers do not directly affect each other's utilities. As a result, the game can be equivalently formulated as a two-player Bayesian Stackelberg game, in which there is just one follower taking on the role of one of the passengers. The type of this passenger is known to the follower but not to the leader and is drawn from a probability distribution such that the probability p_λ of a follower type $\lambda \in \Lambda$ is proportional to its ridership volume.

Furthermore, although in general solving Bayesian Stackelberg games is NP-complete (Conitzer and Sandholm 2006), the utility functions of this game satisfy the zero-sum property, that is, the utility gained by the leader always equals the utility lost by the follower. For such zero-sum Bayesian games, the Stackelberg equilibrium is equivalent to the maximin solution, and the games are solvable by the linear programming formulations of Ponsard and Sorin (1980) or Koller, Megiddo, and von Stengel (1994). However, these LP formulations explicitly enumerate the pure strategies of the leader; since our game of interest has an exponential number of leader pure strategies, even storing these formulations would require exponential space.

Linear Program Formulation

In this section, we introduce our linear-programming-based approach for finding a maximum-revenue (mixed) patrol strategy. As noted above, the leader's space of pure strategies is exponentially large, even with a single patrol unit. We avoid this difficulty by compactly representing mixed patrol strategies by marginal coverage on edges of the transition graph, that is, by the expected numbers of inspections that will occur on these edges. Subsequently, we construct a mixed strategy (that is, a probability distribution over pure strategies) consistent with the marginal coverage.

For expository purposes, we first present a basic formulation of our approach of compactly representing the problem using marginal coverage. This basic formulation also illustrates the key issues that make it difficult for the end user to deploy the patrol strategies computed. We then introduce an

extended formulation to address these issues. We will focus on the high-level ideas and refer interested readers to our technical paper (Yin et al. 2012) for further details.

Basic Formulation

For algorithmic convenience, we add to the transition graph a source v^+ with edges to all possible starting vertices in the transition graph and a sink v^- with edges from all possible ending vertices. We assign these additional dummy edges zero duration and zero effectiveness.

Let x_e be the expected number of inspections that will occur on edge e of this transition graph. We call the vector \mathbf{x} of marginal coverage on all edges the marginal strategy. Since the number of edges in the graph is exponentially fewer than the number of paths, the marginal strategies are a much more compact representation of mixed strategies. A valid vector \mathbf{x} must satisfy the following constraints, which are the constraints defining a fractional flow on the graph when we interpret x_e to be the amount of flow on edge e : the total flows entering and exiting the system are bounded by γ , the number of total patrol units allowed; the total flow entering each intermediate node equals the total flow out of that node. Furthermore, the expected total number of time units spent by the patrols must be bounded by $\gamma \cdot \kappa$.

Recall that our objective is to maximize the leader's expected revenue, which equals the riders' expected total payment. Since the riders are rational players who minimize their expected costs, the expected payment of each rider is the minimum of the ticket price, and his or her expected payment if he or she decides to evade, which equals the fine times the probability of the rider being captured. It turns out that we could upper-bound the probability of capture using a linear expression in terms of \mathbf{x} , leading to a linear program that provides an upper bound on the optimal revenue. Fortunately, once we generate the patrols from the marginals we are able to compute the actual best-response expected utilities of the riders. Our experiments show that the differences between the actual expected utilities and the upper-bounds given by the LP formulation are small. Meanwhile the size of this LP grows only polynomially in the size of the transition graph. Once we solve the LP to find the optimal marginal strategy \mathbf{x} , we can efficiently construct a γ -unit mixed strategy whose marginals match \mathbf{x} .

Issues with the Basic Formulation

There are two fundamental issues with the basic formulation. First, the mixed strategy constructed can fail to satisfy the patrol length limit of κ , notwithstanding the constraint on the sum of the lengths of all patrols, and hence be infeasible. In

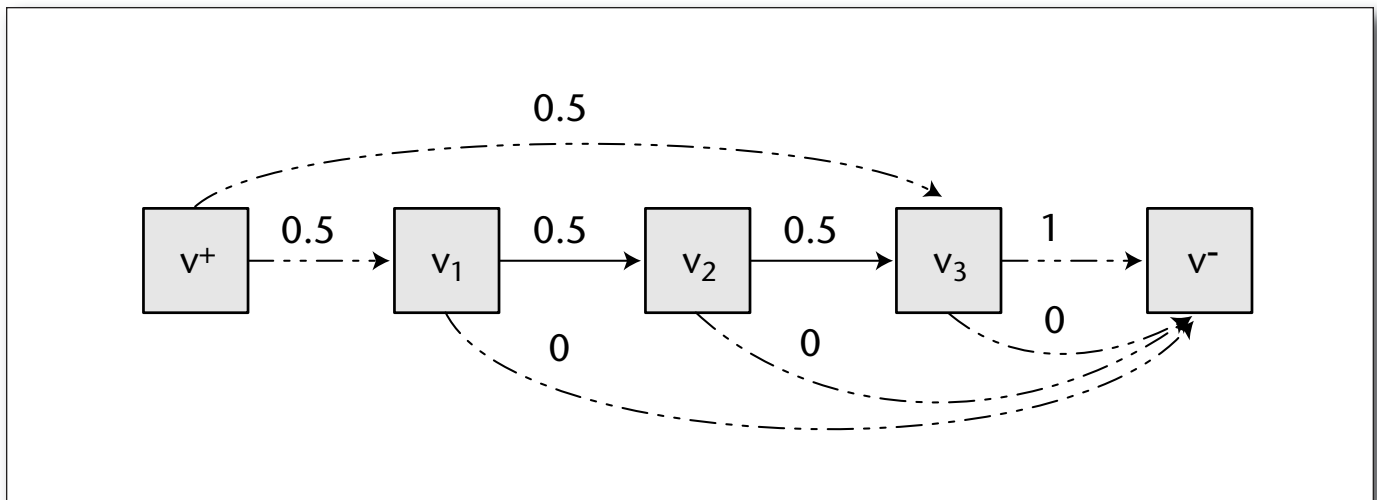


Figure 4. Example of an Infeasible Marginal Strategy.

fact, the marginal strategy computed in the basic formulation may not correspond to any feasible mixed strategy in which all patrols have length at most κ . Consider the counterexample in figure 4. Edges $v_1 \rightarrow v_2$ and $v_2 \rightarrow v_3$ represent two real atomic actions, each with duration 1. Patrols must start from either v_1 or v_3 , but can terminate at any of v_1 , v_2 , and v_3 . This is specified using v^+ and v^- , the dummy source and sink, respectively. We assume $\kappa = 1$ and $\gamma = 1$. It can be verified that the marginal strategy shown in figure 4 satisfies the constraints of our basic LP. However, the only corresponding mixed strategy is to take $v^+ \rightarrow v_3 \rightarrow v^-$ with 50 percent probability and $v^+ \rightarrow v_1 \rightarrow v_2 \rightarrow v_3 \rightarrow v^-$ with 50 percent probability. (To see this, observe that we cannot put any probability mass on the path $v^+ \rightarrow v_1 \rightarrow v_2 \rightarrow v^-$ because the edge $v_2 \rightarrow v^-$ has zero flow; the path $v^+ \rightarrow v_1 \rightarrow v^-$ is excluded by a similar argument.) This mixed strategy is infeasible since its second patrol has duration greater than 1. This patrol length violation arises because the basic formulation only constrains the average patrol length, and therefore allows the use of overlong patrols as long as some short patrols are also used.

Second, the paths selected according to the constructed mixed strategy may switch between trains or between in-station and on-train inspection an impractically large number of times, making the patrol path difficult to implement and error prone. This is an important issue as we want real LASD officers to be able to carry out these strategies. The more switches there are in a patrol strategy, the more instructions the patrol unit has to remember, and the more likely the unit will miss a switch due to imperfections in the train schedule and/or the unit's misexecution of the instructions. For example, in the example, $\langle A, 6pm \rangle \rightarrow \langle B, 7pm \rangle \rightarrow \langle A,$

$8pm \rangle$ and $\langle C, 6pm \rangle \rightarrow \langle B, 7pm \rangle \rightarrow \langle C, 8pm \rangle$ each do one switch while $\langle A, 6pm \rangle \rightarrow \langle B, 7pm \rangle \rightarrow \langle C, 8pm \rangle$ and $\langle C, 6pm \rangle \rightarrow \langle B, 7pm \rangle \rightarrow \langle A, 8pm \rangle$ each do none. Both path pairs cover the same set of edges, making the second preferable because it is easier to implement.

Extended Formulation

Now we present a more sophisticated formulation design to address the two aforementioned issues. The difficulty involved in imposing constraints on the patrol paths (that is, penalizing or forbidding certain paths) in the marginal representation is that paths themselves are not represented, instead being encoded only as marginal coverage.

Hence the key idea is to preserve sufficient path history information within vertices to be able to evaluate our constraints, while avoiding the exponential blowup creating a node for every path would cause. We construct a new graph, called the history-duplicate transition (HDT) graph, by creating multiple copies of the original vertices, each corresponding to different values of history information.

We first explain how to construct the HDT graph from a transition graph G in order to forbid patrol paths longer than κ . The HDT graph is composed of multiple restricted copies of G (that is, subgraphs of G), corresponding to different possible starting time points. For the copy corresponding to starting time point t^* , we only keep the subgraph on vertices $v = \langle s, t \rangle \in V$ where $t^* \leq t \leq t^* + \kappa$. Thus, in each restricted copy of G , the length of any path is guaranteed to be less than or equal to κ . Since there are a finite number of distinct possible starting time points (that is, all distinct discrete time points in V^+), the new graph is a linear expansion

of G . An approximation can be obtained by taking one starting time point every δ time units. In this case, an original vertex (edge) will be kept in at most $\lceil \kappa/\delta \rceil$ copies, implying the new graph is at most $\lceil \kappa/\delta \rceil$ times larger than G .

Figure 5a shows the HDT graph (the shaded portion further explained below) of the example with $\kappa = 2$ and two starting time points, $6pm$ and $7pm$. The HDT graph is thus composed of two restricted copies of the original transition graph. In each vertex, the time shown in parenthesis indicates the starting time point. For example, the original vertex $\langle A, 7pm \rangle$ now has two copies $\langle A, 7pm, (6pm) \rangle$ and $\langle A, 7pm, (7pm) \rangle$ in the HDT graph. For the starting time point of $6pm$, the patrol must end at or before $8pm$, hence we do not need to keep vertices whose discrete time point is $9pm$. For the starting time point of $7pm$, the patrol must start at or after $7pm$, hence we do not need to keep vertices whose discrete time point is $6pm$. The two restricted copies are not two separate graphs but a single graph that will be tied together by the dummy source and sink.

Next, we explain how to further extend the HDT graph to penalize complex patrol paths. The idea is to have each vertex encode the last action occurring prior to it. Specifically, we create multiple copies of a vertex v , each corresponding to a different edge that leads to it. If v is a possible starting vertex, we create an additional copy representing no prior action. If there is an edge from v to v' , we connect all copies of v to the specific copy of v' whose last action was (v, v') . A new edge is called a switching edge if the recorded last actions of its two vertices are of different types (for example, inspecting different trains), unless one of the two vertices is a “no prior action” vertex. As can be verified, the number of switches of a patrol path in the new graph is the number of switching edges it has. To favor simple patrol paths, we impose a cost for using switching edges.

In figure 5b, we show how to apply this extension using the subgraph shown in the shaded box of figure 5a. Since there is only one edge leading to $\langle A, 7pm, (6pm) \rangle$, we create one copy of it representing the action of staying at A . There are three edges leading to $\langle B, 7pm, (6pm) \rangle$, so we create three copies of it representing the actions of taking train from A , staying at B , and taking train from C . The original edges are also duplicated. For example, $\langle B, 7pm, (6pm) \rangle \rightarrow \langle B, 8pm, (6pm) \rangle$ has three copies connecting the three copies of $\langle B, 7pm, (6pm) \rangle$ to the copy of $\langle B, 8pm, (6pm) \rangle$ representing the staying at B action. Among the three copies, only the “Stay” to “Stay” edge is not a switching edge.

Given the final HDT graph $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$, we provide an extended linear program formulation, with variables y_e representing the marginal coverage of an HDT graph edge $e \in \mathcal{E}$ being selected. To penal-

ize switches, we add to the objective function a penalty β for each unit of marginal coverage used by switching edges. Varying the value of β lets us trade off between solution quality (greater revenue) and patrol preference (lower average number of switches). A path in the HDT graph \mathcal{G} trivially corresponds to a path in the transition graph G , since any edge in \mathcal{G} is a duplicate of some edge in G . Because the length of any patrol path in the HDT graph is bounded by κ , the mixed strategy must be feasible.

Real-World Evaluation

We present our evaluation based on real metro schedules and rider traffic data provided by the LASD. We solved the LP in the extended formulation using CPLEX 12.2 on a standard 2.8 GHz machine with 4 GB memory. We first describe the data sets we used, followed by our experimental results.

Data Sets

We created four data sets, each based on a different Los Angeles Metro Rail line: Red (including Purple), Blue, Gold, and Green. For each line, we created its transition graph using the corresponding timetable from www.metro.net. Implementing the LP requires a distribution of types of potential fare evaders (recall that a rider type corresponds to a 4-tuple of boarding station / time and disembarking station / time). We have access to hourly boarding and alighting counts provided by the Los Angeles Sheriff's Department, and from this can estimate the distribution of population on the train system at each hour; but we do not know what percentage of the population are potential fare evaders. In our experiments, we assumed that potential fare evaders were evenly distributed among the general population. Specifically, suppose the percentage of riders boarding in hour i is d_i^+ and the percentage of riders alighting in hour i is d_i^- . Denote the set of rider types that board in hour i by Λ_i^+ and those that alight in hour i by Λ_i^- . Then we would like to compute a fine-grained ridership distribution \mathbf{p} to match the hourly boarding and alighting percentages, that is, to find a point within the following convex region Ω ,

$$\Omega = \{\mathbf{p} \mid \mathbf{p} \succeq 0 \wedge \sum_{\lambda \in \Lambda_i^+} p_\lambda = d_i^+ \wedge \sum_{\lambda \in \Lambda_i^-} p_\lambda = d_i^-, \forall i\}.$$

We estimate the fare-evader distribution by finding the analytic center of Ω , which is efficiently computable.

The inspection effectiveness f_e of an edge is assigned based on the assumption that 10 passengers can be inspected per minute. The inspection effectiveness f_e is capped at 0.5 to capture the fact that the inspector cannot switch between cars

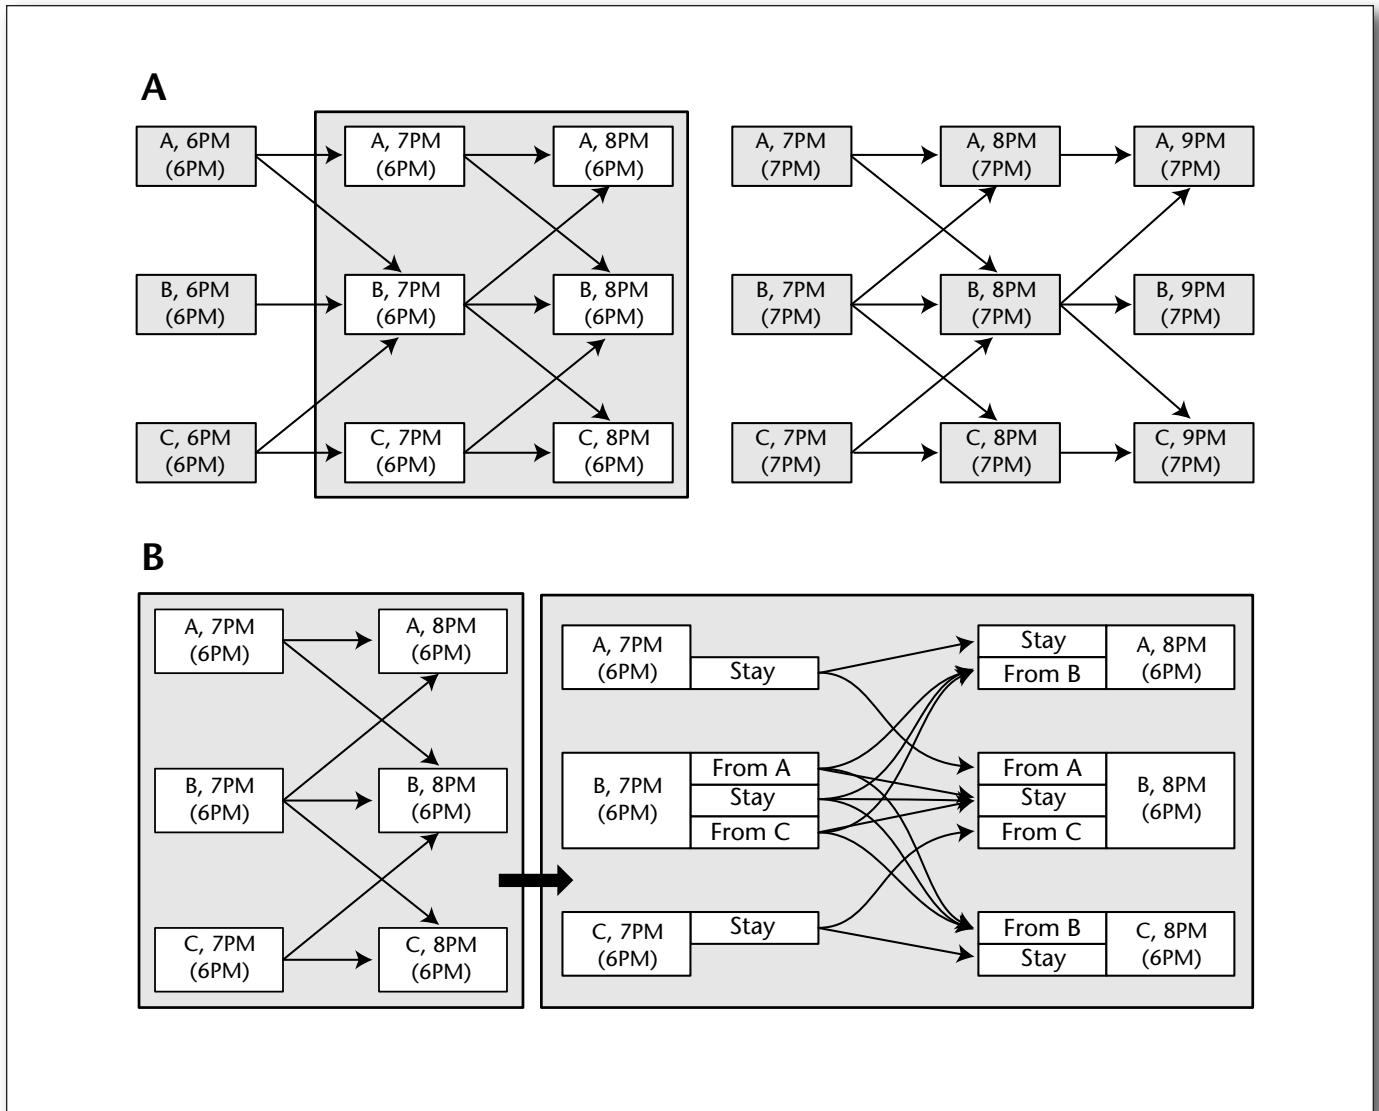


Figure 5. HDT Graph of Example 1.

(a) Two starting time points. (b) Extension storing the last action occurring.

while the train is moving. (Trains contain at least two cars.) The ticket fare was set to \$1.50 (the actual current value) while the fine was set to \$100.00 (Fare evaders in Los Angeles can be fined \$200.00, but they also may be issued warnings.) If we could increase the fine dramatically the riders would have much less incentive for fare evasion, and we could achieve better revenue. However a larger fine is infeasible legally. Table 1 summarizes the detailed statistics for the LA Metro lines.

Line	Stops	Trains	Daily Riders	Types
Red	16	433	149991.5	26033
Blue	22	287	76906.2	46630
Gold	19	280	30940.0	41910
Green	14	217	38442.6	19559

Table 1. Statistics of Los Angeles Metro Lines.

Experimental Results

Throughout our experiments, we fixed γ to 1. In our first set of experiments, we fixed the penalty β for using patrol paths with more switches to 0, and varied the maximum number of hours that an inspector can patrol from 4 to 7 hours. To create

the HDT graph, we took one starting time point every hour.

Figure 6a shows the expected revenue per rider of the mixed patrol strategy we generated, which is the total revenue divided by the number of daily

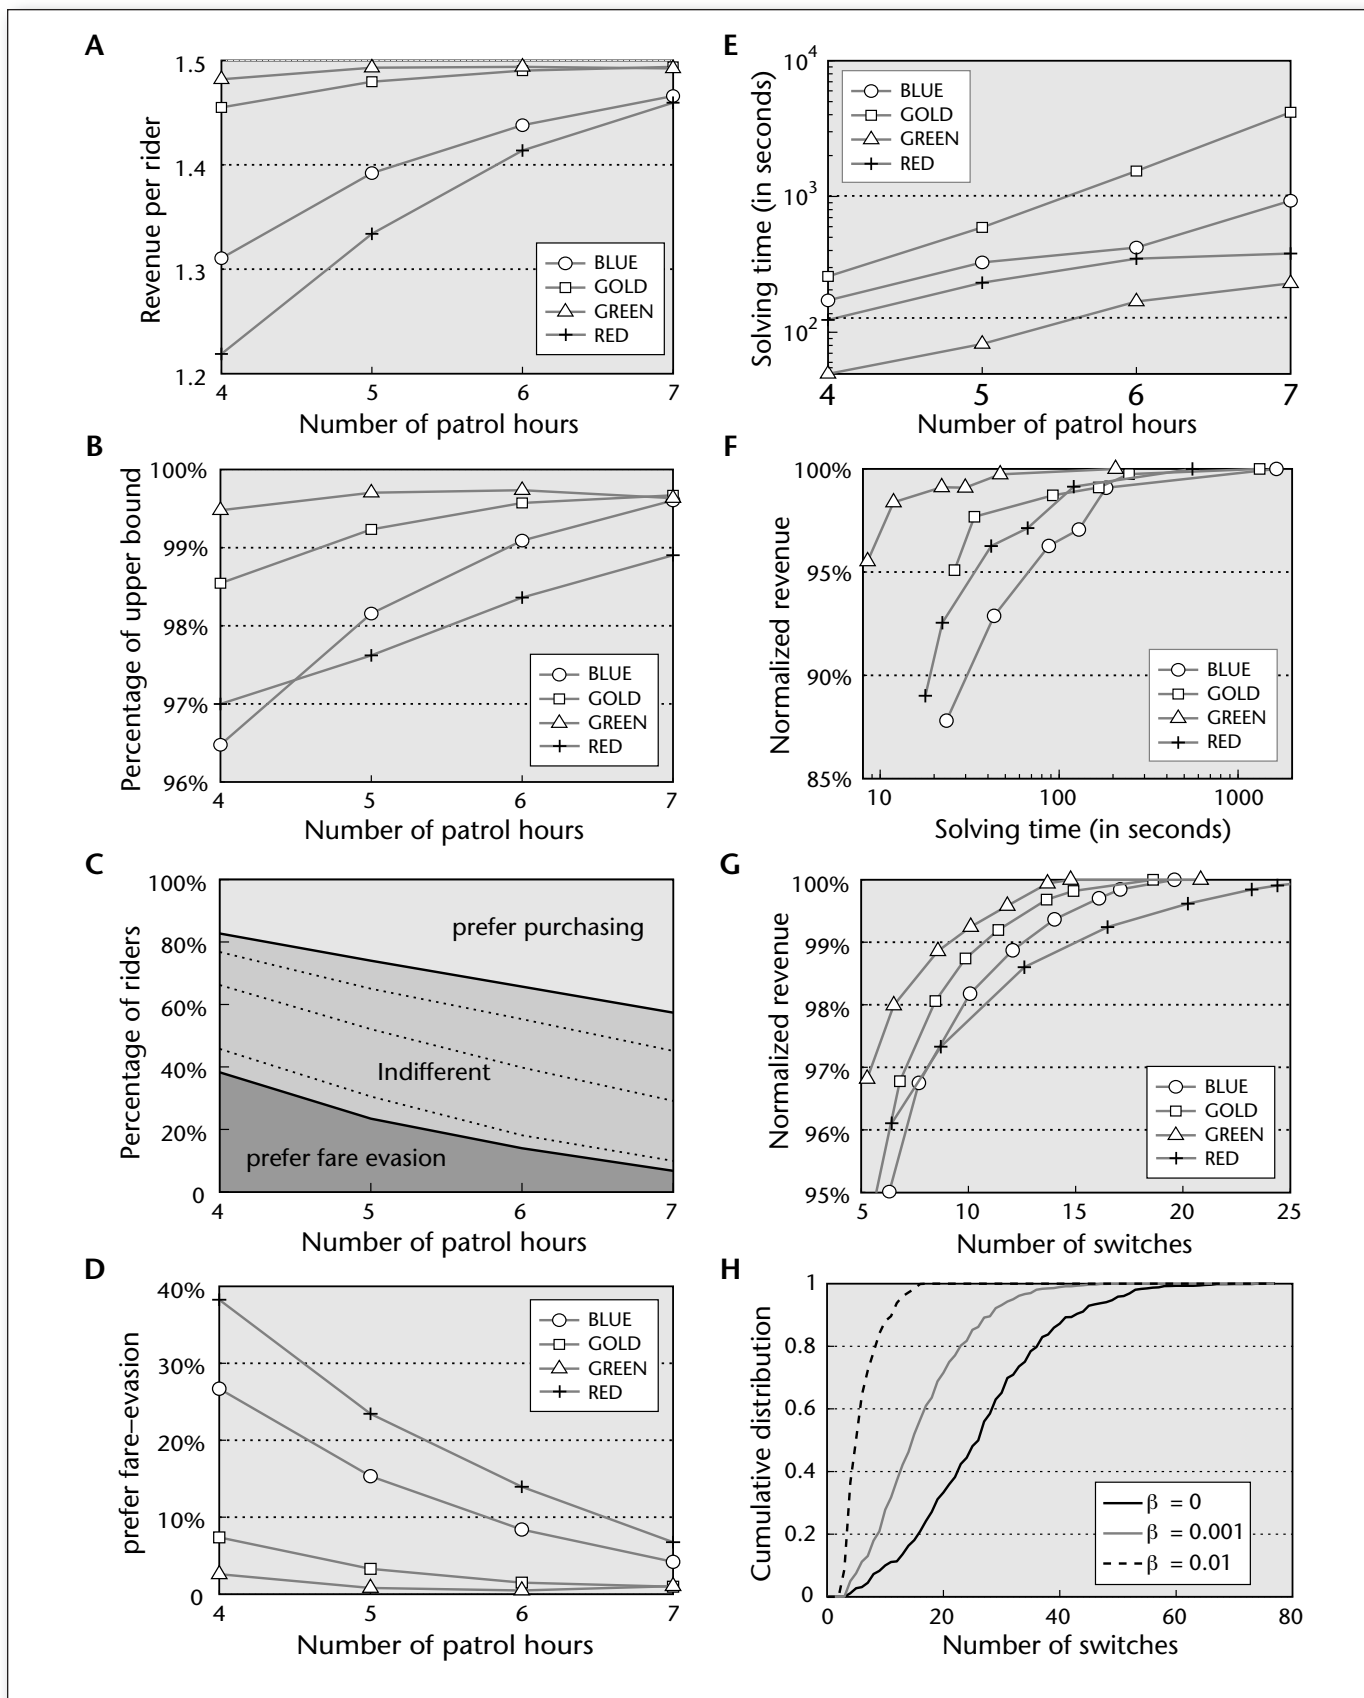


Figure 6. Experimental Results.

riders. Since the LP only returns an upper bound of the attainable revenue, the true expected revenue of the mixed patrol strategy was computed by evaluating the riders' best responses for all rider types. A rider can always pay the ticket price of \$1.50 and will only evade the ticket when the expected fine is lower. Hence the theoretical maximum achievable value is \$1.50, which is achieved when every rider purchases a ticket. As we can see, the per rider revenue increases as the number of patrol hours increases, almost converging to the theoretical upper bound of \$1.50 for the Gold and Green line. Specifically, a 4-hour patrol strategy already provides reasonably good expected value: 1.31 for the Blue line (87.4 percent of the maximum), 1.45 for the Gold line (97.0 percent), 1.48 for the Green line (98.8 percent), and 1.22 for the Red line (81.3 percent). Among the four lines, the Red line has the lowest revenue per rider. This is because the effectiveness of fare inspection decreases as the volume of daily riders increases, and the Red line has a significantly higher number of daily riders than the other lines.

We depict in figure 6b the percentage of the true expected revenue versus the theoretical upper bound returned by the LP. Strategies generated by our method are near optimal; for example, our 4-hour strategies for the Blue, Gold, Green, and Red lines provided expected revenues of 96.5 percent, 98.5 percent, 99.5 percent, and 97.0 percent of the upper bound (and thus at least as much of the optimum), respectively.

To study riders' responses to the computed strategy, we partitioned the entire population of riders into three groups depending on their expected fine if fare evading: riders who prefer purchasing tickets (expected fine is greater than 1.7 — 13.3 percent above the ticket price), riders who prefer fare evasion (expected fine is less than 1.3 — 13.3 percent below the ticket price), and indifferent riders (expected fine is between 1.3 and 1.7). In figure 6c, we show the distribution of the three groups against the strategies computed for the Red line. The three dashed lines inside the region of indifferent riders represent, from top to bottom, the percentages of riders whose expected fine is less than 1.6, 1.5, and 1.4, respectively. As the number of patrol hours increases from 4 to 7, the percentage of riders who prefer fare evasion decreases from 38 percent to 7 percent, the percentage of riders who prefer purchasing tickets increases from 17 percent to 43 percent, and the percentage of indifferent riders remains stable between 45 percent and 50 percent.

Zooming in on the fare evasion, figure 6d shows the percentage of riders who preferred fare evasion against the patrol strategies computed. As we can see, this percentage decreased almost linearly in the number of additional patrol hours beyond 4.

Our 7-hour patrol strategy lowered this percentage to 4.2 percent for the Blue line, 0.01 percent for the Gold line, 0.01 percent for the Green line, and 6.8 percent for the Red line. Again, due to having the highest daily volume, the Red line had the highest percentage of riders who preferred fare evasion.

Finally, figure 6e shows the run time required by CPLEX to solve the LPs we created. As we can see, the run time increased as the number of patrol hours increased for all the metro lines. This is because the size of the HDT graph constructed is roughly proportional to the maximum length of the patrols, and a larger HDT graph requires an LP with more variables and constraints. Among the four lines, the Red and the Green lines have significantly fewer types, and are thus easier to solve than the other two lines.

In our second experiment, we varied the interval δ of taking starting time points, trading off solution quality for efficiency. We fixed the patrol length κ to 4 hours and penalty parameter β to 0. For each line, we tested six interval (δ) settings ranging from 0.5 hour to 4 hours. In figure 6f, the x -axis is the run time (in log-scale) and the y -axis is the normalized revenue against the expected revenue of $\delta = 0.5$ within each line. For each line, a data point from left to right corresponds to $\delta = 4, 3, 2, 1.5, 1,$ and 0.5 respectively. Increasing the run time (by decreasing δ) always led to a better solution; however, the quality gain diminished. For example, for the Blue line, it took 20 seconds of additional run time to increase the solution quality from 87.9 percent ($\delta = 4$ hours) to 92.9 percent ($\delta = 3$ hours), whereas it took 1456 seconds of additional run time to increase the solution quality from 99.1 percent ($\delta = 1$ hour) to 100 percent ($\delta = 0.5$ hour).

In the final experiment, we varied the penalty β for using more switches, trading off between the solution quality and the average number of switches. We fixed the patrol length κ to 4 hours and starting time interval δ to one hour. For each line, we tested seven penalty settings from $\beta = 0$ to $\beta = 0.01$. Figure 6g plots the average number of switches against the normalized revenue against the expected revenue of $\beta = 0$ within each line. For all lines, higher β values led to both lower solution quality and fewer number of switches. For example, the average number of switches in the solution of the highest revenue ($\beta = 0$) ranged from 18.6 (Gold line) to 26.7 (Red line). However, by allowing 3 percent quality loss, this number could be lowered to less than 10 for all the four lines.

To further understand the patrol paths returned in these solutions, we show, in figure 6h, the cumulative probability distributions of the number of switches for the Red line given three settings of β : 0, 0.001, and 0.01. Choosing a lower β tended to lead to more complex patrol paths. For example,



Figure 7. Fare Check at a Station Exit.

the solution of $\beta = 0$ used patrol paths whose number of switches is greater than 20 with 68.9 percent probability; the solution of $\beta = 0.001$ (99.7 percent of the optimum) only used such paths with 31.2 percent probability. And the solution of $\beta = 0.01$ (97.0 percent of the optimum) never used patrol paths that had more than 20 switches.

LASD Evaluation of TRUSTS

Since January 2012, LASD has been testing our generated patrol strategies on the Red Line. For example, in initial test runs on Thursday, January 4, and Friday, January 5, 2012, one patrol unit conducted a 4-hour fare-inspection patrol on each day. A total of 851 fare checks were made, with 41 fare evaders cited and felons felons arrested. The patrols implemented in the two days had four and five switches, respectively, and the officers were able to make the switches we requested. The LASD officers have given us valuable feedback and we have incorporated their suggestions, including inserting breaks in the middle of patrol shifts. Dozens of test patrol runs have been carried out so far, and more tests are scheduled in the future to provide a more thorough evaluation of the effectiveness of our strategies. Figure 7 shows two offi-

cers doing fare checks at a station exit during one of the test patrols.

Summary

In this paper we presented TRUSTS, a novel application for fare-evasion deterrence in urban transit systems. Our development of TRUSTS opens the door to applying game-theoretical randomization beyond previous applications of counterterrorism, to a much broader setting in which common individuals and daily routines are involved. We modeled the domain as a Stackelberg game, providing a novel compact representation of the leader's mixed strategies as flows in the history-duplicate transition graph. We found in our simulations that our method computed close-to-optimal strategies, which effectively deterred fare evasion and ensured high levels of revenue with few patrol hours. We are currently evaluating TRUSTS within the LA Metro Rail system in collaboration with LASD.

We briefly mention a couple of future directions. One of the observations from the test runs is that LASD officers often need to deviate from the patrol schedule due to making felony arrests or being

called to deal with emergencies at other parts of the train system. After dealing with the emergency, which could take a significant amount of time, it is often difficult to catch up to the original patrol schedule. We are currently investigating more efficient ways of utilizing resources in case of such emergencies. Another future direction is to investigate bounded rational models of the passengers, since they are human decision makers and thus might not be perfect optimizers. Incorporating such behavior models of the adversary in security games (Pita et al. 2010; Yang et al. 2011) may potentially increase the robustness of our solutions. And we ultimately want to combine fare checking with LASD's other tasks of crime prevention and counterterrorism. We plan to make use of multiobjective optimization techniques to intelligently allocate resources that makes trade-offs among these different objectives.

Acknowledgements

This research is supported by MURI grant W911NF-11-1-0332. Tuomas Sandholm was funded by the National Science Foundation under grants IIS-0905390, IIS-0964579, and CCF-1101668. Kevin Leyton-Brown was funded by the Natural Sciences and Engineering Research Council of Canada under the Discovery Grants program.

Notes

1. See the Booz Allen Hamilton 2007 report, Faregating Analysis commissioned by LA Metro (boardarchives.metro.net/Items/2007/11/November/20071115EMACItem2%7.pdf)
2. See the 2002 Horizon Research Corporation Metropolitan Transit Authority Fare Evasion Study. (libraryarchives.metro.net/DPGTL/studies/2002/horizon/fare%evasio%n%study.pdf)

References

Alpern, S. 1992. Infiltration Games on Arbitrary Graphs. *Journal of Mathematical Analysis and Applications* 163(1): 286–288.

Becker, G., and Landes, W. 1974. *Essays in the Economics of Crime and Punishment*. New York: Columbia University Press.

Conitzer, V., and Sandholm, T. 2006. Computing the Optimal Strategy to Commit To. In *Proceedings of the 7th ACM Conference on Electronic Commerce*. New York: Association for Computing Machinery.

Dickerson, J. P.; Simari, G. I.; Subrahmanian, V. S.; and Kraus, S. 2010. A Graph-Theoretic Approach to Protect Static and Moving Targets from Adversaries. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Gal, S. 1979. Search Games with Mobile and Immobile Hider. *SIAM Journal on Control and Optimization* 17(1): 99–122.

Halvorson, E.; Conitzer, V.; and Parr, R. 2009. Multi-Step Multi-Sensor Hider-Seeker Games. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*. Menlo Park, CA: AAAI Press.

Jain, M.; Tsai, J.; Pita, J.; Kiekintveld, C.; Rathi, S.; Tambe, M.; and Ordóñez, F. 2010. Software Assistants for Randomized Patrol Planning for the LAX Airport Police and the Federal Air Marshals Service. *Interfaces* 40(4): 267–290.

Koller, D.; Megiddo, N.; and von Stengel, B. 1994. Fast Algorithms for Finding Randomized Strategies in Game Trees. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing*. New York: Association for Computing Machinery.

Pita, J.; Jain, M.; Tambe, M.; Ordóñez, F.; and Kraus, S. 2010. Robust Solutions to Stackelberg Games: Addressing Bounded Rationality and Limited Observations in Human Cognition. *Artificial Intelligence Journal* 174(15): 1142–1171.

Pita, J.; Tambe, M.; Kiekintveld, C.; Cullen, S.; and Steigerwald, E. 2011. Guards — Game Theoretic Security Allocation on a National Scale. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Ponssard, J., and Sorin, S. 1980. The L.P. Formulation of Finite Zero-Sum Games with Incomplete Information. *International Journal of Game Theory* 9(2): 99–105.

Tambe, M. 2011. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. New York: Cambridge University Press.

Tsai, J.; Yin, Z.; Kwak, J.; Kempe, D.; Kiekintveld, C.; and Tambe, M. 2010. Urban Security: Game-Theoretic Resource Allocation in Networked Physical Domains. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*. Menlo Park, CA: AAAI Press.

Wagenaar, W. A. 1972. Generation of Random Sequences by Human Subjects: A Critical Survey of Literature. *Psychological Bulletin* 77(1): 65–72.

Yang, R.; Kiekintveld, C.; Ordóñez, F.; Tambe, M.; and John, R. 2011. Improved Computational Models of Human Behavior in Security Games. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Yin, Z.; Jiang, A. X.; Johnson, M.; Tambe, M.; Kiekintveld, C.; Leyton-Brown, K.; Sandholm, T.; and Sullivan, J. 2012. TRUSTS: Scheduling Randomized Patrols for Fare Inspection in Transit Systems. In *Proceedings of the Twenty-Fourth AAAI Conference on Innovative Applications of Artificial Intelligence (IAAI)*. Menlo Park, CA: AAAI Press.

Zhengyu Yin is a Ph.D. student in the Department of Computer Science at the University of Southern California. His research interests include algorithmic game theory and robust optimization, with applications to security domains. He has published and presented at AAMAS, IJCAI, and AAAI, and others, and published in JAIR and others.

Albert Xin Jiang is a postdoctoral researcher in the Department of Computer Science at the University of Southern California. Much of his research is addressing



AAAI to Colocate with Cognitive Science Society in 2014!

AAAI is pleased to announce that it will colocate with the 2014 Cognitive Science Society Conference in picturesque Québec City, Québec, July 27-31, 2014. The conference will be held at the beautiful Centre des congrès de Québec, and attendees can stay at the adjacent Hilton Québec. More details will be available at the AAAI website soon!

computational problems arising in game theory, including the efficient computation of solution concepts. In 2011 he received the best student paper award at the ACM Conference on Electronic Commerce. He received the 2012 CAIAC Doctoral Dissertation Award and the runner up IFAAMAS-11 Victor Lesser Distinguished Dissertation Award.

Milind Tambe is Helen N. and Emmett H. Jones Professor in Engineering and a professor of computer science and industrial and systems engineering at the University of Southern California (USC). He leads the TEAMCORE Research Group at USC, with research focused on agent-based and multiagent systems. He is a fellow of AAAI and recipient of the ACM/SIGART Autonomous Agents Research Award and the Christopher Columbus Fellowship Foundation Homeland security award. In addition, he is also the recipient of the "influential paper award" from the International Foundation for Agents and Multi-agent Systems.

Christopher Kiekintveld is an assistant professor at the University of Texas at El Paso (UTEP). His research focuses on intelligent agents and decision making, including computational game theory with applications to security and trading agents.

Kevin Leyton-Brown is an associate professor in com-

puter science at the University of British Columbia. He works at the intersection of computer science and micro-economics, addressing computational problems in economic contexts and incentive issues in multiagent systems.

Tuomas Sandholm is a professor in the Computer Science Department at Carnegie Mellon University. He was founder, chairman, and CTO/chief scientist of CombineNet, Inc. from 1997 until its acquisition in 2010. He holds a Ph.D. and MS in computer science and an MS (BS included) with distinction in industrial engineering and management science. He is recipient of the NSF Career Award, the inaugural ACM Autonomous Agents Research Award, the Alfred P. Sloan Foundation Fellowship, the Carnegie Science Center Award for Excellence, and the Computers and Thought Award. He is Fellow of ACM and AAAI.

John P. Sullivan is a career police officer. He currently serves as a lieutenant with the Los Angeles Sheriff's Department. He is coeditor of *Countering Terrorism and WMD: Creating a Global Counter-Terrorism Network* (Routledge, 2006) and *Global Biosecurity: Threats and Responses* (Routledge, 2010). He is coauthor of *Policing Transportation Facilities* (Charles C. Thomas, 1994), and *Emergency Preparedness for Transit Terrorism* (Transportation Research Board, 1997).