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

Paper by Yin et al

Christiana Sabett
Applied Mathematics, Applied Statistics, & Scientific Computation
University of Maryland, College Park

February 27, 2018

Outline

TRUSTS:
Tactical Randomization for Urban Security in Transit Systems

Introduction

TRUSTS:
Tactical Randomization for Urban Security in Transit Systems

Figure: My brother Joey.

Introduction

TRUSTS:
Tactical Randomization for Urban Security in Transit Systems

Figure: A southbound light rail car passing through Linthicum, MD on its way to BWI Airport.

Motivation

- Often transit systems have fare systems where passengers required to buy tickets, but no existing gates/turnstiles; instead, compliance relies on patrols.

Motivation

- Often transit systems have fare systems where passengers required to buy tickets, but no existing gates/turnstiles; instead, compliance relies on patrols.
- Cost-effective measure: thought that revenue loss due to fare evasion less than would otherwise cost to guarantee every passenger buys a ticket.

Motivation

- Often transit systems have fare systems where passengers required to buy tickets, but no existing gates/turnstiles; instead, compliance relies on patrols.
- Cost-effective measure: thought that revenue loss due to fare evasion less than would otherwise cost to guarantee every passenger buys a ticket.
- Need mechanism to assign inspection locations and times, and do so randomly to avoid exploitation due to predictability.

Motivation

- Often transit systems have fare systems where passengers required to buy tickets, but no existing gates/turnstiles; instead, compliance relies on patrols.
- Cost-effective measure: thought that revenue loss due to fare evasion less than would otherwise cost to guarantee every passenger buys a ticket.
- Need mechanism to assign inspection locations and times, and do so randomly to avoid exploitation due to predictability.
- Heavy constraints to consider: train timings, switching between trains, schedule lengths, etc.

Motivation

- Often transit systems have fare systems where passengers required to buy tickets, but no existing gates/turnstiles; instead, compliance relies on patrols.
- Cost-effective measure: thought that revenue loss due to fare evasion less than would otherwise cost to guarantee every passenger buys a ticket.
- Need mechanism to assign inspection locations and times, and do so randomly to avoid exploitation due to predictability.
- Heavy constraints to consider: train timings, switching between trains, schedule lengths, etc.
- TRUSTS is a method for scheduling randomized patrols to inspect transit fares in order to effectively mitigate losses due to fare evasion.

Overview

- TRUSTS currently implemented in LA Metro Rail system.

Overview

- TRUSTS currently implemented in LA Metro Rail system.
- Objective: maximize revenue, including revenue from both ticket sales and fines for fare evasion.

Overview

- TRUSTS currently implemented in LA Metro Rail system.
- Objective: maximize revenue, including revenue from both ticket sales and fines for fare evasion.
- Modeled as a Stackelberg game.
  - One leader: LA Sheriffs Department (LASD).
  - Followers: Metro riders.
  - Leader precommits to mixed patrol strategy; riders decide whether or not to buy a ticket based on that strategy (to minimize their individual cost).

Overview

- TRUSTS currently implemented in LA Metro Rail system.
- Objective: maximize revenue, including revenue from both ticket sales and fines for fare evasion.
- Modeled as a Stackelberg game.
    - One leader: LA Sheriffs Department (LASD).
    - Followers: Metro riders.
    - Leader precommits to mixed patrol strategy; riders decide whether or not to buy a ticket based on that strategy (to minimize their individual cost).
- Problem solved as an LP for optimal flow through a transition graph.

Overview

- TRUSTS currently implemented in LA Metro Rail system.
- Objective: maximize revenue, including revenue from both ticket sales and fines for fare evasion.
- Modeled as a Stackelberg game.
  - One leader: LA Sheriffs Department (LASD).
  - Followers: Metro riders.
  - Leader precommits to mixed patrol strategy; riders decide whether or not to buy a ticket based on that strategy (to minimize their individual cost).
- Problem solved as an LP for optimal flow through a transition graph.
- Added considerations include
  - Length of patrols (avoid patrols that are too long).
  - Train switching frequency (avoid patrols that require difficulty of switching trains).

Problem Setup

- Four LA Metro lines, assumed independent for simplicity.

Problem Setup

- Four LA Metro lines, assumed independent for simplicity.
- Pure leader strategy: sequence of patrol actions of constant bounded duration.

Problem Setup

- Four LA Metro lines, assumed independent for simplicity.
- Pure leader strategy: sequence of patrol actions of constant bounded duration.
- Possible pure follower strategies: buying or not buying.

Problem Setup

- Four LA Metro lines, assumed independent for simplicity.
- Pure leader strategy: sequence of patrol actions of constant bounded duration.
- Possible pure follower strategies: buying or not buying.
- Assumptions:
  1. Train (and rider) paths move in one direction, therefore a train (or rider) does not return to a previous station for a given path duration.
  2. Riders are daily commuters who take a fixed route at a fixed time.
  3. Given (2), riders know the inspection probability perfectly.

Train System

Let

$$G = \langle V, E \rangle \tag{1}$$

be a directed transition graph representing a single train line with discrete time steps.

Train System

Let

$$G = \langle V, E \rangle \tag{1}$$

be a directed transition graph representing a single train line with discrete time steps.

Vertex $v = \langle s, t \rangle$ corresponds to some station/time pair. For edge $e \in E$, e connects two vertices $\langle s, t \rangle$ and $\langle s', t' \rangle$ if a possible train action exists between them, i.e.

1. Traveling action: WLOG, s and s' are adjacent in the station sequence and $\langle s, t \rangle$ and $\langle s', t' \rangle$ are consecutive stops for some train in the schedule
2. Staying action: $s = s', t < t'$ and $\nexists \langle s, t'' \rangle$ such that $t < t'' < t'$

Example

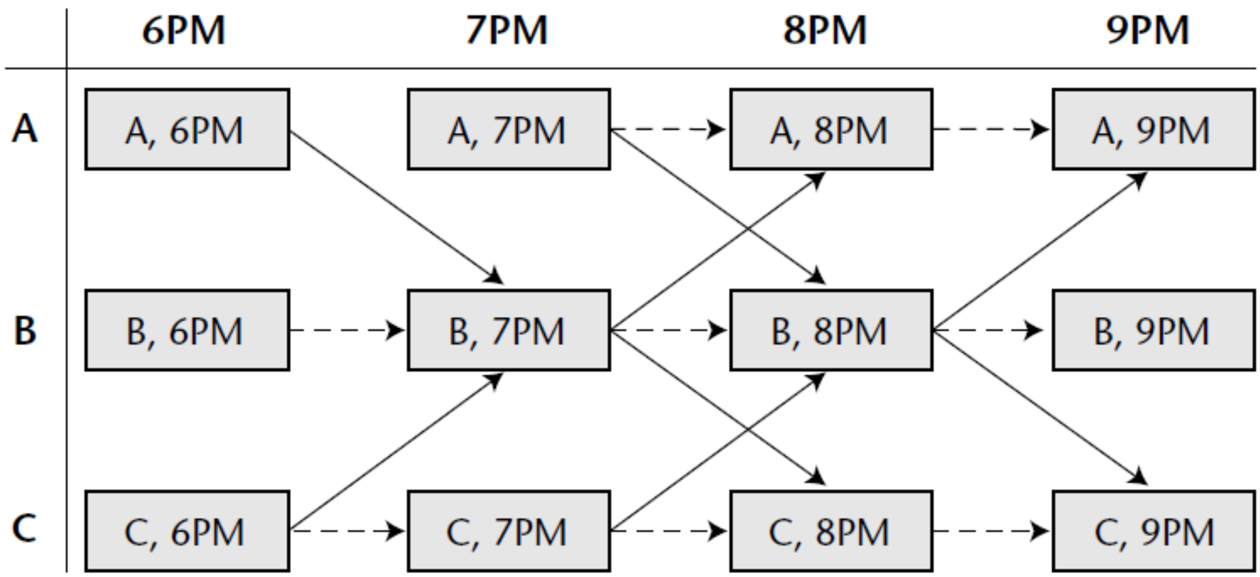

Figure: A train system with three stations and four discrete time points. Dashed lines represent staying actions, solid lines represent traveling actions.

Patrols

Let γ be a (fixed) number of deployable patrol units that can be scheduled for at most κ hours.

Patrols

Let γ be a (fixed) number of deployable patrol units that can be scheduled for at most κ hours.

Patrol units can take one of two actions, occurring over one edge with (duration) length l_e:

1. on-train inspections
2. in-station inspections

Patrols

Let γ be a (fixed) number of deployable patrol units that can be scheduled for at most κ hours.

Patrol units can take one of two actions, occurring over one edge with (duration) length l_e:

1. on-train inspections
2. in-station inspections

Let f_e denote the effectiveness value of edge e, the percentage of riders inspected over that edge.

Patrols

Let γ be a (fixed) number of deployable patrol units that can be scheduled for at most κ hours.

Patrol units can take one of two actions, occurring over one edge with (duration) length l_e:

1. on-train inspections
2. in-station inspections

Let f_e denote the effectiveness value of edge e, the percentage of riders inspected over that edge.

Let $P = [P_1 \ldots P_\gamma]^T$ represent a valid pure patrol strategy, where each path P_i is of size at most κ.

Example Cont'd

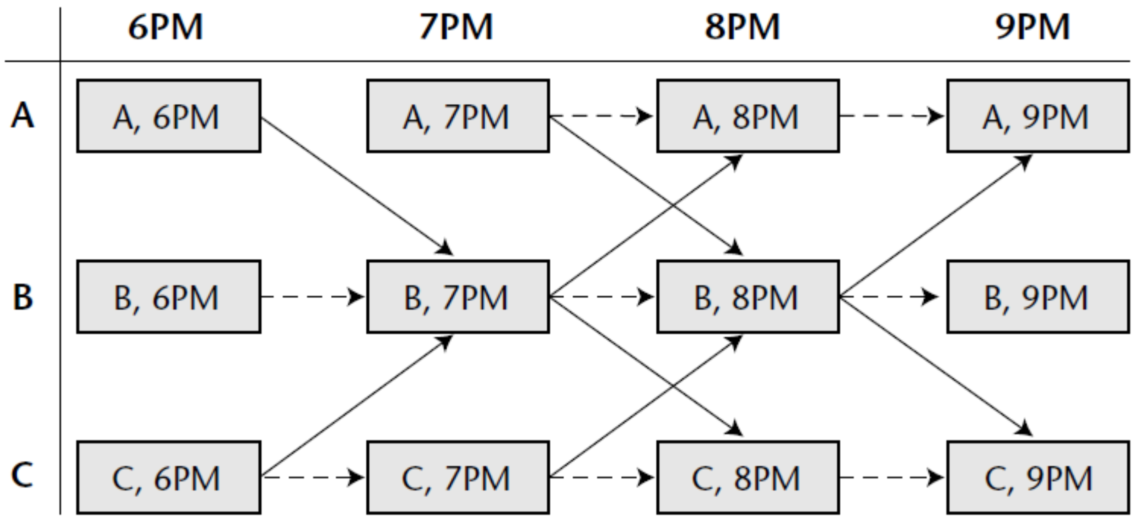

Figure: A train system with three stations and four discrete time points. Dashed lines represent staying actions, solid lines represent traveling actions.

Suppose the number of patrols $\gamma = 1$ with patrol duration $\kappa = 2$.

Example Cont'd

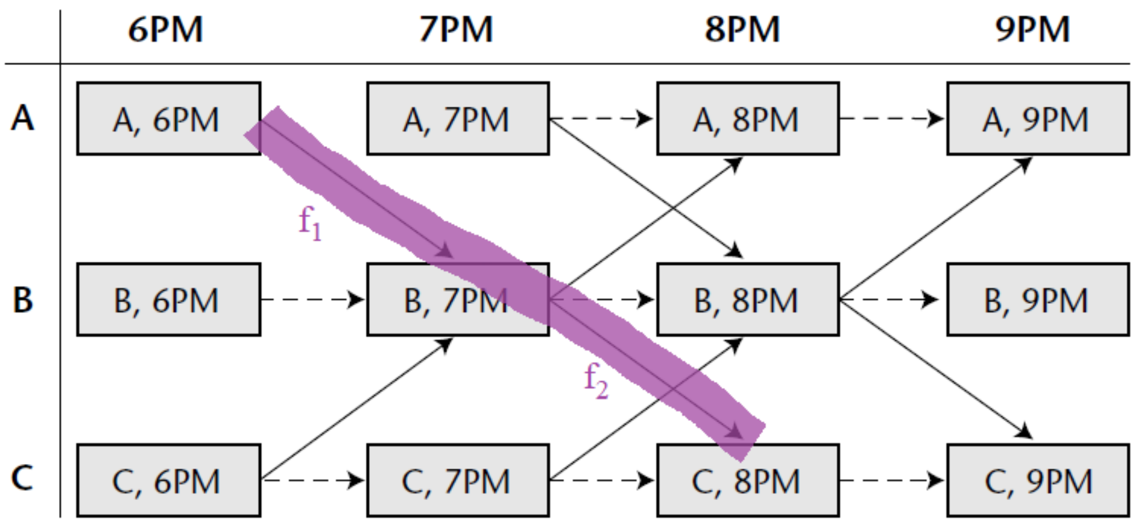

Figure: A train system with three stations and four discrete time points. Dashed lines represent staying actions, solid lines represent traveling actions.

Suppose the number of patrols $\gamma = 1$ with patrol duration $\kappa = 2$. The purple path represents a valid pure patrol strategy. The set of pure patrol strategies consists of all paths of length 2.

Riders

Let ρ be the ticket price, and τ be the fine for fare evasion, $\rho << \tau$.

Riders

Let ρ be the ticket price, and τ be the fine for fare evasion, $\rho << \tau$.

Let rider type λ be the path he or she takes in the graph.

Riders

Let ρ be the ticket price, and τ be the fine for fare evasion, $\rho << \tau$.

Let rider type λ be the path he or she takes in the graph.

Assumptions:

- Riders never follow any “stay” edges mid-ride, because there is only one train line.
- Every rider type ends with a “stay” edge that represents the rider exiting the station (during which they could be inspected).

Riders

Let ρ be the ticket price, and τ be the fine for fare evasion, $\rho << \tau$.

Let rider type λ be the path he or she takes in the graph.

Assumptions:

- Riders never follow any "stay" edges mid-ride, because there is only one train line.
- Every rider type ends with a "stay" edge that represents the rider exiting the station (during which they could be inspected).

The space Λ of rider types corresponds to the set of all subpaths of train paths.

Example Cont'd

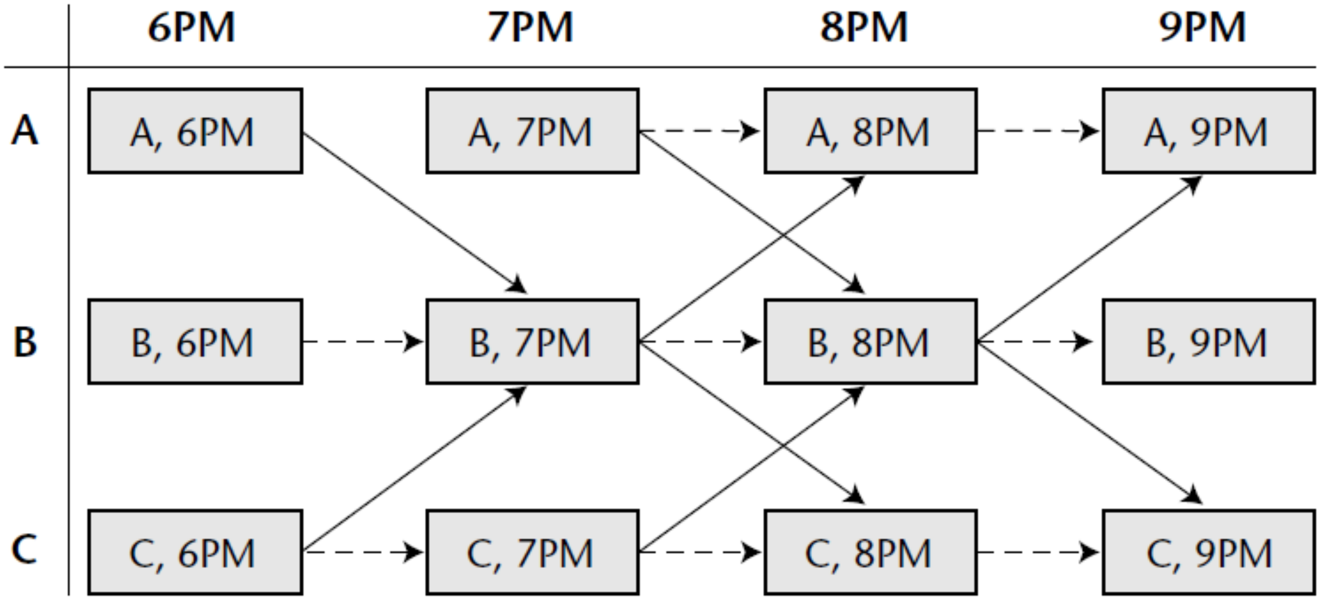

Figure: A train system with four trains, three stations, and four discrete time points. Dashed lines represent staying actions, solid lines represent traveling actions.

Example Cont'd

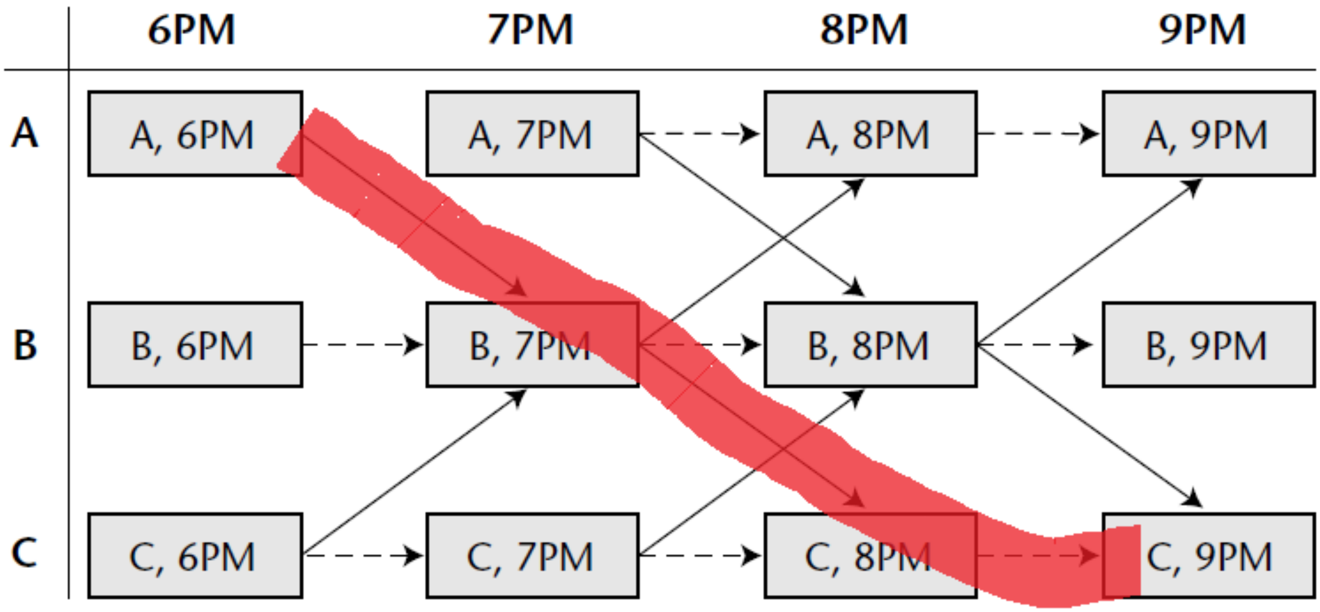

Figure: A train system with four trains, three stations, and four discrete time points. Dashed lines represent staying actions, solid lines represent traveling actions.

Example Cont'd

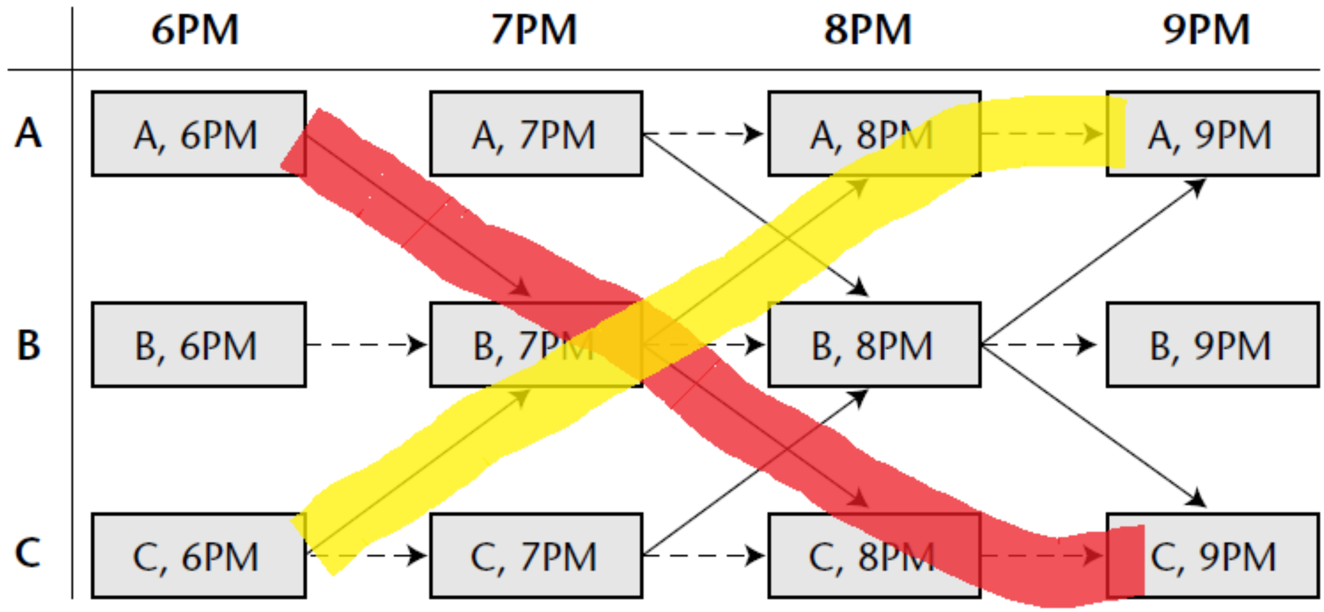

Figure: A train system with four trains, three stations, and four discrete time points. Dashed lines represent staying actions, solid lines represent traveling actions.

Example Cont'd

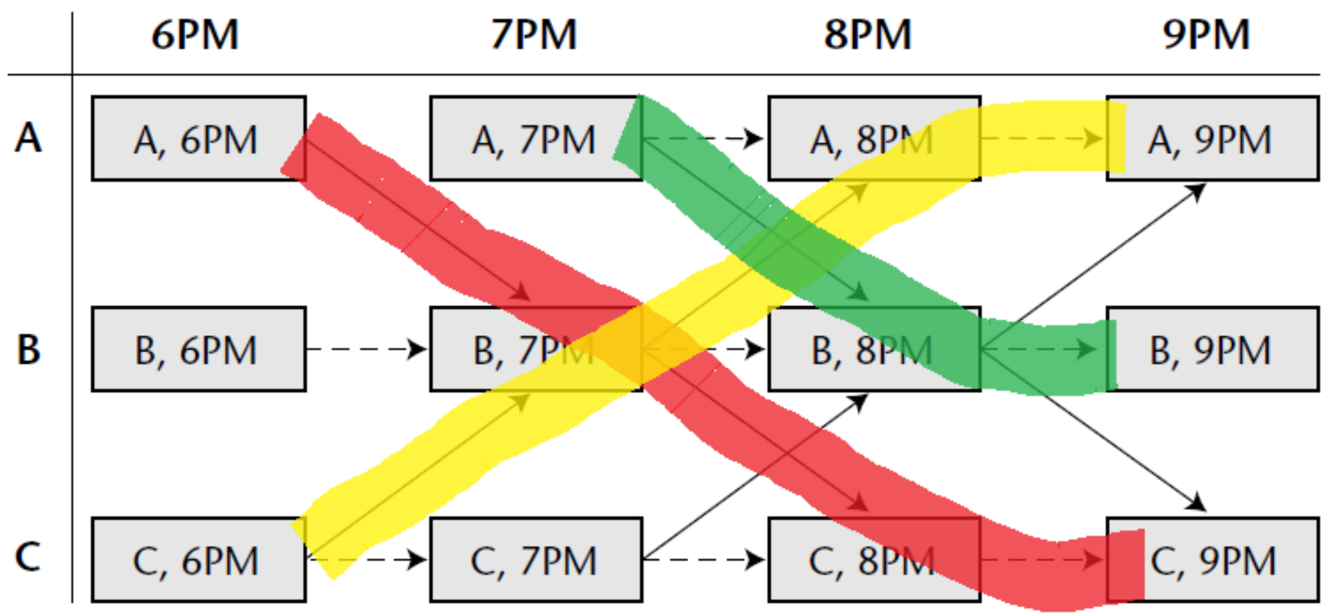

Figure: A train system with four trains, three stations, and four discrete time points. Dashed lines represent staying actions, solid lines represent traveling actions.

Example Cont'd

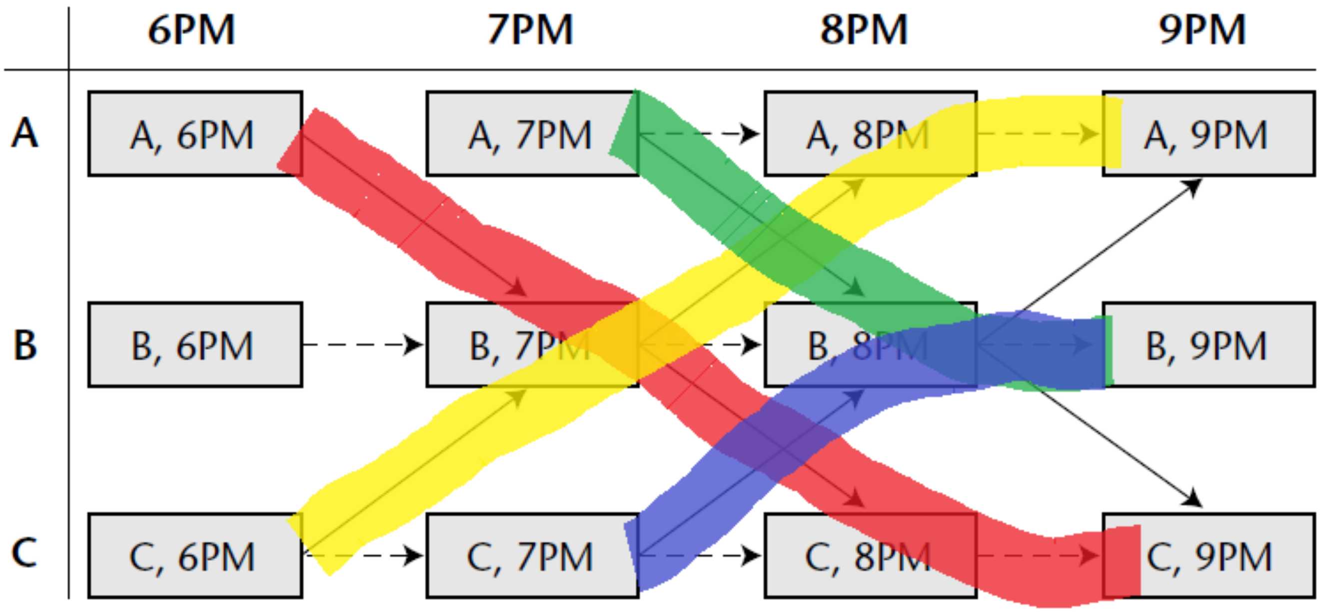

Figure: A train system with four trains, three stations, and four discrete time points. Dashed lines represent staying actions, solid lines represent traveling actions.

Example Cont'd

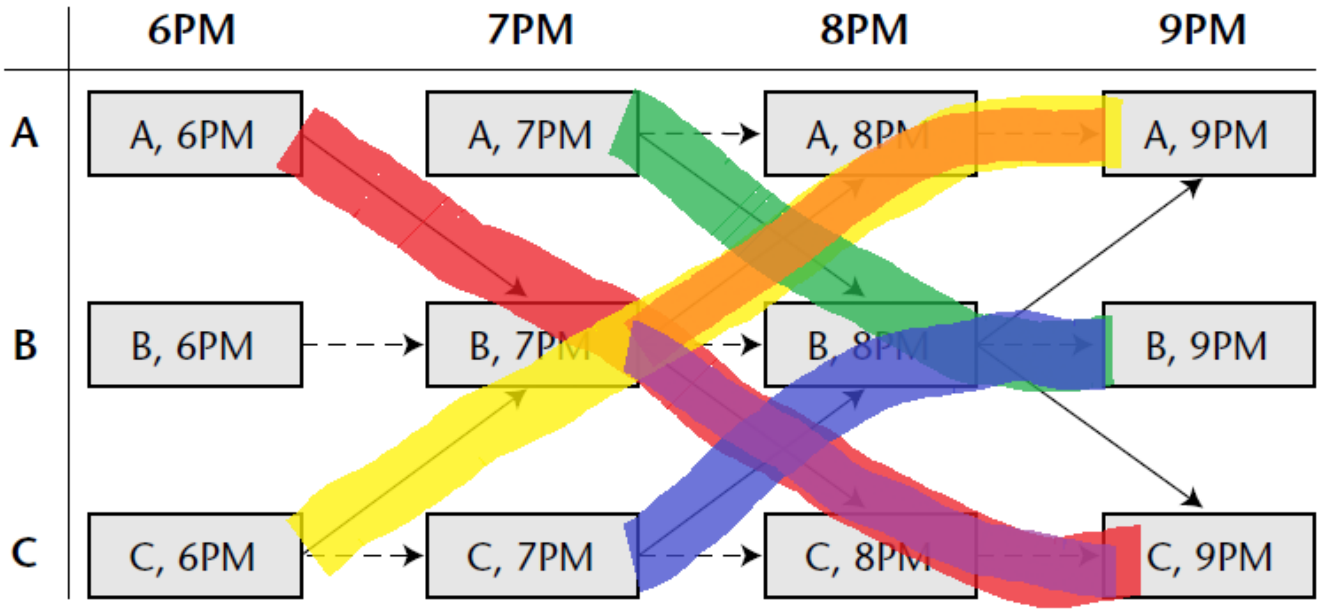

Figure: A train system with four trains, three stations, and four discrete time points. Dashed lines represent staying actions, solid lines represent traveling actions.

Riders Cont'd

Given pure patrol strategy $[P_1 \ldots P_\gamma]^T$, the inspection probability p_l for a rider of type $\lambda \in \Lambda$ is

$$p_l = \min\{1, \sum_{i=1}^{\gamma} \sum_{e \in P_i \cup \lambda} f_e\} \tag{2}$$

Riders Cont'd

Given pure patrol strategy $[P_1 \ldots P_\gamma]^T$, the inspection probability p_l for a rider of type $\lambda \in \Lambda$ is

$$p_l = \min\{1, \sum_{i=1}^{\gamma} \sum_{e \in P_i \cup \lambda} f_e\} \tag{2}$$

The expected utility $U(r)$ for rider r is therefore

$$U(r) = \begin{cases} -\rho, & \text{r buys a ticket} \\ -\tau * p_l, & \text{r is caught not buying a ticket} \end{cases}$$

Riders Cont'd

Given pure patrol strategy $[P_1 \ldots P_\gamma]^T$, the inspection probability p_l for a rider of type $\lambda \in \Lambda$ is

$$p_l = \min\{1, \sum_{i=1}^{\gamma} \sum_{e \in P_i \cup \lambda} f_e\} \quad (2)$$

The expected utility $U(r)$ for rider r is therefore

$$U(r) = \begin{cases} -\rho, & \text{r buys a ticket} \\ -\tau * p_l, & \text{r is caught not buying a ticket} \end{cases}$$

(Note this is the exactly the negative of revenue collected by the leader.)

Objective

- Leader utility $U(\ell)$ is total expected revenue, which can be broken down by interactions with individual riders.

Objective

- Leader utility $U(\ell)$ is total expected revenue, which can be broken down by interactions with individual riders.
- Riders do not affect each others' utilities, and each rider's type is known to the follower but not to the leader.

Objective

- Leader utility $U(\ell)$ is total expected revenue, which can be broken down by interactions with individual riders.
- Riders do not affect each others' utilities, and each rider's type is known to the follower but not to the leader.
- Problem can be reduced to a two-player Bayesian Stackelberg game.

Objective

- Leader utility $U(\ell)$ is total expected revenue, which can be broken down by interactions with individual riders.
- Riders do not affect each others' utilities, and each rider's type is known to the follower but not to the leader.
- Problem can be reduced to a two-player Bayesian Stackelberg game.
- For a zero-sum Bayesian game, the Stackelberg equilibrium is equivalent to the maximum solution.

Objective

- Leader utility $U(\ell)$ is total expected revenue, which can be broken down by interactions with individual riders.
- Riders do not affect each others' utilities, and each rider's type is known to the follower but not to the leader.
- Problem can be reduced to a two-player Bayesian Stackelberg game.
- For a zero-sum Bayesian game, the Stackelberg equilibrium is equivalent to the maximum solution.
- These LPs require explicit enumeration of pure leader strategies; unrealistic for this problem because the space of pure leader strategies is exponentially large.

Basic LP Formulation

For convenience, add source v^+ and sink v^- to G. Additional edges have zero duration and zero effectiveness. Expected total number of time units used by patrols must be bounded by $\gamma \cdot \kappa$.

Basic LP Formulation

For convenience, add source v^+ and sink v^- to G. Additional edges have zero duration and zero effectiveness. Expected total number of time units used by patrols must be bounded by $\gamma \cdot \kappa$.

Let x_e be the expected number of inspections on edge e. Denote the vector $x = [x_e]$ of marginal coverage over every edge $e \in E$ the marginal strategy.

Basic LP Formulation

For convenience, add source v^+ and sink v^- to G. Additional edges have zero duration and zero effectiveness. Expected total number of time units used by patrols must be bounded by $\gamma \cdot \kappa$.

Let x_e be the expected number of inspections on edge e. Denote the vector $x = [x_e]$ of marginal coverage over every edge $e \in E$ the marginal strategy.

Constraints on $\mathbf{x}$:

- Total flow entering and exiting the system bounded by γ.
- Flow into and out of intermediate vertices must be equal.

Basic LP Formulation

$$\max_{\mathbf{x},\mathbf{u}} \sum_{\lambda \in \Lambda} p_\lambda u_\lambda \tag{2}$$

$$\text{s.t. } u_\lambda \leq \min\{\rho,\ \tau \sum_{e \in \lambda} x_e f_e\}, \text{ for all } \lambda \in \Lambda \tag{3}$$

$$\sum_{v \in V^+} x_{(v^+, v)} = \sum_{v \in V^-} x_{(v, v^-)} \leq \gamma \tag{4}$$

$$\sum_{(v', v) \in E} x_{(v', v)} = \sum_{(v, v^\dagger) \in E} x_{(v, v^\dagger)}, \text{ for all } v \in V \tag{5}$$

$$\sum_{e \in E} l_e \cdot x_e \leq \gamma \cdot \kappa, 0 \leq x_e \leq \alpha, \forall e \in E \tag{6}$$

Marginal Representation Example

Problems with the basic formulation:

- Computed strategy x^* may not satisfy patrol length limit κ.
- x^* may switch too often between trains or types of inspections.

Marginal Representation Example

Problems with the basic formulation:

- Computed strategy x^* may not satisfy patrol length limit κ.
- x^* may switch too often between trains or types of inspections.

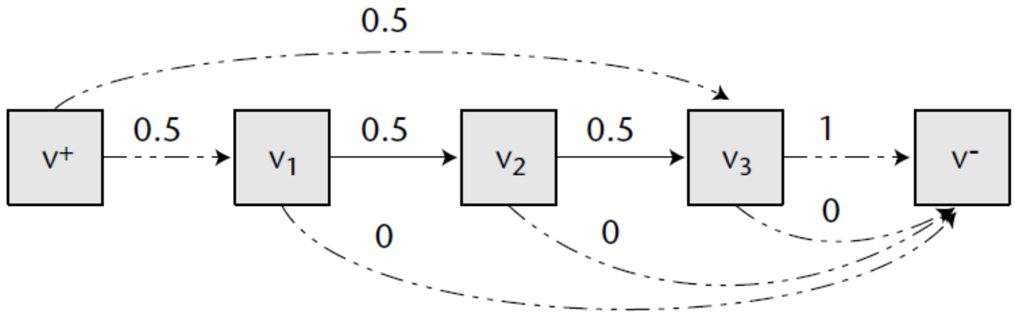

Figure: An infeasible marginal strategy. Each real edge has duration 1. Assume $\gamma = 1$ and $\kappa = 1$. $\mathbf{x} = [0.5\ 0.5]^T$ satisfies the given flow constraints. Corresponding mixed strategy: Take either $v^+ \to v_3 \to v^-$ or $v^+ \to v_1 \to v_2 \to v_3 \to v^-$ with 50% probability. Expected time units spent is 0.5*0 + 0.5*(1+1) = 1, but second patrol strategy has duration $2 > \kappa$.

Extended LP Formulation

Construct a history-duplicate transition (HDT) graph to store path information, in order to impose constraints on the optimal marginal strategy:

1. Create copies of subgraphs of G based on different starting times. For starting time t^*, keep the subgraph on vertices $v = \langle s, t \rangle \in V$ where $t^* \leq t \leq t^* + \kappa$.

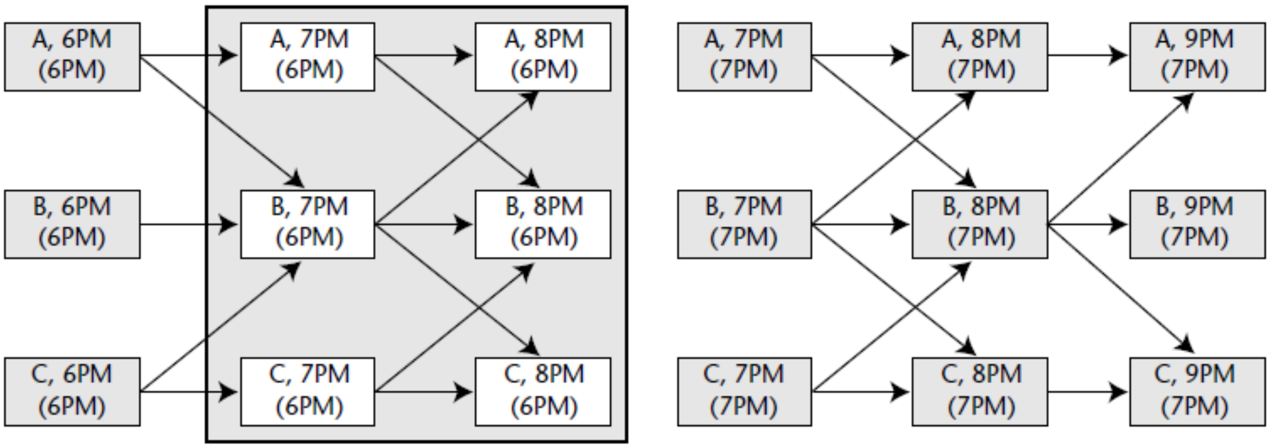

Figure: HDT graph for $\kappa = 2$ with two starting time points, 6pm and 7pm.

Extended LP Formulation

Construct a history-duplicate transition (HDT) graph to store path information, in order to impose constraints on the optimal marginal strategy:

2. For each $v \in V$ with inflow, create a copy of it corresponding to an edge that leads to it. Impose a penalty β for using switching edges in the marginal strategy.

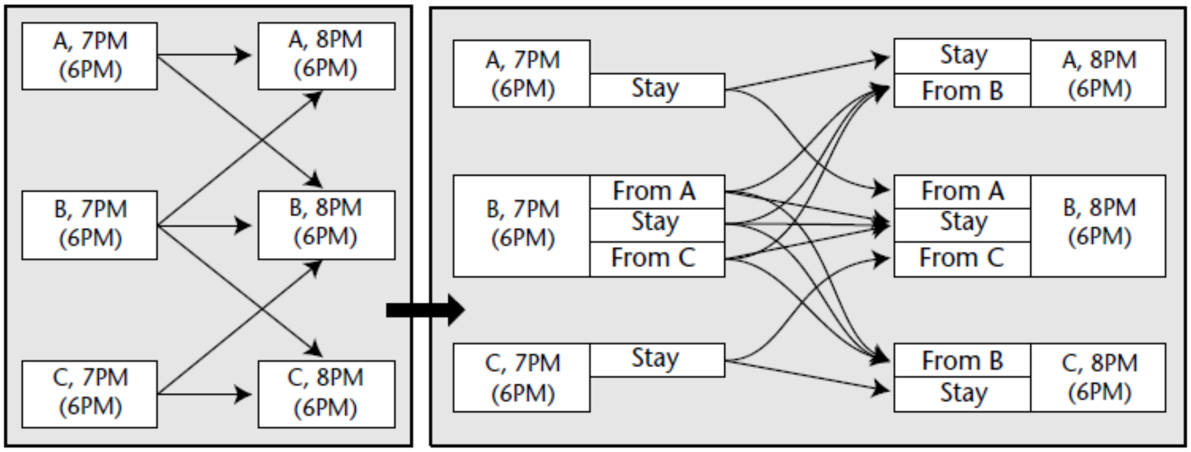

Figure: HDT graph for $\kappa = 2$ with two starting time points, 6pm and 7pm.

Extended LP Formulation

$$\max_{\mathbf{x},\mathbf{y},\mathbf{u}} \sum_{\lambda\in\Lambda} p_\lambda u_\lambda - \beta \sum_{e\in\mathcal{E}} c_e y_e \tag{7}$$

$$\text{s.t. } u_\lambda \le \min\{\rho,\ \tau \sum_{e\in\lambda} x_e f_e\}, \text{ for all } \lambda \in \Lambda \tag{8}$$

$$\sum_{v\in\mathcal{V}^+} y_{(v^+,v)} = \sum_{v\in\mathcal{V}^-} y_{(v,v^-)} \le \gamma \tag{9}$$

$$\sum_{(v',v)\in\mathcal{E}} y_{(v',v)} = \sum_{(v,v^\dagger)\in\mathcal{E}} y_{(v,v^\dagger)}, \text{ for all } v \in \mathcal{V} \tag{10}$$

$$x_e = \sum_{e'\in\Gamma(e)} y_{e'}, \forall e \in E, 0 \le x_e \le \alpha, \forall e \in E \tag{11}$$

Results

- Assumed potential fare evaders evenly distributed among the general population.

Results

- Assumed potential fare evaders evenly distributed among the general population.
- Assigned effectiveness value f_e based on assumption of 10 passengers inspected per minute, with f_e capped at 0.5 because inspectors on trains cannot switch between moving cars.

Results

- Assumed potential fare evaders evenly distributed among the general population.
- Assigned effectiveness value f_e based on assumption of 10 passengers inspected per minute, with f_e capped at 0.5 because inspectors on trains cannot switch between moving cars.
- Ticket fare set to \$1.50, fine set to \$100.

Results

- Assumed potential fare evaders evenly distributed among the general population.
- Assigned effectiveness value f_e based on assumption of 10 passengers inspected per minute, with f_e capped at 0.5 because inspectors on trains cannot switch between moving cars.
- Ticket fare set to \$1.50, fine set to \$100.
- $\gamma = 1$ in all experiments.

Results

- Assumed potential fare evaders evenly distributed among the general population.
- Assigned effectiveness value f_e based on assumption of 10 passengers inspected per minute, with f_e capped at 0.5 because inspectors on trains cannot switch between moving cars.
- Ticket fare set to \$1.50, fine set to \$100.
- $\gamma = 1$ in all experiments.
- First set of experiments: penalty β fixed at 0, κ ranged from four to seven hours. HDT graph had one starting time point every hour.

Results

- Assumed potential fare evaders evenly distributed among the general population.
- Assigned effectiveness value f_e based on assumption of 10 passengers inspected per minute, with f_e capped at 0.5 because inspectors on trains cannot switch between moving cars.
- Ticket fare set to \$1.50, fine set to \$100.
- $\gamma = 1$ in all experiments.
- First set of experiments: penalty β fixed at 0, κ ranged from four to seven hours. HDT graph had one starting time point every hour.
- Second experiments: penalty β fixed at 0, κ fixed at four hours. Intervals of starting time points varied, $0 \leq \delta \leq 4$.

Results

- Assumed potential fare evaders evenly distributed among the general population.
- Assigned effectiveness value f_e based on assumption of 10 passengers inspected per minute, with f_e capped at 0.5 because inspectors on trains cannot switch between moving cars.
- Ticket fare set to \$1.50, fine set to \$100.
- $\gamma = 1$ in all experiments.
- First set of experiments: penalty β fixed at 0, κ ranged from four to seven hours. HDT graph had one starting time point every hour.
- Second experiments: penalty β fixed at 0, κ fixed at four hours. Intervals of starting time points varied, $0 \leq \delta \leq 4$.
- Third experiment: κ fixed at four hours, starting time point interval δ set to one. Penalty β varied.

Results

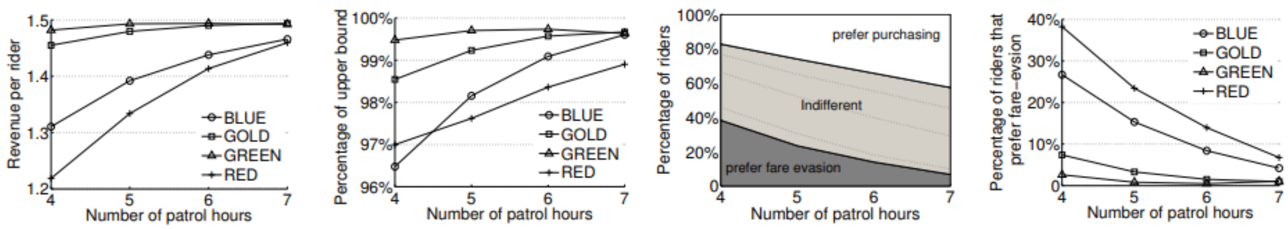

(a) Per passenger revenue of the computed mixed strategy.

(b) Percentage vs. upper bound.

(c) Evasion tendency distribution of Red line.

(d) Percentage of riders that prefer fare evasion

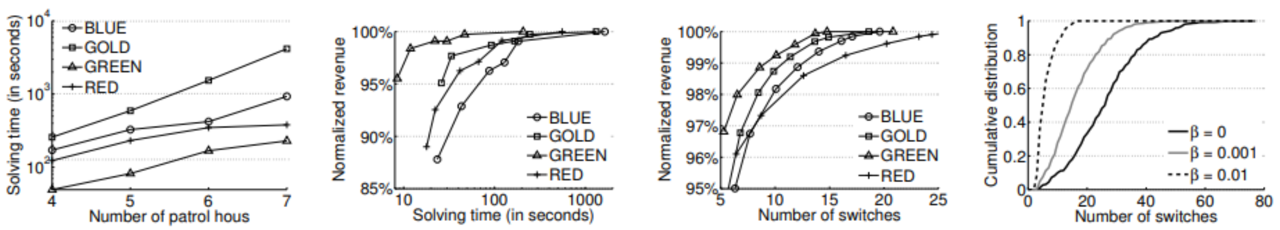

(e) Runtime of solving the LP by CPLEX.

(f) Tradeoffs between optimality and runtime.

(g) Tradeoffs between optimality and patrol preference.

(h) Cumulative probability distribution of the number of switches for the Red line.