Artificial Intelligence Methods for Social Good

Lecture 2-5: Designing Pricing Scheme for Ridesharing Platform

08-537 (9-unit) and 08-737 (12-unit)

Instructor: Fei Fang
feifang@cmu.edu
Wean Hall 4126

Outline

- Domain and Challenges
- Background
  - Linear Program Duality
  - Competitive Equilibrium
  - Minimum-cost Flow Problem
- Spatio-Temporal Pricing for Ridesharing Platforms

Learning Objectives

- Understand the concept of
  - Competitive Equilibrium
  - LP duality, Complementary slackness
- Can write down
  - The dual of any given LP
  - The LP of finding an CE as an LP in various settings
- Can describe
  - Desirable properties of pricing mechanism in ridesharing platforms
  - Key ideas in the design of spatio-temporal pricing mechanism

Pricing Models

- Classic pricing model in Taxi industry
  - $price = w_0 + w_1 l_t + w_2 l_d$
- Surge pricing in ridesharing platform
  - Increase price when demand > supply

Pricing Models

- Why surge pricing?

Evolution of Surge Pricing

- Surge price interface

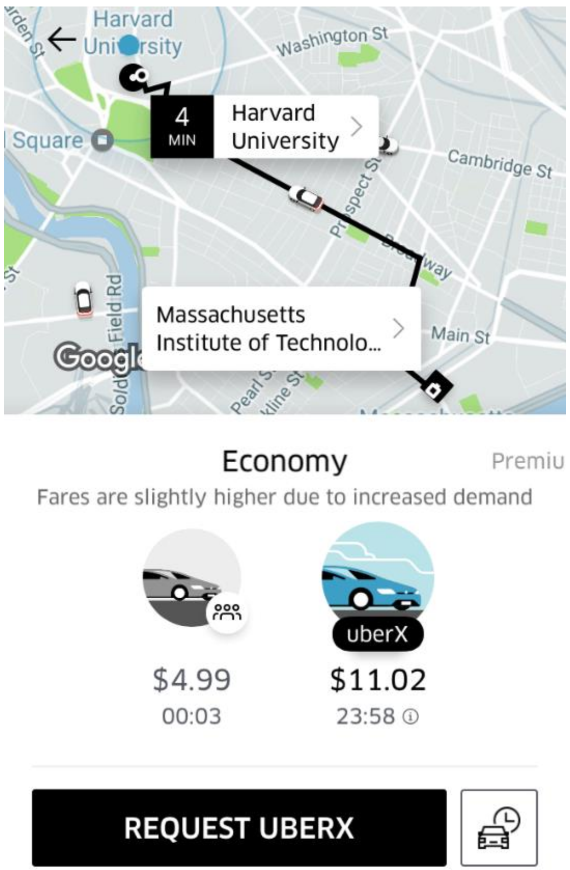

Evolution of Surge Pricing

- Spatial smoothness

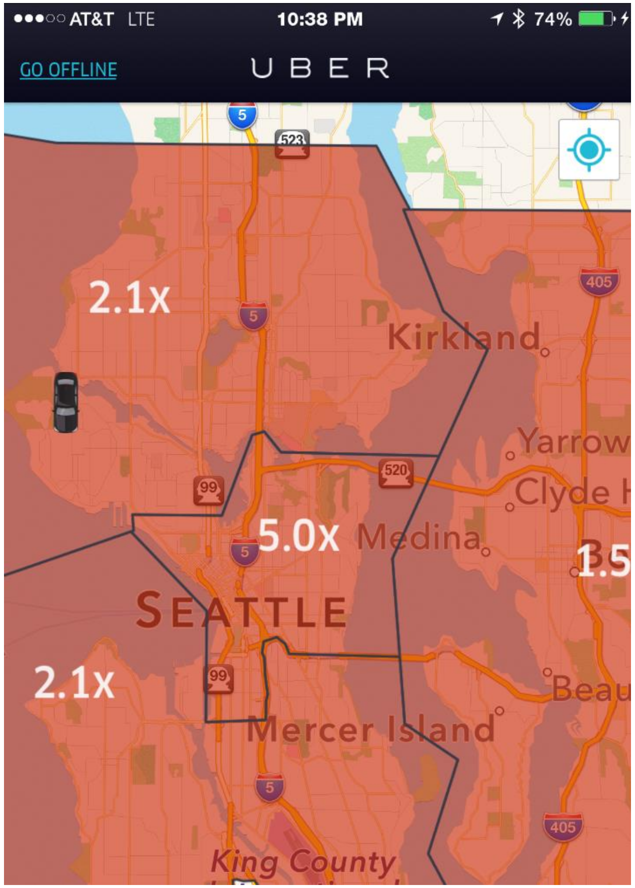

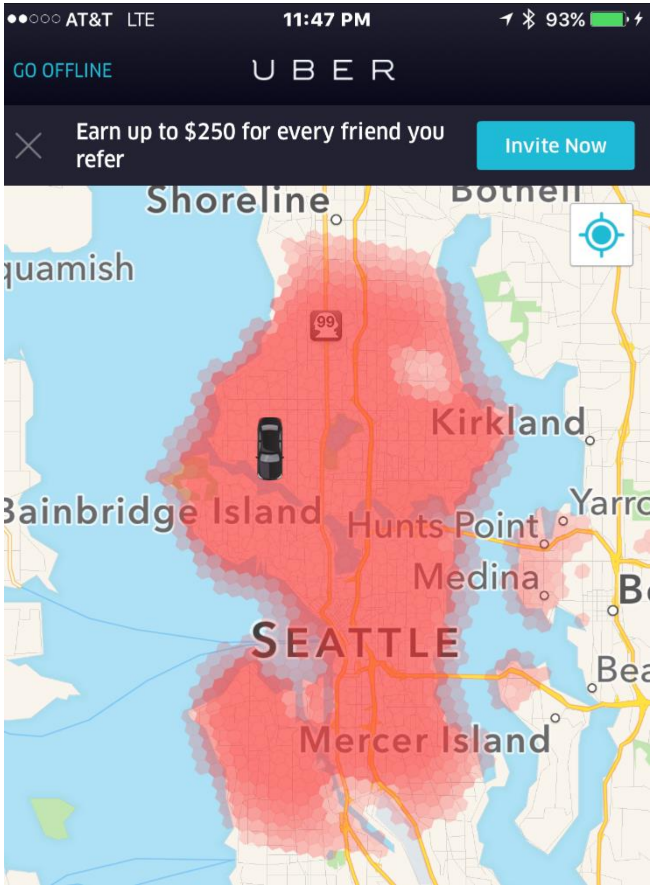

Drivers' Strategic Behavior

- What are the potential strategic behavior of a driver?

Super Bowl Example

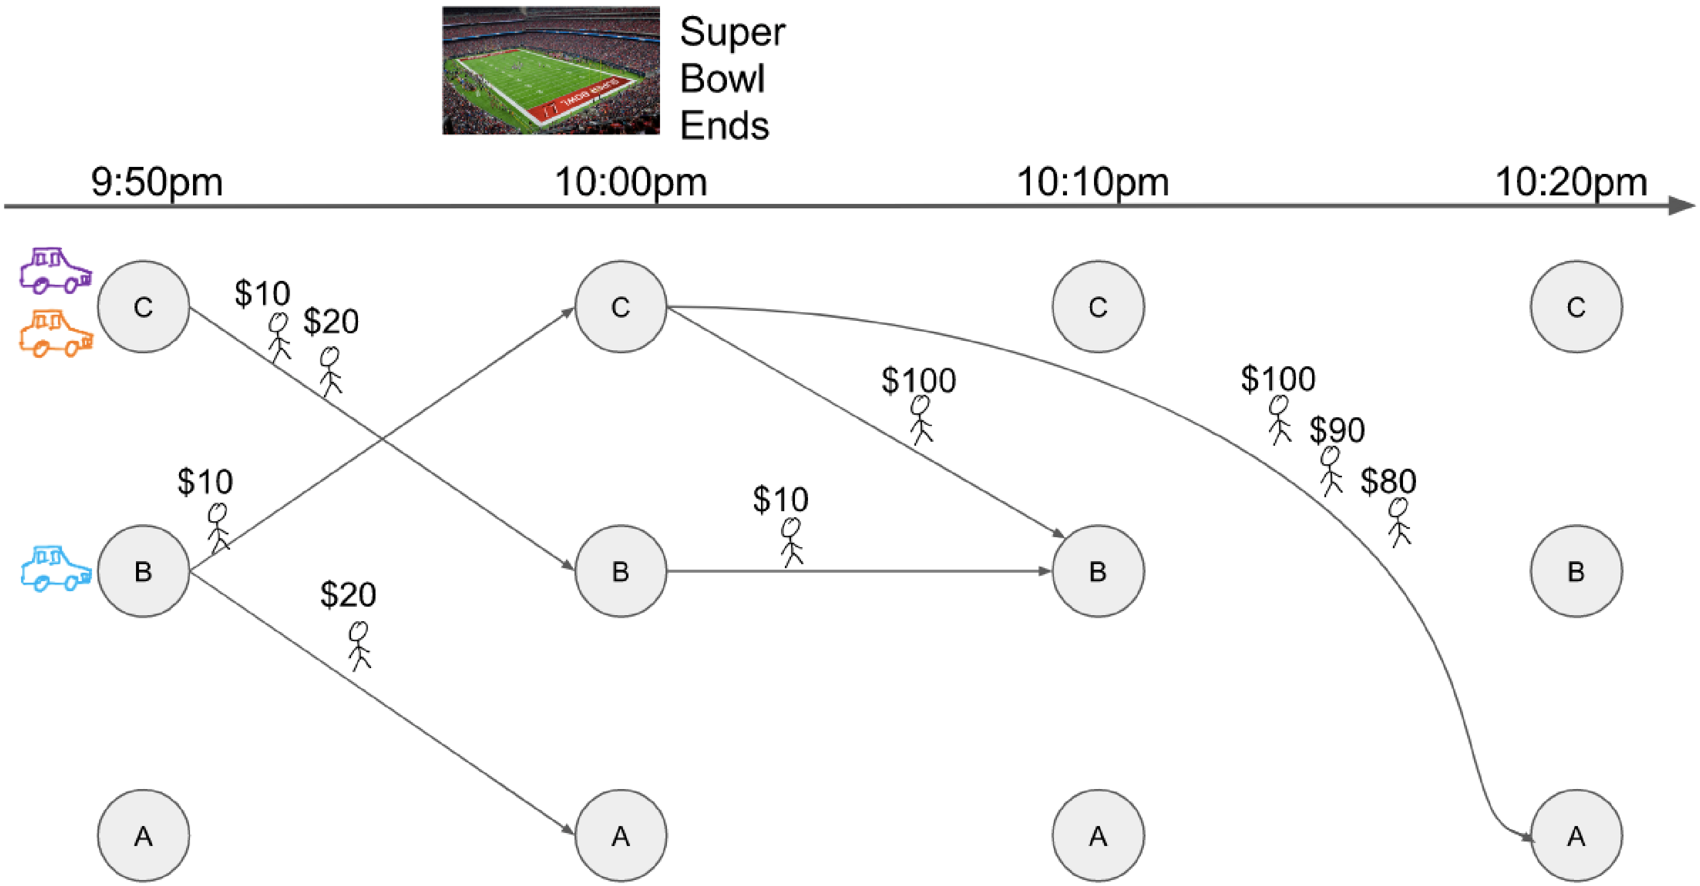

Outline

- Domain and Challenges
- Background
  - Linear Program Duality
  - Competitive Equilibrium
  - Minimum-cost Flow Problem
- Spatio-Temporal Pricing for Ridesharing Platforms

Recall: Linear Program Duality

- LP: An optimization problem whose optimization objective is a linear function and feasible region is defined by linear constraints
- Dual problem of an LP (called primal LP): an LP that is strongly connected to primal LP
- Strong duality theorem: LP and its dual have the same value

Recall: Linear Program Duality

- Exp 3: Maximize Profit

	Price	Labor	Machine
Product 1	$30	0.2 hour	4 hour
Product 2	$30	0.5 hour	2 hour
Total		<=90	<=800

Recall: Linear Program Duality

Maximize	Minimize
ith constraint ≤	ith variable ≥ 0
ith constraint ≥	ith variable ≤ 0
ith constraint =	ith variable unrestricted
jth variable ≥ 0	jth constraint ≥
jth variable ≤ 0	jth constraint ≤
jth variable unrestricted	jth constraint =

Recall: Linear Program Duality

- Complementary slackness
  - Primal constraint is not tight → dual variable=0
  - Primal variable is not zero → dual constraint is tight
  - Dual constraint is not tight → primal variable=0
  - Dual variable is not zero → primal constraint is tight

- Optimality Conditions: If x, y are feasible solutions to the primal and dual problems, respectively, then they are optimal solutions to these problems if, and only if, the complementary-slackness conditions hold for both the primal and the dual problems

Outline

- Domain and Challenges
- Background
  - Linear Program Duality
  - Competitive Equilibrium
  - Minimum-cost Flow Problem
- Spatio-Temporal Pricing for Ridesharing Platforms

Competitive Equilibrium

- Competitive Equilibrium (CE)
  - Also called Walrasian equilibrium
  - Traditional concept in economics
  - Commodity markets with flexible prices and many traders

Competitive Equilibrium

- A very simple setting
  - A set of items $[n] = \{1,2, \dots n\}$
  - A set of buyers $[m] = \{1,2, \dots, m\}$
  - Each buyer i has a valuation for each item j: v_{ij}
  - Given a price vector $p \in \mathbb{R}^n$, agent i's utility is: $u_i(x; p) = v_i \cdot x - p \cdot x$ where $x \in \{0,1\}^n$ indicates which items the agent gets
  - Each agent can get at most one item

Competitive Equilibrium

- A CE consists of:
  - A price vector $p \in \mathbb{R}_+^n$
  - A valid allocation matrix x
    - $x_{ij} \in \{0,1\}$ indicates whether or not item j is allocated to agent i
    - Each item is allocated at most once $\sum_i x_{ij} \leq 1, \forall j$
    - Each buyer can get at most one item $\sum_j x_{ij} \leq 1, \forall i$
    - Use x_i to denote the binary vector for agent i
  - p and x satisfy the following constraints
    - Best response
      - $x_i \in \underset{x: x \in \{0,1\}^n, \sum_j x_j \leq 1}{\operatorname{argmax}} u_i(x; p), \forall i$
    - Market clearance
      - $\forall j, \sum_i x_{ij} = 1$ or $p_j = 0$

Competitive Equilibrium

- Social welfare of an allocation x: $SW(x) = \sum_i \sum_j v_{ij} x_{ij}$
- First Welfare Theorem (Optimality): If x be an equilibrium allocation induced by equilibrium prices p, then x achieves the optimal social welfare
- Second Welfare Theorem (Exchangeability): If p is a set of equilibrium prices and x is an optimal allocation, then (p, x) forms a CE

Competitive Equilibrium

- How to compute CE price when v_i is known?

Quiz 1

- There are three buyers who want to buy a box of chocolate, their valuation of the box of chocolate is \$3, \$5, \$7.5, which of the following (p,x) form a CE? (since there is only one item, p is a scalar, x^j is a scalar)
  - A: $p = \$5.5, x = (1,0,0)$ (meaning agent 1 gets the chocolate)
  - B: $p = \$6, x = (0,1,1)$
  - C: $p = \$7.5, x = (0,0,1)$
  - D: $p = \$y, x = (0,0,1), y \in [5,7.5]$

Competitive Equilibrium

- How to compute CE price when v_i is known?
- Single item case: any price between highest and second highest valuation
  - Buyer optimal price: price = second highest valuation
  - Buyer pessimal price: price = highest valuation
- Multiple item case: Solve by a linear program and its dual
  - Primal solution – allocation x
  - Dual solution – price p
  - Show (x, p) forms a CE using complementary slackness

Competitive Equilibrium

- A slightly generalized setting
  - A set of items $[n] = \{1,2, \dots n\}$
  - A set of agent $[m] = \{1,2, \dots, m\}$
  - Each agent has valuation function $v_i: 2^m \rightarrow \mathbb{R}$
  - Given a price vector $p \in \mathbb{R}^n$, agent's utility is quasilinear: $u_i(x;p) = v_i(x) - p \cdot x$ where $x \in \{0,1\}^n$ indicates which items the agent gets

- A CE consists of:
  - A price vector $p \in \mathbb{R}^n$
  - A valid allocation matrix $x = (x^1, \dots, x^m)$
    - $x_j^i \in \{0,1\}$ indicates whether or not item j is allocated to agent i
    - Each item is allocated at most once $\sum_i x_j^i \leq 1, \forall j$
  - p and x satisfy the following constraints
    - Best response
      - $x^i \in \underset{x}{\operatorname{argmax}}\, u_i(x;p), \forall i$
    - Market clearance
      - $\sum_i x_j^i = 1 \text{ or } p_j = 0, \forall j$

Competitive Equilibrium

- Social welfare of an allocation x: $SW(x) = \sum_i v_i(x^i)$
- First and Second Welfare Theorem still hold (Optimality and Exchangeability)
- Can also be solved by a linear program and its dual
  - Primal solution – allocation x
  - Dual solution – price p
  - Show (x, p) forms a CE using complementary slackness

Competitive Equilibrium

- How is CE connected to mechanism design?

Outline

- Domain and Challenges
- Background
  - Linear Program Duality
  - Competitive Equilibrium
  - Minimum-cost Flow Problem
- Spatio-Temporal Pricing for Ridesharing Platforms

Min-cost Flow Problem

- Given a directed graph $G = (V, E)$ with a source vertex $s \in V$ and a sink vertex $t \in V$, where each edge $(u, v) \in E$ has capacity $c(u, v) \geq 0$ and cost $d(u, v) \geq 0$, a k-unit network flow on the graph is described by the amount of flow on each edge, $f(u, v)$ that satisfies the constraints
  - Capacity constraint: $f(u, v) \leq c(u, v)$
  - Flow conservation: $\sum_{w \in V_{u\rightarrow}} f(u, w) - \sum_{w \in V_{u\leftarrow}} f(w, u) = 0$
  - Required flow:
    - $\sum_{w \in V_{s\rightarrow}} f(s, w) - \sum_{w \in V_{s\leftarrow}} f(w, s) = k$
    - $\sum_{w \in V_{t\rightarrow}} f(t, w) - \sum_{w \in V_{u\leftarrow}} f(w, t) = -k$

Min-cost Flow Problem

Min-cost Flow Problem

- The min-cost flow problem is to find a feasible $s - t$ $k-$ unit flow with minimum cost $\sum_u \sum_v f(u,v)d(u,v)$
- Solution approach: Linear Program
  - Primal: flow
  - Dual: potential (often set $potential(t) = 0$)
- Potential values of the nodes form a lattice
- Lattice
  - If $p \in L$ and $q \in \mathrm{L}$, then $p \vee q \in L, p \wedge q \in L$
  - In this problem, $p \vee q =$ take element-wise maximum, $p \wedge q =$ take element-wise minimum

Outline

- Domain and Challenges
- Background
  - Linear Program Duality
  - Competitive Equilibrium
  - Minimum-cost Flow Problem
- Spatio-Temporal Pricing for Ridesharing Platforms

Super Bowl Example

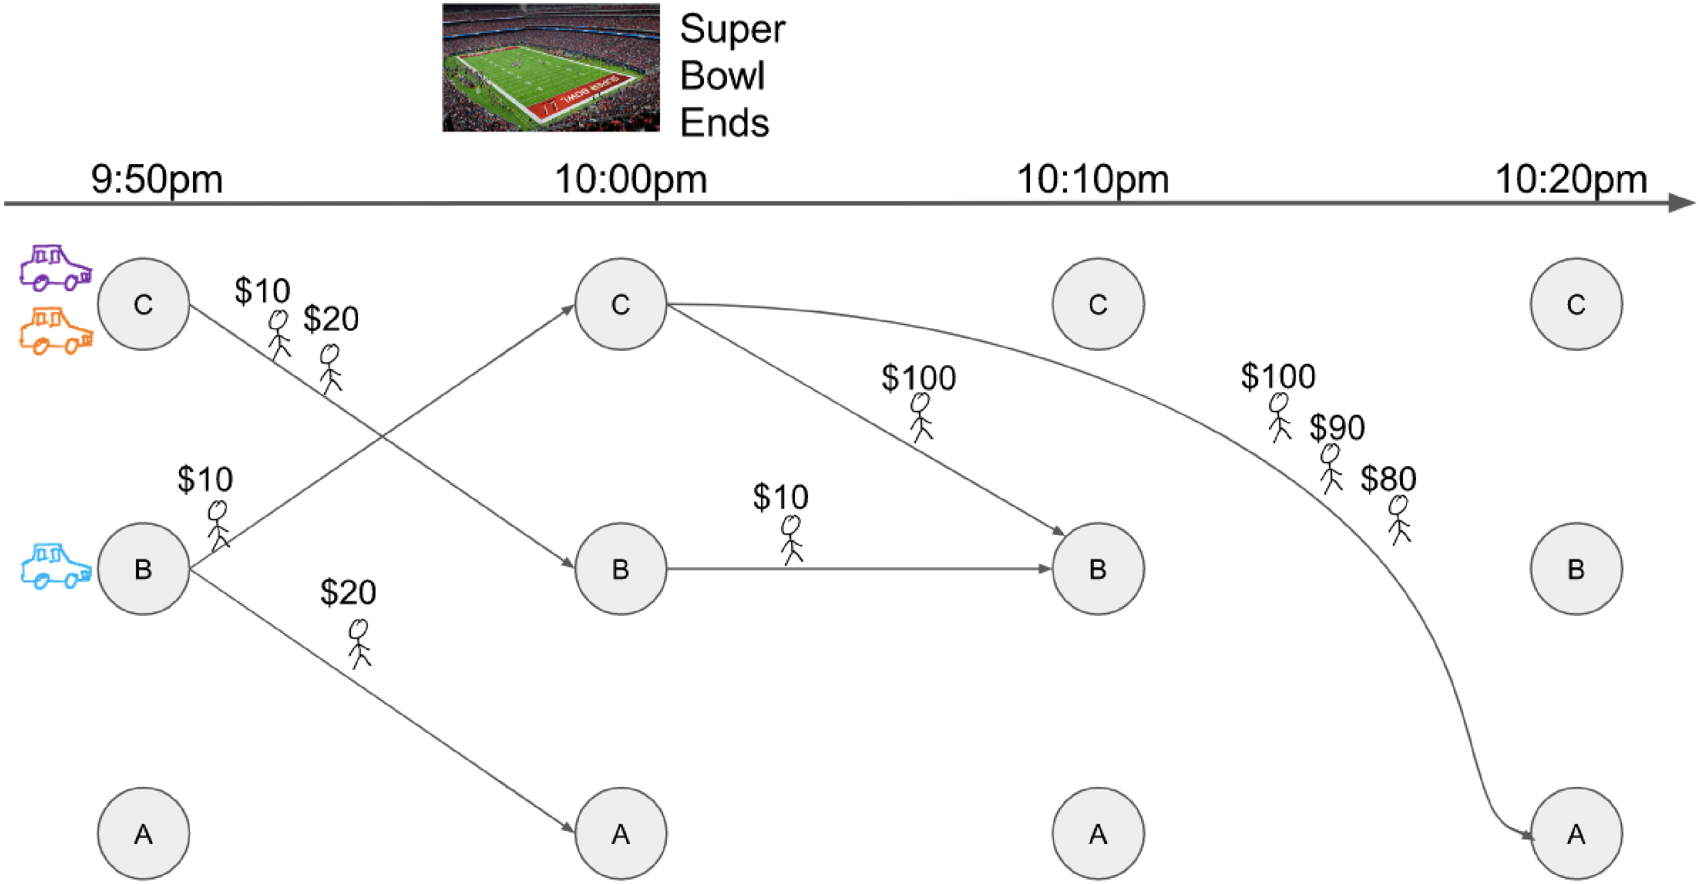

Myopic Pricing

- At current time t, each location has a sub-market
- Allocate cars to the riders with highest valuations
- Driver-pessimal price shown in black

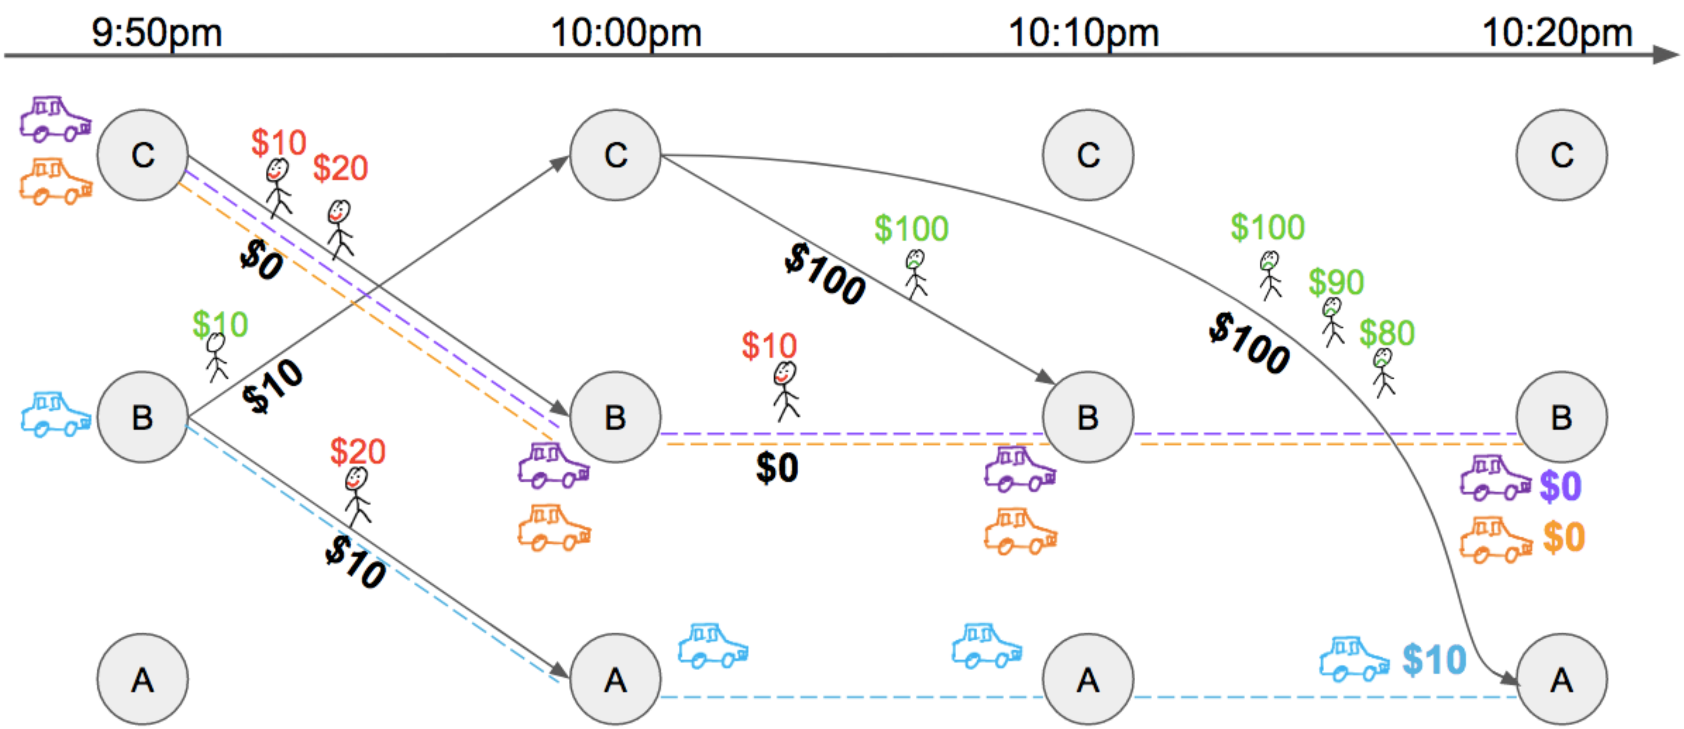

Quiz 2

- With Myopic Pricing, at most, how much more can the purple driver earn if he deviates from the system's assignment and all other drivers always follow the system's assignment? (Options: $100, $90, $80, $0)

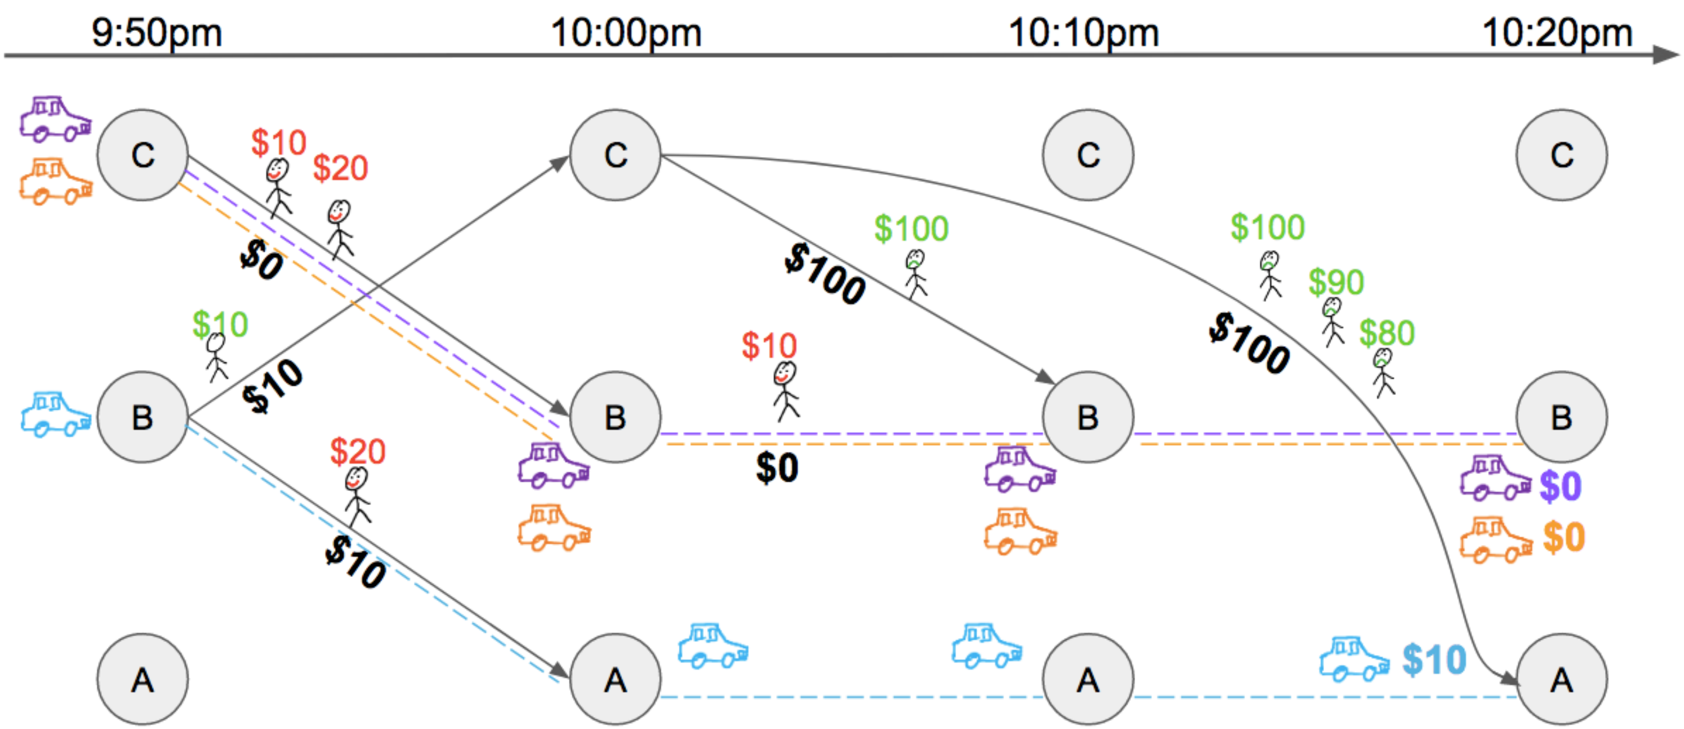

Useful Deviation

- Purple driver rejects the assigned ride at 9:50am to earn more money

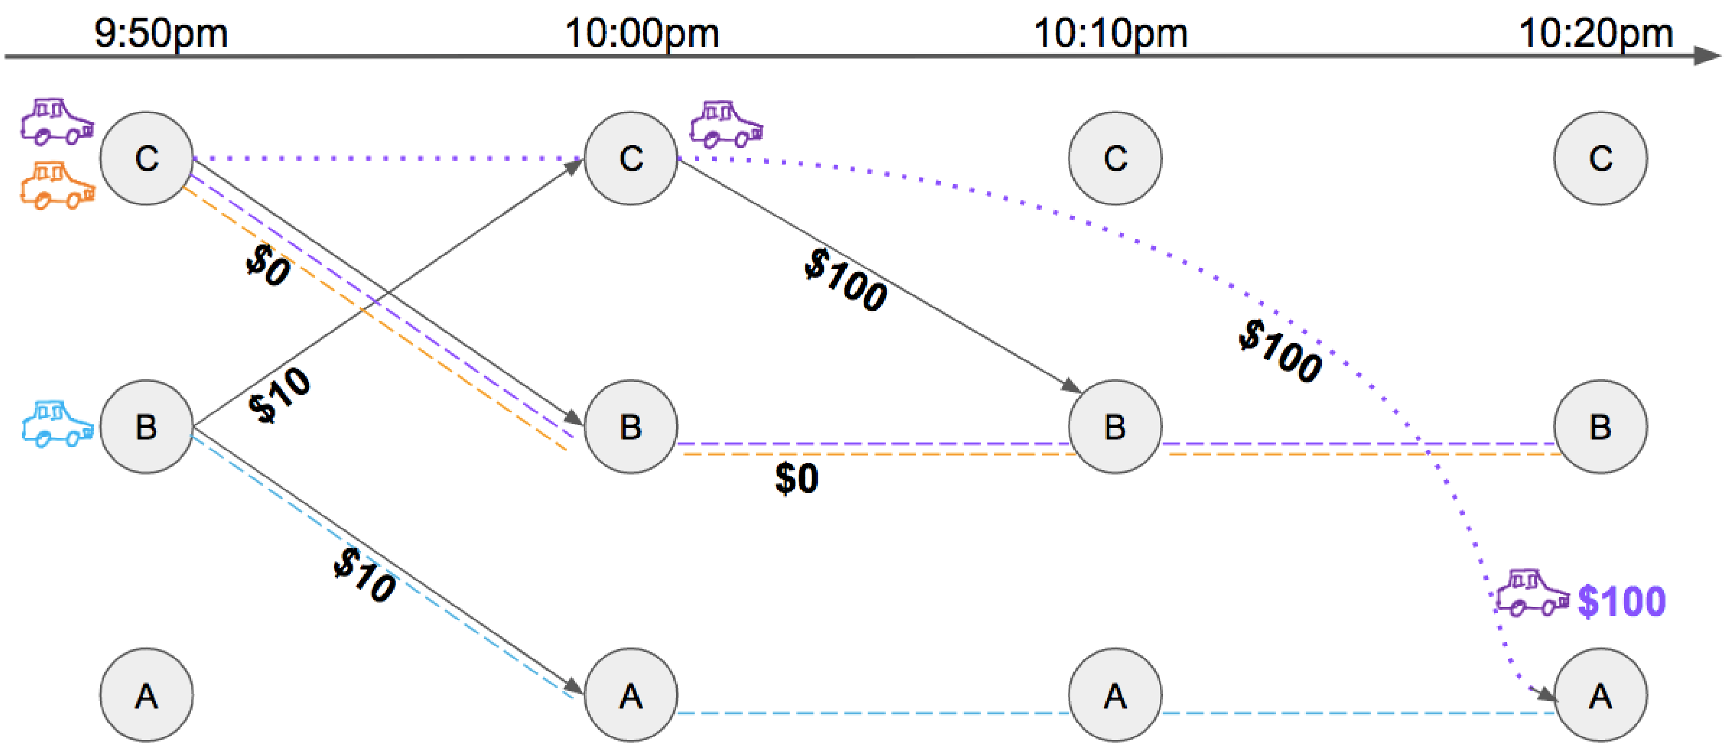

Characteristics of Problem

- Both riders and drivers are customers: autonomous and strategic
- Dynamic: rider and driver arrivals and departures
- Match riders and drivers via priced trips
- Externalities: assigned trips determines when and where the drivers are free to pick up again

Design Goals

- Assign riders to drivers, or tell drivers where to go
- Determine prices for trips depending on the origin, destination, time of the day
- Wish list:
  - Maximize social welfare
  - Align incentives for riders and drivers
  - Envy-free outcomes
  - Budget balance
  - Robust to deviation

Model

- Discrete time, finite planning horizon
- Discrete regions, integer distances

- Driver i entering the system at time $\underline{\tau_i}$ and region e_i, wants to maximize total payment
- Rider j requesting trip from o_j to d_j at time τ_j with value v_j, wants a ride if price is below v_j

- Assumptions:
  - A1. Homogeneous drivers, stay until end of planning horizon
  - A2. Impatient riders
  - A3. Complete and symmetric information

Optimal Plan

- Plans
  - x: Whether or not each rider is picked up
  - $\tilde{z}$: Paths taken by each driver through the planning horizon, and whether or not pick up any rider along the way
  - p: Anonymous trip prices for (o, d, t)

- Optimal Plan $(x, \tilde{z}, p)$
  - Maximize social welfare (total value of riders picked up)
  - Form competitive equilibrium:
    - Rider best response: whoever prefers to be picked gets picked up
    - Driver best response: choose among all possible paths and whether to serve a rider to maximize total payment

Optimal Plan

- Does optimal plan exist?
- Theorem: There exists an optimal plan

- How to compute optimal plan?
- Theorem: An optimal plan can be computed in polynomial time

- Three step argument
  - Step 1: (Similar to commodity market with indivisible goods) Build an LP and write down its dual, prove that the solution of the LP and the solution of the dual LP lead to an optimal plan that maximize social welfare and form an CE
  - Step 2: (Similar to min-cost network flow) Convert the problem into a MCNF problem. Build an LP and write down its dual for the problem. Show that there is a many-to-one mapping between the solution of Step 1 and Step 2
  - Step 3: Argue that we can solve the MCNF (poly time), and use the solution to construct an optimal plan (poly time).

Optimal Plan

- If the price is fixed as shown here, driver is willing to follow the system's assignment

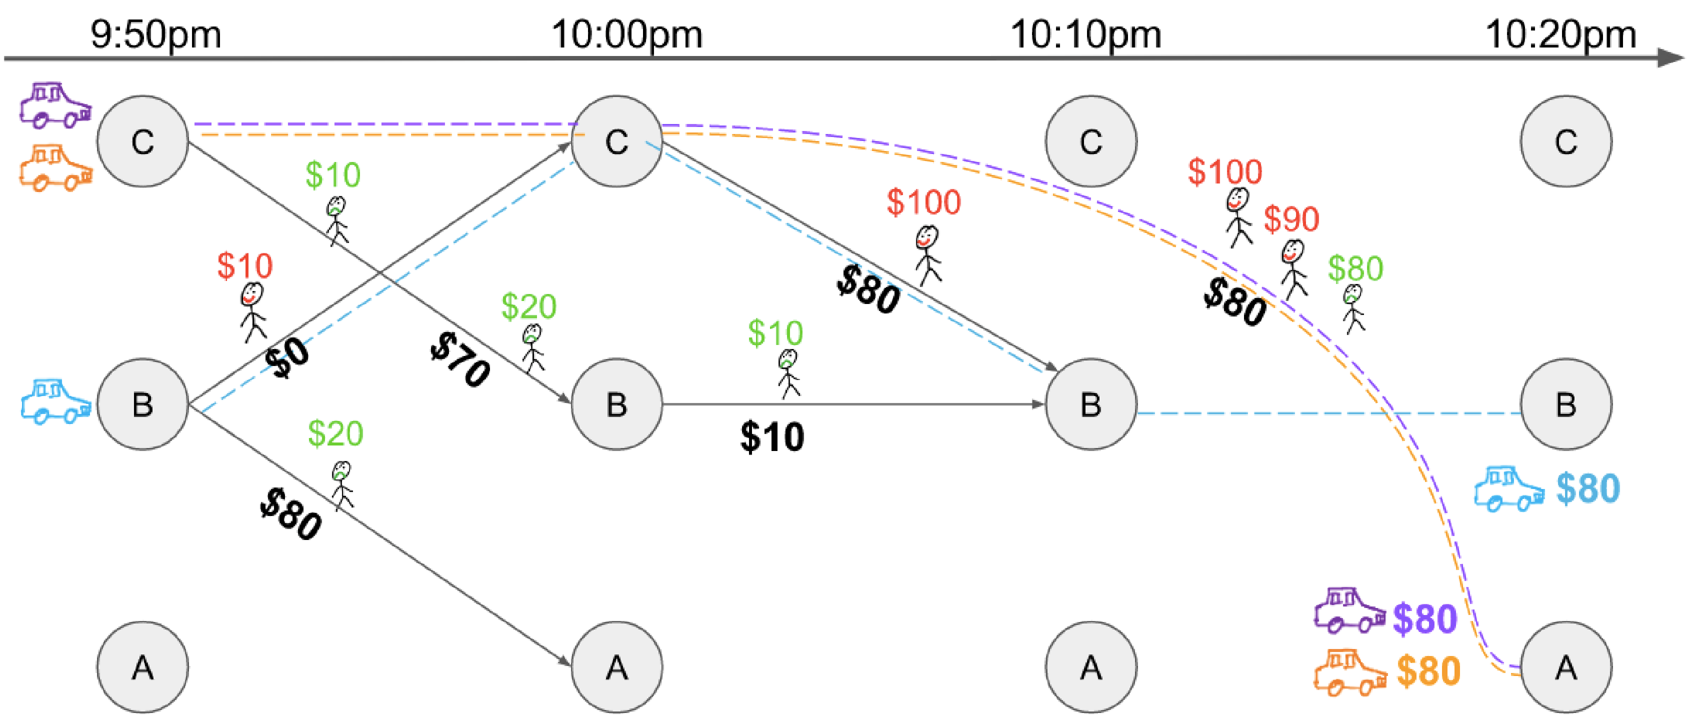

Deviation from Optimal Plan

- If blue car didn't make it to C, should the system stick to the planned price?

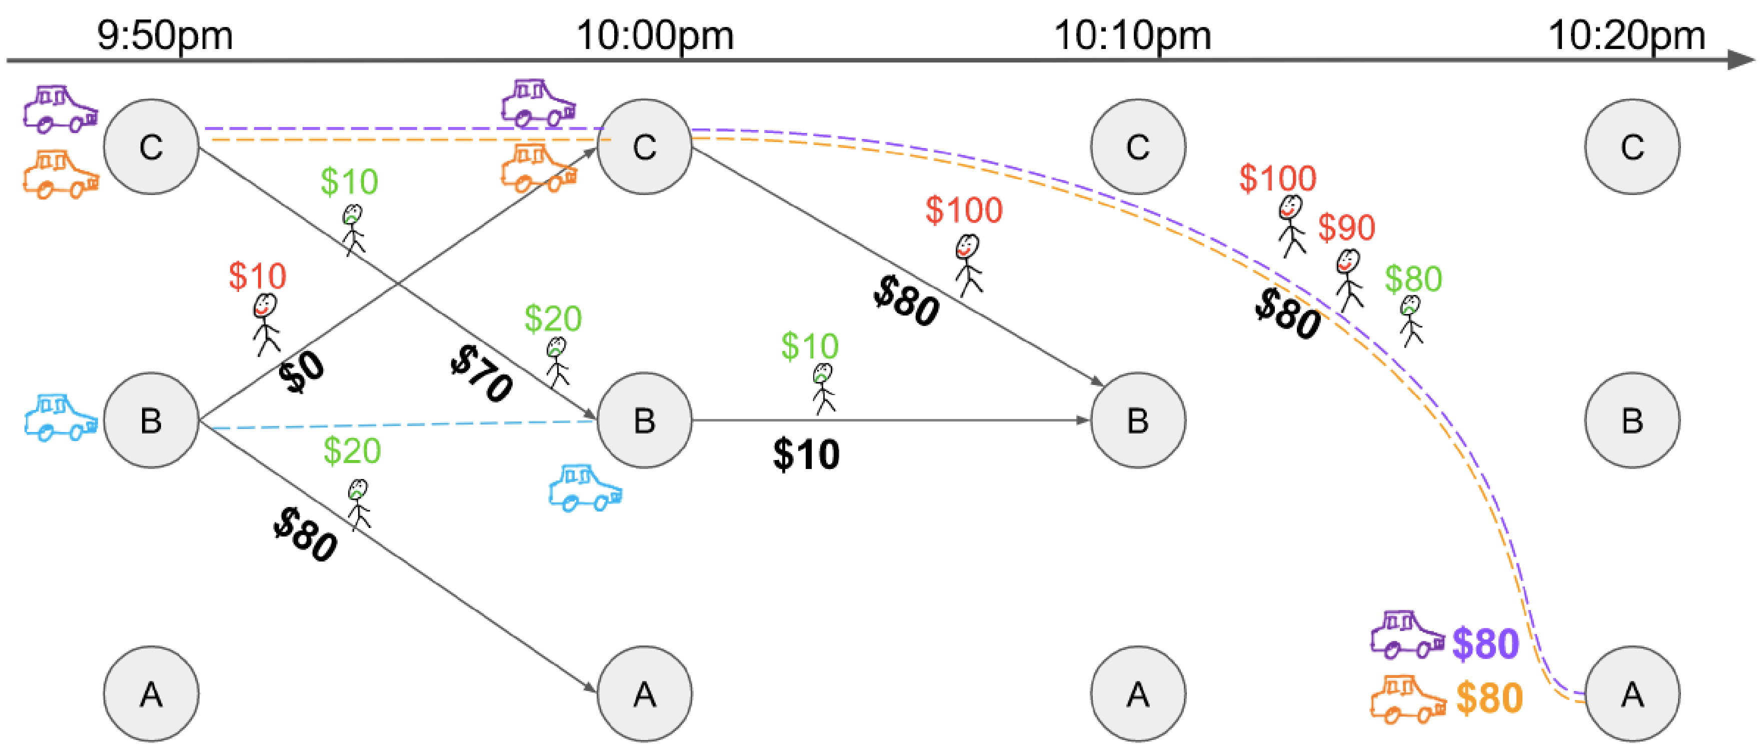

The Spatio-Temporal Pricing Mechanism

- A dynamic mechanism is needed:
  - History independent
  - Time-invariant

- Challenges with the dynamic mechanism:
  - Once the mechanism is announced, the induced game is an extensive-form game
  - An intelligent driver may deliberately deviate from the plan to trigger recomputation trying to earn more money

The Spatio-Temporal Pricing Mechanism

- Spatio-Temporal Pricing:
  - Welfare-optimal assignment
  - Each driver's continuation payment = welfare gain from replicating this driver
  - Trip price = difference in potential value (take bottom element in the lattice)
  - Recompute on off-equilibrium paths

The Spatio-Temporal Pricing Mechanism

- Optimal Plan

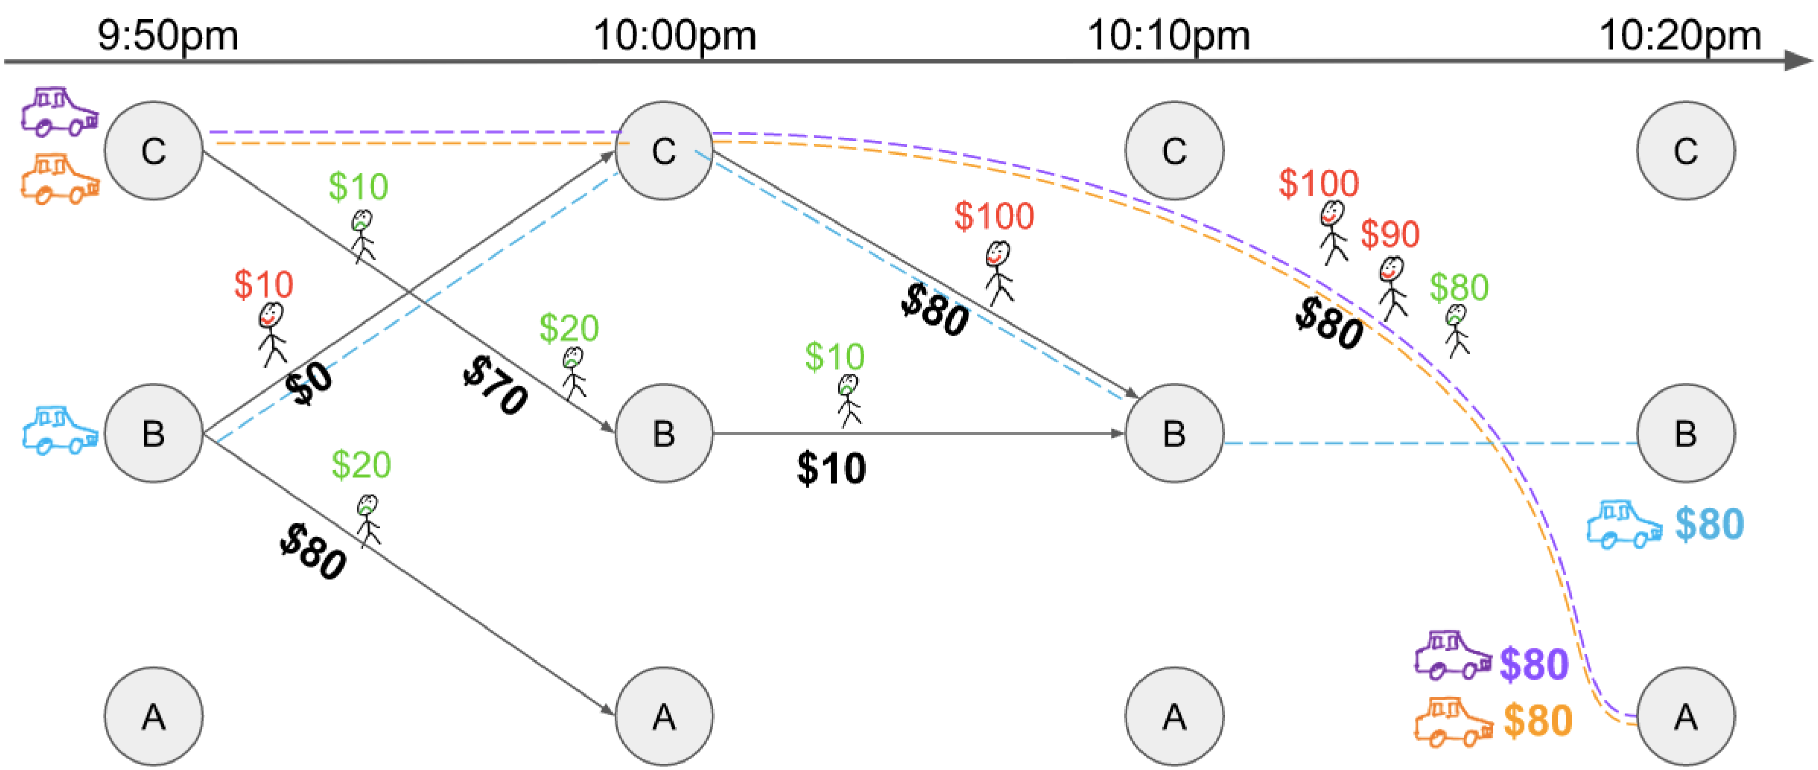

The Spatio-Temporal Pricing Mechanism

- Recompute if there is any deviation from the plan

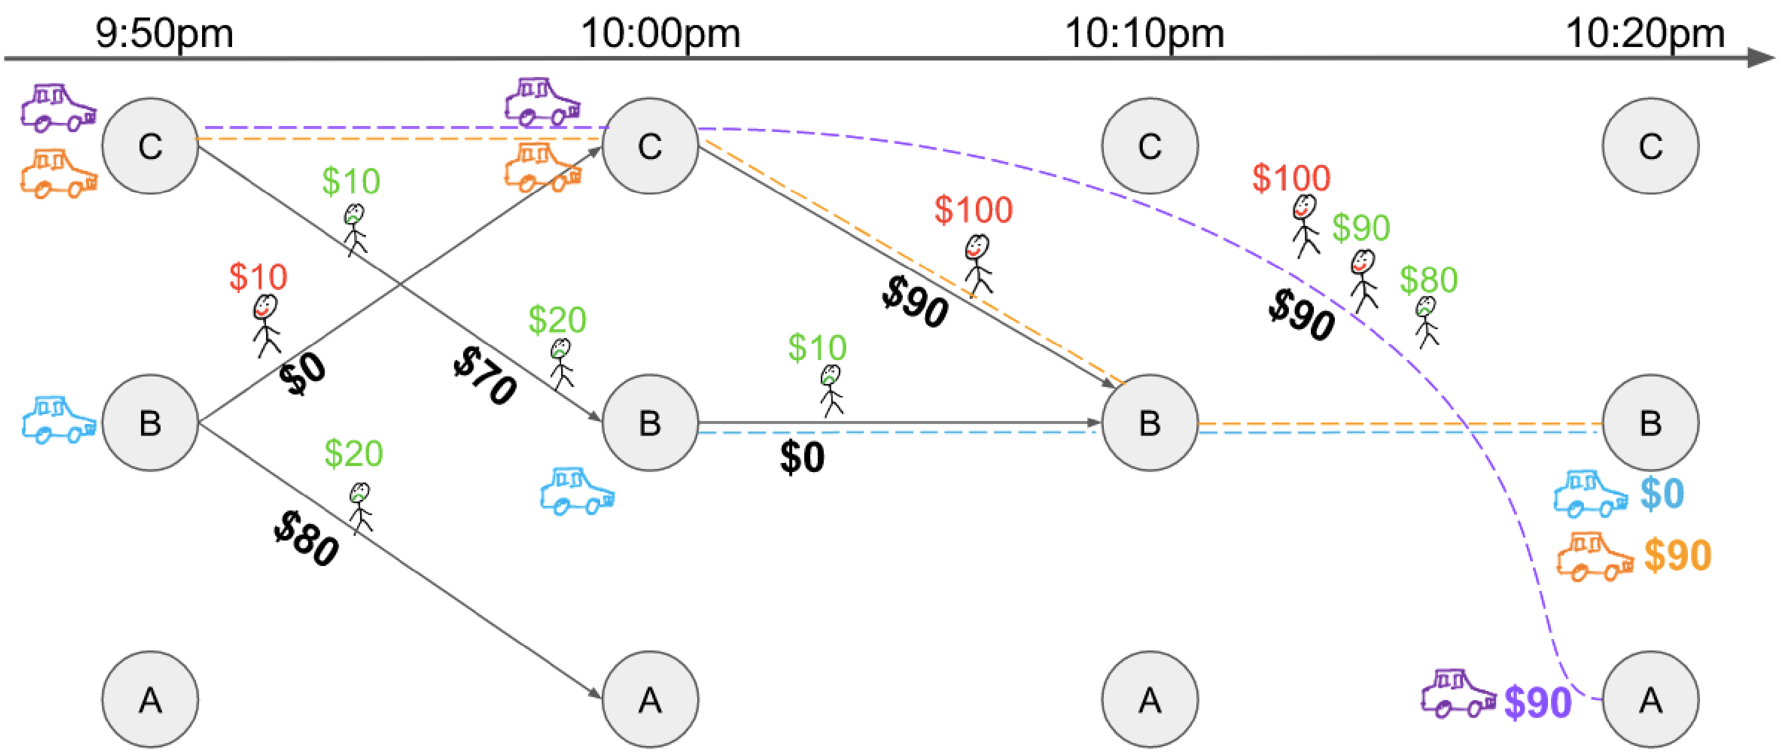

The Spatio-Temporal Pricing Mechanism

- Theorem:
  - Incentive compatible in subgame-perfect equilibrium
  - Envy-free in equilibrium
  - Individually Rational and Budget Balanced for any action profile

Discussion

- Driver collusion?
  - Hard to combat
- Fairness?
  - Subsidies
- Accommodate driver's personal needs?
  - Set "preferred" destination
- Accommodate rider's travel plans?
  - Pre-schedule rides
- Sustainability in areas with less demand?
  - Multimodal ridesharing
- Self-driving cars?
  - Only need to worry about riders and the company

Additional Resources

- Text book
  - *Algorithmic Game Theory 1st Edition, Chapter 11*
  - *by Noam Nisan (Editor), Tim Roughgarden (Editor), Eva Tardos (Editor), Vijay V. Vazirani (Editor)*
  - http://www.cs.cmu.edu/~sandholm/cs15-892F13/algorithmic-game-theory.pdf
- Online course
  - https://www.coursera.org/learn/approximation-algorithms-part-2/lecture/iBEE3/properties-of-lp-duality