Carpool fairness in social networks

Amos Fiat
Anna Karlin
Elias Koutsoupias
Claire Mathieu
Rotem Zach

The carpool problem

The carpool problem

> *Suppose that n people, tired of spending their time and money in gasoline lines, decide to form a carpool… We want a scheduling algorithm that will be perceived as fair by all the members – [Fagin-Williams 1982]*

- every day a subset of n people share a ride
- who should drive?

Objective: everyone drives his fair share

- let $\sigma_1, \ldots, \sigma_T$ be the subsets (requests)
- suppose that driver i has driven $m_{i,T}$ times
- his fair share is $f_{i,T} = \sum_{t:i\in\sigma_t} 1/|\sigma_t|$
- and his unfairness $|m_{i,T} - f_{i,T}|$
- **objective:** minimize $\max_i |m_{i,T} - f_{i,T}|$
- *not in this talk:* one-sided unfairness $(m_{i,T} - f_{i,T})$

Offline vs online

- We consider online algorithms
- An **online algorithm** selects the drivers based only on the past
- The **adversary** selects the sequence of requests, including its length, in advance
- The adversary knows the algorithm, but not the outcome of its random choices (**oblivious adversary**)

Offline algorithms

- There exists an offline algorithm, which
  - can maintain unfairness at most 1
  - for every driver
  - at every time step

for simplicity, we consider requests of size 2

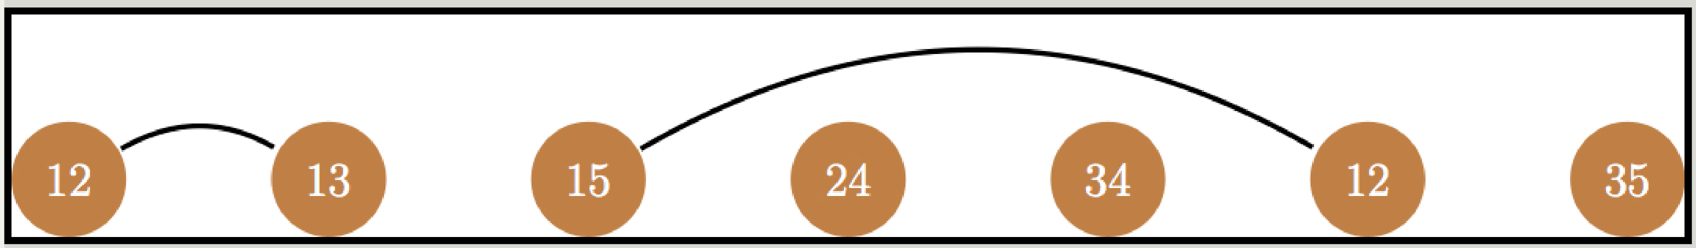

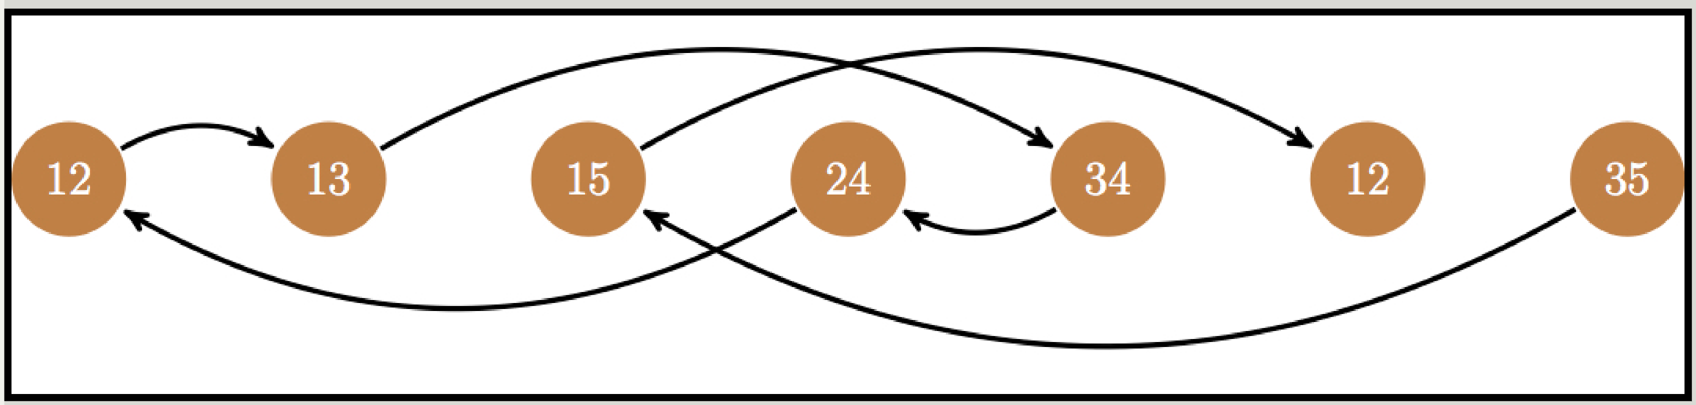

[redacted]

Groups of size 2 suffice

- Every day **only two** people share a ride
- This is without significant loss of generality
  - the general carpool problem reduces to groups of size 2
  - the reduction changes the unfairness only by a factor of 2
- Similar problem: **online edge orientation** of a given graph to minimize (absolute) difference between outdegree and indegree

Basic online algorithms

Random

A random member drives. Unbounded unfairness (proportional to $O(\sqrt{T})$, where $T =$ number of requests)

Local Greedy

In every pair of drivers they drive alternatively. Randomize the first time. Unfairness $O(\sqrt{n \log n})$

Global Greedy

The driver with minimum unfairness drives; in case of a tie, select randomly

- **Conjecture:** Global Greedy has randomized unfairness $\Theta(\log n)$

History

- Fagin and Williams 1982: introduced the problem and the Global Greedy algorithm
- Ajtai, Aspnes, Naor, Rabani, Schulman, and Waarts [AANRSW '96]:
  - reduced the problem to groups of size 2
  - deterministic lower bound $O(n)$
  - randomized algorithms:

    upper bound

    $\Theta(\sqrt{n \log n})$ by the Local Greedy algorithm

    lower bound

    $\Omega(\sqrt[3]{\log n})$ (every algorithm)

This talk: general graphs

- a social network graph
- the request sequence contains only edges of the graph

Overview

	Det	Randomized	Random sequences
Clique	n	$\log^{1/3} n$ .. $\sqrt{n \log n}$	$\Theta(\log \log n)$ (Greedy)
Line	1		$\Omega(\log n / \log \log n)^{1/3}$ (Greedy)
Star	n	static: $\Theta(\sqrt{n})$	1
Planar			$O(\log n)$
All	d	conjecture : $O(\log n)$	

Clique [AANRSW '96]

- deterministic: $\Theta(n)$
  - Adversary requests only pairs of drivers with the same unfairness
  - At the end (after $\Theta(n^3)$ steps), all drivers have distinct unfairness
  - Therefore one of them has unfairness outside the interval $[-n/2 + 1, n/2 - 1]$

- randomized:
  - lower bound $O(\sqrt[3]{\log n})$ (based on the deterministic lower bound)
  - upper bound $\Theta(\sqrt{n \log n})$ (Local Greedy)

- random sequences: Global Greedy has unfairness $\Theta(\log \log n)$

Deterministic algorithms

Overview

	Det	Randomized	Random sequences
Clique	n	$\log^{1/3} n$ .. $\sqrt{n \log n}$	$\Theta(\log \log n)$ (Greedy)
Line	1		$\Omega(\log n / \log \log n)^{1/3}$ (Greedy)
Star	n	static: $\Theta(\sqrt{n})$	1
Planar			$O(\log n)$
All	d	conjecture : $O(\log n)$	

Lower bound for clique

- Tight bound $\Theta(n)$ [AANRSW '96]
- For other graphs? stars?

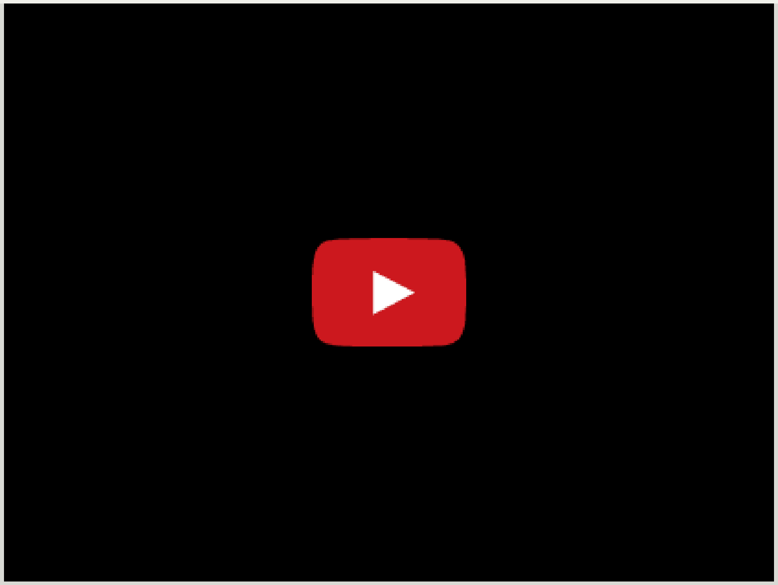

- one-sided fairness? not in this talk

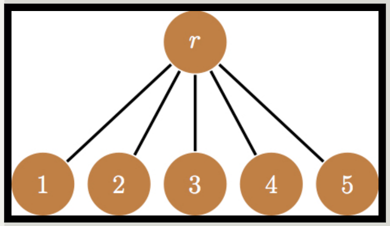

Lower bound on star and arbitrary graphs

- Fix a deterministic algorithm for the star of d leaves
- Reachable states $\vec{x} = (x_1, \ldots, x_d)$, where x_i the unfairness of leaf i
- Define $\varphi(\vec{x}) = x_1 + 2x_2 + \cdots + 2^{d-1}x_d$
- **Claim:** If $\vec{x}$ minimizes $\varphi(\vec{x})$, then $\vec{x} + \vec{1}$ is also reachable
- $|\vec{x} + \vec{1}|_1 - |\vec{x}|_1 = d$
- Therefore either $\vec{x}$ or $\vec{x} + \vec{1}$ has root unfairness $\lceil d/2 \rceil$

Matching upper bound

Theorem: The deterministic unfairness of graphs of degree d is exactly $\lceil d/2 \rceil$.

- Fix an almost balanced orientation (outdegree and indegree differ by at most 1)
- For every oriented edge (i, j), service the odd requests with i and even requests with j

Random sequences

Overview

	Det	**Randomized**	**Random sequences**
Clique	n	$\log^{1/3} n$ .. $\sqrt{n \log n}$	$\Theta(\log\log n)$ (Greedy)
Line	1		$\Omega(\log n / \log\log n)^{1/3}$ (Greedy)
Star	n	static: $\Theta(\sqrt{n})$	1
Planar			$O(\log n)$
All	d	conjecture : $O(\log n)$	

Global Greedy on the line

all pairs of requests are selected uniformly at random

- Global Greedy has expected unfairness $O(\log \log n)$ for the clique [AANRSW'96]
- We show that for sparse networks, Global Greedy does much worse

- In particular, for the line, Global Greedy has expected unfairness $\Omega((\log n / \log\log n)^{1/3})$
- **Proof**
  - There exists a sequence s_n^* of n^3 requests for which Global Greedy with *adversarial tie breaking* has unfairness $\Omega(n)$
  - Break the line into k/n segments of length k
  - Consider a random sequence of length k^3 in one of the segments
  - What is the probability that
    - no request falls into the boundaries and the random sequence is the bad sequence s_k^*? : $1/k^{k^3}$
    - Global Greedy breaks the ties as in the worst-case : $1/2^{k^3}$

- The probability that the unfairness in every segment is less than k is $(1 - 1/(2k)^{k^3})^{n/k}$,
- which is constant when we select $k = (\log n/ \log\log n)^{1/3}$
- It follows that, with constant probability, a random sequence of length $k^3 \cdot n/k$ has unfairness $\Omega((\log n/ \log\log n)^{1/3})$

Random sequences on planar graphs

- Algorithm for stars
  - leaves have unfairness -1, 0, 1
  - when at 0, they help the unfairness of the root
- It has constant unfairness
- Extension to planar graphs
  - partition the edges of the graph into stars
  - every node belongs to at most 6 stars
  - run the algorithm for each star
  - the expected unfairness of every node is constant
  - the maximum unfairness among all nodes is $O(\log n)$

Randomized algorithms

Overview

	Det	Randomized	Random sequences
Clique	n	$\log^{1/3} n$ .. $\sqrt{n \log n}$	$\Theta(\log \log n)$ (Greedy)
Line	1		$\Omega(\log n / \log \log n)^{1/3}$ (Greedy)
Star	n	static: $\Theta(\sqrt{n})$	1
Planar			$O(\log n)$
All	d	conjecture : $O(\log n)$	

Static algorithms

Static randomized online algorithm :: there exists a probability distribution π over the set of states

- the algorithm starts in π
- it remains in π after every possible request.
- **Theorem:** The unfairness of every static algorithm is $\Omega(\sqrt{d})$, where d is the degree of the social graph

- Upper bound (on stars)
  - Balanced Local Greedy algorithm
    - Fix a balanced orientation
    - Every oriented edge (i, j) alternates between unoriented and oriented according to fixed orientation starting at a random state

- Lower bound (on stars)

1. Characterize the stationary distributions of the Markov chain
2. Express the question as linear program
3. Solve the linear program

- Static characterization
- for every x_{-i}: $\sum_{k \in \mathbb{Z}} \pi(k, \vec{x}_{-i})(-1)^k = 0$
- eliminate variables $\pi(\vec{x})$ when $\vec{x}$ has at least one 0
- the value of the linear program is at least equal to

$$\min_{y_i^*} \quad E\left[\left|\sum_i y_i^* X_i\right|\right]$$

  - X_i's are 0-1 unbiased binomial random variables
- this is minimized when half of y_i^*'s are 1 and half are -1
- exactly as in the Local Greedy algorithm

- Complete characterization
- what distributions π are stationary distributions of Markov chains on the line (when $p_{x,x} = \text{const}$)?
  - *answer:* $\pi(k) - \pi(k+1) + \pi(k+2) - \cdots \geq 0$, for every k

Overview

	Det	Randomized	Random sequences
Clique	n	$\log^{1/3} n$ .. $\sqrt{n \log n}$	$\Theta(\log \log n)$ (Greedy)
Line	1		$\Omega(\log n / \log \log n)^{1/3}$ (Greedy)
Star	n	static: $\Theta(\sqrt{n})$	1
Planar			$O(\log n)$
All	d	conjecture : $O(\log n)$	

Thank you

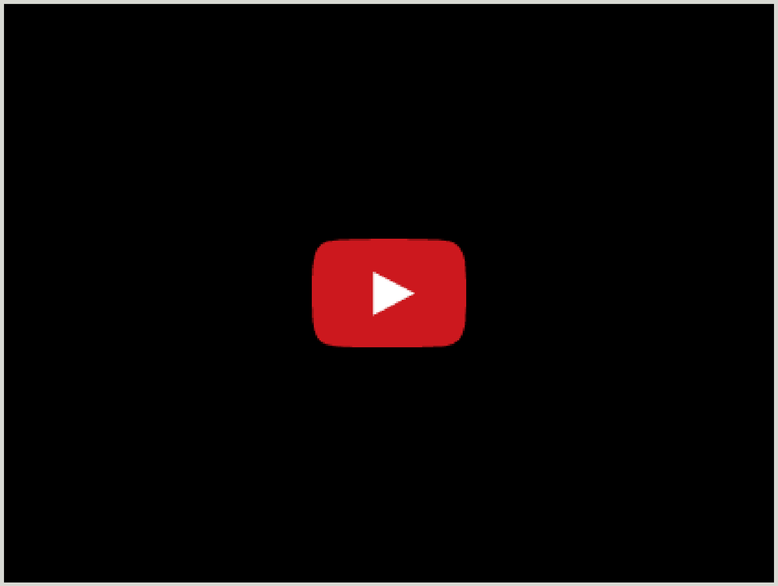

[redacted]

Created by elias.