

Dynamic Traffic Models in Transportation Science

José Correa^{*1}, Carolina Osorio^{*2}, Laura Vargas Koch^{*3},
David Watling^{*4}, and Svenja Griesbach^{†5}

1 Universidad de Chile, CL. correa@uchile.cl

2 HEC Montréal, CA & Google – Mountain View, US. carolina.osorio@hec.ca

3 RWTH Aachen, DE. vargaskoch@gdm.rwth-aachen.de

4 University of Leeds, GB. d.p.watling@its.leeds.ac.uk

5 Universidad de Chile, CL. sgriesbach@cmm.uchile.cl

Abstract

Traffic assignment models are crucial for traffic planners to be able to predict traffic distributions, especially in light of possible changes in the infrastructure, e.g., road constructions, traffic light controls, etc. There is a trend in the transportation community (science and industry) to base such predictions on complex computer-based simulations capable of resolving many elements of a real transportation system. Moreover, cities worldwide, driven by critical sustainability goals, are developing digital twins of their transportation networks to inform the design and the operations of these intricate networks. On the other hand, the theory of dynamic traffic assignments in terms of equilibrium existence, computability, and efficiency, has not matured to the point matching the model complexity inherent in simulations.

The Dagstuhl Seminar was the fourth in a row on this topic and brought together leading scientists in the areas *traffic simulations*, *algorithmic game theory*, and *dynamic traffic assignment*. In this seminar, we tackled important open research problems that were identified in past seminars. Motivated by the increasing importance, in practice, of developing sustainable, flexible, and on-demand mobility services, this seminar identified a new set of important questions and first results in this field.

Seminar July 7–12, 2024 – <https://www.dagstuhl.de/24281>

2012 ACM Subject Classification Networks → Control path algorithms; Networks → Network performance evaluation; Theory of computation → Design and analysis of algorithms

Keywords and phrases Algorithms and Complexity of traffic equilibrium computations, Dynamic traffic assignment models, Simulation and network optimization

Digital Object Identifier 10.4230/DagRep.14.7.1

1 Executive Summary

José Correa (Universidad de Chile, CL)

Carolina Osorio (HEC Montréal, CA & Google – Mountain View, US)

Laura Vargas Koch (RWTH Aachen, DE)

David Watling (University of Leeds, GB)

License © Creative Commons BY 4.0 International license

© José Correa, Carolina Osorio, Laura Vargas Koch, and David Watling

Transportation road networks and the underlying spatio-temporal dynamics of vehicular traffic are increasingly complex. City infrastructure is increasingly real-time and traffic-responsive, urban mobility services are increasingly flexible, multi-modal and tailored to the heterogeneous

* Editor / Organizer

† Editorial Assistant / Collector



Except where otherwise noted, content of this report is licensed under a Creative Commons BY 4.0 International license

Dynamic Traffic Models in Transportation Science, *Dagstuhl Reports*, Vol. 14, Issue 7, pp. 1–16

Editors: José Correa, Carolina Osorio, Laura Vargas Koch, and David Watling



Dagstuhl Reports

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

needs and preferences of the traveling population, road assets are increasingly shared between the transport of individuals and the transport of freight and goods. Importantly, operators and users of the road network increasingly focus on the sustainability of their choices or service offerings. The key sustainability metrics (e.g., greenhouse gas emissions and energy consumption) are tightly linked to the underlying traffic dynamics.

Dynamic traffic assignment models are, therefore, critical tools for transport planners to be able to predict the spatio-temporal distribution of traffic – and thereby the congestion, environmental and social impacts – of transport policies, for example, in the light of possible changes to the infrastructure (e.g. traffic light control), to the transport services offered (e.g., the provision of public transport fleets and frequencies), or to the prices imposed on travelers (e.g., congestion pricing, public transport fares). The prevailing approaches used in the transportation science literature to predict such distributions can be roughly classified into mathematical approaches based on *dynamic traffic assignment models* (DTA) (using the methodology of flows over time) and *simulation-based approaches* (using large-scale microsimulations). The striking advantage of microscopic simulations over DTA models is that the latter usually ignore practically relevant side constraints such as horizontal queueing of vehicles, the feedback of changing network conditions on user behavior, flexible departure time choice, mode choice, and activity schedule choice. Current simulation tools integrate all these dimensions and many more. The increase in model complexity inherent in simulations, however, is not matched by the existing DTA theory. For most simulation software, there is no “mathematical proof of concept”. This means it is not known whether equilibria always exist and the heuristics applied in traffic simulation lead to (approximate) equilibria. Moreover, crucial properties for simulations to give meaningful results, such as uniqueness and continuity of equilibria, are only partially understood.

In the Dagstuhl Seminar, we brought together, for the fourth time, leading researchers from three different communities (Simulations, Dynamic Traffic Assignment, and Algorithmic Game Theory) in order to bridge the gap between complex simulation-based models and the existing theory. In the last four seminars, we made substantial progress by answering questions, e.g., on the controllability of traffic flow or the analysis of equilibria with spatial queues. Still, real-world simulation develops faster than mathematical theory, and at the same time, new challenges in modeling future road networks arise, which are characterized by real-time responsiveness, flexible multi-modal urban mobility offerings, automated and autonomous vehicles, coordinated fleets of vehicles, and environmentally sustainable operations. This seminar focused on these new questions with multiple talks on sustainability in traffic simulation in general and more concrete on ride-hailing, ride-sharing, public transportation, and the control of flow via tolls or information design.

For this edition of the seminar, we again welcomed applied researchers from the industry alongside experts in traffic simulations, algorithmic game theory (AGT), and dynamic traffic assignment (DTA). In this issue of the seminar, we strongly profited from the fact that the community has grown together over the last three Dagstuhl Seminars, and we developed a good mutual understanding and a common language. On the one hand, multiple talks and open problems addressed solutions or progress on questions from previous seminars. On the other hand, at this seminar, we focused on sustainability in traffic planning, established new questions, and extended the core community.

Again, the seminar was a great success. The seminar not only stimulates fruitful collaboration, but by now, we have built a community around this series of seminars with many important research questions and a very supportive and open atmosphere. We got laudatory feedback from many participants, which is also reflected in the survey conducted by Dagstuhl.

2 Table of Contents

Executive Summary

José Correa, Carolina Osorio, Laura Vargas Koch, and David Watling 1

Overview of Talks

Review Talk: Dynamic Traffic Assignment: An Incomplete Review
Hong K. LO 5

Review Talk: Ride-hailing / Ridesharing Operations
Chiwei Yan 5

Review Talk: Combinatorial explorations in the Vickrey bottleneck model
Neil Olver 6

Some open problems in freight train scheduling
Katharina Eickhoff 6

Designing a Public Transport Network that Integrates Fixed Lines and On-Demand
 Mobility
Andrés Salomón Fielbaum Schnitzler 6

Strategic freight transport models with both cooperative and competitive behavior
Gunnar Flötteröd 7

Review Talk: The History and Likely Future of Dynamic User Equilibrium
Terry L. Friesz 7

Oscillating long-term behavior of user equilibria in dynamic traffic models with
 spillback
Theresa Ziemke 7

A Game-theoretical Model of Road Pricing with Endogenized Multi-modal User-
 equilibrium
Gaurav Malik 8

On the Price of Anarchy in Packet Routing Games with FIFO
Martin Strehler 8

A Decomposition Theorem for Dynamic Flows
Lukas Graf 9

Tolls for Dynamic Equilibrium Flows
Tobias Harks 9

Congestion tolling: dollars vs tokens
Ravi Seshadri 9

Establishing cooperation in an evolutionary environment: towards an efficient
 sharing of transport resources in communities
Takamasa Iryo 10

A Recursive Logit Model for Vacant Ride-Sourcing Vehicle Routing Behavior
Song Gao 10

Game-theoretic analysis of user behaviour in reserving transport services
Koki Satsukawa 11

Nash flows over time with tolls <i>Marc Schröder</i>	11
Approximation and Convergence of Large Atomic Congestion Games <i>Marco Scarsini</i>	11
Ordinary and Prophet Planning under Uncertainty in Bernoulli Congestion Games <i>Nicolas Stier-Moses</i>	12
Optimizing Throughput and Makespan of Queuing Systems by Information Design <i>Svenja M. Griesbach</i>	12
Information Design for Congestion Games with Unknown Demand <i>Max Klimm</i>	13
On the Controllability of Combinatorial Problems <i>Jannik Matuschke</i>	13
On cyberattacks in traffic <i>Saif Jabari</i>	14
Open problems	
Bicriteria Nash Flows over Time <i>Tim Oosterwijk</i>	14
Variance of the Travelling Salesman Tour Length for Uniformly Randomly Distributed Points in a Disc <i>Richard Connors</i>	15
Safety distance between trains <i>Katharina Eickhoff</i>	15
Participants	16

3 Overview of Talks

3.1 Review Talk: Dynamic Traffic Assignment: An Incomplete Review

Hong K. LO (The Hong Kong Univ. of Science & Technology, HK)

License © Creative Commons BY 4.0 International license
© Hong K. LO

Traffic Assignment, which predicts the traffic loading on a transportation network, is central to transportation planning and operations. With the advent of real-time information and control, autonomous vehicles, dynamic traffic assignment (DTA) becomes all the more important, for both offline planning and real-time operations. I started my career working on DTA and the related topic dynamic traffic control in the 90's, and then on and off I continued for the last few decades. When I felt the problem got too hard and I could not contribute, I stopped. However, when I found new studies or new approaches to tackle the problem, I felt revived and continued. Hence, this review represents my personal DTA journey, which is incomplete and perhaps biased. It covers travel choice principles expressed as Variational Inequality Problem (VIP), Non-linear Complementarity Problem (NCP), mathematical programming, dynamical systems, touches upon deterministic and stochastic dynamics, attraction domains, stability issues, and equilibrium versus non-equilibrium approaches. As for traffic modeling, this review mainly contrasts static and dynamic traffic models, point-queue versus spatial-queue approaches, and their implications on the modeling results. This review ends with how to extend DTA for bi-level large-scale applications via machine learning techniques, our ongoing work.

3.2 Review Talk: Ride-hailing / Ridesharing Operations


Chiwei Yan (University of California – Berkeley, US)

License © Creative Commons BY 4.0 International license
© Chiwei Yan

In this review talk, I will explore the primary operational challenges faced by ride-hailing and ridesharing platforms, with a bias toward pricing strategies. I will begin by examining ride-hailing services for solo rides and walk through the fundamentals of a simple surge pricing model. Despite being practiced for over a decade, a universally accepted surge pricing algorithm has yet to emerge in either academic literature or practical application. I will talk about a strawman proposal and initiate a discussion on what the future directions could be. Next, I will discuss my recent work on pricing shared rides, allowing multiple riders into the same vehicle. A new challenge arises in this context where the cost of a ride is uncertain beforehand. Contrary to the popular upfront pricing strategies, I argue that platforms should implement a two-tiered pricing policy where prices depend on matching outcomes, demonstrating that this approach results in a win-win-win situation for riders, platforms, and the environment. If time permits, I will conclude with a brief discussion on complementary work involving strategic agents.

3.3 Review Talk: Combinatorial explorations in the Vickrey bottleneck model

Neil Olver (London School of Economics & Political Science, GB)

License  Creative Commons BY 4.0 International license
© Neil Olver

The Vickrey bottleneck model is a simple, yet relevant, model of congestion in dynamic traffic networks. I will survey our current theoretical understanding of the behaviour of equilibria in traffic networks under this model. A focus will be on the very combinatorial structure of the model, and how this gives insight into some of its fundamental properties.

Beyond the most basic setting, there are many orthogonal directions in which the model can be modified or extended to capture more aspects of real-world traffic. I will discuss some interesting questions that arise amongst the many possible combinations of these variations, with an emphasis on those motivated by the perspective of transportation economics.

3.4 Some open problems in freight train scheduling

Katharina Eickhoff (RWTH Aachen, DE)

License  Creative Commons BY 4.0 International license
© Katharina Eickhoff

Nowadays, freight train scheduling becomes an important problem as the traffic volume increases and more goods should be transported by rail instead of road. I present a simplification of this model with extensions and variants. These simplifications should be used to understand the underlying structure and the complexity, which may help in the future to deal with the more realistic, but more complex models.

For a specific scenario, I will give some simple results on complexity, structure and approximability. How these can be extended or related to more complex versions is left as an open problem.

3.5 Designing a Public Transport Network that Integrates Fixed Lines and On-Demand Mobility

Andrés Salomón Fielbaum Schnitzler (University of Sydney, AU)


License  Creative Commons BY 4.0 International license
© Andrés Salomón Fielbaum Schnitzler

With the widespread adoption of on-demand mobility, cities have started to implement on-demand public transport (ODPT). In this talk, we first propose that ODPT can be used as an alternative for passengers whose fixed-line transport option is inconvenient. Focusing on a scenario where ODPT complements bus lines, we propose an incentive-compatible mechanism to split the demand between buses and ODPT.

Second, we show that when ODPT is not door-to-door, vehicles accumulate in main streets, which can effectively be used as a marker for pickup points. We obtain promising results which also might have implications in terms of the fundamental patterns and dynamics underlying people's mobility.

3.6 Strategic freight transport models with both cooperative and competitive behavior


Gunnar Flötteröd (Linköping University, SE)

License  Creative Commons BY 4.0 International license
© Gunnar Flötteröd

Person transport models appear to have been more broadly studied than freight models, arguably because of the relatively less complicated structure of the former. Given the ambition to construct a (national) freight assignment model, this talk attempts to reduce the freight assignment problem to structures known from person transport assignment models. Some interesting differences remain.

3.7 Review Talk: The History and Likely Future of Dynamic User Equilibrium

Terry L. Friesz (Pennsylvania State University – University Park, US)

License  Creative Commons BY 4.0 International license
© Terry L. Friesz

In this presentation we review what has been learned about fixed and elastic demand user equilibrium in a dynamic setting over the last decade. We also discuss the challenges posed by extensions to consider mixed autonomous vehicle (AV) and human-driven vehicle (HDV) flow. We make conjectures about algorithmic innovations and the role of digital twins.

3.8 Oscillating long-term behavior of user equilibria in dynamic traffic models with spillback

Theresa Ziemke (TU Berlin, DE)

License  Creative Commons BY 4.0 International license
© Theresa Ziemke

In this talk, we study the long-term behavior of dynamic traffic equilibria and find that it heavily depends on whether spillback is captured in the traffic model or not. We provide an example for the deterministic queuing model of flows over time (Koch and Skutella, 2011) where no steady state is reached when spillback effects are present. Even though this example comprises a single-commodity instance with a constant inflow rate, the Nash flow over time exhibits an infinite number of phases. This stands in contrast to what has been proven for Nash flows over time without spillback (Cominetti et al., 2021; Olver et al., 2022). In general, this highlights the importance of dynamic transport models: There are effects over time that do not vanish after an initial warm-up phase. Although the importance of dynamic transport models has been apparent for real-world applications with high time dependency, it is still a common motivation for the use of static models that they model the situation beyond the initial warm-up phase, i.e., the steady state. However, our example shows that such a steady state does not necessarily exist. Additionally, we show that similar phase oscillations, as in the Nash flow over time with spillback, can be observed in the agent-based transport simulation MATSim (Horni et al., 2016). This reveals that the non-existence of a steady

state with spillback is not limited to continuous settings (as the simulation models discrete time steps and discrete vehicles). As MATSim is based on a co-evolutionary process which iteratively improves the users' choice, it does not result in exact user equilibria. The fact that the Nash flow phase oscillations can still be observed supports the robustness of this finding. However, the results of the simulation also show that the phase oscillations become rarer and vanish in the long term, the more stochasticity is included in the co-evolutionary process.

3.9 A Game-theoretical Model of Road Pricing with Endogenized Multi-modal User-equilibrium

Gaurav Malik (KU Leuven, BE)

License  Creative Commons BY 4.0 International license
© Gaurav Malik

Modern-day transportation systems involve many stakeholders who influence the overall impacts of each other's decisions. Therefore, it is imperative to model all pertinent stakeholders to reasonably predict scenarios and facilitate informed decision-making among them. We present a Game-theoretical model addressing these motivations. It can be applied to pricing decisions of public and private players/stakeholders in multi-modal transportation systems with different types of travelers.

3.10 On the Price of Anarchy in Packet Routing Games with FIFO


Martin Strehler (Westfälische Hochschule Zwickau, DE)

License  Creative Commons BY 4.0 International license
© Martin Strehler

We investigate packet routing games in which network users selfishly route themselves through a network over time, aiming to reach the destination as quickly as possible. Edges have limited capacities of packets that are forwarded in each discrete time step, and we assume that packets to be forwarded are chosen based on the first-in, first-out (FIFO) principle. Our work builds upon the line of research on packet routing games initiated by Werth, Holzhauser, and Krumke. Interestingly, researchers have established inefficiency results across numerous variations of the model, but for the arguably most natural one, there have been no non-trivial bounds on the price of anarchy. We derive the first non-trivial bounds for packet routing games with the FIFO policy. Specifically, we show that the price of anarchy is at most 2 for the important and well-motivated class of uniformly fastest route equilibria introduced by Scarsini, Schröder, and Tomala. on any linear multigraph. We complement our results with a series of instances on linear multigraphs, where the price of stability converges to at least $\frac{e}{e-1}$. Furthermore, our instances provide a lower bound for the price of anarchy of continuous Nash flows over time on linear multigraphs. This establishes the first lower bound of $\frac{e}{e-1}$ on a graph class where the monotonicity conjecture is proven by Correa, Cristi, and Oosterwijk.

3.11 A Decomposition Theorem for Dynamic Flows

Lukas Graf (Universität Passau, DE)


License  Creative Commons BY 4.0 International license
© Lukas Graf

The famous edge flow decomposition theorem of Gallai (1958) states that any static edge s,d -flow in a directed graph can be decomposed into a linear combination of incidence vectors of paths and cycles. We now study the decomposition problem for the setting of dynamic edge s,d -flows assuming a quite general dynamic flow propagation model. We prove the following decomposition theorem: For any dynamic edge s,d -flow with finite support, there exists a decomposition into a linear combination of s,d -walk inflows and circulations, i.e. edge flows that circulate along cycles with zero transit time. We show that a variant of the classical algorithmic approach of iteratively subtracting walk inflows from the current dynamic edge flow converges to a dynamic circulation. The algorithm terminates in finite time, if there is a lower bound on the minimum edge travel times. We further characterize those dynamic edge flows which can be decomposed purely into linear combinations of s,d -walk inflows.

The proofs rely on the new concept of parameterized network loadings which describe how particles of a different walk flow would hypothetically propagate throughout the network under the fixed travel times induced by the given edge flow. We show several technical properties of this type of network loading and as a byproduct we also derive some general results on dynamic flows which could be of interest outside the context of this paper as well.

3.12 Tolls for Dynamic Equilibrium Flows


Tobias Harks (Universität Passau, DE)

License  Creative Commons BY 4.0 International license
© Tobias Harks

We consider dynamic network flows and study the following question: Which dynamic edge flows can be obtained as tolled dynamic equilibrium flows? In the talk I give a characterization of such flows using strong duality of an associated infinite dimensional linear optimization problem.

3.13 Congestion tolling: dollars vs tokens

Ravi Seshadri (Technical University of Denmark – Lyngby, DK)

License  Creative Commons BY 4.0 International license
© Ravi Seshadri

The case for some form of congestion tolling has long been made given the extent of traffic congestion in most urban transportation networks. However, there is little consensus on whether this tolling should be in the form of dollars (traditional congestion pricing schemes or price regulation) or in the form of tokens (credit-based mobility schemes or quantity regulation). We compare the performance of price and quantity instruments under uncertainty using a simple transportation network consisting of parallel highway routes and a public transport alternative. The demand for each route is determined by a logit mixture model

and the supply consists of static congestion. The comparison is based on the optimum social welfare which is computed for each instrument by solving a non-convex optimization problem involving stochastic user equilibrium constraints. We find that when the tolling system is not day-to-day adaptive and the supply of tokens is fixed, the quantity instrument performs better than the price instrument typically when the marginal external congestion function has a steeper slope than the demand function. The results make a potential case for tolling in tokens.

3.14 Establishing cooperation in an evolutionary environment: towards an efficient sharing of transport resources in communities

Takamasa Iryo (Tohoku University – Sendai, JP)

License  Creative Commons BY 4.0 International license
© Takamasa Iryo

A shared transport service is expected to replace ineffective public transport in rural areas with sparse demand and car transport, which is often not a sustainable transport mode for elderly people. We want to assess how such a shared transport service can emerge through a cooperation process in a community. We first need to assess whether the given setting has a state where users can benefit from a shared ride. This is a cooperative game. We constructed a simple model and showed the existence of the core in a certain condition. On the other hand, establishing a cooperative environment in a community is not always easy. We proposed a simple stochastic day-to-day dynamical model in which each user in a community revises the choice (private mode or shared mode), and the fare is controlled to maintain the long-term neutral budget condition. A preliminary result implies that an appropriate setting of the maximum fare would be beneficial for the successful promotion of the shared ride in a community.

This talk included an initial thought of the joint study with David Watling, Koki Satsukawa, Haruko Nakao, Richard Connors, and Sowa Suzuki.

3.15 A Recursive Logit Model for Vacant Ride-Sourcing Vehicle Routing Behavior


Song Gao (University of Massachusetts – Amherst, US)

License  Creative Commons BY 4.0 International license
© Song Gao

We develop and empirically estimate a recursive logit model to predict vacant ride-sourcing vehicle's routing behavior. A vacant taxi is assumed to follow a Markov Decision Process and the observed trajectories are results from combining a sequential link choice model at each state (node) and state transition probabilities derived from a queuing theory based passenger/driver matching process. The logit model parameters are estimated by maximizing the trajectory likelihood function, and parallel computing is utilized to speed up the estimation. A case study is conducted using taxi GPS data in Shanghai, China with a large-scale network of over 10,000 arcs. Expected fare, expected operating cost, link size (a measure of route overlapping) and number of intersections are found to be significant predictors. Comparison with a base model stopping the MDP at the first matching suggests the advantage of considering multiple pick-up and drop-off cycles.

3.16 Game-theoretic analysis of user behaviour in reserving transport services


Koki Satsukawa (Kanazawa University – Ishikawa-ken, JP)

License  Creative Commons BY 4.0 International license
© Koki Satsukawa

We consider the stability of dynamic user equilibrium (DUE) with atomic users. The DUE problem is formulated as a strategic game wherein atomic users strategically select their routes so as to minimise their route travel times. To analyse the equilibrium, we focus on a particular type of user, called 'the earliest non-assigned user'. We then show that when the existence of the earliest non-assigned user is guaranteed, (i) the existence of a pure Nash equilibrium (i.e. DUE) is guaranteed. (ii) a constructive algorithm is obtained. (iii) the DUE game belongs to a weakly acyclic game, and the stochastic stability of the equilibrium is guaranteed.

3.17 Nash flows over time with tolls

Marc Schröder (Maastricht Univ. School of Business & Economics, NL)

License  Creative Commons BY 4.0 International license
© Marc Schröder

We introduce (fixed) tolls to the deterministic fluid queuing model (also known as Vickrey's bottleneck model) and test whether some of the known properties of dynamic equilibria carry over. Using three examples we show that (1) dynamic equilibria with tolls need not be unique, (2) particles might overtake in dynamic equilibria with tolls, (3) dynamic equilibria with tolls need not reach a steady state, but seem to converge to one.

3.18 Approximation and Convergence of Large Atomic Congestion Games

Marco Scarsini (LUISS University – Rome, IT)


License  Creative Commons BY 4.0 International license
© Marco Scarsini

We consider the question of whether and in what sense, Wardrop equilibria provide a good approximation for Nash equilibria in atomic unsplitable congestion games with a large number of small players. We examine two different definitions of small players. In the first setting, we consider games in which each player's weight is small. We prove that when the number of players goes to infinity and their weights to zero, the random flows in all (mixed) Nash equilibria for the finite games converge in distribution to the set of Wardrop equilibria of the corresponding nonatomic limit game. In the second setting, we consider an increasing number of players with a unit weight that participate in the game with a decreasingly small probability. In this case, the Nash equilibrium flows converge in total variation toward Poisson random variables whose expected values are Wardrop equilibria of a different nonatomic game with suitably defined costs. The latter can be viewed as symmetric equilibria in a Poisson game in the sense of Myerson, establishing a plausible connection

between the Wardrop model for routing games and the stochastic fluctuations observed in real traffic. In both settings, we provide explicit approximation bounds, and we study the convergence of the price of anarchy. Beyond the case of congestion games, we prove a general result on the convergence of large games with random players toward Poisson games.

3.19 Ordinary and Prophet Planning under Uncertainty in Bernoulli Congestion Games


Nicolas Stier-Moses (Meta – Menlo Park, US)

License  Creative Commons BY 4.0 International license
© Nicolas Stier-Moses

We consider an atomic congestion game in which each player i participates in the game with an exogenous and known probability $p_i \in (0, 1]$, independently of everybody else, or stays out and incurs no cost. We compute the parameterized PoA to characterize the impact of demand uncertainty on the efficiency of selfish behavior, considering two different notions of a social planner. A prophet planner knows the realization of the random participation in the game; the ordinary planner does not. As a consequence, a prophet planner can compute an adaptive social optimum that selects different solutions depending on the players that turn out to be active, whereas an ordinary planner faces the same uncertainty as the players and can only minimize the expected social cost according to the player participation distribution. For both type of planners we obtain tight bounds for the PoA, by solving suitable optimization problems parameterized by the maximum participation probability $q = \max_i p_i$. In the case of affine costs, we find an analytic expression for the corresponding bounds.

3.20 Optimizing Throughput and Makespan of Queuing Systems by Information Design

Svenja M. Griesbach (Universidad de Chile, CL)


License  Creative Commons BY 4.0 International license
© Svenja M. Griesbach

We study the optimal provision of information for two natural performance measures of queuing systems: throughput and makespan. A set of parallel links (queues) is equipped with deterministic capacities and stochastic travel times where the latter depend on a realized scenario, and the number of scenarios is assumed to be constant. A continuum of flow particles (users) arrives at the system at a constant rate. A system operator knows the realization of the scenario and may (partially) reveal this information via a public signaling scheme to the flow particles. Upon arrival, the flow particles observe the signal issued by the system operator, form an updated belief about the realized scenario, and decide on a link to use. Inflow into a link exceeding the link's capacity builds up in a queue that increases the travel time on the link. Dynamic inflow rates are in a Bayesian dynamic equilibrium when the expected travel time along all links with positive inflow is equal at every point in time. For a given time horizon T , the throughput induced by a signaling scheme is the total volume of flow that leaves the links in the interval $[0, T]$. We show that the public signaling scheme that maximizes the throughput may involve irrational numbers and provide

an additive polynomial time approximation scheme (PTAS) that approximates the optimal throughput by an arbitrary additive constant $\varepsilon > 0$. The algorithm solves a Lagrangian dual of the signaling problem with the Ellipsoid method whose separation oracle is implemented by a cell decomposition technique. We also provide a multiplicative fully polynomial time approximation scheme (FPTAS) that does not rely on strong duality and, thus, allows to compute also the optimal signals. It uses a different cell decomposition technique together with a piece-wise convex under-estimator of the optimal value function. Finally, we consider the makespan of a Bayesian dynamic equilibrium which is defined as the last point in time when a total given value of flow leaves the system. Using a variational inequality argument, we show that full information revelation is a public signaling scheme that minimizes the makespan.

3.21 Information Design for Congestion Games with Unknown Demand

Max Klimm (TU Berlin, DE)

License  Creative Commons BY 4.0 International license
© Max Klimm

We study a novel approach to information design in the standard traffic model of network congestion games. It captures the natural condition that the demand is unknown to the users of the network. A principal (e.g., a mobility service) commits to a signaling strategy, observes the realized demand and sends a (public) signal to agents (i.e., users of the network). Based on the induced belief about the demand, the users then form an equilibrium. We consider the algorithmic goal of the principal: Compute a signaling scheme that minimizes the expected total cost of the induced equilibrium. We concentrate on single-commodity networks and affine cost functions, for which we obtain the following results. First, we devise a fully polynomial-time approximation scheme (FPTAS) for the case that the demand can only take two values. It relies on several structural properties of the cost of the induced equilibrium as a function of the updated belief about the distribution of demands. We show that this function is piecewise linear for any number of demands, and monotonic for two demands. Second, we give a complete characterization of the graph structures for which it is optimal to fully reveal the information about the realized demand. This signaling scheme turns out to be optimal for all cost functions and probability distributions over demands if and only if the graph is series-parallel. Third, we propose an algorithm that computes the optimal signaling scheme for any number of demands whose time complexity is polynomial in the number of supports that occur in a Wardrop equilibrium for some demand. Finally, we conduct a computational study that tests this algorithm on real-world instances.

3.22 On the Controllability of Combinatorial Problems

Jannik Matuschke (KU Leuven, BE)

License  Creative Commons BY 4.0 International license
© Jannik Matuschke

Consider a finite ground set E , a set of solutions $X \subseteq \mathbb{R}^E$ and a class of objective functions \mathcal{C} on X . We are interested in subsets $S \subseteq E$ that can *control* X in the sense that we can induce any given solution $x \in X$ as an optimum for any given objective function $c \in \mathcal{C}$ by

adding linear terms to c on the coordinates corresponding to S . We observe that if X is either a convex set or consists of binary vectors, a set S controls X if and only if it enables us to *identify* any given solution by its coordinates on S . For the case that X is convex, we moreover show that the sets controlling X induce a matroid. As a consequence, min-weight controlling sets can be computed efficiently from a representation of the affine hull of X in this case.

While the aforementioned result extends to the case where X is the set of bases of a matroid, other natural discrete structures are much less tractable: In particular, when X is the set of s - t -paths in a directed graph, deciding whether an identifying set of a certain cardinality exists is Σ_2^P -complete. The problem remains NP-hard even when the underlying graph is acyclic, but we derive an approximation for this case by deriving a tight bound on the gap between the size identifying solutions for X and the size of identifying solutions for its convex hull.

3.23 On cyberattacks in traffic

Saif Jabari (New York University – Abu Dhabi, AE)

License  Creative Commons BY 4.0 International license
© Saif Jabari

Self-driving vehicle deployment is on the rise in many cities worldwide, mainly as robotaxis (at the time of this writing). Autonomous vehicles (AVs) use deep learning to perform various tasks, from environment perception (e.g., identifying traffic signs, pedestrians, and lane markings) to executing control decisions (e.g., braking, acceleration, and lane changing). Deep neural networks (DNNs) are highly vulnerable to structured perturbations to inputs, so-called “adversarial attacks.” They are also susceptible to adversarial training samples that trigger poor behavior. This presentation first provides a crash course in adversarial training, followed by two types of perturbations that affect the DNNs that AVs use. The first perturbation targets the graph inputs commonly used in object detection tools for AVs. I explore using an error-correcting scheme to reconstruct potentially perturbed graphs and share preliminary work that produces a performance guarantee for the error-correcting scheme. The second perturbation I will present takes place at the time of training, where small “adversarial” samples, added during training, create hidden vulnerabilities that an adversary can trigger post-training. I present an example of such trojans crafted using dynamic traffic flow modeling and show that such trojans can be very stealthy, affecting the operation of an AV-based congestion control system.

4 Open problems

4.1 Bicriteria Nash Flows over Time

Tim Oosterwijk (VU Amsterdam, NL)

License  Creative Commons BY 4.0 International license
© Tim Oosterwijk

Can these models with particles that are not so singleminded be extended, to for example other objective functions, linear tardiness costs, tolls or heterogeneous deadlines?

4.2 Variance of the Travelling Salesman Tour Length for Uniformly Randomly Distributed Points in a Disc

Richard Connors (University of Luxembourg, LU)

License  Creative Commons BY 4.0 International license
© Richard Connors

Consider a disc of area A centered at the origin, within which we generate N points uniformly randomly. We solve the traveling salesman problem (TSP) to find the tour of minimum length, L , that visits these N points plus the origin (including the origin is not crucial but is the problem of interest here). Clearly, the optimal tour length tends to increase with A and N . A well-known result from Beardwood et al (1959) “The shortest path through many points” is that the optimal tour length scales like $L(A, N) \sim \beta A^{\frac{1}{2}} N^{\frac{1}{2}}$, asymptotically where β is a constant. However, for any fixed A and N , each randomly generated instance will result in a different optimal TSP tour, with different lengths. Open Problem: When N points are generated randomly uniformly in a disc of area A , what is the variance of the optimal TSP tour lengths $\text{Var}[L(A, N)]$ as a function of A and N ?

4.3 Safety distance between trains

Katharina Eickhoff (RWTH Aachen, DE)

License  Creative Commons BY 4.0 International license
© Katharina Eickhoff

We have given a directed acyclic network $G = (V, A)$ with a source $s \in V$ and a sink $t \in V$. There are travel times $\tau_a \in \mathbb{Z}_{>0}$ for each arc $a \in A$ (the same for all trains). We want to embed k trains in this network, where these trains have to keep a safety distance D in between two consecutive trains, i.e., a train can enter an arc after at least D time units from the point in time the train before it enters this arc. Note that the length is also included in D , i.e., D is actually the sum of the length and the safety distance. The goal is to minimize the time when the last train arrives at t (aka makespan).

It is known that this problem is weakly NP-hard (already for two trains). For three trains, we were able to show that there is an optimal solution where the trains travel in convoys, i.e., any two trains either travel along the same, or along disjoint paths. This is achieved by a rearranging an optimal solution without changing the objective value.

Question (“Convoy Conjecture”): For at least four trains, is there always an optimal solution where the trains travel in convoys?

The problem with rearranging more than three trains in this setting is, that it may be necessary to do simultaneous swaps to keep the solution optimal.

If the “convoy conjecture” is true, we know that intermediate waiting is not necessary, and that we can reduce the problem to finding disjoint paths in the underlying network G , which is much faster than computing *very* disjoint paths in the time-expanded network. This can be solved by an FPTAS for a fixed number of trains. However, this algorithm would grow exponentially with the number k of trains, and the question remains whether this can be improved.

Partial answers from the seminar: During the seminar we were able to construct a D -approximation algorithm for maximizing the number of trains in a given time horizon T and a D -approximation algorithm for minimizing the time horizon T (aka makespan) for sending k trains. Moreover, we showed that both of these problems are strongly NP-hard via a reduction from numerical 3D matching.

Participants

- Umang Bhaskar
TIFR Mumbai, IN
- Roberto Cominetti
Adolfo Ibáñez University –
Santiago, CL
- Richard Connors
University of Luxembourg, LU
- José R. Correa
University of Chile –
Santiago de Chile, CL
- Katharina Eickhoff
RWTH Aachen, DE
- Andrés Salomón Fielbaum
Schnitzler
University of Sydney, AU
- Gunnar Flötteröd
Linköping University, SE
- Martin Gairing
University of Liverpool, GB
- Song Gao
University of Massachusetts –
Amherst, US
- Lukas Graf
Universität Passau, DE
- Svenja M. Griesbach
University of Chile –
Santiago de Chile, CL
- Tobias Harks
Universität Passau, DE
- Martin Hofer
RWTH Aachen, DE
- Takamasa Iryo
Tohoku University – Sendai, JP
- Saif Jabari
New York University –
Abu Dhabi, AE
- Max Klimm
TU Berlin, DE
- Ekkehard Köhler
BTU Cottbus-Senftenberg, DE
- Hong K. Lo
The Hong Kong Univ. of Science
& Technology, HK
- Gaurav Malik
KU Leuven, BE
- Jannik Matuschke
KU Leuven, BE
- Haruko Nakao
University of Luxembourg –
Esch-sur-Alzette, LU
- Neil Olver
London School of Economics &
Political Science, GB
- Tim Oosterwijk
VU Amsterdam, NL
- Carolina Osorio
HEC – Montréal, CA & Google
Research – Mountain View, US
- Dario Paccagnan
Imperial College London, GB
- Britta Peis
RWTH Aachen, DE
- Koki Satsukawa
Kanazawa University –
Ishikawa-ken, JP
- Marco Scarsini
LUISS University – Rome, IT
- Daniel Schmand
Universität Bremen, DE
- Marc Schröder
Maastricht Univ. School of
Business & Economics, NL
- Ravi Seshadri
Technical University of Denmark
– Lyngby, DK
- Alexander Skopalik
University of Twente –
Enschede, NL
- Nicolas Stier-Moses
Meta – Menlo Park, US
- Martin Strehler
Westfälische Hochschule
Zwickau, DE
- Laura Vargas Koch
RWTH Aachen, DE
- Bernhard von Stengel
London School of Economics &
Political Science, GB
- Peter Wagner
DLR – Berlin, DE
- David Watling
University of Leeds, GB
- Chiwei Yan
University of California –
Berkeley, US
- Theresa Ziemke
TU Berlin, DE

