

Dynamic Traffic Models in Transportation Science

Martin Gairing^{*1}, Carolina Osorio^{*2}, Britta Peis^{*3}, David Watling^{*4},
and Katharina Eickhoff^{†5}

1 University of Liverpool, GB. gairing@liverpool.ac.uk

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

3 RWTH Aachen, DE. britta.peis@oms.rwth-aachen.de

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

5 RWTH Aachen, DE. katharina.eickhoff@oms.rwth-aachen.de

Abstract

Traffic assignment models are crucial for transport planners to be able to predict the congestion, environmental and social impacts of transport policies, for example in the light of possible changes to the infrastructure, to the transport services offered, or to the prices charged to travellers. The motivation for this series of seminars – of which this seminar was the third – is the prevalence in the transportation community of basing such predictions on complex computer-based simulations that are capable of resolving many elements of a real systems, while on the other hand, the theory of dynamic traffic assignments (in terms of equilibrium existence, computability and efficiency) had not matured to the point matching the model complexity inherent in simulations.

Progress has been made on this issue in the first two seminars (Dagstuhl Seminar 15412 and 18102), by bringing together leading scientists in the areas of *traffic simulation*, *algorithmic game theory* and *dynamic traffic assignment*. We continued this process this seminar. Moreover, we started to address the growing real-life challenge of new kinds of 'mobility service' emerging, before the tools are available to incorporate them in such planning models. These services include intelligent/dynamic ride-sharing and car-sharing, through to fully autonomous vehicles, provided potentially by a variety of competing operators.

Seminar May 8–13, 2022 – <http://www.dagstuhl.de/22192>

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.12.5.92

1 Executive Summary

Martin Gairing (University of Liverpool, GB)

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

Britta Peis (RWTH Aachen, DE)

David Watling (University of Leeds, GB)

License  Creative Commons BY 4.0 International license

© Martin Gairing, Carolina Osorio, Britta Peis, and David Watling

Traffic assignment models are crucial for transport planners to be able to predict the traffic distributions – 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

* 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. 12, Issue 5, pp. 92–111

Editors: Martin Gairing, Carolina Osorio, Britta Peis, and David Watling



DAGSTUHL
REPORTS

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

controls), to the transport services offered (e.g. the provision of public transport fleets and frequencies), or to the prices imposed on travellers (e.g. road tolls, 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, activity schedule choice, and such. 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” such as formal proofs for the following fundamental properties:

- a (dynamic) equilibrium always exists,
- an equilibrium is unique,
- an equilibrium is efficiently computable,
- there is a smooth function of equilibria with respect to parameters (such as street capacities, tolls, etc.)

In this seminar we brought together, for the third 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 doing so, we build on the progress made in the first two seminars, and at the same time addressed the emerging real-life challenge of modeling the new kinds of ‘mobility service’ that are increasingly envisaged for the future, such as intelligent/dynamic ride-sharing and car-sharing, through to fully autonomous vehicles, provided potentially by a variety of competing operators.

The purpose of this third seminar was to build on the momentum, collaborations and new directions identified at the previous two seminars, following on from the impacts identified earlier. Similarly to the first two seminars, one objective was to make new advances on identifying, formulating and solving existing open problems by combining insights from the different academic fields. Thus, we continued working on themes, initiated at the previous Dagstuhl seminars, related to the representation of queueing in dynamic models and the identification of new requirements and new open problems arising from future forms of mobility and travellers’ behaviour, for example, systems of autonomous vehicles, and coordinated/shared routing.

Again, the seminar was a big success. Beside forster collaboration of the last seminars of this series, the seminar stimulates new and very fruitful ones. We got laudatory feedback from many participants which is also reflected in the survey conducted by Dagstuhl.

2 Table of Contents

Executive Summary

<i>Martin Gairing, Carolina Osorio, Britta Peis, and David Watling</i>	92
--	----

Overview of Talks

Towards day-to-day assignment game models for mobility ecosystems <i>Joseph Chow</i>	96
Tutorial Talk: Integrated travel behavior/network modeling <i>Gunnar Flötteröd</i>	96
Existence and Complexity of Approximate Equilibria in Weighted Congestion Games <i>Yiannis Giannakopoulos</i>	97
Dynamic Traffic Assignment for Electric Vehicles <i>Lukas Graf</i>	97
Tutorial Talk: The role of Information in DTA Models <i>Tobias Harks</i>	98
Signaling in Network Congestion Games <i>Martin Hoefer</i>	98
Stability and Instability in DTA <i>Takamasa Iryo</i>	98
Equilibria in Multiclass and Multidimensional Atomic Congestion Games <i>Max Klimm</i>	99
Something about traffic lights and bicycles <i>Ekkehard Köhler and Martin Strehler</i>	99
Hitting paths with random sets in abstract networks <i>Jannik Matuschke</i>	99
Tutorial Talk: Continuity, Uniqueness and Long-Term Behaviour of Nash Flows Over Time <i>Neil Olver</i>	100
Combining Bayesian optimization with analytical dynamic traffic models to tackle dynamic simulation-based transportation optimization problems <i>Carolina Osorio</i>	100
In Congestion Games, Taxes Achieve Optimal Approximation <i>Dario Paccagnan and Martin Gairing</i>	101
Stackelberg Max Closure <i>Britta Peis</i>	101
The Complexity of Gradient Descent <i>Rahul Savani</i>	102
Learning Augmented Shortest Paths: Dijkstra with Predictions <i>Guido Schäfer</i>	102
Bicriteria Nash Flows over Time <i>Daniel Schmand</i>	102

Something about Complexity of Equilibria or Bicycles <i>Alexander Skopalik</i>	103
Tutorial Talk: A Faster Algorithm for Quickest Transshipments via an Extended Discrete Newton Method <i>Martin Skutella</i>	103
Stackelberg Max Closure <i>Bernhard von Stengel</i>	104
Open problems	
Controllability of Traffic Networks <i>Richard Connors</i>	104
Embedding multiple trains in a given network <i>Katharina Eickhoff</i>	105
Concepts for driver's route choice behaviour <i>Takamasa Iryo</i>	106
Hitting Sets, Randomly <i>Jannik Matuschke</i>	106
The Expected Price of Anarchy in Wardrop Routing <i>Daniel Schmand</i>	107
Complexity of network loading in the fluid-queue model <i>Martin Skutella</i>	108
Towards defining a set of open problems for transportation analysis <i>David Watling</i>	109
Identifying paths within a given bound <i>David Watling</i>	110
Participants	111

3 Overview of Talks

3.1 Towards day-to-day assignment game models for mobility ecosystems

Joseph Chow (New York University, US)

License  Creative Commons BY 4.0 International license
 Joseph Chow

While there are many methods in within-day dynamics of on-demand mobility services, the literature on day-to-day dynamics is scarce. In this open problem session, I will present a steady state assignment game model for analyzing and designing platforms of mobility ecosystems that output passenger flows, stable pricing ranges, and platform subsidies. I will discuss challenges and opportunities to identifying an equivalent day-to-day adjustment model corresponding to that model: stable outcome space becoming an attractor, learning functions for user and operators that include cost allocations, using the path-dependency of the dynamic system to tackle the nonuniqueness of stable outcome spaces corresponding to a matching solution, deterministic and stochastic variants, and introducing policy interventions to guide ecosystems toward desired attractors.

3.2 Tutorial Talk: Integrated travel behavior/network modeling

Gunnar Flötteröd (Linköping University, SE & Swedish National Road and Transport Research Institute (VTI) – Stockholm SE)

License  Creative Commons BY 4.0 International license
 Gunnar Flötteröd

Strategic transport planning models that predict travel demand and behavior in consideration of a congested transport system have been used for many decades. Historically, these models are aggregate and continuum flow based, meaning that human behavior and physical vehicle propagation are abstracted into a limited number of representative traveler and vehicle flows. The last few decades have seen the convergence of network flow meso/microsimulations and activity-based travel demand models into “agent-based” model systems that simulate individual synthetic travelers and vehicles instead of largely anonymous flows. Greater realism may be expected from agent-based models, given their potentially higher level of detail. However, careful structural modeling is needed to leverage these capabilities. The talk presents the possible advantages that arise from an agent-based approach and emphasizes the accompanying modeling challenges.

3.3 Existence and Complexity of Approximate Equilibria in Weighted Congestion Games

Yiannis Giannakopoulos (Universität Erlangen-Nürnberg, DE)

License  Creative Commons BY 4.0 International license
© Yiannis Giannakopoulos

Joint work of Giorgos Christodoulou, Martin Gairing, Yiannis Giannakopoulos, Dogo Poças, Clara Waldmann

We study the existence of approximate pure Nash equilibria (α -PNE) in weighted atomic congestion games with polynomial cost functions of maximum degree d . Previously it was known that d -approximate equilibria always exist, while nonexistence was established only for small constants, namely for 1.153-PNE. We improve significantly upon this gap, proving that such games in general do not have $\tilde{\Theta}(\sqrt{d})$ -approximate PNE, which provides the first super-constant lower bound.

Furthermore, we provide a black-box gap-introducing method of combining such nonexistence results with a specific circuit gadget, in order to derive NP-completeness of the decision version of the problem. In particular, deploying this technique we are able to show that deciding whether a weighted congestion game has an $\tilde{O}(\sqrt{d})$ -PNE is NP-complete. Previous hardness results were known only for the special case of exact equilibria and arbitrary cost functions.

The circuit gadget is of independent interest and it allows us to also prove hardness for a variety of problems related to the complexity of PNE in congestion games. For example, we demonstrate that the question of existence of α -PNE in which a certain set of players plays a specific strategy profile is NP-hard for any $\alpha < 3^{d/2}$, even for unweighted congestion games.

Finally, we study the existence of approximate equilibria in weighted congestion games with general (nondecreasing) costs, as a function of the number of players n . We show that n -PNE always exist, matched by an almost tight nonexistence bound of $\tilde{\Theta}(n)$, which we can again transform into an NP-completeness proof for the decision problem.

3.4 Dynamic Traffic Assignment for Electric Vehicles

Lukas Graf (Universität Augsburg, DE)

License  Creative Commons BY 4.0 International license
© Lukas Graf

Joint work of Lukas Graf, Tobias Harks, Prashant Palkar

Electric vehicles (EVs) are a great promise for the coming decades in order to allow for mobility but at the same time take measures against the climate change by reducing the emissions of classical combustion engines.

However, those new types of vehicles also come with new challenges like the limited range and the more complex and time consuming recharging process. Modelling traffic involving EVs, thus, requires a more involved model than the one used for traditional vehicles. To accommodate this, we augment the deterministic queuing model of Vickrey with energy constraints and study capacitated dynamic equilibria in this model. Here, in contrast to the traditional dynamic equilibria, we can neither assume that particles only take simple paths (leading to an infinite dimensional strategy space) nor synchronized intermediate arrival times (i.e. particles starting at the same time from the same source towards the same destination might still arrive at intermediate nodes at different times).

We will explore these difficulties and show how we can overcome them to show existence of equilibria for EVs as well as an even more general traffic models. We complement our theoretical results by a computational study in which we apply a fixed-point algorithm to compute energy- dynamic equilibria.

3.5 Tutorial Talk: The role of Information in DTA Models

Tobias Harks (Universität Augsburg, DE)

License  Creative Commons BY 4.0 International license
© Tobias Harks

I will first give a brief overview on some existing DTA models assuming i) perfect and full information on travel times, flows etc. and the other extreme (ii) assuming only instantaneous (incomplete) information. After this overview I introduce a DTA model, where agents base their instantaneous routing decisions on real-time delay predictions. I formulate a mathematically concise model and derive properties of the predictors that ensure a dynamic prediction equilibrium exists. I will show that the framework subsumes the well-known full information and instantaneous information models, in addition to admitting further realistic predictors as special cases.

3.6 Signaling in Network Congestion Games

Martin Hoefer (Goethe-Universität – Frankfurt am Main, DE)

License  Creative Commons BY 4.0 International license
© Martin Hoefer
Joint work of Svenja Griesbach, Martin Hoefer, Max Klimm, Tim Koglin

A largely untapped potential for improving traffic flows is rooted in the inherent uncertainty of travel times. Large mobility services have an informational advantage over single network users as they are able to learn traffic conditions from data. A benevolent provider may use this informational advantage in order to steer a traffic equilibrium into a favorable direction. The resulting optimization problem is a task commonly referred to as signaling or Bayesian persuasion.

In this talk, we discuss our recent progress on this signaling problem. We characterize cases in which optimal signaling schemes are structurally simple and/or can be computed in polynomial time. Our results reveal interesting connections between optimal signaling, network structure, and support enumeration for equilibria.

3.7 Stability and Instability in DTA

Takamasa Iryo (Tohoku University – Sendai, JP)

License  Creative Commons BY 4.0 International license
© Takamasa Iryo

Stability and instability issues in DTA have been studied actively in recent years but still many things are not well understood. This talk will give a brief introduction of existing studies dealing with stability issues in DTA, especially those on for DUE (dynamic user-equilibrium) solutions, and raise possible future research directions to tackle with this problem.

3.8 Equilibria in Multiclass and Multidimensional Atomic Congestion Games

Max Klimm (TU Berlin, DE)

License  Creative Commons BY 4.0 International license
© Max Klimm

We study the existence of pure Nash equilibria in atomic congestion games with different user classes where the cost of each resource depends on the aggregated demand of each class. A set of cost functions is called consistent for this class if all games with cost functions from the set have a pure Nash equilibrium. We give a complete characterization of consistent sets of cost functions showing that the only consistent sets of cost functions are sets of certain affine functions and sets of certain exponential functions. This characterization is also extended to a larger class of games where each atomic player may control flow that belongs to different classes.

3.9 Something about traffic lights and bicycles

Ekkehard Köhler (BTU Cottbus-Senftenberg, DE) and Martin Strehler (Westfälische Hochschule Zwickau, DE)

License  Creative Commons BY 4.0 International license
© Ekkehard Köhler and Martin Strehler

Joint work of Ekkehard Köhler, Markus Rogge, Robert Scheffler, Martin Strehler

We report on work in progress on the question of finding good paths for bicycles in a traffic network with (possibly coordinated) traffic lights. While the standard dynamic shortest path problem in time-dependent networks with cyclic time windows is known to be easy, we study the related problem of finding routes that keep a small number of stops. We show that the problem is strongly NP-hard for paths and trails, whereas it is weakly NP-hard for walks. We give a pseudo-polynomial algorithm for this third option. Further we report about ongoing work on designing algorithms for the case of variable speeds and the employment of appropriate power consumption and recovery models for bicycles for the above formulated problem.

3.10 Hitting paths with random sets in abstract networks

Jannik Matuschke (KU Leuven, BE)

License  Creative Commons BY 4.0 International license
© Jannik Matuschke

Consider a set family \mathcal{P} on a common ground set E , together with a requirement $r(P) \in [0, 1]$ for each $P \in \mathcal{P}$. Each element $e \in E$ is equipped with a marginal probability $p(e)$ such that $\sum_{e \in P} p(e) \geq r(P)$ for all $P \in \mathcal{P}$. In this talk, we investigate the question under which conditions these marginal probabilities can be turned into a probability distribution on the subsets of E such that each $P \in \mathcal{P}$ is intersected by the resulting random set with probability at least $r(P)$.

Motivated by a network security game, Dahan et al. (2001) showed the existence of such a distribution for the case that \mathcal{P} is the family of maximal chains in a partial order (equivalently: the set of s - t -paths in a directed acyclic graph) and the requirements r fulfill

a natural conservation property. We extend this result to the case that \mathcal{P} is an abstract network fulfilling a certain ordering assumption. We also give a significantly simplified argument for the original result, using a classic idea from the theory of flows over time. In particular, our results imply that equilibria for the aforementioned security game can be efficiently computed even when the underlying digraph contains cycles.

3.11 Tutorial Talk: Continuity, Uniqueness and Long-Term Behaviour of Nash Flows Over Time

Neil Olver (London School of Economics, GB)

License  Creative Commons BY 4.0 International license
© Neil Olver

Joint work of Neil Olver, Leon Sering, Laura Vargas Koch

We consider a dynamic model of traffic that has received a lot of attention in the past few years. Users control infinitesimal flow particles aiming to travel from a source to destination as quickly as possible. Flow patterns vary over time, and congestion effects are modelled via queues, which form whenever the inflow into a link exceeds its capacity. We answer some rather basic questions about equilibria in this model: in particular uniqueness (in an appropriate sense), and continuity: small perturbations to the instance or to the traffic situation at some moment cannot lead to wildly different equilibrium evolutions.

To prove these results, we make a surprising connection to another question: whether, assuming constant inflow into the network at the source, do equilibria always eventually settle into a “steady state” where all queue delays change linearly forever more? Cominetti et al. proved this under an assumption that the inflow rate is not larger than the capacity of the network – eventually, queues remain constant forever. We resolve the more general question positively.

3.12 Combining Bayesian optimization with analytical dynamic traffic models to tackle dynamic simulation-based transportation optimization problems

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

License  Creative Commons BY 4.0 International license
© Carolina Osorio

Cities and companies alike often resort to the use of detailed, dynamic, and stochastic traffic simulators to tackle their transportation optimization problems. Methods from the field of simulation-based optimization are often used. However, their use to tackle problems with time-dependent decision variables is limited. In this talk, we present recent work that formulates an analytical, differentiable, and compute efficient dynamic traffic model, and integrates it within a Bayesian optimization (BO) framework to tackle high-dimensional dynamic transportation problems. We present validation results on small network problems, as well as a large-scale traffic signal control case study for Midtown Manhattan, New York City. We discuss how the proposed approach enhances the scalability of BO methods.

3.13 In Congestion Games, Taxes Achieve Optimal Approximation

Dario Paccagnan (Imperial College London, GB) and Martin Gairing (University of Liverpool, GB)

License © Creative Commons BY 4.0 International license
© Dario Paccagnan and Martin Gairing

We consider the problem of minimising social cost in atomic congestion games and show, perhaps surprisingly, that efficiently computed taxation mechanisms yield the *same* performance achievable by the best polynomial time algorithm, even when the latter has full control over the players' actions. It follows that no other tractable approach geared at incentivising desirable system behaviour can improve upon this result, regardless of whether it is based on taxations, coordination mechanisms, information provision, or any other principle. In short: Judiciously chosen taxes achieve optimal approximation.

Three technical contributions underpin this conclusion. First, we show that computing the minimum social cost is NP-hard to approximate within a given factor depending solely on the admissible resource costs. Second, we design a tractable taxation mechanism whose efficiency (price of anarchy) matches this hardness factor, and thus is optimal. As these results extend to coarse correlated equilibria, any no-regret algorithm inherits the same performances, allowing us to devise polynomial time algorithms with optimal approximation.

3.14 Stackelberg Max Closure

Britta Peis (RWTH Aachen, DE)

License © Creative Commons BY 4.0 International license
© Britta Peis
Joint work of Karsten Jungnitsch, Marc Schröder, Britta Peis

In a Stackelberg max closure game, we are given a digraph whose vertices correspond to projects from which firms can choose and whose arcs represent precedence constraints. Some projects are under the control of a leader who sets prices in the first stage of the game, while in the second stage, the firms choose a feasible subset of projects of maximum value. For a single follower, the leader's problem of finding revenue-maximizing prices can be solved in strongly polynomial time. In this paper, we focus on the setting with multiple followers and distinguish two situations. In the case in which only one copy of each project is available (limited supply), we show that the two-follower problem is solvable in strongly polynomial time, whereas the problem with three or more followers is NP-hard. In the case of unlimited supply, that is, when sufficient copies of each project are available, we show that the two-follower problem is already APX-hard. As a side result, we prove that Stackelberg min vertex cover on bipartite graphs with a single follower is APX-hard.

3.15 The Complexity of Gradient Descent

Rahul Savani (RWTH Aachen, DE)

License  Creative Commons BY 4.0 International license
© Rahul Savani

Joint work of John Fearnley, Paul Goldberg, Alexandros Hollender, Rahul Savani

PPAD and PLS are successful classes that each capture the complexity of important game-theoretic problems: finding a mixed Nash equilibrium in a bimatrix game is PPAD-complete; and finding a pure Nash equilibrium in a congestion game is PLS-complete. Many important problems, such as solving a Simple Stochastic Game or finding a mixed Nash equilibrium of a congestion game, lie in both classes. However, it was strongly believed that their intersection does not have natural complete problems. We show that it does: any problem that lies in both classes can be reduced in polynomial time to the problem of finding a stationary point of a function. Our result has been used to show that computing a mixed equilibrium of a congestion game is also complete for the intersection of PPAD and PLS.

3.16 Learning Augmented Shortest Paths: Dijkstra with Predictions

Guido Schäfer (CWI – Amsterdam, NL)

License  Creative Commons BY 4.0 International license
© Guido Schäfer

We study the use of machine learning techniques to solve a fundamental shortest path problem, known as the *single-source many-targets shortest path problem (SSMTSP)*. Basically, our idea is to equip an adapted version of Dijkstra’s algorithm with machine learning predictions to solve this problem: Based on the trace of Dijkstra’s algorithm, we design a neural network that predicts the shortest path distance after a few iterations. The prediction is then used to prune the search space explored by Dijkstra’s algorithm, which significantly reduces the number of operations on the underlying priority queue. Crucially, our algorithms always compute the exact shortest path distances (even if the prediction is inaccurate) and never use more queue operations than the standard algorithm. We provide a probabilistic analysis of our new algorithm which proves a lower bound on the number of saved queue operations. Our bound quantifies these savings depending on the accuracy of the prediction and applies to arbitrary graphs with (partial) random weights. We also present extensive experimental results on random instances showing that the actual savings are oftentimes significantly larger.

3.17 Bicriteria Nash Flows over Time

Daniel Schmand (Universität Bremen, DE)

License  Creative Commons BY 4.0 International license
© Daniel Schmand

Joint work of Tim Oosterwijk, Daniel Schmand, Marc Schröder

Flows over time are a natural way to incorporate flow dynamics that arise in various applications such as traffic networks. In this work we introduce a natural variant of the deterministic fluid queuing model in which users aim to minimize their costs subject to arrival

at their destination before a pre-specified deadline. We determine the existence and the structure of Nash flows over time and fully characterize the price of anarchy for this model. The price of anarchy measures the ratio of the quality of the equilibrium and the quality of the optimum flow, where we evaluate the quality using two different natural performance measures: the throughput for a given deadline and the makespan for a given amount of flow. While it turns out that both prices of anarchy can be unbounded in general, we provide tight bounds for the important subclass of parallel path networks. Surprisingly, the two performance measures yield different results here.

3.18 Something about Complexity of Equilibria or Bicycles

Alexander Skopalik (University of Twente – Enschede, NL)

License  Creative Commons BY 4.0 International license
© Alexander Skopalik

In this talk we revisit the complexity of computing a(n) (approximate) pure equilibrium in atomic congestion games. While our understanding of the complexity of pure equilibria is reasonably well, i.e, it is PLS-hard in general but we have a good picture of the landscape of tractable spacial cases, the same is not true of approximate pure equilibria. We discuss the gap between general inapproximability result of Vöcking and S. and the approximation algorithm of Caragiannis et al. for certain cost function. We briefly outline a possible approach towards closing this gap.

3.19 Tutorial Talk: A Faster Algorithm for Quickest Transshipments via an Extended Discrete Newton Method

Martin Skutella (TU Berlin, DE)

License  Creative Commons BY 4.0 International license
© Martin Skutella
Joint work of Miriam Schlöter, Martin Skutella, Khai Van Tran

The Quickest Transshipment Problem is to route flow as quickly as possible from sources with supplies to sinks with demands in a network with capacities and transit times on the arcs. It is of fundamental importance for numerous applications in areas such as logistics, production, traffic, evacuation, and finance. More than 25 years ago, Hoppe and Tardos presented the first (strongly) polynomial-time algorithm for this problem. Their approach, as well as subsequently derived algorithms with strongly polynomial running time, are hardly practical as they rely on parametric submodular function minimization via Megiddo's method of parametric search. The main contribution of this paper is a considerably faster algorithm for the Quickest Transshipment Problem that instead employs a subtle extension of the Discrete Newton Method. This improves the previously best known running time of $\tilde{O}(m^4k^{14})$ to $\tilde{O}(m^2k^5 + m^3k^3 + m^3n)$, where n is the number of nodes, m the number of arcs, and k the number of sources and sinks.

3.20 Stackelberg Max Closure

Bernhard von Stengel (London School of Economics, GB)

License  Creative Commons BY 4.0 International license
© Bernhard von Stengel

LP duality (the strong duality theorem of linear programming) and the minimax theorem for zero-sum games are considered “equivalent” in the sense that one can easily be proved from the other. However, the classic proof by Dantzig (1951) of LP duality from the minimax theorem is flawed. It needs an additional assumption of strict complementarity. We show that this assumption amounts to assuming the Lemma of Farkas, which implies LP duality directly. We fix this with a new, different proof. We also mention a beautiful existing direct proof of the Lemma of Farkas based on minimally infeasible systems of inequalities.

4 Open problems

4.1 Controllability of Traffic Networks

Richard Connors (University of Luxembourg, LU)

License  Creative Commons BY 4.0 International license
© Richard Connors

Let $G = (N, E)$ be a graph, where N is the set of nodes, and A is the set of arcs. This graph can be taken to represent a transportation network where the nodes are intersections and the arcs are links between intersections. On this network, there is a set of origin-destination (OD) pairs W and a set of paths P which connect OD pairs. The cost, c_a , of travelling on a particular link, a , is given as a sum of the travel time $t_a(v_a)$ and a monetary cost or toll τ_a . The travel time is a function of the flow v_a on the link a , and is assumed to be strictly monotonically increasing. The demand for OD pair w is q^w . The following constraints apply over the system:

$$f_p^w \geq 0 \quad \text{for all } p \in P, w \in W, \quad q^w = \sum_p f_p^w \quad \text{for all } w \in W, \quad v_a = \sum_w \sum_p f_p^w \delta_{ap}^w$$

Here f_p^w is the path flow on path p between OD pair w and δ_{ap}^w is an indicator equal to 1 if link a is on path p between OD pair w , and zero otherwise. We will write $\Omega(q)$ as the feasible space of link flows defined in this way.

The link costs are given by strictly increasing functions $c_a(v_a) = t_a(v_a) + \tau_a$. User Equilibrium (UE) wherein the OD flows satisfy Wardrop’s first principle, is a solution to $\min_v \int_0^{v_a} c_a(z) dz$ for demand-feasible link flow vectors $v \in \Omega(q)$.

It is well-known that that UE is typically inefficient, in that it does not minimise the total network travel time. The flow pattern corresponding to the minimum network travel time satisfies $\min_v \sum_a t_a(v_a)v_a$ again constrained to demand-feasible flows.

The UE path flows equilibrate the total OD cost $\sum_a c_a \delta_{ap}^w$ on each path p connecting OD w , including the cost arising from any link tolls. Applying network tolls therefore enables us to achieve the system optimum (minimum travel time) flow pattern, while satisfying UE with respect to the total costs.

We can apply tolls to achieve many other demand-feasible flow patterns e.g. for a single OD network we could ensure non-zero flow on only one path by applying sufficiently high tolls on the links we wish to remain unused. In this case the toll levels are not unique, and the set of tolled links may be non-unique. This is particularly the case if we allow “negative tolls”.

For a given network we say that we have **full controllability** if we can push the UE solution to be ANY feasible link flow solution $v \in \Omega(q)$.

Problem 1: Minimum toll set for full controllability under fixed demand

For a given network with fixed demand, using positive or negative tolls, determine the minimum number of links to toll in order to achieve full controllability (i.e. the toll levels can be varied to achieve the different flow patterns, but the set of tolled links is fixed).

Problem 2: Minimum toll set to achieve system optimal under any demand

For a given network, determine the minimum number of links to toll (positively or negatively) so that UE gives the system optimum flow pattern at ANY level of OD demand(s).

Notes: In the simplest case there is only one OD pair. Certain network topologies may be trivial e.g. spanning tree. For which network topologies can bounds be established?

4.2 Embedding multiple trains in a given network

Katharina Eickhoff (RWTH Aachen, DE)

License © Creative Commons BY 4.0 International license
© Katharina Eickhoff

Let $G = (V, E)$ be a directed graph which corresponds to the infrastructure of a train network. We like to embed two (or multiple) trains with given origins and destinations and given departure and arrival times. Some vertices/edges may be blocked in a specific time window, for example since another fixed train uses the track at this moment.

The easier case is to embed a single train in the network which takes the times the infrastructure can be used into account. One approach to solve the problem is to construct a time expanded network, delete all vertices/edges which cannot be used and find a path in the time expanded network.

For two (or multiple) trains we have to find disjoint path so that the trains do not collide. Shiloach and Perl [1] provide an algorithm which solves this problem in directed acyclic graphs. But disjoint path in general are not enough: Due to safety reasons, the trains will block the nodes/edges longer than the time they are on them. That raises the question how to find “very disjoint” paths in a directed graph.

The discussion during the seminar leads to promising ideas, e.g. how to adapt the algorithm by Shiloach and Perl [1]. However, constructing the time expanded network blows up the instance. Thus, this idea is very time consuming and the question arise how to speed up.

References

- 1 Shiloach, Y., and Y. Perl. *Finding two disjoint paths between two pairs of vertices in a graph*. Journal of the ACM (JACM) 25(1), pp. 1-9, 1978.

4.3 Concepts for driver's route choice behaviour

Takamasa Iryo (Tohoku University – Sendai, JP)

License  Creative Commons BY 4.0 International license
© Takamasa Iryo

Dynamic traffic assignment problems considering atomic vehicles can be described as strategic games. On the other hand, finding a pure Nash equilibrium in DTA games is not always easy. Some of them can be NP-hard. Even the existence of a pure Nash equilibrium solution is not a general property. Considering a mixed Nash equilibrium solution is a way to resolve the issue of existence at least, but it may not be natural to assume that drivers stochastically choose routes day-to-day.

Therefore it would be an open question to investigate which concept is more suitable to describe drivers' route choice behaviour other than pure/mixed Nash equilibria. I have asked participants to say some candidates, and the following have been raised:

- Approximate pure nash equilibrium
- Evolutionary dynamics
- CURB (Closed Under Rational Behaviour)
- Preparation set (Voorneveld, 2004)
- Sink equilibria
- Level-K

4.4 Hitting Sets, Randomly

Jannik Matuschke (KU Leuven, BE)

License  Creative Commons BY 4.0 International license
© Jannik Matuschke

Let \mathcal{P} be a set system on a ground set E . Let $r : \mathcal{P} \rightarrow [0, 1]$ be a requirement function on the sets and let $p : E \rightarrow [0, 1]$ be a vector of marginals. We are interested in characterizing marginals for which there is a probability distribution on the subsets of E such that the resulting random set intersects each $P \in \mathcal{P}$ with probability at least $r(P)$. We call such a distribution a *feasible decomposition* of p .

It is easy to see that the following condition is necessary for the existence of a feasible decomposition:

$$\sum_{e \in P} p(e) \geq r(P) \quad \forall P \in \mathcal{P}. \quad (\star)$$

The above setting was first studied by Dahan et al. [1], motivated by an application in a network security game. They showed that (\star) is also sufficient for the existence of a feasible decomposition, if \mathcal{P} is the set of maximal chains of a partial order \preceq on E (or, equivalently, the set of s - t -paths in a directed acyclic graph) and the requirements r fulfill the following *conservation law*:

$$r(P) + r(Q) = r(P \times_e Q) + r(Q \times_e P) \quad \forall P, Q \in \mathcal{P}, e \in P \cap Q,$$

where $P \times_e Q = \{f \in P : f \preceq e\} \cup \{f \in Q : e \preceq f\}$. In addition, they provide an efficient algorithm for computing a feasible decomposition in the *affine setting*, where the requirements are of the form $r(P) = 1 - \sum_{e \in P} \mu(e)$ for some $\mu(e) \in [0, 1]^E$.

For this affine setting, both the existence result and the algorithm by Dahan et al. [1] can be extended to the case where \mathcal{P} is an abstract network [3]. An abstract network, introduced by [2], is a system (E, \mathcal{P}) , where the elements of $\mathcal{P} \subseteq 2^E$ are called *paths*. Each path $P \in \mathcal{P}$ is equipped with a complete order \preceq_P of the elements in P , and the following switching axiom is fulfilled: For each $P, Q \in \mathcal{P}$ and each $e \in P \cap Q$, there is a path $P \times_e Q \in \mathcal{P}$ with $P \times_e Q \subseteq \{f \in P : f \preceq_P e\} \cup \{f \in Q : e \preceq_Q f\}$. In particular, the set of s - t -paths in a digraph is an abstract network. As a consequence of the aforementioned generalization, equilibria for the security game studied by Dahan et al. [1] can be efficiently computed even when the underlying digraph contains cycles.

It is also possible to show that when \mathcal{P} is the set of s - t -paths in a directed acyclic graph, the conservation law is actually equivalent to the affine case, and a corresponding weight vector μ can be found efficiently [3]. As a result, a feasible distribution consistent with the marginals can also be computed efficiently in this case.

The conservation law can be extended to general digraphs (with cycles) and even abstract networks. In fact, it resembles Hoffman's [2] supermodularity condition for abstract networks:

$$r(P) + r(Q) \leq r(P \times_e Q) + r(Q \times_e P) \quad \forall P, Q \in \mathcal{P}, e \in P \cap Q. \quad (1)$$

This leads to the following open question.

Open Problem 1: Assuming that \mathcal{P} is the set of s - t -path in a digraph and that r fulfills (1), is (\star) a sufficient condition for the existence of a feasible decomposition?

One can also investigate the existence of feasible decompositions other set systems. However, it seems that, when going beyond paths (in abstract networks), (\star) is no longer a sufficient condition – simple counter-examples exist, e.g., when \mathcal{P} is the set of bases of a matroid or the set of perfect matchings in a graph [3]. These counter-examples all have a similar underlying structure. So it appears that stronger assumptions on r and \mathcal{P} are necessary here, leading to the second open question.

Open Problem 2: Is there a way to characterize the sufficiency of (\star) by means of forbidden substructures in r and \mathcal{P} ?

References

- 1 Mathieu Dahan, Saurabh Amin, and Patrick Jaillet. *Probability distributions on partially ordered sets and network interdiction games*. Mathematics of Operations Research, 47:458–484, 2022.
- 2 A. J. Hoffman. *A generalization of max flow-min cut*. Mathematical Programming, 6(1):352–359, 1974.
- 3 Jannik Matuschke. *Hitting paths with random sets in abstract networks*. Unplushed note. Presented at Dagstuhl Seminar “Dynamic Traffic Models in Transportation Science”.

4.5 The Expected Price of Anarchy in Wardrop Routing

Daniel Schmand (Universität Bremen, DE)

License © Creative Commons BY 4.0 International license
© Daniel Schmand

The price of anarchy (PoA) is a standard concept for analyzing the inefficiency of equilibria in strategic games. Researchers have mainly focused on worst-case bounds for a specific class of games, e.g. Wardrop routing games. It is well known that the PoA is at most $4/3$

in Wardrop routing games with affine-linear cost functions and that this bound is tight [1]. If cost functions are polynomials of maximal degree k , this bound increases in the order of $\Theta(k/\log k)$ [2].

However, the known worst case examples do not only need a bad network structure but also a demand that is tailored to the given network. In real-life traffic networks the demand is typically not fixed, but varying over time and due to uncertainties. These facts motivate a study of the Price of Anarchy dependent on the demand. Numerical studies suggest that in many networks the price of anarchy is only large for a small range of demands [3]. Recently, it was shown that the price of anarchy tends to 1 if the demand goes to 0 or infinity [5]. Additionally, Cominetti, Dose, and Scarsini analyzed the behavior and properties of the PoA-function in dependence of the demand [4].

The question of establishing a bound on the expected price of anarchy for a demand drawn from some given distribution is still open. Specifically, for a fixed network let W_d be the cost of a Wardrop equilibrium, and O_d the cost of an optimal flow, both with demand d . We are interested in establishing a bound on $\mathbb{E} \left[\frac{W_d}{O_d} \right]$, where the expectation is taken over the demand drawn from a given probability distribution. During and after the 2015 Dagstuhl seminar, we have analyzed the Pigou network with costs 1 and x^k . We have shown that the expected price of anarchy for a demand drawn from a uniform distribution $U[0,D]$ for $D \in \mathbb{R}$ is in the order of $\Theta(\log k)$. It is still open to extend this result to other network structures.

References

- 1 Roughgarden, Tardos: *How bad is selfish routing*
- 2 Roughgarden: *The price of anarchy is independent of the network topology*
- 3 O'Hare, Connors, Watling: *Mechanisms that Govern how the Price of Anarchy varies with Travel Demand*
- 4 Cominetti, Dose, Scarsini: *The price of anarchy in routing games as a function of the demand*
- 5 Colini-Baldeschi, Cominetti, Mertikopoulos, Scarsini: *When is selfish routing bad? The price of anarchy in light and heavy traffic*

4.6 Complexity of network loading in the fluid-queue model

Martin Skutella (TU Berlin, DE)

License  Creative Commons BY 4.0 International license
© Martin Skutella

Our open problem is to prove results on the complexity of the network loading problem in the fluid-queue model. The most basic yet interesting setting is as follows: We consider a directed graph $D = (V, A)$ with capacities $u \in \mathbb{Z}_{>0}^A$ and transit times $\tau \in \mathbb{Z}_{\geq 0}^A$ on the arcs. Moreover, we are given directed paths P_1, \dots, P_k along with fixed flow rates r_1, \dots, r_k . Assuming that we start to send flow at rates r_i into every path P_i at time 0, the network loading problem is to compute the resulting flow rate functions on the arcs of the network. Notice that the time at which a flow particle traveling along path P_i arrives on some arc $a \in P_i$ not only depends on its departure time at the start node of P_i but also on the latencies experienced on the predecessor edges on path P_i . This fact induces involved interdependencies among the flow rate functions on the arcs of the network. Notice that we can assume without loss of generality that all paths share the same start- and end-node.

We are interested in the computational complexity of decision problems such as, for example: For a given arc $a \in A$, a given point in time θ , and a given threshold value $s \in \mathbb{Q}_{>0}$,

- is the flow rate on arc a at time θ at least s ?
- is there a point in time $\theta' \geq \theta$ such that the flow rate on arc a at time θ' is at most s ?

The described problem setting can also be modified, for example by considering more general inflow patterns on the given paths (e.g., piecewise constant with a limited number of break points), or by allowing for time-dependent arc capacities etc.

The open problem is to come up with NP-hardness, PSPACE-hardness, or maybe even undecidability results in this context.

4.7 Towards defining a set of open problems for transportation analysis

David Watling (University of Leeds, GB)

License  Creative Commons BY 4.0 International license
© David Watling

Many branches of mathematics have benefited from having well-defined, unsolved problems, yet to my knowledge there has never been a clear set of such problems defined in transportation analysis. Such problems could serve, for example, as a focus for theoretical work to prove some property, or to extend the sufficient conditions under which a property is known to hold. Alternatively they could be a focus for numerical simulation work, developing a body of evidence that might lead to a more well-defined conjecture, or finding counter-examples that we then know must be ruled out by any sufficient conditions of a theoretical proof. If an unproven conjecture amasses sufficient numerical evidence in support of it, then it might be used as a starting point from which to prove other results.

The purpose of the presentation is to develop this suggestion further, with the intention that it could lead to the start of a technical discussion group that would aim to develop an initial specification of such a set of open problems. The motivation for selecting problems could be diverse. For example, it could be in order to bring together a wide body of knowledge in order to make a clear statement of what is currently known and not known; for example, in dynamic traffic assignment, for which kinds of component sub-models, network structure, etc. are we able to guarantee existence, uniqueness or stability of equilibrium, and in which cases do we know of counter-examples? Alternatively, the selection of such problems could be also be based on what the community believe would be useful results, even if extremely difficult to prove theoretically, as a basis for developing further work. A third possibility is to consider the way in which real transportation systems are changing, and are expected to change in the future, and the new demands they place on the modelling of strategic agents; for example, automated vehicles and more shared transport options give rise to very different kinds of problem that will need pressing attention. The talk will not actually aim to precisely define any such problems, but is intended as a stimulus for more discussion at Dagstuhl and beyond.

4.8 Identifying paths within a given bound

David Watling (University of Leeds, GB)

License  Creative Commons BY 4.0 International license
© David Watling

A kind of “behavioural” route choice model has been defined in the transportation literature where drivers choose routes based on random utility theory, but where the random error terms are bounded in a special way, such that the difference in random utility between any given route and a minimum cost route for that movement is a random variable on a closed interval. The consequence of this is that travellers use all routes with a cost no greater than $x\%$ more than minimum cost. In order to implement such models in an equilibrium framework it would be desirable to be able to solve sub-problems which, for fixed (i.e. flow-independent) arc costs, identify the set of all routes that are within such a bound. The discussions discussed the behavioural aspects of this model and the complexity of the problem. As part of the discussions ensuing, an outline of a potential solution approach was also proposed (by Martin Strehler):

1. Do shortest path from origin to all vertices.
2. Do shortest path from destination to all vertices with edges reversed (so gives shortest paths from all vertices to destination).
3. Check for vertices which can be deleted by seeing which routes via the vertex violates bound.
4. Identify edges by the same kind of way.
5. What remains is a minimal subnetwork containing all paths (but not all paths in the subnetwork will satisfy bound).

Participants

- Joseph Chow
New York University, US
- Richard Connors
University of Luxembourg, LU
- Katharina Eickhoff
RWTH Aachen, DE
- Miriam Fischer
Imperial College London, GB
- Gunnar Flötteröd
Linköping University, SE &
Swedish National Road and
Transport Research Institute
(VTI) – Stockholm SE
- Martin Gairing
University of Liverpool, GB
- Yiannis Giannakopoulos
Universität Erlangen-
Nürnberg, DE
- Lukas Graf
Universität Augsburg, DE
- Tobias Harks
Universität Augsburg, DE
- Martin Hofer
Goethe-Universität –
Frankfurt am Main, DE
- Takamasa Iryo
Tohoku University – Sendai, JP
- Max Klimm
TU Berlin, DE
- Ekkehard Köhler
BTU Cottbus-Senftenberg, DE
- Jannik Matuschke
KU Leuven, BE
- Haruko Nakao
University of Luxembourg, LU
- Neil Olver
London School of Economics, GB
- Carolina Osorio
HEC Montréal, CA & Google –
Mountain View, US
- Dario Paccagnan
Imperial College London, GB
- Britta Peis
RWTH Aachen, DE
- Rahul Savani
University of Liverpool, GB
- Guido Schäfer
CWI – Amsterdam, NL
- Daniel Schmand
Universität Bremen, DE
- Alexander Skopalik
University of Twente –
Enschede, NL
- Martin Skutella
TU Berlin, DE
- Martin Strehler
Westfälische Hochschule
Zwickau, DE
- Bernhard von Stengel
London School of Economics, GB
- David Watling
University of Leeds, GB

