



Synergistic Traffic Assignment

Thomas Bläsius
Karlsruhe Institute of Technology
Karlsruhe, Germany
thomas.blaesius@kit.edu

Adrian Feilhauer
Karlsruhe Institute of Technology
Karlsruhe, Germany
adrian.feilhauer@kit.edu

Markus Jung
Karlsruhe Institute of Technology
Karlsruhe, Germany
markus.jung@student.kit.edu

Moritz Laupichler
Karlsruhe Institute of Technology
Karlsruhe, Germany
moritz.laupichler@kit.edu

Peter Sanders
Karlsruhe Institute of Technology
Karlsruhe, Germany
sanders@kit.edu

Michael Zündorf
Karlsruhe Institute of Technology
Karlsruhe, Germany
michael.zuendorf@alumni.kit.edu

ABSTRACT

Traffic assignment analyzes traffic flows in road networks that emerge due to traveler interaction. Traditionally, travelers are assumed to use private cars, so road costs grow with the number of users due to congestion. However, in sustainable transit systems, travelers share vehicles s.t. more users on a road lead to higher sharing potential and reduced cost per user. Thus, we invert the usual *avoidant traffic assignment* (ATA) and instead consider *synergistic traffic assignment* (STA) where road costs *decrease* with use.

We find that STA is significantly different from ATA from a game-theoretical point of view. We show that a simple iterative best-response method with simultaneous updates converges to an equilibrium state. This enables efficient computation of equilibria using optimized speedup techniques for shortest-path queries. In contrast, ATA requires slower sequential updates or more complicated iteration schemes that only approximate an equilibrium. Experiments with a realistic scenario for the city of Stuttgart indicate that STA indeed quickly converges to an equilibrium. We envision STA as a part of *software-defined transportation* systems that dynamically adapt to current travel demand. As a first demonstration, we show that an STA equilibrium can be used to incorporate traveler synergism in a simple bus line planning algorithm to potentially greatly reduce the required vehicle resources.

CCS CONCEPTS

• **Theory of computation** → **Algorithmic game theory**; **Exact and approximate computation of equilibria**; **Network games**.

KEYWORDS

traffic assignment; practical applications of game theory; emergent behavior; equilibrium concepts; software-defined mobility

ACM Reference Format:

Thomas Bläsius, Adrian Feilhauer, Markus Jung, Moritz Laupichler, Peter Sanders, and Michael Zündorf. 2025. Synergistic Traffic Assignment. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*, Detroit, Michigan, USA, May 19 – 23, 2025, IFAAMAS, 9 pages.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

1 INTRODUCTION

Passenger transportation is in a crisis. Established systems often predominantly rely on the use of private cars, which leads to congestion, pollution, and bad usage of space in urban areas. Traditional, statically planned public transportation is expensive, slow and inflexible. Thus, there is a lot of interest in more flexible systems that adapt to actual demand using a carefully orchestrated combination of individual transportation, shared taxis, and larger shared vehicles. We are interested in evaluating the potential of such *software-defined transportation* systems.

To this end, previous work evaluated the usage of large fleets of shared taxis in urban areas [1, 8, 22]. This alone has limited sharing potential which does not suffice to service the traveler demand in entire metropolitan areas. Using walking/bicycles/scooters to reach pickup and dropoff points improves upon shared taxis but the overall savings are moderate (about 20% less energy and slightly improved travel times) [15, 25]. Thus, an essential component of efficient software-defined transportation is a way to identify roads with high sharing potential solely based on the (previously measured or simulated) demands of travelers (*agents*). We model this using a *traffic assignment* congestion game where agents' strategies are paths in a road network (a graph), resulting in loads on road segments (edges – the resources of the game). We are interested in *equilibria*, i.e., strategies (path choices) for each agent such that no agent can reduce their cost by unilaterally changing their path.

Traffic assignment has been intensively studied but so far exclusively for the *avoidant* case (ATA) where costs increase with load due to congestion. This is important when considering individual transportation. For our application, load acts as a measure of sharing potential, which means costs *decrease* with load. We call this *synergistic traffic assignment* (STA). The focus of this paper is to analyze best response for STA, where agents react to the current situation with a strategy that is optimal for themselves. We are interested in whether this converges to an equilibrium, with the ultimate goal to efficiently compute such a stable state.

Motivated by the fact that shortest-path computations can be substantially accelerated if the edge costs do not change too often, we consider different best-response variants for STA. In *sequential* best response, each step consists of only one agent changing their strategy. In *simultaneous* best response, all agents choose their new strategy for the next iteration at the same time. As an additional variation orthogonal to this, we distinguish between *impact-aware* or *impact-blind* best response. In the impact-aware variant (typically

considered in game theory), each agent takes their own impact on the cost function into account when considering to change their strategy. In the impact-blind variant, agents make their decision on the current cost of edges, not on the cost after their own change. We note that this makes little difference in practical scenarios with many agents where the impact of each individual agent on the cost is negligible. This yields four variants for STA for all combinations of sequential vs. simultaneous and impact-aware vs. impact blind; see Section 3 for formal definitions and Table 1 for an overview.

For the impact-aware case, it follows from the more general setting of congestion games (CG) that sequential best response converges to an equilibrium [32]. For simultaneous best response, however, we observe that there can be best-response cycles. We additionally give an example of a best-response cycle in the case where a group of only two equivalent agents change their strategy simultaneously; see Section 4.1. For the impact-blind case, we show that best response converges to an equilibrium, even in the simultaneous setting; see Section 4.2. This difference between impact-blind and impact-aware is somewhat surprising, as they seem similar. Moreover, it is in stark contrast to ATA, where impact-blindness may cause best-response cycles even in the sequential setting.

Beyond the theoretical result, the convergence of the simultaneous and impact-blind variant allows fast computation of an equilibrium. In each round, all agents change simultaneously and base their choice on the same edge costs. Thus, we can use sophisticated preprocessing techniques for path planning that are several orders of magnitude faster than individual applications of Dijkstra’s algorithm; see Section 5. In contrast, ATA has to use expensive agent-by-agent updates or more complex iteration schemes that can use simultaneous updates but only approximate an equilibrium.

In Section 6, we perform an experimental evaluation using realistic demand data for Stuttgart, Germany. Fewer than 20 iterations suffice to find an equilibrium. In contrast, established ATA algorithms are considered far from an equilibrium at this number of iterations [7, 31]. Moreover, the employed shortest-path speedup technique of *customizable contraction hierarchies* benefits from the fact that STA reinforces a hierarchy of roads, leading to faster queries in later iterations of STA. Using the right parametrization of the cost function, we find routes that have only moderately (less than 25%) longer travel times than free-flow shortest paths while attaining considerably (more than two times) higher sharing potential. We also look at a simple model for deriving bus lines from traffic flows. We find that basing the bus lines on an STA equilibrium rather than free-flow traffic reduces the total vehicle operation time significantly. Though the model is too simplistic to draw solid conclusions for real-world settings, it shows that STA can be a useful tool for such applications.

Summary of Contributions.

- Introduction of synergistic traffic assignment (STA) as a way to extract sharing potential for a set of traveler demands
- Theoretical analysis of the best-response process for different variants of STA
- Practically efficient algorithm for finding an STA equilibrium based on the theoretical insights
- Prototypical demonstration that STA equilibria identify high sharing potential in realistic inputs with uses in line planning

2 RELATED WORK

To the best of our knowledge, there is no prior work on synergistic traffic assignment, i.e., traffic assignment with decreasing edge cost functions. Thus, here we consider research on avoidant traffic assignment (ATA) that may also be relevant for synergistic traffic assignment. ATA is a well-studied problem with decades of research in the operations and traffic planning communities. Additionally, ATA is often considered in game theory, as a prototypical example of a non-cooperative congestion game in pure strategies [30, 32].

User Equilibrium. Wardrop [36] states that travelers in a road network naturally act selfishly, changing their path if a different path with a better travel time for themselves exists, even to the detriment of others. An assignment in which no traveler can change their path to their benefit is called a *user equilibrium* (UE). Beckmann et al. [4] find that a UE always exists for ATA, which matches a result from game theory about the existence of pure strategy Nash equilibria in congestion games [32]. As traffic without central control naturally tends to a UE, it is an important issue of traffic analysis to compute UEs for a given network and traveler demand.

Simple Approaches to Finding a UE. A simple approach to finding a UE is a process called iterated *sequential best response*, in which one traveler at a time may change their path while other travelers’ paths are kept fixed. Edge costs are updated according to the change of one traveler before continuing with the next traveler. Sequential best response always converges to a UE [28, 32]. However, the number of iterations until convergence can be exponential in the number of travelers [13, 14].

In iterated *simultaneous best response*, each traveler chooses their next path simultaneously, while edge costs are kept fixed. Then, all paths are changed at once, before updating edge costs and repeating the process. While the simultaneous method needs fewer edge cost updates than sequential best response, it does not always converge. Instead, it may encounter best-response cycles where travelers revert to a previous state in a cyclical manner [34]. Thus, simultaneous best response is also not suited to find a UE for ATA.

Finding a UE in Practice. In practice, more sophisticated methods find a UE by applying convex optimization to Beckmann’s *transformation* [4] of ATA into a convex program whose minimum is a UE. A usual categorization [17, 31, 37] divides these algorithms into *link-based* [18, 34], *path-based* [9, 16, 23, 24], and *bush-based* [2, 11, 29] approaches. These algorithms perform simultaneous path updates but they utilize the convexity of the problem to ensure progress towards the UE and avoid best-response cycles. Note two limitations common to these approaches: First, travelers are treated as non-atomic, i.e., the load of a single O-D pair can be split among multiple paths. Second, the algorithms iteratively approach the UE in increasingly smaller steps but do not actually reach it.

Each algorithm mentioned here requires an edge cost update and a re-computation of shortest paths in every iteration. The overhead for these updates can be expected to dominate the total running time. Buchhold et al. [7] show that a recent shortest-path speedup techniques called *customizable contraction hierarchies* (CCH) can adequately reduce running times for cost updates and shortest-path queries to tenths of seconds per iteration for millions of O-D pairs.

3 MODELS AND NOTATION

We define traffic assignment as a congestion game. Let $[k] = \{1, \dots, k\}$ be a set of k agents. In a *congestion game* there is a finite set E of available *resources* that are used by the agents. Each agent $i \in [k]$ has a set of *possible strategies* P_i where each strategy $p \in P_i$ is a subset of resources, i.e., $p \subseteq E$. A *strategy profile* is a vector $S = (p_1, \dots, p_k)$ of strategies such that $p_i \in P_i$ for $i \in [k]$. The *load* $\ell_e(S)$ of a resource $e \in E$ with respect to a strategy profile S is the number of agents using e , i.e., $\ell_e(S) = |\{p_i \in S \mid e \in p_i\}|$. Each resource $e \in E$ has a *cost function* $c_e: \mathbb{N} \rightarrow \mathbb{R}$, where $c_e(\ell)$ is the cost of resource e given that it has load ℓ . The resulting cost for an agent is the sum of the costs of all resources used by the agent. More formally, we define the *cost* of a strategy $p \subseteq E$ given a strategy vector S as $\text{cost}(p, S) = \sum_{e \in p} c_e(\ell_e(S))$. Thus, for $S = (p_1, \dots, p_k)$, the cost for agent $i \in [k]$ is $\text{cost}(p_i, S)$.

With this, one can use the common definitions for the Nash equilibrium and best response. Before defining them formally, we introduce the traffic assignment game, which also motivates less common variants of these concepts. Colloquially speaking, in traffic assignment, each agent chooses a path in a graph that gets them from their origin to their destination, using the edges as resources.

More formally, let $G = (V, E)$ be a directed graph with a cost function $c_e: \mathbb{N} \rightarrow \mathbb{R}$ for each edge. An *origin-destination pair* (O-D pair) is a pair $(s, t) \in V \times V$ with $s \neq t$. For $s, t \in V$, an *st-path* is a set of edges $p \subseteq E$ such that the subgraph induced by p contains s and t , all vertices except s have exactly one incoming edge, and all vertices except t have exactly one outgoing edge. For a set of O-D pairs $X = \{(s_1, t_1), \dots, (s_k, t_k)\}$, the *traffic assignment game*, is a congestion game where the edge set E forms the resources and the possible strategies for agent $i \in [k]$ are the $s_i t_i$ -paths, i.e., $P_i = \{p \subseteq E \mid p \text{ is an } s_i t_i\text{-path}\}$. A traffic assignment game is *avoidant* if the cost functions c_e are non-decreasing, i.e., if more traffic only increases cost. Conversely, it is *synergistic* if the cost functions c_e are non-increasing. We also use the terms *avoidant* and *synergistic* for the more general case of congestion games.

3.1 Best Response

Consider a congestion game and let $S = (p_1, \dots, p_k)$ be a strategy profile. For an agent $i \in [k]$, we use S_{-i} to denote the omission of i 's strategy. Moreover, for any alternative strategy $p'_i \in P_i$ of agent i , we use (S_{-i}, p'_i) to denote the replacement of p_i with p'_i in S , i.e., $(S_{-i}, p'_i) = (p_1, \dots, p_{i-1}, p'_i, p_{i+1}, \dots, p_k)$. Note that $S = (S_{-i}, p_i)$.

A *best response* of agent i to a strategy profile $S = (p_1, \dots, p_k)$ is a strategy $p'_i \in P_i$ that minimizes $\text{cost}(p'_i, (S_{-i}, p'_i))$, which is the cost for agent i after they change their strategy from p_i to p'_i . We say that agent i is *content* with S if p_i is a best response to S . If all agents are content with S , then S is a *Nash equilibrium*. With *sequential best response*, we refer to the process of changing the strategy of one agent at a time to their best response with respect to the current strategy profile. If this reaches an equilibrium, we say that the process *converges*. Otherwise, we obtain a so-called *best-response cycle*, in which the process indefinitely cycles through a repeating sequence of strategy profiles (these are the only two options as the state space is finite). The following theorem by Rosenthal [32] states that the sequential best response defined above converges.

THEOREM 3.1 (ROSENTHAL [32]). *In a congestion game, sequential best response converges to a Nash equilibrium.*

As traffic assignment games are a special case of congestion games, Theorem 3.1 also holds for our setting. This is algorithmically useful as it provides a way to compute an equilibrium. In each step, one has to compute a shortest path in G for the O-D pair of the currently considered agent. Unfortunately, this has the effect that the costs of the edges change after every shortest-path computation, obstructing the use of algorithmic shortest-path techniques that are based on pre-computations. We overcome this by incorporating two types of adjustments to the above best-response variant.

3.2 Simultaneous Best Response

Simultaneous best response is a process where, in every step, all agents change their strategy to their best response with respect to the current strategy profile. While this feels like a step in the right direction in the sense that the strategy profile changes less frequently in relation to the number of shortest-path computations, there are two major downsides. First, we show in Section 4.1 that this leads to a best-response cycle. This is not very surprising as, in the synergistic setting, two agents may attempt to join each other and thereby avoid each other due to simultaneously changing their strategy. We extend this by also providing a best-response cycle in a setting that is almost sequential in the sense that only two equivalent agents choose their new path together; see Section 4.1.

The second downside is that the simultaneous setting alone is not enough to let us use pre-computation techniques for the shortest-path queries. The reason for this is that the cost of an edge e for agent i depends on whether i is currently using it. To be more precise, if $S = (p_1, \dots, p_k)$ is the current strategy profile and $e \in p_i$, then the cost of e is $c_e(\ell(S))$. If, however, $e \notin p_i$, then the cost of e is $c_e(\ell(S) + 1)$ as i switching to a strategy that uses e would mean that the load of e increases by 1. In other words, the best response aims to minimize $\text{cost}(p'_i, (S_{-i}, p'_i))$, i.e., the cost of the edges is not based on S but on the updated profile (S_{-i}, p'_i) . To resolve this, we introduce a best-response variant in which the agents ignore the impact of their own change. In ATA this setting of simultaneous updates based on the costs after the last round is sometimes done under the premise of *capacity restraint* [21, 34]. For a resource selection game, Harder et al. [20] considered a similar concept and used the terms *impact-aware* and *impact-blind*. Following this, we also call the best-response variant defined above and the resulting Nash equilibrium *impact-aware*.

3.3 Impact-Blind Best Response

An *impact-blind best response* of agent i to a strategy profile $S = (p_1, \dots, p_k)$ is a strategy $p'_i \in P_i$ that minimizes $\text{cost}(p'_i, S)$. All definitions for the impact-aware setting can be directly translated (impact-blind Nash equilibrium, sequential/simultaneous, convergence, best-response cycle). We note that in practical settings, the impact of a single agent changing is usually not very high. Fortunately, the change from impact-aware to impact-blind has the effect that even simultaneous best response converges. We note that this is specific to the synergistic setting, which makes intuitive sense for the following reason; see Section 4.2 for details. In the synergistic setting, the own impact of changing the strategy only reinforces the

Table 1: Overview of our results on the best-response variants of congestion games (CG) and traffic assignment (TA).

| | sequential | simultaneous |
|---|---|-------------------------------|
| impact-aware $\text{cost}(p'_i, (S_{-i}, p'_i))$ | converges (CG) Theorem 3.1 [32] | cycles (STA) Observation 1 |
| impact-blind $\text{cost}(p'_i, S)$ | converges (synergistic CG) Theorem 4.2 | |

decision, which leads to slightly higher stability of the impact-blind variant (the converse is true for the avoidant setting). Nonetheless, we find it quite surprising that this seemingly minor effect is sufficient to obtain convergence rather than a best-response cycle.

4 GAME ANALYSIS

Here we analyze the different best-response variants for STA; see Table 1 for an overview. We start with the first row (impact-aware) in Table 1. For the sequential case (first column), convergence already follows from Rosenthal’s results [32] on congestion games; recall Theorem 3.1. For the simultaneous case (second column), we give a best-response cycle in Section 4.1. We additionally give a best-response cycle for a variant that lies between the simultaneous and the sequential setting. For the second row (impact-blind), we show in Section 4.2 that for synergistic congestion games (including STA) best response converges to an equilibrium, even in the simultaneous setting. This notably distinguishes STA from ATA, which has best-response cycles in all cases except for the sequential impact-aware setting covered by Theorem 3.1.

4.1 Best-Response Cycles

Simultaneous best response does not converge for synergistic traffic assignment, as agents simultaneously attempting to swap towards each other makes it so that they essentially swap past each other and never meet. The example in Figure 1 yields the following observation; see full paper [6] for a proof.

OBSERVATION 1. *Simultaneous and impact-aware best response for synergistic traffic assignment has a best-response cycle.*

In the following, we additionally consider a best-response variant in which only agents with the same O-D pair change their strategy simultaneously. We additionally assume that the agents know of each other, i.e., in their response valuation, they account for the fact that all agents with the same O-D pair will make the same choice.

Formally, we consider two agents to be equivalent if they have the same O-D pair and we call the resulting equivalence classes *groups*. Then *group-simultaneous best response* refers to a process in which all agents from one group simultaneously change their strategy, while the different groups are considered sequentially. Consider a step in which one group changes and assume without loss of generality that this group consists of the first a agents, i.e., the group is the set $[a] \subseteq [k]$. Given a strategy profile $S = (p_1, \dots, p_k)$ and a strategy $p' \in P_1$ that is valid for the group $[a]$, let $(S_{-[a]}, p') = (p', \dots, p', p_{a+1}, \dots, p_k)$ be the strategy profile




Figure 1: Graph (left) with a best-response cycle (right) for simultaneous impact-aware best response. The two bold edges have cost 1 for load below 2 and cost 0 for load at least 2. Two edges have cost ϵ . All other edges have cost 0.

where all agents in $[a]$ change their strategy to p' . The *group-impact-aware best response* for agents in $[a]$ is the strategy $p' \in P_1$ for which $\text{cost}(p', (S_{-[a]}, p'))$ is minimized. We get the following.

THEOREM 4.1. *Group-simultaneous and group-impact-aware best response for synergistic traffic assignment has a best-response cycle, even if the largest group has size 2.*

PROOF SKETCH. Consider the graph in Figure 2 (left) with the O-D pairs (s_i, t_i) for $i \in [4]$. Note that the agents 1 and 2 have the same O-D pair and thus form a group. Most edge costs are constant (indicated by the numbers, 0 if there is no number), except for the three bold edges. Here, a cost of $a \xrightarrow{\ell} b$ indicates that the cost is a for loads below ℓ and the cost is decreased to b for load at least ℓ . Following the figure, we refer to the agents as blue (1 and 2), red (3), and orange (4). Also, we refer to the path from s_1 to t_1 with the three bold edges as *horizontal path*.

The best-response process cycles through the four configurations shown in Figure 2 (right). The argument is roughly as follows. Red and orange can either go to the horizontal path early (configuration A and B) or late (configuration C and D). In the early case, blue has an incentive to use the horizontal path. However, blue using the horizontal path incentivizes red and orange to go the horizontal path late. See full paper [6] for a full proof. \square

Concerning the best-response variants in Table 1, the group-simultaneous and group-impact-aware variant lies between sequential and simultaneous. Moreover, it lies above impact-aware in the sense that the agents are not only aware of their own impact but also of that of equivalent agents. We note that we find Theorem 4.1 somewhat surprising as the grouped setting is quite close to the sequential setting, in which best response converges. We only need groups of size 2, which makes the setting almost sequential. Moreover, using impact-aware but not group-impact-aware best response would also make it equivalent to the sequential setting for the following reason: In the impact-aware variant, the group as a whole would swap if and only if the first agent of the group would swap individually. Moreover, for later agents of the group, the previously swapped agents only reinforce the choice.

4.2 Convergence of Impact-Blind Best Response

Here, we show that impact-blind best response converges to an equilibrium for synergistic congestion games. This is particularly interesting for the simultaneous variant, but our proof also holds for




Figure 2: Graph (left) with a best-response cycle (right) for Theorem 4.1. Edges are implicitly oriented from left to right.

the sequential setting. Note that the requirement of the congestion game being synergistic is crucial, i.e., the cost of a resource must not increase with a higher load. For non-synergistic congestion games like ATA the impact blindness is in fact detrimental. Thus, their impact-blind formulations can run into best-response cycles.

Our proof uses a potential function $\phi: P_1 \times \dots \times P_k \rightarrow \mathbb{R}$ that maps a strategy profile S to a number. Assume agents change their strategy using impact-blind best response, yielding a new strategy profile S' . We show that in this case $\phi(S') < \phi(S)$ holds. As the potential cannot decrease indefinitely (the strategy space is finite), the process must reach an equilibrium.

We use the potential function introduced by Rosenthal [32] for the sequential and impact-aware setting. It sums for each resource the cost of all loads up to the actual load. To make this more formal, recall that $c_e(\ell)$ is the cost of resource e for load ℓ and that $\ell_e(S)$ refers to the load of e given the strategy profile S . The *edge potential* $\phi_e(S)$ of e is defined as $\phi_e(S) = \sum_{\ell=0}^{\ell_e(S)} c_e(\ell)$. The *potential* ϕ is the sum over all resources, i.e., $\phi(S) = \sum_{e \in E} \phi_e(S)$.

Our proof that ϕ decreases with each round of impact-blind best response works roughly as follows. Given strategy profile $S = (p_1, \dots, p_k)$, each agent i changes their strategy from p_i to p'_i such that $\text{cost}(p'_i, S)$ is minimized. With this change, agent i anticipates that their cost decreases by $\Delta_i = \text{cost}(p_i, S) - \text{cost}(p'_i, S)$. Note that the actual cost decrease can be different as the agent did not factor in their own impact, or even the impact of the other agents changing at the same time in the simultaneous setting. However, we can show that the potential decreases at least by the total anticipated cost of all agents. As $\Delta_i \geq 0$ (otherwise the agent would not have changed strategy), this results in a potential decrease.

THEOREM 4.2. *In a synergistic congestion game, impact-blind best response (simultaneous or sequential) converges to an equilibrium.*

PROOF. We consider the general setting where each round of best response consists of any non-empty subset of agents changing their strategy. This in particular includes the simultaneous and sequential variants. Let $S = (p_1, \dots, p_k)$ be the current strategy profile and let $S' = (p'_1, \dots, p'_k)$ be the resulting profile after the change. Let Δ_i be the anticipated cost decrease for agent i as defined above and let $\Delta = \sum_{i=1}^k \Delta_i$. As $\Delta_i > 0$ for some i , we also have $\Delta > 0$.

We now change perspective on the anticipated cost decrease from the agents to the resources. For $e \in p_i \setminus p'_i$, agent i changed from using e to not using e . Thus, the resource e contributes $c_e(\ell_e(S))$ positively to Δ_i . Conversely, if $e \in p'_i \setminus p_i$, then e contributes $-c_e(\ell_e(S))$ to Δ_i . To define the overall contribution of e to Δ , note that the load of e decreased by $\ell_e(S) - \ell_e(S')$. Each of these load decreases corresponds to an agent anticipating a cost decrease of $c_e(\ell_e(S))$, thus we define the anticipated cost decrease of an edge e as $\Delta_e = c_e(\ell_e(S)) \cdot (\ell_e(S) - \ell_e(S'))$. It directly follows that $\sum_{e \in E} \Delta_e = \Delta > 0$.

Recall that our goal is to show that $\phi(S) - \phi(S') > 0$. For this, we show that for every resource $e \in E$ it holds that $\phi_e(S) - \phi_e(S') \geq \Delta_e$, which implies $\phi(S) - \phi(S') \geq \Delta > 0$. To show $\phi_e(S) - \phi_e(S') \geq \Delta_e$, note that the potential difference is

$$\phi_e(S) - \phi_e(S') = \sum_{\ell=0}^{\ell_e(S)} c_e(\ell) - \sum_{\ell=0}^{\ell_e(S')} c_e(\ell).$$

First consider the case that $\ell_e(S) \geq \ell_e(S')$, i.e., decreasing load. Then we get the following, where the inequality holds as the sum has $\ell_e(S) - \ell_e(S')$ terms and c_e is non-increasing

$$\phi_e(S) - \phi_e(S') = \sum_{\ell=\ell_e(S')+1}^{\ell_e(S)} c_e(\ell) \geq (\ell_e(S) - \ell_e(S')) \cdot c_e(\ell_e(S)) = \Delta_e.$$

For the case $\ell_e(S) \leq \ell_e(S')$, we analogously obtain

$$\phi_e(S) - \phi_e(S') = - \sum_{\ell=\ell_e(S)+1}^{\ell_e(S')} c_e(\ell) \geq (\ell_e(S) - \ell_e(S')) \cdot c_e(\ell_e(S)) = \Delta_e,$$

which concludes the proof. \square

As STA is a synergistic congestion game, the above theorem holds for STA. In the remainder of the paper, we utilize the fact that the simultaneous variant of impact-blind best response converges.

5 STA ALGORITHM

In a traffic assignment game with converging best response, computing an equilibrium essentially boils down to iterated shortest-path computations. To make this efficient, one can potentially apply speed-up techniques that have been developed in the context of route planning [3]. These techniques exploit the fact that the graph

does not change by pre-computing auxiliary information once in the beginning, which is then used to enable highly efficient queries.

For traffic assignment, the graph topology does not change between shortest-path queries, but the edge costs change depending on the load. Fortunately, there are speed-up techniques that support such cost changes [10, 12], which work in three phases. The first phase is a so-called *metric-independent preprocessing*, where pre-computations are done based only on the topology of the graph, ignoring the edge costs. The second phase is called *customization* and performs pre-computations based on the edge costs. The third phase are the shortest-path *queries*, which are highly efficient due to the data structures computed in the first two phases.

Concerning the different variants for STA, impact-blind and simultaneous best response is particularly suited for this algorithmic approach as all agents do shortest-path computations with respect to the same cost function once before updating it due to the changed load. This is the reason why we focus on impact-blind and simultaneous best response for computing STA equilibria. We note that this is viable as this best-response variant converges for STA due to Theorem 4.2, which is not true for ATA.

Algorithm Description. With these preliminary considerations, we are now ready to state our algorithm. We have a *load vector* storing the load of each edge in the current traffic assignment. It is initialized with all 0s. Moreover, we do the metric independent preprocessing once in the beginning. Then, in each iteration, we first do the customization based on the edge costs defined by the current load vector. Afterwards, the resulting data structure is used to compute a shortest path for every O-D pair. Finally, the load vector is updated according to the resulting paths. This is iterated until no agent changes their path, which is equivalent to checking whether the load vector changed, due to the proof of Theorem 4.2.

Implementation Details. There are different variants of the above-mentioned three-phase approach with different trade-offs between the running times of the different phases. In our setting, the bottleneck is the query. We thus use a so-called *customizable contraction hierarchy* (CCH) [12], which provides faster queries at the cost of slower customization compared to CRP [10]. Explaining CCH is beyond the scope of this paper and we refer the reader to the recent survey [5]. For readers familiar with CCH, we want to briefly mention the details of our implementation. We compute the nested dissection order with InertialFlowCutter [19]. We use perfect customization, elimination tree queries, and all improvements stated by Bläsius et al. [5]. We do not use parallelization or SIMD operations.

6 EXPERIMENTAL EVALUATION

We provide experimental results on the convergence and resulting traffic flows of our STA algorithm using real-world data and an implementation in C++17.¹ We compile with gcc 13.3.1 using -O3. We use a machine with Fedora 39 (kernel 6.8.11), 32 GiB of DDR4-4266 RAM, and an AMD Ryzen 7 PRO 5850U CPU clocked at 4.40 Ghz.

6.1 Inputs and Methodology

We evaluate our algorithm on the road network of the city of Stuttgart, Germany, and the surrounding region. We use a road

¹The code is available at <https://github.com/molaupi/synergistic-traffic-assignment>.

Table 2: Four traffic Scenarios used for the evaluation of STA.

| Scenario | Analysis period | O-D pairs |
|----------|-------------------|-----------|
| S-morn | Tue., 7:30-8:30am | 32 034 |
| S-even | Tue., 4:30-5:30pm | 38 021 |
| S-day | a whole Tuesday | 453 926 |
| S-week | a whole week | 2 946 810 |

network based on the OSM network of Germany² and the OSM relation of the Stuttgart city limits.³ We only consider travel demand within these city limits but allow travelers to also use major streets in the surrounding area. For this, we include roads of OSM rank at least *tertiary link*⁴ in a rectangular box around the city that is 3 times the width and height of the bounding box of the city limits. The resulting graph contains 70 915 vertices and 138 862 edges. We retrieve traversal times $d(e)$ for each edge e using the length and speed limit of each road segment encoded in the OSM data.

Our demand data [33] models travel demand in Stuttgart for a whole week and was originally forecast using mobiTopp [26, 27], which was calibrated from a household travel survey conducted in 2009/2010 [35]. We use the same test scenarios that Buchhold et al. [7] use to evaluate ATA algorithms, i.e., a typical morning peak, a typical evening peak, a typical day and a whole week; see Table 2.

The cost function for STA may significantly depend on the particular application of STA, e.g., the type of vehicle used. Here we use a simple function with a single tuning parameter $r \in [0, 1]$, which we call the *selfishness parameter*. This allows us to interpolate between ignoring sharing ($r = 1$) and aggressively optimistic sharing ($r = 0$). Let $d(e)$ and ℓ_e be the travel time (assuming free flow) and the load of edge e . We define the cost $c_e^r(\ell)$ of e as

$$c_e^r(\ell_e) = r \cdot d(e) + (1 - r) \cdot \frac{d(e)}{\ell_e + 1}.$$

6.2 Sharing Evaluation

Due to our simple cost function, we evaluate sharing using quantities that are independent of the concrete definition of c_e^r . Let $\text{dist}(u, v)$ be the length of the shortest path between u and v with respect to the travel time d . Further, we define D as a natural extension of d on sets of edges, i.e., $D(p) = \sum_{e \in p} d(e)$. Recall that p_i denotes the path of agent $i \in [k]$ and ℓ_e is the load on edge e .

We define the *average stretch* as $\frac{1}{k} \sum_{i \in [k]} \frac{D(p_i)}{\text{dist}(s, t)}$.

The *average sharing* is $\frac{1}{k} \sum_{i \in [k]} \frac{\sum_{e \in p_i} d(e) \cdot (\ell_e - 1)}{D(p_i)}$,

i.e., the sharing of each agent $i \in [k]$ is the number of other agents they share their path with, averaged over the whole path. In other words, if you pick a random point during i 's journey, then i 's sharing value is the expected number of other agents on i 's current edge. As there is also sharing without STA (i.e., for $r = 1$), we define the *normalized average sharing* as the average sharing relative to the average sharing obtained for $r = 1$.

²Available at <https://download.geofabrik.de/europe/germany.html>.

³See <https://www.openstreetmap.org/relation/2793104>.

⁴See <https://wiki.openstreetmap.org/wiki/Key:highway>.




Figure 3: The normalized average sharing in relation to the average stretch in the morning scenario for different values of r ; the numbers at the points indicate r .




Figure 4: Fraction of agents that share a fraction of at least x of their travel time with at least ℓ other agents. We use $\ell \in \{1, 10, 100\}$ and show values for $r = 0$ (high sharing) and $r = 1$ (free flow) in the S-morn scenario.

Figure 3 shows the trade-off between average stretch and normalized average sharing for different choices of r . Stretch and sharing both increase with lower r , seemingly in a linear relation. However, the slope is substantially below 1, indicating that slightly longer paths can result in higher sharing. For instance, at $r \approx 0.0075$, the sharing can be doubled at the cost of only 1.25 times longer paths.

Figure 4 analyzes the distribution of sharing potential along riders' paths. The plot has one curve per pair of selfishness $r \in \{0, 1\}$ and load threshold $\ell \in \{1, 10, 100\}$. Each curve shows the fraction of agents (y-axis) for whom at least a fraction x of their path (x-axis) is shared with at least ℓ others in the r STA flow. The curves for STA with strong sharing ($r = 0$) stay high even for large values of ℓ , particularly in comparison to just using individual shortest paths ($r = 1$). Thus, most agents achieve high sharing on a large portion of their path, which is important for applications like ride sharing.

6.3 Performance Evaluation

Our theoretical results prove convergence, but give no guarantees for how many iterations are necessary. Thus, in this section, we evaluate the number of iterations and running times in practice.

Figure 5 shows the number of iterations depending on r . Smaller values of r require more iterations but even for $r = 0$, fewer than 20 iterations suffice. The number of agents appears to have no effect.




Figure 5: The number of iterations for the STA best response to converge. The reported numbers include the last iteration that did not result in a change.




Figure 6: Time spent for shortest-path computations and load accumulation per iteration in the morning scenario.

To give an impression of the overall running times, the entire process with $r = 0$ takes about one second for S-morn and S-even, ten seconds for S-day, and a minute for S-week, respectively.

In Figure 6, we can see that in later iterations, the shortest-path queries become noticeable faster. For $r = 0$, the time spent on queries almost halves after only 6 iterations. As the number of O-D pairs stays the same, this implies that the CCH (with perfect customization) can leverage the increased sharing and smaller number of relevant edges for faster queries.

As claimed previously, we observe that shortest-path computations are indeed the bottleneck compared to customization. For instance, for $r = 0$, the smallest scenario (S-morn), and the fastest iteration, the time spent on shortest-path computations is roughly 37 ms. Meanwhile, an average of only 10.5 ms is spent on customization per iteration. It thus makes sense to opt for a CCH variant with more efficient queries at the cost of slower customization. In the future, shortest-path queries can be optimized further, e.g., by employing thread- and instruction-level parallelism (cf. [7]).

6.4 Bus Line Planning

As a proof of concept for the merits of STA, we use STA in a simple bus line planning problem. We propose a greedy path-based line planning algorithm to find trunk bus lines for a feeder-trunk system. We show that bus lines based on STA paths reduce total vehicle travel time compared to lines based on free-flow (FF) paths. Note that this proof-of-concept model is not optimized for realism but for simplicity while grasping the effects of sharing between travelers.

Bus Line Problem. We consider a line planning problem where we are given a road network and travel demands. The goal is to find bus lines subject to a bus operation time budget such that the total vehicle travel time (for buses and feeder vehicles) is minimized. We assume a road network $G = (V, E)$ with travel times $d(e)$ for every $e \in E$, a set of O-D pairs $X = \{(s_x, t_x)\}$, a bus operation time budget B , a fixed service frequency f for every line, and an operation time window $[0, T]$ s.t. every journey fits within that window.

We consider a bus line L to be a simple path in G . Buses travel the length of L from start to finish. Let δ_L be the travel time from start to finish. The total vehicle time needed to service L during $[0, T]$ with frequency f is $T_L = T \cdot f \cdot \delta_L$.

Travelers may enter and exit a bus anywhere on the line. For any part of a journey that is not covered by bus lines, the traveler has to use a feeder (e.g., a taxi). A feeder is immediately available between any two places but it can never be shared with other travelers.

Our goal is to minimize *total vehicle operation time (TVOT)*, i.e., the sum of operation times of all buses and feeder vehicles. Without any buses, every traveler uses the feeder system to go directly from their origin to their destination without any sharing. This gives us a baseline TVOT that we then aim to reduce by introducing buses. More specifically, we want to find a set of bus lines \mathcal{L} such that the TVOT is minimized but the sum of bus operation times fits into the bus operation time budget B , i.e., $\sum_{L \in \mathcal{L}} T_L \leq B$.

Algorithm. We use a simple path-based bus line planning algorithm to compare lines based on STA paths to lines based on FF paths. As input, the algorithm takes a set of paths P , containing a path p_x from s_x to t_x for every O-D pair $x \in X$. Let $N_P(e)$ be the number of paths in P that contain e . We greedily construct lines:

A new line L is started at the edge $e \in E$ with the largest $N_P(e)$. We extend L forwards by iteratively appending edges at the end, always choosing the edge e' with the largest $N_P(e')$. We stop extending L forwards if there is no next edge with $N_P(e') > 0$. We then extend the line backward analogously.

Then, for each path $p \in P$ that overlaps with L , we remove p from P , slice the part covered by L out of p , and add the remaining subpaths of p back to P . We update $N_P(e)$ for $e \in E$ accordingly. Note that we do this for only a subset of paths such that a fixed seating capacity of buses is not exceeded anywhere on the line.

We keep generating new lines until all traveler demand is covered by bus lines, i.e., until P becomes empty. We use a knapsack solver to choose a subset of lines \mathcal{L} that fits into the operation time budget B and maximizes traveler coverage.

Evaluation. We evaluate bus lines for the S-morn instance (cf. Table 2). We assume buses with a capacity of 80 and a fixed service frequency of $f = 1/10\text{min}$. For the input paths P , we use STA paths with different values of the selfishness parameter r (see Section 6.1), with $r = 1$ being equivalent to using FF paths. For simplicity's sake, we restrict each traveler to their input path and allow them to only use bus lines with which they overlap during line construction.

In Figure 7, we show the TVOT for different bus budgets and different selfishness $r \in \{0, 0.01, 1\}$. Note that with a budget of $B = 0$ h and selfishness of $r = 1$, travelers use only feeder vehicles on shortest paths, giving us the baseline TVOT of about 6600 h.

Unsurprisingly, bus lines reduce the TVOT for any value of r . Thus, we consider the minimal TVOT for different values of r .




Figure 7: Total vehicle operation time (in hours) for varying bus budgets and $r \in \{0, 0.01, 1\}$.

As previously stated in Section 6.2, STA paths are longer with smaller selfishness. Since we restrict travelers to their input paths even when using feeder vehicles, these longer paths result in a larger TVOT with smaller selfishness r if the bus budget is small. This effect is particularly strong for $r = 0$, which is why we focus on the comparison between $r = 0.01$ and $r = 1$ (FF) in the following.

Given a large enough budget, the improved sharing with STA leads to better bus lines which make up for the longer paths and reduce the TVOT to smaller values than FF-based lines. At a bus budget of $B = 1800$ h, STA-based lines with $r = 0.01$ lead to a minimum TVOT of 2140 h while FF-based lines require at least 2641 h. Thus, the FF-based lines require more than twice as much additional feeder service (841 h) compared to the STA-based lines (340 h). This reduction of about 500 h is equivalent to 7.5% of the baseline TVOT for using only feeder vehicles.

7 CONCLUSIONS AND FUTURE WORK

We have demonstrated that synergistic traffic assignment is an interesting game-theoretic concept that allows fast computations of equilibria. These equilibria seem relevant for applications in the context of envisioned software-defined transportation systems. A lot remains to be done there. Cost functions could try to better model costs and behavior of agents (perhaps using simulation or machine learning). More sophisticated algorithms to interpret the equilibria would take into account actually available vehicles and constraints on using them. Perhaps more interestingly, agents could use equilibrium information (based on historical data) to look for routes and vehicles worth sharing. This brings us back to the interesting algorithmic question of how STA equilibria can be computed/approximated in a dynamic fashion.

ACKNOWLEDGMENTS

This work was supported by funding from the pilot program Core Informatics of the Helmholtz Association (HGF). This paper was created in the “Country 2 City - Bridge” project of the „German Center for Future Mobility“, which is funded by the German Federal Ministry for Digital and Transport.

REFERENCES

- [1] Javier Alonso-Mora, Samitha Samaranyake, Alex Wallar, Emilio Frazzoli, and Daniela Rus. 2017. On-demand High-Capacity Ride-sharing via Dynamic Trip-Vehicle Assignment. *Proceedings of the National Academy of Sciences* 114, 3 (2017), 462–467. <https://doi.org/10.1073/pnas.1611675114>

- [2] Hillel Bar-Gera. 2002. Origin-Based Algorithm for the Traffic Assignment Problem. *Transportation Science* 36, 4 (11 2002), 398–417. <https://doi.org/10.1287/trsc.36.4.398.549>
- [3] Hannah Bast, Daniel Delling, Andrew V. Goldberg, Matthias Müller-Hannemann, Thomas Pajor, Peter Sanders, Dorothea Wagner, and Renato F. Werneck. 2016. Route Planning in Transportation Networks. In *Algorithm Engineering - Selected Results and Surveys*, Lasse Kliemann and Peter Sanders (Eds.). Lecture Notes in Computer Science, Vol. 9220. Springer, 19–80. https://doi.org/10.1007/978-3-319-49487-6_2
- [4] Martin Beckmann, Charles B. McGuire, and Christopher B. Winsten. 1955. *Studies in the Economics of Transportation*. Technical Report. The Rand Corporation.
- [5] Thomas Bläsius, Valentin Buchhold, Dorothea Wagner, Tim Zeitz, and Michael Zündorf. 2025. Customizable Contraction Hierarchies – A Survey. [arXiv:2502.10519](https://arxiv.org/abs/2502.10519) [cs.DS]
- [6] Thomas Bläsius, Adrian Feilhauer, Markus Jung, Moritz Laupichler, Peter Sanders, and Michael Zündorf. 2025. Synergistic Traffic Assignment. [arXiv:2502.04343](https://arxiv.org/abs/2502.04343) [cs.GT]
- [7] Valentin Buchhold, Peter Sanders, and Dorothea Wagner. 2019. Real-time Traffic Assignment Using Engineered Customizable Contraction Hierarchies. *ACM Journal of Experimental Algorithmics* 24, 1 (2019), 2.4:1–2.4:28. <https://doi.org/10.1145/3362693>
- [8] Valentin Buchhold, Peter Sanders, and Dorothea Wagner. 2021. Fast, Exact and Scalable Dynamic Ridesharing. In *2021 Proceedings of the Symposium on Algorithm Engineering and Experiments (ALENEX)* (Philadelphia, PA). SIAM, 98–112. <https://doi.org/10.1137/1.9781611976472.8>
- [9] Styliani-Stella Constantine Dafermos. 1968. *Traffic Assignment and Resource Allocation in Transportation Networks*. Ph.D. thesis. Johns Hopkins University.
- [10] Daniel Delling, Andrew V. Goldberg, Thomas Pajor, and Renato F. Werneck. 2017. Customizable Route Planning in Road Networks. *Transportation Science* 51, 2 (2017), 566–591. <https://doi.org/10.1287/trsc.2014.0579>
- [11] Robert B. Dial. 2006. A Path-based User-Equilibrium Traffic Assignment Algorithm that Obviates Path Storage and Enumeration. *Transportation Research Part B: Methodological* 40, 10 (2006), 917–936. <https://doi.org/10.1016/j.trb.2006.02.008>
- [12] Julian Dibbelt, Ben Strasser, and Dorothea Wagner. 2016. Customizable Contraction Hierarchies. *ACM Journal of Experimental Algorithmics* 21, 1 (2016), 1.5:1–1.5:49. <https://doi.org/10.1145/2886843>
- [13] Stéphane Durand and Bruno Gaujal. 2016. Complexity and Optimality of the Best Response Algorithm in Random Potential Games. *Lecture Notes in Computer Science* 9928 (2016), 40–51. https://doi.org/10.1007/978-3-662-53354-3_4
- [14] Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. 2004. The Complexity of Pure Nash Equilibria. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing* (Chicago, IL). ACM, 604–612. <https://doi.org/10.1145/1007352.1007445>
- [15] Andres Fielbaum, Xiaoshan Bai, and Javier Alonso-Mora. 2021. On-demand Ridesharing with Optimized Pick-up and Drop-off Walking Locations. *Transportation Research Part C: Emerging Technologies* 126 (2021). <https://doi.org/10.1016/j.trc.2021.103061>
- [16] Michael Florian, Isabelle Constantin, and Dan Florian. 2009. A New Look at Projected Gradient Method for Equilibrium Assignment. *Transportation Research Record* 2090, 1 (2009), 10–16. <https://doi.org/10.3141/2090-02>
- [17] Michael Florian and Donald Hearn. 1995. Network Equilibrium Models and Algorithms. In *Network Routing*. Handbooks in Operations Research and Management Science, Vol. 8. Elsevier, Chapter 6, 485–550. [https://doi.org/10.1016/S0927-0507\(05\)80110-0](https://doi.org/10.1016/S0927-0507(05)80110-0)
- [18] Marguerite Frank and Philip Wolfe. 1956. An Algorithm for Quadratic Programming. *Naval research logistics quarterly* 3, 1-2 (1956), 95–110.
- [19] Lars Gottesbüren, Michael Hamann, Tim Niklas Uhl, and Dorothea Wagner. 2019. Faster and Better Nested Dissection Orders for Customizable Contraction Hierarchies. *Algorithms* 12, 9 (2019). <https://doi.org/10.3390/a12090196>
- [20] Jonathan Gadea Harder, Simon Krogmann, Pascal Lenzner, and Alexander Skopalik. 2023. Strategic Resource Selection with Homophilic Agents. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, Edith Elkind (Ed.). IJCAI, 2701–2709. <https://doi.org/10.24963/ijcai.2023/301>
- [21] Alan M. Hershdorfer. 1966. Predicting the Equilibrium of Supply and Demand: Location Theory and Transportation Network Flow Models. In *Papers - Seventh Annual Meeting "Transportation - An Emerging Science"* (San Francisco). Transportation Research Forum.
- [22] Sin C. Ho, W.Y. Szeto, Yong-Hong Kuo, Janny M.Y. Leung, Matthew Petering, and Terence W.H. Tou. 2018. A Survey of Dial-a-Ride Problems: Literature Review and Recent Developments. *Transportation Research Part B: Methodological* 111 (2018), 395–421. <https://doi.org/10.1016/j.trb.2018.02.001>
- [23] R. Jayakrishnan, Wei K. Tsai, Joseph N. Prashker, and Subodh Rajadhyaksha. 1994. Faster Path-based Algorithm for Traffic Assignment. In *73rd Annual Meeting of the Transportation Research Board* (Washington, DC). SAGE, 75–83.
- [24] Amit Kumar and Srinivas Peeta. 2011. An Improved Social Pressure Algorithm for Static Deterministic User Equilibrium Traffic Assignment Problem. In *90th Annual Meeting of the Transportation Research Board* (Washington, DC). SAGE. <https://trid.trb.org/View/1091267>
- [25] Moritz Laupichler and Peter Sanders. 2024. Fast Many-to-Many Routing for Dynamic Taxi Sharing with Meeting Points. In *Symposium on Algorithm Engineering and Experiments (ALENEX)* (Alexandria, VA). SIAM, 74–90. <https://doi.org/10.1137/1.9781611977929.6>
- [26] Nicolai Mallig, Martin Kagerbauer, and Peter Vortisch. 2013. mobiTopp—A Modular Agent-based Travel Demand Modelling Framework. *Procedia Computer Science* 19 (2013), 854–859. <https://doi.org/10.1016/j.procs.2013.06.114>
- [27] Nicolai Mallig and Peter Vortisch. 2015. Modeling Car Passenger Trips in mobiTopp. *Procedia Computer Science* 52 (2015), 938–943. <https://doi.org/10.1016/j.procs.2015.05.169>
- [28] Dov Monderer and Lloyd S. Shapley. 1996. Potential Games. *Games and Economic Behavior* 14 (5 1996), 124–143. Issue 1. <https://doi.org/10.1006/game.1996.0044>
- [29] Yu Marco Nie. 2010. A Class of Bush-based Algorithms for the Traffic Assignment Problem. *Transportation Research Part B: Methodological* 44, 1 (2010), 73–89. <https://doi.org/10.1016/j.trb.2009.06.005>
- [30] Guillermo Owen. 2013. *Game Theory* (4th ed.). Emerald Group.
- [31] Olga Perederieieva, Matthias Ehrgott, Andrea Raith, and Judith Y.T. Wang. 2015. A Framework for and Empirical Study of Algorithms for Traffic Assignment. *Computers & Operations Research* 54 (2015), 90–107. <https://doi.org/10.1016/j.cor.2014.08.024>
- [32] Robert W. Rosenthal. 1973. A Class of Games Possessing Pure-strategy Nash Equilibria. *International Journal of Game Theory* 2 (1973), 65–67. <https://doi.org/10.1007/BF01737559>
- [33] Johannes Schlaich, Udo Heidl, and Regine Pohlner. 2011. Verkehrsmodellierung für die Region Stuttgart—Schlussbericht. *Unpublished Manuscript* (2011).
- [34] Yosef Sheffi. 1985. *Urban Transportation Networks*. Prentice-Hall.
- [35] Verband Region Stuttgart. 2011. Mobilität und Verkehr in der Region Stuttgart 2009/2010: Regionale Haushaltsbefragung zum Verkehrsverhalten. *Schriftenreihe Verband Region Stuttgart* 29 (2011), 1–138.
- [36] John Glen Wardrop. 1952. Some Theoretical Aspects of Road Traffic Research. *Proceedings of the Institution of Civil Engineers* 1, 3 (1952), 325–362. <https://doi.org/10.1680/ipeds.1952.11259>
- [37] Zhong Zhou and Matthew Martimo. 2010. Computational Study of Alternative Methods for Static Traffic Equilibrium Assignment. In *Proceedings of the 12th world conference on transport research (WCTRS)* (Lisbon, Portugal). 1–15.