



King's Research Portal

DOI: 10.1007/978-3-030-29911-8_37

Document Version Peer reviewed version

Link to publication record in King's Research Portal

Citation for published version (APA):

Serafino, P., Ventre, C., Tran-thanh, L., Zhang, J., An, B., & Jennings, N. (2019). Social Cost Guarantees in Smart Route Guidance. In A. C. Nayak, & A. Sharma (Eds.), *Pacific Rim International Conference on Artificial Intelligence: Trends in Artificial Intelligence - 16th Pacific Rim International Conference on Artificial Intelligence, Proceedings* (pp. 482-495). Article Chapter 37 (Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); Vol. 11671 LNAI). Advance online publication. https://doi.org/10.1007/978-3-030-29911-8_37

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
You may not further distribute the material or use it for any profit-making activity or commercial gain
You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Social Cost Guarantees in Smart Route Guidance *

Paolo Serafino¹, Carmine Ventre², Long Tran-Thanh³, Jie Zhang³, Bo An⁴, and Nick Jennings⁵

¹ Gran Sasso Science Institute, Italy paolo.serafino@gssi.it
 ² University of Essex, UK, c.ventre@essex.ac.uk
 ³ University of Southampton, UK, {ltt08r,jie.zhang}@soton.ac.uk
 ⁴ Nanyang Technological University, Singapore, boan@ntu.edu.sg
 ⁵ Imperial College London, UK, n.jennings@imperial.ac.uk

Abstract. We model and study the problem of assigning traffic in an urban road network infrastructure. In our model, each driver submits their intended destination and is assigned a route to follow that minimizes the social cost (i.e., travel distance of all the drivers). We assume drivers are strategic and try to manipulate the system (i.e., misreport their intended destination and/or deviate from the assigned route) if they can reduce their travel distance by doing so. Such strategic behavior is highly undesirable as it can lead to an overall suboptimal traffic assignment and cause congestion. To alleviate this problem, we develop moneyless mechanisms that are resilient to manipulation by the agents and offer provable approximation guarantees on the social cost obtained by the solution. We then empirically test the mechanisms studied in the paper, showing that they can be effectively used in practice in order to compute manipulation resistant traffic allocations.

1 Introduction

Recent years have witnessed increasing interest in the development of efficient traffic control systems [15,9]. This is motivated by the significant negative impact on the quality of life of both road users and residents caused by heavy traffic congestion levels in large cities such as London, Beijing, and Los Angeles. Indeed, heavy congestion is known to be a major cause of air and noise pollution, which are widely recognized as the main cause of many health issues [14,22]. Adding to this is the economic cost associated with the large amount of time spent in traffic jams, which reduces the productivity of the economy [13]. Moreover, the situation is expected to become significantly worse in the future when the population, and thus the traffic flow, in large cities will be much bigger than at present. Unfortunately, conventional traffic control systems have proven unable to efficiently decrease congestion levels, as they are not designed to be

^{*}This work has been partially supported by the Italian MIUR PRIN 2017 Project ALGADIMAR "Algorithms, Games, and Digital Markets".

adaptive to the dynamics of city traffic, which changes over space and time. On the other hand, it has been shown [20,16] that by putting some sort of intelligence/smartness into traffic control systems, we can make them adapt to the changes of the traffic flow. A key objective within these smart traffic control systems is to address the so-called *traffic assignment problem* (TAP), in which mobile agents (i.e., typically drivers) declare their intended destination to the system, perhaps via their satellite navigation systems, and are then assigned a route to follow, in such a way that some objective function of the overall traffic flow in the system is optimized (i.e., minimizing the total traveled distance or maintaining an efficient traffic load balance). As these agents are typically selfinterested and strategic (i.e., they try to maximize their own utility, disregarding whether this is detrimental to the global optimization goal), they may manipulate the system whenever they can benefit from doing so [16,24]. This kind of opportunistic behavior is highly undesirable as it will increase the total social cost (i.e., decreasing the total load balance or increasing the total congestion level). As such, incentivizing agents not to be strategic is a key design objective of these traffic assignment systems [20,16,24]. Given this, we focus on strategyproof TAP mechanisms, which guarantee that it is in the agent's best interest to always report her true destination and follow the assigned route. Furthermore, we assume that money transfers between the mechanism and the agents are not available. This is a common assumption in many domains [19] that will facilitate the likely real-world deployment of the system by lowering set up costs (i.e., avoiding the construction of tolling booths). The remainder of the paper is organized as follows. In Section 2 we discuss related works. In Section 3 we introduce our model for TAP and prove that Pareto optimal allocations theoretically guarantee that agents will follow their assigned paths (Theorem 1). We then move to study deterministic (Section 4) and randomized (Section 5) Pareto optimal mechanisms for our problem. We show that the approximation ratio of *deterministic* strategyproof mechanisms is lower bounded by 3 (Theorem 2), while the Serial Dictatorship mechanism can achieve an upper bound of $2^n - 1$ and it is Pareto-optimal and non-bossy (Theorems 4 and 5), where n is the number of agents (Theorems 4 and 5). Furthermore, if we require non-bossiness and Pareto optimality, we are able to close this approximation ratio gap by showing that the Bipolar Serial Dictatorship mechanism is the *only* strategyproof mechanism. For randomized mechanisms, we show that the approximation ratio is lower bounded by $\frac{11}{10}$ (Theorem 7). In addition, the Random Serial Dictatorship mechanism can achieve an *n*-approximation (Theorems 8 and 9), while still preserving the desired properties of Pareto-optimality and non-bossyness. In addition to these theoretical results, we present an extensive experimental evaluation on traffic networks generated from real road network data, which show how the mechanisms studied in the paper provide good performance in practice, despite the high theoretical worst case approximation guarantee.

2 Related Work

There is a large body of literature on traffic network modelling and assignment [2,21,7,8]. However, these works typically ignore the strategic behaviour of participating agents. Nevertheless, they can be useful to model the underlying traffic network in our work. In particular, we follow the widely used traffic model proposed in [2]. To tackle the strategic behaviour of the agents, several researchers have suggested employing mechanism design with money and auction theory for traffic control [20,16,24,4]. These works typically rely on the computation of the VCG auction in order to assign vehicles to paths. However, they require monetary incentives, and typically focus on a local control level, such as intersection management (as VCG is typically computationally hard, and thus, not readily scalable [6]). A number of researchers have focused on mechanism design without money [19,5]. However, none of these mechanisms can be easily applied to the traffic assignment problem, as they do not take into account the features of the underlying traffic network structure. As we will show, TAP bears some resemblance to the problem of assigning indivisible objects [3,23,10], although these results are not directly applicable to our scenario. Indeed TAP has a much more complex structure (mainly due to the underlying traffic network topology) which traditional assignment mechanisms fail to address.

3 Model and Preliminary Definitions

A traffic assignment problem (TAP) consists of a set of agents $A = \{a_1, \ldots, a_n\}$ and a road network infrastructure, represented as a directed graph G = (V, E), where: (i) $V = \{v_1, \ldots, v_{|V|}\}$ is the set of nodes representing the junctions of the road network infrastructure; and (ii) $E \subseteq V \times V$ is the set of directed edges representing one-way road segments. Each edge $e \in E$ has a *capacity* $c: E \to \mathbb{N}^+$, which determines the maximum number of agents that can travel through the edge at any given time, and a weight function $w: E \to \mathbb{R}^+$ which represents the cost incurred by the agent traveling through the edge (i.e., travel distance). Furthermore, each edge is associated to a transit time $\tau : E \to \mathbb{Z}^+$ which represents the free travel time of the edge (i.e., the minimum travel time needed to travel through the road at maximum allowed speed). This means that agent a_i setting off at time t from node v_o and heading to node v_d through the edge (v_o, v_d) will reach node v_d at time $t + \tau(v_o, v_d)$, and will occupy edge (v_o, v_d) only in the time interval $[t, t+\tau(v_o, v_d)]$. Unless stated otherwise, we assume that edges (u, v) and (v, u) are symmetrical: for all $(u, v), (v, u) \in E$ $c(u, v) = c(v, u), (v, u) \in E$ w(u, v) = w(v, u) and $\tau(u, v) = \tau(v, u)$.

As in [17], we assume that if the flow of traffic through an edge does not exceed its capacity, then no congestion occurs and the traveling time equals the free travel time. Initially, at time t = 0, agents reside on a (publicly known) set

 $O \subseteq V$ of nodes⁶ of the graph, O_i being the initial location of agent a_i . Each agent $a_i \in A$ wants to reach an intended destination $D_i \in V$, which is the agent's private information and will be referred to in the remainder as her *type*.

Agents submit (or *bid*) a destination to an *allocation mechanism*, which then assigns each agent a path in order to optimize a certain objective function. More formally, let \mathcal{P} be the set of all possible simple paths between any two nodes in G. Let $\mathbf{D} = (D_1, \ldots, D_n) \in V^n$ be a vector of declarations (also referred to as *bids*) by the agents and \mathbf{D}_{-i} be the vector of declarations of all agents but a_i . A mechanism $M^{G,O}: V^n \to \mathcal{P}^n$ maps a vector of declarations to *feasible paths* (i.e., not exceeding the capacity of the edges at any given time) on G, given the initial locations O of the agents. We write $M(\mathbf{D})$ instead of $M^{G,O}(\mathbf{D})$ when G and O can be deduced from the context. The path associated to agent a_i is denoted as $M_i(\mathbf{D})$.

A traffic assignment $S = M(\mathbf{D})$ induces a flow over time⁷ $f_S : E \times \mathcal{T} \to \mathbb{N}^+$, where \mathcal{T} is a suitable discretization of time w.r.t. the transit times of the edges of G (for simplicity we will assume that $\mathcal{T} = \{0, 1, \ldots, T\}$, where T is a time horizon sufficient for the network to clear. Thus, $f_S(u, v; t) = |\{a_i \in A | (u, v) \in S_i\}|$ is the number of agents that are assigned a path that contains edge (u, v) at time $t \in \mathcal{T}$. Feasibility constraints imply that $f_S(u, v; t) \leq c(u, v)$ for all $t \in \mathcal{T}$.

In the remainder, without loss of generality, we will study the problem on the *time-expanded network* [11,12] of G and consider the *static* flow through it (i.e, the transit of an agent over and edge is instantaneous). A time-expanded network is a properly constructed directed graph with cost and capacity functions on the edges just like G, but no transit time (i.e. travel time is instantaneous through all the edges). This is without loss of generality from the point of view of SP, Pareto-optimality, non-bossines and approximation guarantee since it is well known (see [11,12]) that a flow over time is equivalent to a static flow on the corresponding time-expanded network.

Let $f_S^{-i}: E \to \mathbb{N}$ be the flow induced by traffic assignment S generated by agents $A \setminus \{a_i\}$, formally for all $e \in E$, $f_S^{-i}(e) = |\{a_j \in A : e \in S_j, j \neq i\}|$. The residual graph G_f^{-i} is a graph such that: (i) G_f^{-i} has the same nodes and edges as G; (ii) each edge $e \in E$ of G_f^{-i} has capacity $c(e) - f_S^{-i}(e)$. For any two nodes $u, v \in V$, let $\mathcal{P}_{u,v}$ denote the set of simple paths in G connecting u to v. Furthermore, for all traffic assignments $S = M(\mathbf{D})$ and all agents a_i , let $\mathcal{P}_{u,v}^i(S) = \{P \in \mathcal{P}_{u,v} | \forall e \in P, c(e) > f_S^{-i}(e)\}$. Informally, $\mathcal{P}_{u,v}^i(S)$ is the set of paths connecting u and v that have spare capacity from the perspective of agent a_i (i.e., they can be used by agent a_i) when the other agents implement

⁶Restricting origins/destinations of journeys to road junctions is without loss of generality since fictitious nodes that serve the sole purpose of acting as starting/ending point of a journey can always be created by edge splitting operations.

⁷Sometimes also referred to as *dynamic flow* in the literature. We prefer the term *flow* over time as the adjective *dynamic* has often been used for settings where the input data arrive online or change over time. We assume that all the agents are present at time t = 0 and the network is cleared after the last agent reaches their destination.

S. Then, the set of reactions available to agent a_i having type D_i at allocation S is defined as $R_i(S) = \mathcal{P}^i_{O_i,D_i}(S)$.

Agents are not constrained to follow their assigned path but can choose a different one, subject to capacity constraints⁸. To model this, as per [18], we assume that, after the mechanism computes a traffic allocation, the agents can react by choosing an action from a set $R_i \subseteq \mathcal{P}$. Hence, the actual cost function of an agent depends on: (i) her true type D_i ; (ii) the allocation S chosen by the mechanism on input the bids reported by the agents; and (iii) the reactions chosen by the agents.

We can now formally define the cost function of an agent. Given an allocation $S' = M(D'_i, \mathbf{D}_{-i})$, the cost of an agent of type D_i with respect to S' is defined as: $cost_i(S', D_i) = \min_{P \in R_i(S')} w(P)$ where $w(P) = \sum_{(u,v) \in P} w(u,v)$ denotes the cost of P. We assume that agents are risk-neutral. In what follows, we define a set of desiderata for our allocation mechanism, namely: (i) strategyproofness, (ii) Pareto optimality and (iii) non-bossiness.

A deterministic mechanism M is *strategyproof* (SP for short) if, for all agents a_i , for all declarations D_i and D'_i and all declarations of the other agents \mathbf{D}_{-i} , agent a_i cannot decrease her cost by misreporting her true type, namely:

$$cost_i(M(\mathbf{D}), D_i) \le cost_i(M(D'_i, \mathbf{D}_{-i}), D_i)$$
 (1)

A randomized mechanism is *strategyproof in expectation* if (1) holds in expectation (i.e., over the random choices of the mechanism). A randomized mechanism is *universally strategyproof* if agents cannot gain by lying regardless of the random choices made by the mechanism, i.e., the output of the mechanism is a distribution over strategyproof deterministic allocations.

The social cost of an allocation S is defined as $SC(S, \mathbf{D}) = \sum_{a_i \in A} cost_i(S, D_i)$. A mechanism OPT is optimal for TAP if $OPT(\mathbf{D}) \in \arg\min_{S \in \mathcal{P}^n} SC(S, \mathbf{D})$ for all **D**. A mechanism M is an α -approximation (w.r.t the optimal social cost) with $\alpha \in \mathbb{R}, \alpha \geq 1$, being referred to as the approximation ratio of M, if, for all **D**, $SC(M(\mathbf{D}), \mathbf{D}) \leq \alpha \cdot SC(OPT(\mathbf{D}), \mathbf{D})$.

A traffic allocation $S \in \mathcal{P}^n$ is *Pareto optimal* if there exists no other feasible traffic allocation S' such that $cost_j(S', D_j) \leq cost_j(S, D_j)$ for all a_j , and $cost_k(S', D_k) < cost_k(S, D_k)$ for some a_k . Pareto optimal allocations are of particular interest in our scenario, because, as proven in Theorem 1, they are a min-cost response in the available reactions $R_i(S)$ of an agent. This gives us a theoretical guarantee that agents will actually implement Pareto optimal solutions returned by the mechanism.

Theorem 1. Let $S = M(\mathbf{D})$ be a traffic assignment and let $R_i(S)$ be the set of reactions available to a_i at S. If S is Pareto optimal, then $M_i(\mathbf{D}) \in \arg\min_{P \in R_i(S)} w(P)$.

⁸Agents are not prevented from using edges not belonging to their assigned paths, as this would result in a waste of public resources (i.e., road capacity). To avoid congestion, we disincentivize agents from using edges that, according to the scheduled traffic, are filled to capacity. This can be easily implemented through the use of traffic cameras that check cars' number plates.

Finally, a mechanism M is non-bossy if $M_i(\mathbf{D}) = M_i(D'_i, \mathbf{D}_{-i})$ implies that $M_j(\mathbf{D}) = M_j(D'_i, \mathbf{D}_{-i})$, for all $a_i, a_j \in N$ and all \mathbf{D} and D'_i . In other words, non-bossyness excludes (arguably undesirable) mechanisms that allow one agent to change the allocation of other agents without changing her own too. In the remainder of this paper, we focus on strategyproof mechanisms for TAP that approximately achieve the optimal social cost. In particular, we are interested in mechanisms that are also Pareto-optimal and non-bossy.

4 Deterministic Mechanisms

In this section, we discuss deterministic mechanisms for TAP. We first provide a lower bound on the approximation ratio of SP deterministic mechanisms.

Theorem 2. There is no α -approximate deterministic SP mechanism for the traffic assignment problem with $\alpha < 3 - \varepsilon$, for any $\varepsilon > 0$.

These impossibility results suggest that in order to achieve strategyproofness we have to give up on optimality. This naturally leads to asking to what extent can we approximate the optimal social welfare while satisfy the desired properties. As a first step to answer this question, we examine the well-known Serial Dictatorship mechanism that is deterministic and notoriously satisfies our three desiderata (i.e., strategyproofness, Pareto optimality and non-bossiness).

Definition 1. Mechanism Serial Dictatorship (SD), given an ordering $a_1 \prec \ldots, \prec a_n$ of the agents, allocates paths to agents in n stages such that at stage i agent a_i is allocated her minimum cost path in the residual graph $G_f^{-\{a_1,\ldots,a_{i-1}\}}$.

The following theorem proves that SD is feasible under some mild conditions:

Theorem 3. If G is K-edge-connected⁹, mechanism SD is feasible for K agents.

Next we provide an upper bound on the approximation ratio of SD, and thus, on its worst case performance. In order to prove our result, we make the following assumption:

Definition 2. The deviation on capacious path assumption (DoCP) assumes that whenever the SD mechanism allocates to an agent a path that is different from the one that the optimal mechanism would allocate, the assigned path has sufficient capacity to potentially be allocated to all the remaining agents.

To better understand this assumption, consider the following example. With reference to Figure 1, let a_i be an agent and P_i^* be the path she is assigned in the optimal allocation (i.e., $OPT_i = P_i^*$). If agent a_i is not assigned P_i^* by SD, there must be an agent a_j , where $j \prec i$ in the ordering used by SD, such that: (i) $SD_j = P_j \neq OPT_j$ and (ii) $P_j \cap P_i^* \neq \emptyset$ and (iii) at least one edge of P_i^*

 $^{^{9}\}mathrm{A}$ graph is K-edge-connected if it remains connected when strictly fewer than K edges are removed.

is saturated after a_j is assigned P_j . In such a situation, we say that agent a_i is blocked by agent a_j . Let $\alpha_i \in P_j \cap P_i^*$ ($\beta_i \in P_j \cap P_i^*$, respectively) be the first (last, respectively) node of P_i^* in P_j . The DoCP assumption postulates that if a_j blocks a_i , then the alternative path of blocked agent a_i through blocking agent $a_j \ \Gamma_i^j = (O_i, \alpha_i, O_j, D_j, \beta_i, D_i)$ has at least capacity $n - |\{a_k \in A | a_j \prec a_k\}|$ in the residual graph $G_f^{-\{a_1,\ldots,a_j\}}$. That is, all agents yet to be assigned by SD after a_j is blocked by agent a_j then path Γ_i^j always exists, although unless we assume DoCP, it might not have spare capacity to be assigned to agent a_i . It is not



Fig. 1. Deviation on capacious paths

difficult to see that if we relax the DoCP assumption, then the approximation ratio of SD is not bounded by any function of the number of agents on certain pathological TAP instances.

Theorem 4. Under the DoCP assumption, SD is at most $(2^n - 1)$ -approximate.

Proof (Proof sketch). We prove the claim by induction on the number of players. Let OPT_i denote the cost and solution (with a slight abuse of notation) of the optimal allocation that only considers bids of agents $j \leq i$. Similarly, let SD_i denote the cost and solution of SD on input all the bids of agents $j \leq i$. Base of the induction (i = 1): trivially $OPT_1 = SD_1$. Now assume that the claim is true for i - 1 and, for $j \leq i$, let P_j^* (P_j , respectively) be the path assigned to agent j by OPT_i (SD_i , respectively). For a path P, we let w(P) denote the cost of the path in the given graph G. We want to prove that under the DoCP assumption, the following holds:

$$w(P_i) \le OPT_i + SD_{i-1}.$$
(2)

If $P_i^* = P_i$ then we are done. Therefore, we can assume that $P_i^* \neq P_i$. This means that the paths P_j allocated to agents j < i by SD_i saturate some of the edges of P_i^* . Now, for at least one of these agents, say $\overline{j}, P_{\overline{j}}^* \neq P_{\overline{j}}$ for otherwise also in OPT_i path P_i^* would be unavailable to i. But then $w(P_i) \leq w(\Gamma_i^{\overline{j}}), \Gamma_i^{\overline{j}}$ being the path that connects O_i to D_i through $O_{\overline{j}}$, as per the definition of DoCP. Note that, under the DoCP assumption, $\Gamma_i^{\overline{j}}$ is always feasible. Since $\Gamma_i^{\overline{j}}$ uses only edges in $OPT_i \cup SD_{i-1}$ (i.e. P_i^* and P_j^* are in OPT_i , paths (O_i, α_i) and (β_i, D_j)

belong to SD_{i-1}), (2) is proven. Finally, (2) and the inductive hypothesis yield:

$$SD_{i} = SD_{i-1} + w(\Gamma_{i}^{j}) \le 2SD_{i-1} + OPT_{i}$$

$$\le 2((2^{i-1} - 1)OPT_{i-1}) + OPT_{i} \le (2^{i} - 1)OPT_{i}$$

As the $(2^n - 1)$ -approximation ratio can be prohibitively large for large n, we ask ourselves whether we can further improve this upper bound. Unfortunately, the following theorem answers this question in the negative.

Theorem 5. Under the DoCP assumption, the bound of Theorem 4 is tight.

We now provide a characterization of SP, Pareto-optimal, and non-bossy mechanisms for a subset of instances of TAP, named TAP⁺ and we prove that the family of all mechanisms satisfying the above properties is comprised by a generalization of SD, namely *Bi-polar Serial Dictatorship* (BSD). Such a characterization extends naturally to TAP instances. TAP⁺ is subset of instances of TAP having a peculiar structure: (*i*) every agent has the same source node O; (*ii*) O has outgoing edges with unitary capacity and no ingoing edges, let $E_O = \{(O, v_1), \ldots, (O, v_m)\}$ denote the set of outgoing edges of O; and (*iii*) the set of possible destinations that the agents can declare is restricted to a given subset $\mathcal{D} \subset V$.

Definition 3. Given an ordering of the agents $\{i_1, i_2\} \prec i_3 \prec \ldots \prec i_n$ and a bipartition $\{X_1, X_2\}$ of the set of alternatives X (i.e., paths in the case of TAP) such that $X_1 \cap X_2 = \emptyset$ and $X_1 \cup X_2 = X$, a BSD mechanism executes SD with ordering $i_2 \prec i_1 \prec \ldots \prec i_n$ if $\min_{x \in X} cost_1(x) = \min_{x \in X} cost_2(x) = x \in X_2$; otherwise SD with ordering $i_1 \prec i_2 \prec \ldots \prec i_n$ is executed.

Theorem 6. A traffic allocation mechanism for TAP^+ is Pareto-optimal, SP and non-bossy if and only if it is a Bi-polar Serially Dictatorial Rule.

Proof (Proof sketch). We reduce an instance of the problem of assigning indivis*ible objects* with general ordinal preferences [3] (AIO for short) to TAP⁺. In an instance of AIO, a set of objects $X = \{x_1, \ldots, x_m\}$ has to be assigned to a set of agents $A = \{a_1, \ldots, a_n\}$, such that every agent receives at most one object and no agent is left without an object if there are objects still available. Agents have ordinal general preferences \succeq_i , where $x \succeq_i y$ for $x, y \in X$ means that agent i (weakly) prefers object x to object y. From an instance of AIO, we build an instance of TAP⁺ as follows. TAP⁺ has the same set of agents A as AIO. Graph G of TAP⁺ has a node O such that $O_i = O$ for all $a_i \in A$. For every object $x_i \in X$ we construct in G a node v_i and an edge (O, v_i) such that $c(O, v_i) = 1$ and $w(O, v_i) = \varepsilon$ for $0 < \varepsilon \ll 1$. Let Ψ be the set of all possible preference relations over X. We construct $|\Psi|$ destination nodes D_k , one for each preference relation $\succeq \in \Psi$ and for each $k \in 1, \ldots, |\Psi|$. For each $j \in \{1, \ldots, m\}$ we add an edge (v_i, D_k) having capacity 1 and weight $w(v_i, D_k)$ equal to the ranking of x_i according to \succeq . We can now transform an instance of the so-constructed TAP⁺ problem to an instance of the AIO problem, and vice versa. In [3] it is proved that BSD is the only Pareto optimal, SP and non-bossy mechanism for AIO. This characterization transfers to TAP⁺ due to the reduction sketched above.

Next, we investigate the performance of BSD and show that it does not asymptotically perform better than SD. In particular, we state that:

Lemma 1. BSD cannot achieve an approximation ratio lower than $\Omega(2^n)$.

5 Randomized Mechanisms

Given the undesirable approximation guarantees of deterministic mechanisms, we now turn to randomization. Randomized mechanisms can often be interpreted as fractional mechanisms for the deterministic solutions, under mild conditions. We start by proving the following inapproximability lower bound:

Theorem 7. There is no α -approximate universally truthful randomized mechanism for the traffic assignment problem with $\alpha < 11/10$.

In the remainder of this section, we study the randomized version of SD for TAP, which is universally strategyproof, (ex-post) Pareto optimal and non-bossy.

Definition 4. The Randomized Serial Dictatorship (RSD) mechanism computes uniformly at random an ordering σ over the agents and returns the output of SD over ordering σ .

The following results gives a tight bound on the approximation ratio of RSD.

Theorem 8. Under the DoCP assumption, RSD is at most n-approximate.

Proof (Proof sketch). We are going to prove the claim by induction on the number of agents. As above, let OPT_i denote the cost of the optimal solution with paths assigned only to agents a_j , with $j \leq i$. With a slight abuse of notation we also let OPT_i denote the solution itself. Similarly, RSD_i denotes the expected cost of RSD on input all the bids of agents $a_j, j \leq i$. For the base of the induction with i = 1, it is clear that RSD_1 is the optimal solution. Now assume that the claim is true for i - 1 and consider an instance with i agents. Let $I_{-k}(P)$, P being a path from O_k to D_k , be the instance of the problem without agent a_k and with the capacity of the directed edges in P diminished by one (i.e., as if the path P were used by a_k). Note that by the DoCP assumption, one of the agents a_j , with $j \neq k$, is guaranteed to be able to use the edges of P in the opposite direction than a_k . We now let $OPT_{-k,P}$ and $RSD_{-k,P}$ be the cost of the optimum and expected cost of RSD on $I_{-k}(P)$, respectively. Moreover, let π_j be the path minimizing the cost of agent a_j (i.e., the path that SD would assign to a_j if she was the first to choose). We then have

$$\begin{split} RSD_i &= \frac{1}{i} \sum_{k=1}^{i} \left(w(\pi_k) + RSD_{-k,\pi_k} \right) \le \frac{1}{i} \sum_{k=1}^{i} \left(w(\pi_k) + (i-1)OPT_{-k,\pi_k} \right) \\ &\le \frac{1}{i} \sum_{k=1}^{i} w(\pi_k) + \frac{1}{i} \sum_{k=1}^{i} \left((i-1)(OPT_i + w(\pi_k)) \right) \le \frac{1}{i} OPT_i + (i-1)OPT_i + \frac{i-1}{i} OPT_i \\ &= i \cdot OPT_i \end{split}$$

where the first equality follows from the definition of RSD, i.e., with probability 1/i each agent k will have the first choice. As for the inequalities, we note that the first follows from the inductive hypothesis whilst the last from the observation that $OPT_i \geq \sum_{k=1}^{i} w(\pi_k)$. We are left with the second inequality. That is, we prove that under the DoCP $OPT_{-k,\pi_k} \leq OPT_i + w(\pi_k)$. If OPT_i allocates π_k to agent a_k then we are done. Otherwise, let P_k be the path that a_k gets in OPT_i and note that the paths P_j allocated to agents $a_j \ j \neq k$ by OPT_i saturates some of the edges of P_k ; let $a_{\bar{j}}$ be one of these agents. Consider now the solution S to $I_{-k}(\pi_k)$ where all agents but $a_{\bar{j}}$ are allocated the same path as in OPT_i and agent $a_{\bar{j}}$ is given, instead of $P_{\bar{j}}$, the alternative path Γ_j^k through agent a_k . Observe that $\Gamma_{\bar{j}}^k$ uses the same directed edges of $P_{\bar{j}}$ and P_k and the edges of π_k in opposite direction and, as observed above, under the DoCP assumption, is a feasible path for $a_{\bar{j}}$ and S a feasible solution to $I_k(\pi_k)$, whose social cost is denoted SC(S). But then:

$$OPT_{-k,\pi_k} \le SC(S) = OPT_i - w(P_j) - w(P_k) + w(P) \le OPT_i + w(\pi_k)$$

where the last inequality follows from the fact that the edges in $P \setminus (P_k \cup P_j)$ are a subset of the edges in π_k .

Theorem 9. The approximation ratio of RSD is $\Omega(n)$.

This means that by allowing randomness in the allocation mechanism, we can reduce the exponential approximation of the deterministic case to a linear one.



Fig. 4. Experimental results on NY-4000 Fig. 5. Experimental results on NY-10000

6 Experimental Results

In this section we present the results of the experimental evaluation we conducted in order to assess whether the theoretical inapproximability lower bounds impose a high approximation cost on real-life instances. In short, we will show that they do not. In particular, we have measured the approximation ratio obtained by SD and RSD on three real-life graphs extracted from the DIMCAS 99 shortest path implementation challenge benchmark datasets [1]. In particular, Rome99 represents a large portion of the directed road network of the city of Rome, Italy, from 1999. The graph contains 3353 vertices and 8870 edges. Vertices correspond to intersections between roads and edges correspond to roads or road segments. NY-4000 and NY-10000 are two subgraphs extracted from NY-d, a larger distance graph (with 264,346 nodes and 733,846 edges) representing a large portion the road network infrastructure of New York City, USA. The two graphs were obtained by taking a subset, respectively, of the first 4000 and 10000 nodes of the graph while ensuring that the connectivity was preserved by adding edges representing paths through nodes of the original graph not included in the subgraph. In Table 3 some statistics related to the structural characteristics of our test graphs are reported, where δ^+_{AVG} represents the average outdegree of a node (i.e. the average number of edges originating from a node) and c_{AVG} is the average capacity of the outgoing edges of a node. In our experimental assessment, we studied the variation of the approximation ratio of SD and RSD on the test graphs while varying the resource augmentation factor. The resource augmentation factor is the key parameter of the resource augmentation framework [5], a novel comparison framework where a truthful mechanism that allocates "scarce resources" is evaluated by its worst-case performance on an instance where such "scarce resources" are augmented, against the optimal mechanism on the same instance with the original amount of resources. In [5] it is argued that this is a fairer comparison framework than the traditional approximation ratio, which compares the performance of a mechanism that is severely limited by the requirement of truthfulness to that of an omnipotent mechanism that operates under no restrictions and has access to the real inputs of the agents. An equivalent resource augmentation framework is often also used in the analysis of online algorithms. In the TAP scenario, the natural resource to be augmented is the capacity of the existing edges, modelled by the augmentation factor γ , which in our framework is defined as the factor by which the average capacity of the edges departing from a node is multiplied, spreading the excess capacity evenly among the outgoing edges of the node. More formally, if $c_{AVG}(v)$ is the average capacity of node v, then the augmented average capacity $c_{AVG}^{\gamma}(v) = \gamma \cdot c_{AVG}$, and the capacity of each outgoing edge is set as $\frac{c_{AVG}(v)}{\delta^+(v)}$, where $\delta^+(v)$ is the outdegree of v. In our experiments we ranged the augmentation factor γ in the interval [1,2], which means increasing the initial capacity until it is doubled. To run our experiments, we generated three separate populations of agent-origin-destination triplets, one population for each test graph, each comprising a number of triplets roughly equal to 1/3 of the nodes of the graph. The size of the population of triplets was empirically tailored to let the competition for popular links arise without making the allocation problem unfeasible. For each agent-origin-destination triplet in the population, both the origin and the destination were independently drawn

uniformly at random from the set of the nodes of the graph, with replacement (i.e. the same node can be the origin/destination of multiple triplets).

Figures 2, 4 and 5 show the results of our experimental analysis, respectively on graph Rome99, NY-4000 and NY-10000. In particular, the left hand side plot represents the absolute value of the social cost for the optimal mechanism, expressed in kilometers, for SD and for RSD, whereas the right hand side plot represents the approximation ratio for SD and RSD. From our experimental analysis we can see that the actual approximation ratio of both SD and RSD is much lower than the predicted theoretical worst-case approximation. In particular, our experiments show that the approximation ratios of SD and RSD are quite similar and strongly o(n) on the investigated road networks. This is due to the fact that such theoretical approximation lower bounds rely on pathological instances that are quite unlikely to occur in real life graphs. It is also worth noting the beneficial effect that augmenting the capacity of existing roads has on the approximation ratio: increasing the augmentation factor steadily decreases the approximation ratio on both Rome99 and NY-4000. On the other hand a marked decrease is noticeable only if we increase the augmentation factor to 1.8 in the case of NY-10000. This phenomenon is due to the already reach topological structure of NY-10000, which necessitates less augmentation to yield good performances.

7 Conclusions

In this paper we investigate the problem of strategyproof traffic assignment without monetary incentives. We study two SP mechanism for our problem, namely Serial Dictatorship and its randomized counterpart Random Serial Dictatorships. For deterministic mechanisms we prove that Serial Dictatorship is $2^n - 1$ under some mild assumptions, and characterize Bipolar Serial Dictatorship as the only SP, Pareto optimal and non-bossy deterministic mechanism for our problem. In the randomized case, we prove that Random Serial Dictatorship is *n*-approximate. Finally we assess the performance of Serial Dictatorship and Random Serial Dictatorship on real road network infrastructure, and show that they exhibit good approximation guarantees. In particular, RSD is almost indistinguishable from SD, which means that the instances giving rise to the inapproximability results rarely occur in practice. Note that our work is the first that addresses the problem of moneyless strategyproof traffic assignment. Although it ignores a number of properties that occur in real-world scenarios (e.g., dynamic network behavior, or asynchronous bid submissions), it still serves as a proof of concept for moneyless strategyproof assignment mechanisms.

References

1. 9th DIMACS Implementation Challenge - Shortest Paths. http://www.diag. uniroma1.it/challenge9/download.shtml (1999)

- Beckmann, M., McGuire, C., Winsten, C.B.: Studies in the economics of transportation. Techn. report (1956)
- Bogomolnaia, A., Deb, R., Ehlers, L.: Strategy-proof assignment on the full preference domain. J. Economic Theory 123(2), 161–186 (2005)
- 4. Brenner, J., Schäfer, G.: Online Cooperative Cost Sharing, pp. 252–263 (2010)
- Caragiannis, I., Filos-Ratsikas, A., Frederiksen, S.K., Hansen, K.A., Tan, Z.: Truthful facility assignment with resource augmentation: An exact analysis of serial dictatorship. In: WINE'16. pp. 236–250 (2016)
- Conitzer, V., Sandholm, T.: Failures of the VCG mechanism in combinatorial auctions and exchanges. In: AAMAS'06. pp. 521–528. ACM (2006)
- Coogan, S., Arcak, M.: A compartmental model for traffic networks and its dynamical behavior. IEEE Trans. on Automatic Control 60(10), 2698–2703 (2015)
- Daganzo, C.F.: The cell transmission model: A dynamic representation of highway traffic consistent with the hydrodynamic theory. Transportation Research Part B: Methodological 28(4), 269–287 (1994)
- Djahel, S., Salehie, M., Tal, I., Jamshidi, P.: Adaptive traffic management for secure and efficient emergency services in smart cities. In: PERCOM'13. pp. 340– 343. IEEE (2013)
- Filos-Ratsikas, A., Frederiksen, S.K.S., Zhang, J.: Social welfare in one-sided matchings: Random priority and beyond. In: SAGT'14. pp. 1–12 (2014)
- Ford, L.R., Fulkerson, D.R.: Constructing maximal dynamic flows from static flows. Oper. Res. 6(3), 419–433 (Jun 1958)
- 12. Ford, L.R., Fulkerson, D.R.: Flows in Networks. Princeton University Press (1962)
- 13. Goodwin, P.: The economic costs of road traffic congestion (2004)
- Krzyzanowski, M., Kuna-Dibbert, B., Schneider, J.: Health effects of transportrelated air pollution. WHO Regional Office Europe pp. 1–190 (2005)
- Leontiadis, I., Marfia, G., Mack, D., Pau, G., Mascolo, C., Gerla, M.: On the effectiveness of an opportunistic traffic management system for vehicular networks. IEEE Trans. on Int. Transportation Sys. 12(4), 1537–1548 (2011)
- Levin, M.W., Fritz, H., Boyles, S.D.: On optimizing reservation-based intersection controls. IEEE Trans. on Int. Transportation Sys. 18(3), 505–515 (2017)
- Nesterov, Y., de Palma, A.: Stationary dynamic solutions in congested transportation networks: Summary and perspectives. Networks and Spatial Economics 3(3), 371–395 (Sep 2003)
- Nissim, K., Smorodinsky, R., Tennenholtz, M.: Approximately optimal mechanism design via differential privacy. In: ITCS'12. pp. 203–213 (2012)
- Procaccia, A.D., Tennenholtz, M.: Approximate mechanism design without money. ACM TEAC. 1(4), 18:1–18:26 (2013)
- Raphael, J., Maskell, S., Sklar, E.: From goods to traffic: first steps toward an auction-based traffic signal controller. In: PAAMS'15. pp. 187–198. Springer (2015)
- Skabardonis, A., Dowling, R.: Improved speed-flow relationships for planning applications. Transportation Research Record: Journal of the Transportation Research Board (1572), 18–23 (1997)
- Stansfeld, S.A., Matheson, M.P.: Noise pollution: non-auditory effects on health. Oxford Journals, Medicine and Health, British Medical Bulletin 68(1), 243–257 (2003)
- Svensson, L.G.: Strategy-proof allocation of indivisible goods. Social Choice and Welfare 16(4), 557–567 (Sep 1999)
- Vasirani, M., Ossowski, S.: A market-inspired approach for intersection management in urban road traffic networks. JAIR 43, 621–659 (2012)