

Urgency-aware Optimal Routing in Repeated Games through Artificial Currencies

Mauro Salazar¹, Dario Paccagnan², Andrea Agazzi³, W.P.M.H. (Maurice) Heemels¹

Abstract—When people choose routes minimizing their individual delay, the aggregate congestion can be much higher compared to that experienced by a centrally-imposed routing. Yet centralized routing is incompatible with the presence of self-interested agents. How can we reconcile the two? In this paper we address this question within a *repeated game* framework and propose a fair incentive mechanism based on artificial currencies that routes selfish agents in a system-optimal fashion, while accounting for their temporal preferences. We instantiate the framework in a parallel-network whereby agents commute repeatedly (e.g., daily) from a common start node to the end node. Thereafter, we focus on the specific two-arcs case whereby, based on an artificial currency, the agents are charged when traveling on the first, fast arc, whilst they are rewarded when traveling on the second, slower arc. We assume the agents to be rational and model their choices through a game where each agent aims at minimizing a combination of today’s discomfort, weighted by their urgency, and the average discomfort encountered for the rest of the period (e.g., a week). We show that, if prices of artificial currencies are judiciously chosen, the routing pattern converges to a system-optimal solution, while accommodating the agents’ urgency. We complement our study through numerical simulations. Our results show that it is possible to achieve a system-optimal solution whilst reducing the agents’ perceived discomfort by 14-20% when compared to a centralized optimal but urgency-unaware policy.

I. INTRODUCTION

MOBILITY systems are currently facing significant challenges due to users’ dissatisfaction, traffic congestion and environmental pollution [1], [2]. At the same time, the advent of driving technologies, the internet of things, the concept of sharing economies and new automotive technologies is leading to structural transformations in the way we conceive mobility, providing us with unprecedented opportunities to handle the aforementioned challenges. For instance, intermodal autonomous mobility-on-demand systems—fleets of robotaxis servicing travel demand jointly with public transit—are a promising solution for urban scenarios, as they combine the high-efficiency of public transportation systems with the point-to-point mobility service provided by fleets of connected autonomous vehicles. Whilst the possibility of routing customers through socially-optimal intermodal routes can significantly improve transportation efficiency [3], this requires users to sacrifice their individual welfare for the




Fig. 1. Network with one origin and one destination node connected by two arcs. Each arc j is traversed by x_j people per unit time, resulting in a societal cost $c_j(x_j)x_j$ and a discomfort cost $d_j(x_j)$ perceived by the agents. Each arc is assigned a price p_j . In this particular case, the agents pay $p_1 > 0$ to take the first, fast route, and receive $r_2 = -p_2 > 0$ when traveling the second, slower route.

“greater good” [4]. The central issue revolves around the inherent tension between the drivers’ individual objective (e.g., minimizing the travel-time from A to B) and the societal goal (e.g., minimizing the city-wide congestion). In this respect, it is well known that traffic patterns arising from self-interested decision making are often inefficient [5], and this is, to a high degree, what we experience in every densely populated city.

A promising solution to alleviate these issues is to employ monetary tolls [6]–[9]. However, approaches based on congestion pricing are associated with two fundamental drawbacks: i) they discriminate users with lower incomes; ii) they do not account for individual preferences such as the users’ temporary urgency and needs. Against this backdrop, this paper presents an incentive scheme based on an artificial currency—here called *Karma*, borrowing the terminology from [10]—to align the routing of self-interested users with the system-optimum allocation, whilst accounting for their temporal individual needs. Specifically, our framework is based on a currency that can neither be bought nor exchanged, but only spent or gained when traveling. On the microscopic level, it gives each user equal possibilities to choose when to be self-interested—selecting the fastest path for a toll (e.g., the upper one in Fig. 1)—and when to be altruistic—traveling the slower path for a reward (e.g., the lower one in Fig. 1)—whilst, on the mesoscopic level, it aligns the average population behavior to the system optimum.

Related work: The interest in the design of tolls dates back to the work of Pigou [6]. Since then, a large body of literature

¹Control Systems Technology group, Department of Mechanical Engineering, Eindhoven University of Technology, The Netherlands, e-mail: {m.r.u.salazar,w.p.m.h.heemels}@tue.nl.

²Computational Optimization group, Department of Computing, Imperial College London, United Kingdom, e-mail: d.paccagnan@imperial.ac.uk.

³Department of Mathematics, Duke University, United States, e-mail: andrea.agazzi@duke.edu.

pertaining to the fields of transportation and economics has developed this approach [6]–[8]. When considering the continuous flow approximation of the classical congestion game model [11] (typically employed to describe mesoscopic traffic patterns), the marginal cost mechanism [12] produces local tolls ensuring that every equilibrium allocation coincides with the system optimum. The problem of optimal tolling is much more challenging when considering the original congestion game model. In this context, the recent result in [13] derives optimal local tolls that minimize the system inefficiency and can be considered the counterpart of marginal cost tolls for discrete congestion games. Designing tolling mechanisms that account for the agents’ sensitivity distribution is particularly challenging even if considering congestion games with a continuum of agents, unless the exact sensitivity of each agent is fully known to the designer. In this respect, the results are limited to providing efficiency bounds for existing tolling mechanisms [14], to deriving optimal tolls when the sensitivity distribution is a piecewise constant function, or to proving the existence of optimal tolls for general distributions [15]. Our work significantly departs from all these studies for two fundamental reasons: First, we utilize an artificial currency that can neither be bought nor exchanged; second, we account for the agents’ sensitivity in a repeated game scenario with the objective of achieving the system optimum in the long run.

Whilst the use of artificial currency as a means to coordinate self-interested decision makers has recently attracted significant attention, see e.g., [10], [16], [17], most of the existing works are based on auction mechanisms and have not been applied to managing traffic routing. Perhaps closest in spirit to our work is [10], whereby the authors introduce an artificial currency called Karma to coordinate the behavior of competitive decision makers and allocate shared resources (e.g., intersection coordination). Their central idea consists in allowing for agents to pass on using resources today, in exchange for Karma, which will allow them to claim the resource in the future. While our work is based on a similar philosophy, it departs significantly from the former in that we envision its application to mobility systems (e.g., traffic routing). As such and in contrast to [10], each resource is associated to a measure of its quality (e.g., travel time), and our model accounts for the presence of more than two decision makers. Moreover, agents are not required to submit Karma bids, whilst their routing is controlled through simple payment transactions.

Statement of contributions: This paper introduces a repeated game framework and leverages an artificial currency to route agents in a system-optimal fashion while accounting for their temporal preferences. In particular, we focus on repetitive events, e.g., daily commutes, whereby traveling agents choose between two possible routes, as shown in Fig. 1, depending on their daily sensitivity (e.g., urgency) to the discomfort incurred when traveling (e.g., travel time). Considering a central operator who wants to align the mesoscopic agents’ behavior with the societal optimum, we first instantiate a simple pricing framework that is not relying

on complex auction mechanisms, but gives the agents full freedom of choice. Second, we propose a decision-model for rational agents for which we compute the best response strategy in closed form, demonstrating it to exhibit a clear structure. Third, we show that the only possible static pricing policy resulting from Karma-conservation arguments makes the system optimum the only and globally asymptotically stable equilibrium of the system, in the sense that the population behavior will converge to it for any initial Karma allocation. It is important to highlight how our mechanism does not require any information on the underlying sensitivity distribution, i.e., it is distribution-free. Finally, we simulate our scheme in different case studies, validating our theoretical findings numerically.

Organization: The remainder of this paper is structured as follows: We instantiate our framework in Section II, where we introduce the repeated game problem on a parallel arcs network for a central operator and for an individual agent. From Section III onward we focus on a two-arcs parallel network, by first deriving the optimal static prices via necessary Karma-conservation conditions. Section IV computes the best response strategy for the agent analytically, which is used in Section V to study the average aggregate behavior of a population of agents playing the best response strategy, showing that it converges to the desired optimal flows. We validate our findings with numerical simulations for different case studies in Section VI, and conclude the paper in Section VII with a discussion and an outlook on future research.

II. PROBLEM FORMULATION

This section introduces the routing problem from both the mesoscopic perspective of a central operator who wants to minimize the societal costs, and from the individual perspective of an agent aiming at minimizing the discomfort incurred when traveling under the proposed scheme.

Our framework combines three ingredients: i) a transportation network represented by a digraph, ii) cost functions representing the societal and personal costs (a measure of discomfort such as travel time) resulting from the aggregate route choices of the agents, and iii) a static pricing policy based on the artificial currency Karma.

A. Central Operator’s Problem

Consider a parallel road-network, consisting of a single origin and destination node but connected by $n \in \{1, 2, \dots\}$ directed arcs, e.g., the example in Fig. 1, whereby $n = 2$. Such a digraph can be used to describe a daily commute, but our framework lends itself to model more general resource-allocation problems whereby the quality of each resource depends on its usage-level.

From a mesoscopic perspective, the agents want to reach the destination from the origin traveling through one of the n arcs in the digraph at each (discrete) time $t \in \mathbb{N}$. Thereby, the component $x_j(t)$ of vector $x(t) \in [0, 1]^n$ represents the fraction of users crossing arc $j \in \{1, 2, \dots, n\}$ at time t . Every day, each agent has a probability $P_{\text{home}} \in [0, 1]$ not

to travel. This way, the probability for an agent to travel is $P_{\text{go}} = 1 - P_{\text{home}}$ and $\sum_{j=1}^n x_j(t) = P_{\text{go}}$ people will travel per day. Crossing each arc j entails a specific societal cost per person $c_j(x_j(t))$ dependent on the number of people traversing it $x_j(t)$. Moreover, it also causes each person a discomfort $d_j(x_j(t))$ (e.g., travel time). We assume $c : [0, 1]^n \rightarrow \mathbb{R}_+^n$ and $d : [0, 1]^n \rightarrow \mathbb{R}_+^n$ to be monotonically increasing functions for each arc j , where $\mathbb{R}_+ = [0, \infty)$ is the set of real non-negative numbers. In general, the societal cost $c(x)$ and the discomfort $d(x)$ need not be aligned—for instance, the former might represent energy consumption and the latter travel time.

We assume the presence of a central operator (e.g., a municipal authority) who needs to design incentives so that the aggregate flows converge to the minimizer of the total societal cost $c(x)^\top x$.

Problem 1 (Central Operator’s Problem). *The central operator aims at routing customers so the aggregate route choice $x(t)$ converges to*

$$x^* \in \arg \min_{x \in [0, 1]^n} c(x)^\top x \quad (1a)$$

$$\text{s.t. } \mathbb{1}^\top x = P_{\text{go}}. \quad (1b)$$

In order to steer peoples’ behavior towards a social optimum x^* , the central operator endows each agent with the artificial currency Karma and sets a price $p_j \in \mathbb{R}$ to cross each arc j . Agents are not permitted to buy or exchange Karma, and they can only select the arcs keeping their Karma-level non-negative. Crucially, some of the arcs have negative prices, so that agents will be rewarded when crossing them.

From a microscopic perspective, we model the routing choice for an individual agent i at time t as $y^i(t) \in \{0, 1\}^n$ being a vector with $y_j^i(t) = 1$ if person i decides to cross arc j at t , and 0 otherwise. Thereby, a non-traveling agent is modeled as $y^i(t) = \mathbb{0}$. Given a scenario with M agents, their individual decisions are linked to the mesoscopic flows as $x(t) = \frac{1}{M} \sum_{i=1}^M y^i(t)$. In particular, assuming that at time t each agent i owns an amount $k^i(t) \in \mathbb{R}_+$ of Karma, given a routing choice $y^i(t)$, her Karma will be updated as $k^i(t+1) = k^i(t) - p^\top y^i(t)$.

B. Individual Agent’s Problem

On the microscopic level, we assume individual agents to make choices in order to minimize their traveling discomfort without reaching a negative level of Karma. In contrast to conventional monetary tolling schemes, the individual agent’s problem cannot be captured in a static setting: In our framework, the agents are playing against their future selves, and must account for their future preferences to decide when to use Karma and when to gain it. Specifically, on some days they might be more sensitive to discomfort than on other days—for instance, when going to an important meeting, travel time may be more important than when simply commuting to the working place. To this end, we define the sensitivity that a person i might have at time t w.r.t. discomfort as $s^i(t) \in \mathbb{R}_+$ and use it as a weighting factor. This way, the cost perceived by agent i when crossing arc j is $s^i(t) \cdot d_j(x_j(t))$. Hereby, we assume $s^i(t)$

to be i.i.d. extractions (w.r.t. i and t) from an underlying common probability distribution with probability density function $\rho : [s_{\min}, s_{\max}] \rightarrow \mathbb{R}_+$, support set $[s_{\min}, s_{\max}] \subseteq \mathbb{R}_+$ and average value \bar{s} . For the sake of simplicity, from now on we drop dependence on t , x and i whenever it is clear from the context.

We assume that each agent will choose the route associated with the least discomfort perceived on the respective day of traveling and the one expected in the T days to follow, whilst accounting for her Karma allowance. Specifically, we model the decision of each agent in the present moment as $y^i \in \{0, 1\}^n$, whilst we condense her future behavior into the fractional variable $\bar{y}^i \in [0, 1]^n$. We assume the agents to consider today’s discomfort on the arcs $d(x(t))$ to be constant for their whole planning horizon (e.g., a week). Thereby, the agents minimize the discomfort perceived today with a sensitivity $s^i(t)$ and the discomfort perceived in the remaining T steps on the horizon with an average sensitivity \bar{s} . Finally, we assume agents to be conservative in terms of Karma: Given an initial Karma level $k^i(t)$, they will constrain their choices so that their Karma at the end of the horizon will not fall below a non-negative reference value k_{ref}^i . We explicitly account for path constraints (positive Karma at every time step) on the agent’s Karma level only for the current decision, whilst assuming that in the remaining T time-steps captured by the average behavior \bar{y}^i they will be satisfied. Formally, we get the following individual agent’s problem:

Problem 2 (Individual Agent’s Problem). *At time t , given the flows x and prices p , respectively, a traveling agent with Karma level $k \geq 0$ and reference k_{ref} , and sensitivity s will choose her route as y^* resulting from*

$$(y^*, \bar{y}^*) \in \arg \min_{y \in \mathcal{Y}, \bar{y} \in \bar{\mathcal{Y}}} s \cdot d(x)^\top y + T \cdot \bar{s} \cdot d(x)^\top \bar{y} \quad (2a)$$

$$\text{s.t. } k - p^\top y - T \cdot p^\top \bar{y} \geq k_{\text{ref}} \quad (2b)$$

$$p^\top y \leq k, \quad (2c)$$

with $\mathcal{Y} = \{y \in \{0, 1\}^n : \mathbb{1}^\top y = 1\}$ and $\bar{\mathcal{Y}} = \{\bar{y} \in [0, 1]^n : \mathbb{1}^\top \bar{y} = 1\}$. We define the set containing all points y^* solving (2) as $\mathcal{Y}^*(x, s, k, k_{\text{ref}}) \subseteq \mathcal{Y}$. Non-traveling agents have $y^* = \mathbb{0}$.

Note that a discount factor can be readily included by simply scaling the average sensitivity \bar{s} in the objective.

C. Infinitely Many Agents Setting and Wardrop Equilibrium

In this paper, we consider the limit case whereby agents form a continuum with $M \rightarrow \infty$. To describe the population, we use a notation similar to [17]. Thereby, we describe the distribution of the Karma level k and reference k_{ref} in the population with the density function $\eta : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$. This way, we can describe an infinite-agents population with $\rho(s)$ and $\eta(k, k_{\text{ref}})$. For the infinite-agents setting, the Nash and Wardrop Equilibrium (WE) are identical [18] and can be defined as follows:

Definition II.1 (Wardrop Equilibrium). *$x^{\text{WE}} \in [0, 1]^n$ satisfying $\mathbb{1}^\top x^{\text{WE}} = P_{\text{go}}$ is a WE if and only if there exist*

$y^*(x^{\text{WE}}, s, k, k_{\text{ref}}) \in \mathcal{Y}^*(x^{\text{WE}}, s, k, k_{\text{ref}})$ so that

$$\int_{s_{\min}}^{s_{\max}} \int_0^{\infty} \int_0^{\infty} y^*(x^{\text{WE}}, s, k, k_{\text{ref}}) \cdot \rho(s) \cdot \eta(k, k_{\text{ref}}) ds dk dk_{\text{ref}} = x^{\text{WE}}.$$

Informally, x^{WE} is a WE if the aggregate best-response based on assuming $x = x^{\text{WE}}$ reconstructs the same vector x^{WE} . Observe that the agents' decision process can be interpreted as a model predictive control algorithm that implements solely the optimal decision for today y^* , whilst discarding the optimal future decisions \bar{y}^* .

Taking a static planning perspective, at each time-step t , we model the aggregate choices of the population with the WE $x^{\text{WE}}(t)$. The central operator's problem then is to choose the prices so that the daily WE $x^{\text{WE}}(t)$ will converge to the system optimum x^* .

Problem 3 (Pricing Problem). *Given a desired system optimum x^* , the pricing problem consists of finding $p \in \mathbb{R}^n$ so that $\lim_{t \rightarrow \infty} x^{\text{WE}}(t) = x^*$.*

D. Discussion

A few remarks are in order. First, we use a static setting whereby on each day we compute the WE to model the users' behavior. This is in line with the mesoscopic perspective of our study. Second, we assume the agents to be rational and to share the discomfort function $d(x)$ and sensitivity probability density function $\rho(s)$, but differentiate them with regard to their daily level of sensitivity $s^i(t)$ and their reference Karma-level k_{ref}^i .

III. TWO-ARCS PRICING PROBLEM

For the remainder of this paper, we will focus on the parallel network with two arcs shown in Fig. 1. This model can represent a daily commute routine where users can choose between two options, a fast (or comfortable) route, and one that is slower (or less comfortable). We leave the study of parallel networks with more arcs and more general transportation networks to future research endeavors. We assume, without loss of generality, that at the desired socially-optimal solution x^* it holds that $d_1(x_1^*) < d_2(x_2^*)$.¹ Therefore, we introduce a price $p_1 > 0$ for arc 1 and a reward $r_2 = -p_2 > 0$ for arc 2. Crucially, given a desired solution x^* , a necessary condition required to solve Problem 3, is that the total Karma must be conserved in the steady state. Hence, the steady-state prices must satisfy $p^\top x^* = 0$, implying that

$$p_1 = r_2 \cdot \frac{x_2^*}{x_1^*}. \quad (3)$$

Since scaling the prices' (p_1, p_2) with a common factor would not change the set of equilibria, (3) automatically delivers the only possible prices that could solve Problem 3. Therefore, in the two-arcs setting under consideration, there is no control freedom when choosing the steady-state prices. However, we will show in the remainder of this paper that choosing the prices as in (3) will make the aggregate behavior of the agents described by Problem 2 converge to the desired system optimum resulting from Problem 1.

¹If it were to hold $d_2(x_2^*) < d_1(x_1^*)$, given the monotonicity properties of $d(x)$, we could simply swap the arcs' labeling and recover this problem.




Fig. 2. The best response strategy (4) resulting from the solution of Problem 2 for $d_1(x) < d_2(x)$.

IV. BEST RESPONSE STRATEGY

This section derives the best response strategy for an agent with Karma-level k , reference level k_{ref} and daily sensitivity s . We begin with a technical assumption ensuring feasibility of the problem at hand.

Assumption 1 (Feasible Desired Flow). *The desired mesoscopic flow satisfies $x^* > 0$ and $x_1^*/x_2^* \in [1/T, T]$, so that $r_2/p_1 \in [1/T, T]$.*

This assumption is needed because, if the ratio between price p_1 and reward r_2 was too small or too large, there would not always exist a $\bar{y} \in \bar{\mathcal{Y}}$ so that $p^\top (y + T \cdot \bar{y})$ is zero for all $y \in \mathcal{Y}$ and, therefore, a *Karma-sustained* planning would be infeasible. However, it is reasonable to consider a long-enough horizon T , so that the latter requirement would not be limiting.

The solution of Problem 2, namely the agents' best response strategy, can be computed in closed-form as shown in the following theorem.

Theorem IV.1 (Best Response Strategy). *We define $k_{\text{inf}} = \max\{0, k_{\text{ref}} - (T + 1) \cdot r_2\}$, $k_{\text{poor}} = \max\{p_1, k_{\text{ref}} + p_1 - T \cdot r_2\}$, $k_{\text{rich}} = k_{\text{ref}} + T \cdot p_1 - r_2$, and $k_{\text{wealthy}} = k_{\text{ref}} + (T + 1) \cdot p_1$. Consider a player with Karma $k \geq k_{\text{inf}}$, sensitivity s , Karma reference k_{ref} , and given prices p . Given x so that for the discomfort it holds $d_1(x_1) < d_2(x_2)$, the best response strategy of a traveling agent in terms of which arc j to choose, so that $y_{j^*}^* = 1$ and 0 otherwise, is*

$$j^* = \begin{cases} 2 & \text{if } k \in [k_{\text{inf}}, k_{\text{poor}}) \\ \begin{cases} 1 & \text{if } s > \bar{s} \\ 2 & \text{if } s < \bar{s} \end{cases} & \text{if } k \in [k_{\text{poor}}, k_{\text{rich}}) \\ \begin{cases} 1 & \text{if } s > \bar{s} \cdot \frac{k_{\text{wealthy}} - k}{p_1 + r_2} \\ 2 & \text{if } s < \bar{s} \cdot \frac{k_{\text{wealthy}} - k}{p_1 + r_2} \end{cases} & \text{if } k \in [k_{\text{rich}}, k_{\text{wealthy}}) \\ 1 & \text{if } k \geq k_{\text{wealthy}}. \end{cases} \quad (4)$$

Moreover, given x so that $d_1(x_1) = d_2(x_2)$, the best response strategy of a traveling agent is

$$j^* \begin{cases} = 2 & \text{if } k \in [k_{\text{inf}}, k_{\text{poor}}) \\ \in \{1, 2\} & \text{if } k \geq k_{\text{poor}}. \end{cases} \quad (5)$$

Finally, given x so that $d_1(x_1) > d_2(x_2)$, the best response strategy of a traveling agent is

$$j^* = 2 \quad \forall k \geq k_{\text{inf}}. \quad (6)$$

The proof can be found in the extended version of this paper [19]. Fig. 2 depicts the best response strategy (4). As shown in the proof, Problem 2 is infeasible for $k < k_{\text{inf}}$. Therefore, we will always consider the case whereby the Karma of any agent is initialized above k_{inf} . Moreover, under policy (4) the set $[k_{\text{inf}}, k_{\text{wealthy}} + r_2)$ is positively invariant and attractive from above, so that it will be $k \geq k_{\text{inf}}$ for all future times. Finally, we observe that the best response strategies (4)–(6) are independent of the quantitative discomfort level; they depend only on the prices, the current level of Karma k , and the sensitivity s .

A. Wardrop Equilibrium

One can verify that, when prices are chosen according to (3), there exists no WE satisfying $d_1(x_1) > d_2(x_2)$. This is due to the fact that, when $d_1(x_1) > d_2(x_2)$, the agent's best response would entail choosing arc 2, thus leading to $d_1(0) < d_2(P_{\text{go}})$. Finally, we define \bar{x} so that $d_1(\bar{x}_1) = d_2(\bar{x}_2)$, whereby, given the monotonicity properties of $d(x)$, we have that $d_1(x_1) < d_2(x_2)$ for all $x_1 < \bar{x}_1$. In the following lemma, we show that a WE always exists, and is characterized by $d_1(x_1^{\text{WE}}) \leq d_2(x_2^{\text{WE}})$.

Lemma IV.2 (Wardrop Equilibrium for Two Arcs). *Let $\bar{x}_1 < \bar{x}_1$, i.e., $d_1(\bar{x}_1) < d_2(\bar{x}_2)$, and define the mesoscopic flow resulting from the best response (4) $y^*(\bar{x}, s, k, k_{\text{ref}}) \in \mathcal{Y}^*(\bar{x}, s, k, k_{\text{ref}})$ as*

$$x^{\text{BR}} = \int_{s_{\text{min}}}^{s_{\text{max}}} \int_0^\infty \int_0^\infty y^*(\bar{x}, s, k, k_{\text{ref}}) \cdot \rho(s) \cdot \eta(k, k_{\text{ref}}) ds dk dk_{\text{ref}}. \quad (7)$$

A WE exists and is characterized by $x^{\text{WE}} = x^{\text{BR}}$ if $x_1^{\text{BR}} < \bar{x}_1$, and/or $x^{\text{WE}} = \bar{x}$.

The proof can be found in the extended version of this paper [19]. We define the WE at \bar{x} as *uncontrolled*, since it corresponds to the WE resulting when no pricing is applied. In fact, it exists when the population has too much Karma and can therefore easily afford the first arc.

For the remainder of this paper we will assume that for the discomfort functions it holds $d_1(x_1) < d_2(x_2)$ for all admissible x . This way, in addition to Lemma IV.2, we can prove the existence of a unique WE with $d_1(x_1^{\text{WE}}) < d_2(x_2^{\text{WE}})$ for any Karma distribution $\eta(k, k_{\text{ref}})$, directly stemming from the best response strategy (4). This assumption enables the formal analysis of the mesoscopic aggregate behavior presented in Section V below. However, we believe this assumption can be lifted, since by the prices definition (3)—i.e., $p_1 x_1^* - r_2 x_2^* = 0$ —and the fact that $x_1^* < \bar{x}_1$, it holds that $p^\top \bar{x} > 0$. Therefore, the uncontrolled WE at \bar{x} is not sustainable in time, as the population would continuously lose Karma, until eventually reaching a distribution with WE $d_1(x_1^{\text{WE}}) < d_2(x_2^{\text{WE}})$ converging to x^* . We corroborate our conjecture with numerical simulations in Section VI, whilst leaving a formal mathematical proof to an extended version of this paper.




Fig. 3. The steady-state distribution in relative Karma levels for an exponentially distributed s .

V. MESOSCOPIC AVERAGE BEHAVIOR

This section studies the dynamics of the aggregate behavior of the population. In the case under consideration, the WE always results from the individual agent's best response strategy (4), i.e., it corresponds to x^{BR} from (7) in Lemma IV.2. Since the positively invariant set of (4) $k \in [k_{\text{inf}}, k_{\text{wealthy}} + r_2)$ is attractive, we assume each agent to start with a level of Karma already in it. We first define the fraction of agents with a certain level of Karma:

$$\begin{aligned} P_{\text{poor}} &= P(k \in [k_{\text{inf}}, k_{\text{poor}})) \\ P_{\text{ok}} &= P(k \in [k_{\text{poor}}, k_{\text{rich}})) \\ P_{\text{rich}} &= P(k \in [k_{\text{rich}}, k_{\text{wealthy}})) \\ P_{\text{wealthy}} &= P(k \in [k_{\text{wealthy}}, k_{\text{wealthy}} + r_2)) \end{aligned} \quad (8)$$

Next, given the probability density function of the sensitivity $\rho(s)$ and $\hat{s} \in [s_{\text{min}}, s_{\text{max}}]$, we define

$$\begin{aligned} P_{\text{chill}}(\hat{s}) &= P(s < \hat{s}) = \int_{s_{\text{min}}}^{\hat{s}} \rho(s) ds \\ P_{\text{rush}}(\hat{s}) &= P(s > \hat{s}) = 1 - P_{\text{chill}}(\hat{s}) = \int_{\hat{s}}^{s_{\text{max}}} \rho(s) ds. \end{aligned} \quad (9)$$

Overall, these six measures represent the fraction of agents present in one of the regions shown in Fig. 2, and can be used to compute the average population response. For the sake of simplicity, but without loss of generality, from now on we assume $r_2 = -p_2 \geq p_1 > 0$, and $k_{\text{ref}} + p_1 - T \cdot r_2 \geq p_1$, i.e., $k_{\text{poor}} = k_{\text{ref}} + p_1 - T \cdot r_2$. The cases with $k_{\text{poor}} = p_1$ and/or $r_2 \leq p_1$ can be studied with identical arguments and lead to the same result. Furthermore, we assume p_1 and r_2 to be commensurate, and scale them to be integer and co-prime (in practice, it is possible to round the prices after a suitable scaling without appreciably affecting the results, as shown in Section VI below). We define $i = k - k_{\text{ref}} + T \cdot r_2 + 1$, quantize the probability distributions (8) as

$$\begin{aligned} P_{\text{poor}} &= \sum_{i=1}^{p_1} \mathbf{P}_{\text{poor}}(i), & P_{\text{ok}} &= \sum_{i=1}^{(T-1) \cdot (p_1 + r_2)} \mathbf{P}_{\text{ok}}(i), \\ P_{\text{rich}} &= \sum_{i=1}^{p_1 + r_2} \mathbf{P}_{\text{rich}}(i), & P_{\text{wealthy}} &= \sum_{i=1}^{r_2} \mathbf{P}_{\text{wealthy}}(i), \end{aligned} \quad (10)$$

and stack them in the vector $\mathbf{P} = (\mathbf{P}_{\text{poor}}^\top, \mathbf{P}_{\text{ok}}^\top, \mathbf{P}_{\text{rich}}^\top, \mathbf{P}_{\text{wealthy}}^\top)^\top$, as shown in Fig. 3. The probability vector \mathbf{P} is the quantized version of $\eta(k, k_{\text{ref}})$, and evolves in line with the best response strategy (4) as detailed in (11) on page 6. In matrix form, it can be written as

$$\mathbf{P}^+ = \mathbf{A}\mathbf{P}, \quad (13)$$

$$\begin{aligned}
\mathbf{P}(i)^+ &= P_{\text{home}} \cdot \mathbf{P}(i) + P_{\text{go}} \cdot P_{\text{rush}}(\bar{s}) \cdot \mathbf{P}(i + p_1) & \forall i \in [1, r_2] \\
\mathbf{P}(i)^+ &= P_{\text{home}} \cdot \mathbf{P}(i) + P_{\text{go}} \cdot (P_{\text{chill}}(\bar{s}) \cdot \mathbf{P}(i - r_2) + P_{\text{rush}}(\bar{s}) \cdot \mathbf{P}(i + p_1)) & \forall i \in [r_2 + 1, r_2 + p_1] \\
\mathbf{P}(i)^+ &= P_{\text{home}} \cdot \mathbf{P}(i) + P_{\text{go}} \cdot (P_{\text{chill}}(\bar{s}) \cdot \mathbf{P}(i - r_2) + P_{\text{rush}}(\bar{s}) \cdot \mathbf{P}(i + p_1)) & \forall i \in [r_2 + p_1 + 1, (T - 1) \cdot (p_1 + r_2)] \\
\mathbf{P}(i)^+ &= P_{\text{home}} \cdot \mathbf{P}(i) + P_{\text{go}} \cdot (P_{\text{chill}}(\bar{s}) \cdot \mathbf{P}(i - r_2) + P_{\text{rush}}(\bar{s}) \cdot \frac{T \cdot (p_1 + r_2) + 1 - i}{p_1 + r_2} \cdot \mathbf{P}(i + p_1)) & \forall i \in [(T - 1) \cdot (p_1 + r_2) + 1, T \cdot (p_1 + r_2)] \\
\mathbf{P}(i)^+ &= P_{\text{home}} \cdot \mathbf{P}(i) + P_{\text{go}} \cdot (P_{\text{chill}}(\bar{s}) \cdot \frac{(T + 1) \cdot (p_1 + r_2) + 1 - i}{p_1 + r_2} \cdot \mathbf{P}(i - r_2) + \mathbf{P}(i + p_1)) & \forall i \in [T \cdot (p_1 + r_2) + 1, T \cdot p_1 + (T + 1) \cdot r_2] \\
\mathbf{P}(i)^+ &= P_{\text{home}} \cdot \mathbf{P}(i) + P_{\text{go}} \cdot P_{\text{chill}}(\bar{s}) \cdot \frac{(T + 1) \cdot (p_1 + r_2) + 1 - i}{p_1 + r_2} \cdot \mathbf{P}(i - r_2) & \forall i \in [T \cdot p_1 + (T + 1) \cdot r_2 + 1, (T + 1) \cdot (p_1 + r_2)]
\end{aligned} \tag{11}$$

$$A_{\text{chill}}(i, j) = \begin{cases} 1 & \text{if } i \in [r_2 + 1, r_2 + p_1] \wedge j = i - r_2 \\ P_{\text{chill}}(\bar{s}) & \text{if } i \in [p_1 + r_2 + 1, T \cdot (p_1 + r_2)] \wedge j = i - r_2 \\ P_{\text{chill}}(\bar{s}) \cdot \frac{(T + 1) \cdot (p_1 + r_2) + 1 - i}{p_1 + r_2} & \text{if } i \in [T \cdot (p_1 + r_2) + 1, (T + 1) \cdot (p_1 + r_2)] \wedge j = i - r_2 \\ 0 & \text{otherwise} \end{cases} \tag{12a}$$

$$A_{\text{rush}}(i, j) = \begin{cases} P_{\text{rush}}(\bar{s}) & \text{if } i \in [1, (T - 1) \cdot (p_1 + r_2)] \wedge j = i + p_1 \\ P_{\text{rush}}(\bar{s}) \cdot \frac{T \cdot (p_1 + r_2) + 1 - i}{p_1 + r_2} & \text{if } i \in [(T - 1) \cdot (p_1 + r_2) + 1, T \cdot (p_1 + r_2)] \wedge j = i + p_1 \\ 1 & \text{if } i \in [T \cdot (p_1 + r_2) + 1, T \cdot p_1 + (T + 1) \cdot r_2] \wedge j = i + p_1 \\ 0 & \text{otherwise} \end{cases} \tag{12b}$$




Fig. 4. Sparsity pattern of the matrix A_{go} . The probability of choosing the first arc paying p_1 is denoted in blue, whilst the probability of choosing the second arc receiving r_2 in red. In this example, we chose $p_1 = 2$, $r_2 = 3$ and $T = 3$. Full dots represent 1-entries, whilst empty dots denote $P_{\text{rush}}(\bar{s})$ (blue) and $P_{\text{chill}}(\bar{s})$ (red). Stars represent $P_{\text{rush}}(\cdot)$ (blue) and $P_{\text{chill}}(\cdot)$ (red) evaluated as in the third and second condition of (12a) and (12b), respectively. Since $P_{\text{rush}} + P_{\text{chill}} = 1$, column-stochasticity of A_{go} can be determined by inspection.

where $A = P_{\text{home}} \cdot I + P_{\text{go}} \cdot A_{\text{go}}$ is a non-negative, square and column-stochastic matrix, i.e., $A^\top \mathbf{1} = \mathbf{1}$. Furthermore, A is primitive if $P_{\text{home}} > 0$. Since A is non-negative, for any $\mathbf{P} \geq \mathbf{0}$, it holds that $\mathbf{P}^+ \geq \mathbf{0}$. Moreover, being A column-stochastic, we have that $\mathbf{1}^\top \mathbf{P}^+ = \mathbf{1}^\top \mathbf{P} = \mathbf{1}^\top \mathbf{P}$ [20]. The sparsity pattern of $A_{\text{go}} = A_{\text{chill}} + A_{\text{rush}}$ is shown in Fig. 4, whereby the red dots represent the probability for a traveling agent with Karma deviation $k = i + k_{\text{ref}} + T \cdot r_2 - 1$ to choose the slow arc and are defined as in (12a), whilst the blue dots represent the probability to choose the fast arc and are defined as in (12b). The system has discrete linear time-invariant dynamics. Its ω -limit sets are either fixed points or limit cycles. Such ω -limit sets can be found by eigenvalue

analysis as we do below, showing that limit cycles can only appear if $P_{\text{home}} = 0$. Since $\mathbf{P} \geq \mathbf{0}$ is a probability distribution, it must hold $\mathbf{1}^\top \mathbf{P} = 1$. Hence, the trivial equilibrium $\mathbf{P} = \mathbf{0}$ is not admissible. As the equilibrium is defined as $\mathbf{P}^e = A\mathbf{P}^e$, it must be spanned by the eigenvector related to the 1-eigenvalue of A .

With this in mind, we can now show that in the large population limit, the equilibrium in Karma distribution \mathbf{P}^e is globally asymptotically stable and corresponds to the desired system optimum x^* .

Theorem V.1 (Globally Asymptotically Stable and Optimal Equilibrium). *Given the prices (3) and a population of agents acting in line with (4) with $P_{\text{home}} > 0$, the equilibrium of the Karma distribution dynamics (13) $\mathbf{P}^e \geq \mathbf{0}$ with $\mathbf{1}^\top \mathbf{P}^e = 1$ is globally asymptotically stable and its resulting flows correspond to the system optimum x^* .*

The proof can be found in the extended version of this paper [19]. This theorem shows that the prices resulting from Karma conservation arguments indeed solve Problem 3. Note that equivalent results can be obtained by leveraging the geometric ergodicity properties of the Markov chain \mathbf{P} through straightforward application of Doeblin's theorem [21].

VI. NUMERICAL RESULTS

This section presents the results obtained via numerical simulations for the following case-studies: First, a scenario where the social cost corresponds to the sum of the agents' cost, i.e., $d(x)^\top x$ as in standard routing settings [9]. There we consider the case whereby on average $P_{\text{home}} = 5\%$ of the population does not travel and the limit case where all agents travel every day, i.e., $P_{\text{home}} = 0$. Second, we study the case where the societal cost does not correspond to the sum of the users' costs. For all the scenarios we consider $M = 1000$ agents with a horizon of a week with $T = 6$, and sample their daily sensitivity from an exponential distribution

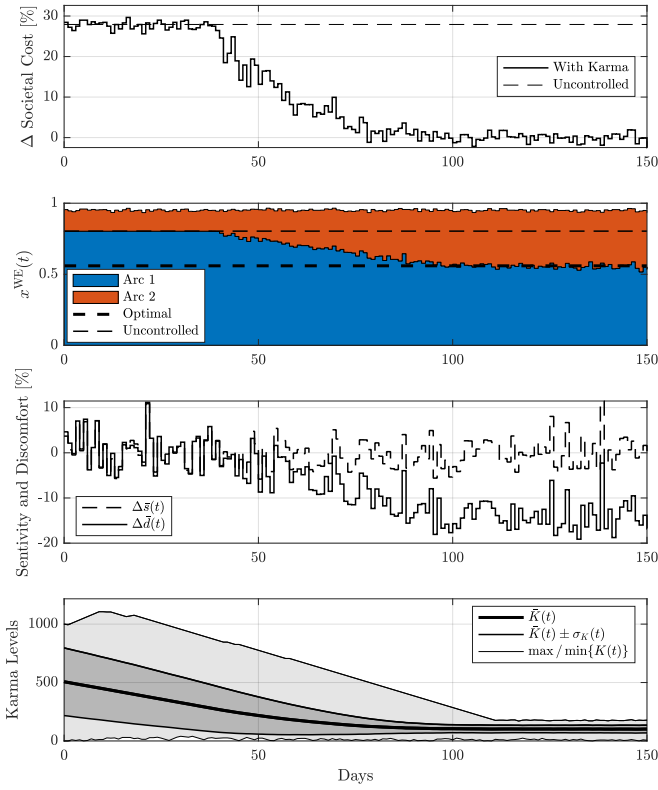


Fig. 5. Societal cost, mesoscopic flows, relative average sensitivity and discomfort deviation, and Karma distribution for a population with $P_{\text{home}} = 5\%$ and very large initial Karma levels.

defined on \mathbb{R}_+ . In line with our theoretical findings, the case where the discomfort functions satisfy $d_1(x_1) < d_2(x_2)$ for any admissible x resulted in the expected convergence of the aggregate behavior to the system optimum for any of the scenarios mentioned. For the sake of brevity, we omit such results and focus on the more interesting case whereby $d_1(x_1) > d_2(x_2)$ for $x_1 > \bar{x}_1$, and the uncontrolled WE at \bar{x} exists. Specifically, we model the discomfort as a travel-time Bureau of Public Roads (BPR) function [22]

$$d(x_j) = d_{0,j} \cdot \left(1 + \alpha \cdot (x_j/\kappa_j)^\beta\right), \quad (14)$$

with $d_0 = (1, 2)^\top$, $\kappa = (1/2, 2/3)^\top$, $\alpha = 0.15$ and $\beta = 4$, for which $\bar{x}_1 = 0.80$ for both values of P_{home} . We simulate each day by iteratively computing the Nash equilibrium (approximating the infinite-agents WE) resulting when each agent is solving Problem 2 (which can be efficiently solved as a linear program). The computation of the Nash equilibrium always needed only a couple of iterations.

A. The Social Cost is the Sum of the Agents' Cost

Setting $P_{\text{home}} = 5\%$, the system optimum is $x^* = (0.56, 0.39)^\top$, for which we design prices according to (3) as $p = (10, -14)^\top$, rounding them as mentioned in Section V above. We initialize the Karma reference levels k_{ref}^i from a uniform distribution defined between 0 and 100, whilst initializing the Karma levels $k^i(0)$ between 0 and 500, so that an extremely large fraction of agents starts with k above k_{wealthy} and thus \bar{x} is the only possible WE. As expected,




Fig. 6. Societal cost and mesoscopic flows for $P_{\text{home}} = 0$.

Fig. 5 shows that providing the agents with too much Karma will indeed result in the uncontrolled WE. However, as the average Karma level of the population $\bar{K}(t)$ shown in the fourth subplot is depleted, the system-level behavior and cost converge very close to the system optimum with an average relative societal cost difference below 0.1%. We measure the perceived discomfort of the single agents $s^i(t) \cdot d(x(t))^\top y^i(t)$ and average it over the population. The third subplot of Fig. 5 compares it with the average discomfort that would be perceived by the agents if they were allocated in a random and sensitivity-unaware fashion to the same flows as

$$\Delta \bar{d}(t) = \frac{\sum_{i=1}^{1000} s^i(t) \cdot d(x(t))^\top y^i(t) - \bar{s} \cdot d(x(t))^\top y^i(t)}{\sum_{i=1}^{1000} \bar{s} \cdot d(x(t))^\top y^i(t)}, \quad (15)$$

which exhibits a behavior that is similar to the relative deviation of the daily average sensitivity from the distribution's mean $\Delta \bar{s}(t) = \frac{\sum_{i=1}^{1000} s^i(t) - \bar{s}}{1000 \cdot \bar{s}}$. In line with the goal of our framework, while converging to the system optimum, the agents perceive a discomfort about 14% lower compared to the case where the agents would be randomly allocated in a system-optimal but urgency-unaware fashion.

Next, we study the limit case where all agents are traveling, i.e., $P_{\text{home}} = 0$. Fig. 6 shows the results obtained in this scenario with system-optimal solution $x^* = (0.57, 0.43)^\top$ for which we design again the prices via (3) and round them to $p = (10, -13)^\top$. In this case, we initialize the Karma initial and reference values between 0 and 100. Interestingly, the proposed scheme seems to work well also for this limit case that could admit periodic solutions—as the matrix A would no longer be primitive but only row-stochastic. Again, the population behavior converges very close to the system-optimum with a relative deviation of about 0.2% and an average $\Delta \bar{d}$ of about -14%. Since in this scenario the number of users traveling every day is constant, the optimal societal cost is never outperformed. Overall, these results prompt us to study in more detail also the convergence properties of this limit case with $P_{\text{home}} = 0$.

B. The Social Cost is not the Sum of the Agents' Cost

The proposed framework enables to steer the behavior of the population to any average choice satisfying Assumption 1. In this case, we set the societal cost as $c(x) = x$,




Fig. 7. Societal cost and mesoscopic flows for $c(x) = x$ and $P_{\text{home}} = 5\%$.

leading to the symmetric optimal flows $x^* = (0.5, 0.5)^\top$ for which the prices resulting from (3) can be chosen as $p = (10, -10)^\top$. Fig. 7 shows that, despite different societal and users' objectives and a very high price of anarchy for the uncontrolled case, the mesoscopic flows converge to the system optimum. Finally, also in this scenario, our framework enables agents to significantly improve the perceived discomfort by about 20% w.r.t. a system-optimal but random allocation, whilst aligning their behavior with the desired system optimum.

VII. CONCLUSION

This paper explored the application of artificial currencies to route self-interested agents in a system-optimal fashion whilst accounting for their temporal needs. Specifically, we instantiated a repeated game in a static setting whereby each day traveling agents choose whether to cross the most comfortable route for a price or to receive a reward for traveling a less comfortable route. For a two-parallel-arcs network we showed that a static pricing choice resulting from simple Karma-conservation arguments guarantees the mesoscopic average behavior to converge to the system optimum, significantly reducing the societal costs w.r.t. the uncontrolled equilibrium. What is more, our scheme is fairly simple and does not rely on any auction mechanism for allocation, but leaves full freedom of choice to the agents as long as their Karma level is non-negative. As a result, it enables a considerable reduction of the perceived discomfort with respect to a random but optimal allocation. In practice, our numerical results were in line with our findings in all the scenarios studied: Our scheme reached a societal cost less than 1% higher than the system optimum, while significantly reducing the average perceived discomfort by 14-20%.

This work can be extended as follows. We would like to study more realistic network structures, such as more general transportation graphs with different origin-destination pairs. Furthermore, we would like to devise learning-based control algorithms to adapt the prices in real-time and implement them within high-fidelity simulation environments.

ACKNOWLEDGMENTS

We thank Dr. Ilse New for proofreading this paper. Furthermore, we are grateful to Dr. Marco Pavone for the fruitful discussions and advice. This paper is dedicated to the loving memory of Dr. Kilian Schindler.

REFERENCES

- [1] B. Tuttle and T. Cowles, "Traffic jams cost americans \$124 billion in 2013," *Time - Money*, 2014.
- [2] EPA, "Inventory of us greenhouse gas emissions and sinks: 1990-2016," *United States Environmental Protection Agency*, 2018.
- [3] M. Salazar, N. Lanzetti, F. Rossi, M. Schiffer, and M. Pavone, "Intermodal autonomous mobility-on-demand," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 9, pp. 3946–3960, 2020.
- [4] S. Wollenstein-Betech, A. Houshmand, M. Salazar, M. Pavone, C. G. Cassandras, and I. C. Paschalidis, "Congestion-aware routing and rebalancing of autonomous mobility-on-demand systems in mixed traffic," in *Proc. IEEE Int. Conf. on Intelligent Transportation Systems*, 2020, in press.
- [5] T. Roughgarden and E. Tardos, "How bad is selfish routing?" *Journal of the Association for Computing Machinery*, vol. 49, pp. 236–259, 2002.
- [6] A. C. Pigou, *The Economics of Welfare*. Macmillan, 1920.
- [7] P. Bergendorff, D. W. Hearn, and M. V. Ramana, "Congestion toll pricing of traffic networks," *Network Optimization*, pp. 51–71, 1997.
- [8] S. A. Morrison, "A survey of road pricing," *Transportation Research Part A: Policy and Practice*, vol. 20, no. 2, pp. 87–97, 1986.
- [9] P. N. Brown and J. R. Marden, "Studies on robust social influence mechanisms – incentives for efficient network routing in uncertain settings," *IEEE Control Systems Magazine*, vol. 37, no. 1, pp. 98–115, 2017.
- [10] A. Censi, S. Bolognani, J. G. Zilly, S. S. Mousavi, and E. Frazzoli, "Today me, tomorrow thee: Efficient resource allocation in competitive settings using karma games," in *Proc. IEEE Int. Conf. on Intelligent Transportation Systems*, 2019.
- [11] R. W. Rosenthal, "A class of games possessing pure-strategy nash equilibria," *International Journal of Game Theory*, vol. 2, no. 1, pp. 65–67, 1973.
- [12] M. J. Beckmann, C. B. McGuire, and C. B. Winsten, "Studies in the economics of transportation," *Rand Corporation: Research Memoranda*, 1955.
- [13] D. Paccagnan, R. Chandan, B. L. Ferguson, and J. R. Marden, "Incentivizing efficient use of shared infrastructure: Optimal tolls in congestion games," *arXiv preprint arXiv:1911.09806*, 2019.
- [14] R. Meir and D. C. Parkes, "When are marginal congestion tolls optimal?" in *ATT@ IJCAI*, R. Meir and D. C. Parkes, Eds., 2016.
- [15] L. Fleischer, K. Jain, and M. Mahdian, "Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games," in *IEEE Symp. on Foundations of Computer Science*, 2004.
- [16] C. Prendergast, "The allocation of food to food banks," *EAI Endorsed Trans. Serious Games*, vol. 3, no. 10, 2016.
- [17] A. Gorokh, S. Banerjee, and K. Iyer, "From monetary to non-monetary mechanism design via artificial currencies," *Available at SSRN 2964082*, 2019.
- [18] D. Paccagnan, B. Gentile, F. Parise, M. Kamgarpour, and J. Lygeros, "Nash and wardrop equilibria in aggregative games with coupling constraints," *IEEE Transactions on Automatic Control*, vol. 64, no. 4, pp. 1373–1388, 2018.
- [19] M. Salazar, D. Paccagnan, A. Agazzi, and W. P. M. H. Heemels, "Urgency-aware optimal routing in repeated games through artificial currencies," in *European Control Conference*, 2021, submitted. Extended version available at <https://arxiv.org/abs/2011.11595>.
- [20] F. Bullo, *Lectures on Network Systems*, 1st ed. Kindle Direct Publishing, 2020, with contributions by J. Cortes, F. Dorfler, and S. Martinez. [Online]. Available: <http://motion.me.ucsb.edu/book-Ins>
- [21] W. Doeblin, "Sur les proprietes asymptotiques de mouvements régis par certains types de chaînes simples (suite et fin)," *Bulletin mathématique de la Société Roumaine des Sciences*, vol. 39, no. 2, pp. 3–61, 1937.
- [22] Bureau of Public Roads, "Traffic assignment manual," U.S. Dept. of Commerce, Urban Planning Division, Tech. Rep., 1964.