Scaling Marginal Cost Tolling to Address Heterogeneity under Imperfect Information in Routing Games

Andries Rosseau
Vrije Universiteit Brussel
Brussels, Belgium
andries.rosseau@vub.be

Paolo Turrini
Warwick University
Warwick, United Kingdom
p.turrini@warwick.ac.uk

Marjon Blondeel
Vrije Universiteit Brussel
Brussels, Belgium
marjon.blondeel@vub.be

Ann Nowé
Vrije Universiteit Brussel
Brussels, Belgium
ann.nowe@vub.be

ABSTRACT

In routing games, agents select routes in a network in order to minimize their individual latency. Resulting Nash equilibria are known generally not to minimize the total latency across the system, which often requires further coordination. A well-known method that addresses the inefficiency caused by self-interested decision making is *marginal cost tolling* (MCT). Under the traditional assumption of homogeneous agents that trade off time (latency) and money (tolls) equally, marginal cost tolling induces optimal behavior and minimizes total latency. However, how should agents be tolled when their preferences are heterogeneous? We introduce *μ-MCT*, a tolling mechanism that scales marginal cost tolls for routing networks with unknown heterogeneous preferences. In contrast to previous work on heterogeneous routing games, μ-MCT does not assume knowledge of the agents' preferences, thereby respecting privacy concerns, nor does it require knowledge of the network structure. Moreover, an equal amount is tolled to agents that travel the same route, which addresses fairness concerns as well. μ-MCT only has a single parameter, μ, which scales marginal cost tolls and creates a spectrum of tolling mechanisms. We show the properties of μ-MCT for several heterogeneous populations in a set of benchmark networks with high inefficiency. Our results indicate that μ-MCT can considerably improve total latency for a broad range of μ values (and even for surprisingly small tolls). We further ask what μ value should be chosen when optimization is limited and discuss sample-efficient gradient-free learning. μ-MCT is easy to compute, requiring only a derivative of the latency, and can be an elegant tolling mechanism for routing networks when working under imperfect information.

KEYWORDS

Congestion, Routing, Price of Anarchy, Multi-Agent, Learning, Mechanism, Marginal Cost, Heterogeneous

1 INTRODUCTION

In routing games [37, 48], a population of agents – also referred to as users or players – must each select a route on a network from their origin to their destination. Every edge in the network admits a latency function, which is the time it takes to traverse the edge[1]. Some applications of routing games are urban traffic management [42], internet traffic routing [20], supply chain and logistics [43] and aircraft routing and scheduling [6]. If we assume users act rationally and choose their routes according to their own self-interest, a Nash equilibrium [26] can be obtained where no individual is able to decrease their latency by deviating from its chosen path [24, 35, 39]. However, in many routing games, this results in a suboptimal outcome where the total latency of all users could still be smaller if the traffic were routed through a central planner or another coordination mechanism. This discrepancy is expressed as the *Price of Anarchy* (PoA) [20, 27, 39], which is the ratio of the total latency under Nash equilibrium to the minimal possible total latency, which is called the *system optimum* (SO). The closer this ratio is to 1, the less inefficiency is caused by moving from centralized route planning to decentralized (selfish) decision making.

One way to reduce the Price of Anarchy in routing games is to apply tolling. Traditionally, tolls are modeled simply by adding them directly to the latency, creating a new cost function for the agents that induces a different traffic flow. A well-known tolling method that is optimal (i.e., minimizes the total latency) under known homogeneous preferences is *marginal cost tolling* (MCT) [5]. Marginal cost tolling assumes that all agents trade off time (latency) and monetary incentives (tolls) equally. In reality, agents may have heterogeneous preferences over latency and tolls.

To address heterogeneity, previous work has proposed to scale the marginal cost toll for each individual agent with respect to its specific preference value [31, 32, 45]. However, this approach requires agents to specify precise numerical values for their preferences, which can be a difficult task for agents (e.g., humans). In addition, all these preference values need to be communicated to the tolling agent, which can become impractical in large-scale networks with many agents. Moreover, the acquisition of this information might raise concerns about privacy. Finally, the assumption that users can be tolled a different amount for using the same route can again be impractical, but more importantly might be perceived as unfair. Our approach addresses all of these concerns in an adaptation of marginal cost tolls.

Proc. of the Adaptive and Learning Agents Workshop (ALA 2025), Avalos, Aydeniz, Müller, Mohammedalamen (eds.), May 19 – 20, 2025, Detroit, Michigan, USA, ala-workshop.github.io. 2025.

[1] In many domains, latency can mean something different or more abstract than a time delay.

Our contribution: We present **μ-MCT**, a tolling mechanism that generalizes marginal cost tolls for heterogeneous users. It has a single tunable parameter μ that scales the tolling imposed on the agents. We apply μ-MCT in routing games with a high Price of Anarchy, such as generalized Pigou [30], Braess graphs [8], and the well-known Sioux Falls benchmark [17], while varying the degree of heterogeneity over agent preferences. Our results show a broad range of μ values yielding equilibria that considerably improve the total latency over no tolling, even when tolls are small (compared to classical marginal cost tolls). While these results are promising and an indicator for real world application potential, we also show the possibility to create adversarial networks where μ-MCT actually increases total latency for certain μ values. If ample μ values can be evaluated, this negative result should be of no practical concern, and μ-MCT can be a powerful tolling mechanism that is able to deal with unknown heterogeneity, respects privacy, and tolls agents equally without knowing the network structure.

Our benchmarks demonstrate that a broad range of μ values perform well from the outset. However, we also explore the potential of integrating μ-MCT with a gradient-free learning approach in cases where an initial parameter selection proves insufficient. We mainly discuss Bayesian optimization since it prioritizes sample-efficiency and therefore performs well with limited computational budget.

In summary, μ-MCT:

- Provides the same toll per edge for all agents (meaning easier practical implementation and addressing fairness concerns);
- Requires no direct information about individual agent preferences, thereby respecting privacy;
- Scales to any network structure, network size, or heterogeneity distribution;
- Performs well in practice, even when considering surprisingly small tolls;
- Is easy to compute (for any given μ), requiring only a derivative of the latency.

Related work

A different approach that sets one (global) toll per edge for heterogeneous users is presented in [7, 12, 15]. Here, tolls are static constants. However, given some plausible assumptions, they can be set so that the system optimum is always achieved. This is a powerful result, but in practice, computing an optimal set of static tolls becomes infeasible when the number of players and the network size increase. Moreover, when tolls do not scale dynamically with the flow, static tolls can toll agents large amounts at times when this is not required. When possible, the flow-dependent nature of tolls is a beneficial property for real world settings.

More relevant to our work is Δ-*tolling*, proposed in [41]. Δ-tolling also tolls edges dynamically by assuming that all latency functions are a BPR function [19, 42]. This is often a suitable model for real-world traffic networks. Based on this assumption, Δ-tolling only needs evaluations of the latencies for every edge at every time step. However, it has two parameters that need to be tuned to the problem setting; three if one counts the average of the preference distribution (which the authors acknowledge). In contrast, our approach assumes knowledge of the latency functions (which can also be gathered from real-world data), but only has one tuneable parameter.

The setting of μ-MCT is one of imperfect information regarding the users' preferences. The role of information in congestion games is well established [4, 33], but the literature so far has mainly studied this from the point of view of whether routes are known to the agents participating in the congestion games. In contrast, this work studies imperfect information from the point of view of mechanism design.

Brown and Marden [9] show that any network-agnostic tolling mechanism that improves the total latency on one or more networks, will necessarily degrade performance on another specific network (except when restricted to parallel-link networks). This implies that no single μ value can always do at least as well as no tolling in all networks (except the trivial $\mu \to +\infty$ setting of no tolling).

2 BACKGROUND

2.1 Routing games

We model a routing game as a directed graph $G \equiv (V, E)$, where V represents the finite set of vertices v and E the finite set of directed edges e. An edge is an ordered pair of distinct vertices. It is possible to have multiple edges between the same two vertices, which are called parallel edges. In a routing game, traffic must travel from a source vertex to a destination vertex along the edges of G. Each routing game has at least one *source-destination* (SD) pair. A path p between a source s and a destination d is an ordered set of edges for which the end vertex of one edge is the start vertex of the subsequent edge in the set, and which at no point circles back on itself[2]. A flow f is defined on the set E of all edges e, $\{f_e\}_{e\in E}$, or the set P of all paths p, $\{f_p\}_{p\in P}$, and indicates the amount of traffic on each edge or path. For any edge flow, $f_e = \sum_{p\in P: e\in P} f_p$, meaning that the flow on edge e is equal to the sum of the flow on paths that contain edge e. The total flow on all paths equals the traffic rate r:

$$r = \sum_{p\in P} f_p. \tag{1}$$

A routing game has a latency function l defined on all edges $e \in E$. The edge latency function $l_e(f_e)$, is defined as a function of the flow f_e per edge e, which represents the delay an agent experiences to traverse the edge. Latency functions are assumed to be well-behaved in the sense that they are continuous, non-negative, monotonically increasing with f_e, continuously differentiable, and convex. The latency of a path p with respect to the flow f is the sum of the latencies of its edges:

$$l_p(f_p) = \sum_{e\in p} l_e(f_e). \tag{2}$$

The total latency in the system produced by flow f is then:

$$L(f) \equiv \sum_{p\in P} l_p(f_p) \cdot f_p = \sum_{e\in E} l_e(f_e) \cdot f_e, \tag{3}$$

[2] Routing games can in principle contain cyclic paths, but they will never be considered by rational agents, since latency and tolls can only be non-negative and therefore only add to the experienced cost. This also avoids infinite strategy spaces.

(Proof of the second equality can be found in [37].)

We consider unweighted, non-atomic routing games, where agents each control an infinitesimal fraction of flow [35]. If we assume users act rationally, a Nash equilibrium [26] will be obtained where no individual can decrease their latency by deviating from its chosen path [24]. A permutation of indistinguishable agents will also lead to a Nash equilibrium that has the same flow on all edges; therefore, the flow that corresponds to a Nash equilibrium is called a *Nash flow* f^{NF}. Non-atomic routing games have a unique Nash flow [5].

2.2 The Price of Anarchy

In many routing games, the total travel time of all users can still be reduced if traffic was routed through a central planner or some other coordination mechanism [34, 36]. A system optimum (SO) flow f^* captures this notion and minimizes the total system latency:

$$f^* = \arg\min_{f} L(f) \tag{4}$$

The ratio of the total travel time under the Nash flow to the total travel time under the system optimum is called the Price of Anarchy (PoA) [27], defined as:

$$PoA(G, r, l) = \frac{L(f^{NF})}{L(f^*)}, \tag{5}$$

The closer this ratio is to 1, the less the network suffers from the inefficiency caused by moving from centralized route planning to decentralized (self-interested) decision making.

2.3 Tolling

To improve the inefficiency of traffic networks, represented by the Price of Anarchy, tolls can be imposed on the network edges, and are denoted as τ_e. We consider the case of dynamic tolling, where tolls are a function of the amount of flow (i.e., $\tau_e(f_e)$). Many successful real-world applications of dynamic pricing exist [1–3, 44]. We also assume the standard description of tolls that are continuous, non-negative, monotonically increasing with f_e, continuously differentiable, and convex.

Traditionally, tolls are simply added to the latency, defining a *cost function* $c_e(f_e)$ per edge e:

$$c_e(f_e) = l_e(f_e) + \tau_e(f_e). \tag{6}$$

Tolls are implemented with the purpose of reducing the total latency in the system, and therefore also reducing the Price of Anarchy. New Nash flows can be induced by the edge *costs* instead of only the latencies. The system optimum flow is still defined by the minimal possible total latency in the system, irregardless of tolls. To clearly distinguish between the Price of Anarchy with and without tolls, we define the *Induced Price of Anarchy* (IPoA), as:

$$IPoA(G, r, c) = \frac{L(f_c^{NF})}{L(f^*)}, \tag{7}$$

where we indicate that the Nash flow is induced by the cost function c instead of solely the latency l.

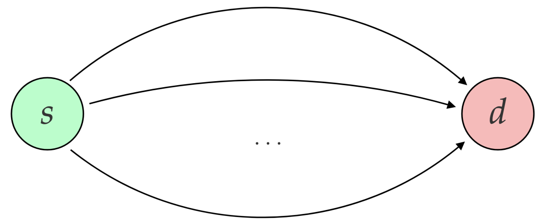

Figure 1: General representation of a parallel edge network. A famous example of a parallel edge network is the Pigou network [30].

2.4 Marginal Cost Tolling

A well-known way of establishing tolls is through *marginal cost tolling* (MCT) [30, 36]. The idea behind marginal cost tolling is to toll each agent according to the increase in latency its presence or actions cause for the whole system. The marginal cost toll $\tau_e^*(f_e)$ is defined as:

$$\tau_e^*(f_e) = f_e \cdot l_e'(f_e). \tag{8}$$

where $l_e'(f_e) = \frac{dl_e(f_e)}{df_e}$.

The *marginal cost function* c^* is then defined as:

$$c_e^*(f_e) = l_e(f_e) + f_e \cdot l_e'(f_e). \tag{9}$$

If the marginal cost function is imposed as the cost function on any routing problem with known homogeneous preferences instead of its original latency function, the induced Nash flow $f_{c^*}^{NF}$ becomes the same as the system optimal flow f^* under the original latency function l_e [5, 37]. This means that $IPoA(G, r, c^*) = 1$.

3 HETEROGENEOUS PREFERENCES

Traditionally, tolls are simply added to latencies, assuming that all agents value latencies and tolls equally. This is a strong assumption that does not represent the complexity of real world scenarios, where heterogeneity over preferences is the norm rather than the exception. Therefore, under the more general assumption that agents have heterogeneous preferences, the edge cost function becomes:

$$c_{e,i}(f_e) = (1 - \eta_i)\; l_e(f_e) + \eta_i\; \tau_e(f_e), \tag{10}$$

where agent $i \in [0, r]$ has preference $\eta_i \in [0, 1)$ over tolls, and $1 - \eta_i$ is then the preference over latency. Since agents are represented on a continuous spectrum (and an agent $i \in [0, r]$ controls an infinitesimal portion dr of all traffic r), we can order the continuum of agents according to increasing η_i, without loss of generality. Therefore, η is a monotonically increasing function mapping $[0, r]$ to $[0, 1)$. The special case of homogeneous agent preferences is given by a constant function, and we denote its value simply by $\bar{\eta}$.

Note that a routing game with potentially different cost functions per agent (a player-specific game) no longer implies a unique equilibrium based on the classical potential-game argument [35], which validated Equation 7. Uniqueness can instead be shown under reasonable assumptions through Brouwer's fixed-point theorem [10].

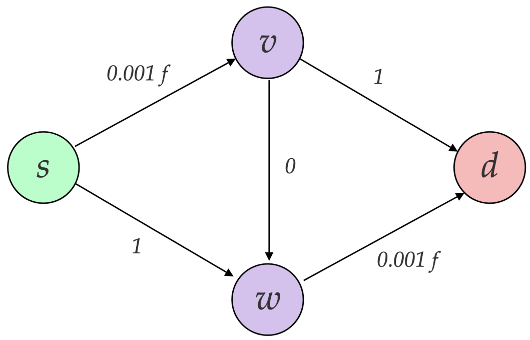

Figure 2: An example of a first-order Braess network for $N = 1000$. The Nash flow consists of all agents selecting $s - v - w - d$, leading to an individual latency of 2, and a total latency of 2000. The system optimum flow is that half of the agents select $s - v - d$ and the other half $s - w - d$, leading to an individual cost of 3/2 and a total latency of 1500. This produces a Price of Anarchy of 4/3. The system optimum is not a Nash flow, since each agent is incentivized to switch to $s - v - w - d$ which has an individual cost of 1.

3.1 Individually tolling heterogeneous users

How can one toll the system to reduce the Price of Anarchy in the case of heterogeneous preferences? We know that marginal cost tolling in the case of homogeneous preferences leads to an induced Price of Anarchy of 1. An intuitive solution in the heterogeneous case is to scale the marginal cost toll (Equation 8) with respect to the preference value η_i of each agent i, and add latency $l_e(f_e)$:

$$\tau_{e,i}(f_e) = l_e(f_e) + \frac{f_e \cdot l'_e(f_e)}{\eta_i}, \tag{11}$$

This causes the total cost $c_{e,i}$ to reduce once again to the homogeneous marginal cost function, which induces the system optimal flow:

$$\begin{aligned} c_{e,i}(f_e) &= (1 - \eta_i)\ l_e(f_e) + \eta_i\ \tau_{e,i}(f_e) \\ &= (1 - \eta_i)\ l_e(f_e) + \eta_i\ \left(l_e(f_e) + \frac{f_e \cdot l'_e(f_e)}{\eta_i} \right) \\ &= l_e(f_e) + f_e \cdot l'_e(f_e) \\ &= c^*_e(f_e), \end{aligned} \tag{12}$$

While the marginal cost function c^*_e is independent of i, the toll in equation 11 is not, and requires knowledge of all agent preferences. This is the approach taken in the works of [13, 31, 32, 45].

4 μ-MCT

In many applications, agents (e.g., humans) might not know their exact numerical preferences. Moreover, the logistics of communicating these values to the central tolling agent might be hard to implement in practice, and acquiring this information per agent might violate privacy concerns. Finally, it could be deemed unfair when agents pay different tolls, even though they traverse the same route. To address these concerns, this work studies the situation of imperfect information, meaning we do not know η_i, and can only impose one toll per edge for all agents. We present **μ-MCT**, a mechanism that tolls an edge according to:

$$\tau_{e,\mu}(f_e) = l_e(f_e) + \frac{f_e \cdot l'_e(f_e)}{\mu}, \tag{13}$$

The difference with Equation 11 is that μ-MCT is not tailored to each individual in the population (which would assume knowledge of all preferences η_i). Instead, apart from the flow and the derivative of the latency, $\tau_{e,\mu}$ is now determined by a global and fixed μ, which is always the same for every agent and every edge in the network. The cost function on edge e for agent i then becomes:

$$\begin{aligned} c_{e,i,\mu}(f_e) &= (1 - \eta_i)\ l_e(f_e) + \eta_i\ \tau_{e,\mu}(f_e) \\ &= (1 - \eta_i)\ l_e(f_e) + \eta_i\ \left(l_e(f_e) + \frac{f_e \cdot l'_e(f_e)}{\mu} \right) \\ &= l_e(f_e) + \frac{\eta_i}{\mu}\ \left(f_e \cdot l'_e(f_e) \right). \end{aligned} \tag{14}$$

The μ-MCT tolling function depends on the latency, the derivative of the latency, the flow, and μ. Only μ can be directly controlled by the tolling agent, but agents still select routes according to their experienced cost function (Equation 14), which also depends on the intrinsic η_i values of the agents. Given the latencies and the η_i values of the population, each value for μ then induces a specific flow, which corresponds to a total latency $L(f)$ (Equation 3) and an induced Price of Anarchy.

Note that an agent's η_i and the tolling agent's μ (generally) do not cancel out to reduce to homogeneous marginal cost tolling (as was the case for previous works [31, 32, 45]). From now on, we define the IPoA as:

$$IPoA(G, r, \eta, \mu) = \frac{L(f^{NF}_{\eta,\mu})}{L(f^*)}, \tag{15}$$

where $L(f^{NF}_{\eta,\mu})$ is the total latency for the Nash flow induced by the combination of the preference function η and the tolling parameter μ, which together create the cost functions that the agents experience (Equation 14).

4.1 Lower and upper bounds on the IPoA

The IPoA induced by μ-MCT depends on η and μ (Equation 14). The lower bound is given by the homogeneous case, where $\forall i \in [0, r] : \eta_i = \bar{\eta}$, and where μ is set equal to $\bar{\eta}$, thereby obtaining the system optimum and an IPoA of 1. No algorithm can do better than this.

However, under μ-MCT, the IPoA has no upper bound. To show this, consider the adversarial network with two parallel edges connecting one source-destination pair, for which $l_1(f_1) = f_1$ and $l_2(f_2) = C$. If we consider the traffic rate r, and an arbitrary large constant $C \gg r$, the untolled Nash flow and the system optimum flow coincide: all agents select the first edge, giving a total latency of r^2 and a PoA of 1. However, given any η, in the limit of $\mu \rightarrow 0$, the cost for edge 1 (cf. Equation 14) will become larger than C for all agents, steering them towards the second edge. This then leads to a total latency of $r \cdot C$, and therefore an IPoA of $\frac{C}{r}$, showing that the IPoA does not have an upper bound, since C can be arbitrarily large.

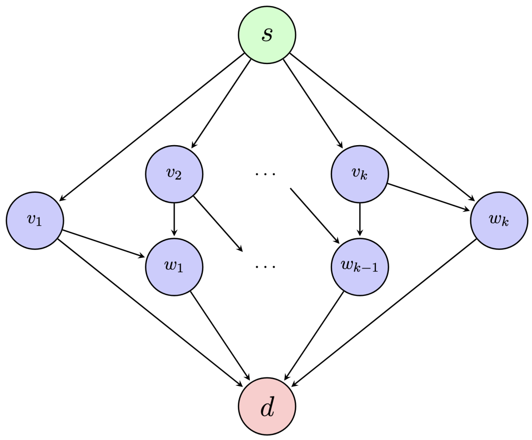

Figure 3: General representation of a (k-th order) Braess network.

5 EXPERIMENTS

How does tolling with μ-MCT fare when applied to well-known benchmarks with high inefficiency and several heterogeneous populations? We compare with regards to the outcome without tolling and to the system optimum.

Note that to be able to run numerical best-response dynamics for the evaluation of the total latency induced by a μ value, we discretize (fixed) non-atomic networks and scale the arguments of the latency functions with the now finite number of players N, where player $i \in (1, 2, ..., N)$ has preference η_i. In a strict sense, this makes our experiments atomic, and atomic routing games have different bounds on the Price of Anarchy, with and without marginal cost tolling. For a comprehensive overview of marginal cost tolls and the Price of Anarchy in atomic congestion games, we refer to [29]. However, since our networks are *fixed* and we increase the number of players to a large N, we are effectively in the regime of *large games*, where properties such as a unique Nash flow and PoA are recovered from non-atomic games in the limit [11, 14, 23, 25, 38]. This means that our numerical results will converge to the continuous, non-atomic solution with arbitrary precision, given a large enough N.

5.1 Networks

5.1.1 Pigou networks. The first kind of networks considered are Pigou networks. The classical Pigou network has two parallel edges connecting a source-destination pair (cf. Figure 1). One edge has a latency function that scales positively with the flow, and the other has a constant latency equal to the maximal latency on the first edge. The Nash flow is for all agents to select the first edge, while the system optimal flow averts a portion of the flow to the constant latency edge. The simplest Pigou network has a latency function $l_1(f_1) = f_1/N$ on the first edge, and a constant latency of $l_2(f_2) = N$ for the second edge. The Nash flow without tolling consists of all agents choosing the first edge, and the system optimum flow is half-half over the two edges, creating a PoA of 4/3 [39]. We can extend the simple Pigou network by adding more edges with varying (non-linear) functions of the flow, again increasing the size of the problem (and therefore the amount of options agents have) and thereby also varying the initial PoA. Another possibility is to replace an edge with a Braess network [8]. In general, a wide range of constructions are possible which we will categorize under the umbrella of Pigou networks.

5.1.2 Braess networks. We also consider Braess networks, where in each network, a version of the Braess' paradox is present [8]. The Braess' paradox is a counterintuitive phenomenon in routing networks where adding extra edges can worsen the total latency rather than improve it (i.e., increase the PoA). The paradox occurs because agents act selfishly to minimize their own travel time, shifting the equilibrium in a way that increases the total latency. An example of the classic Braess network is given in Figure 2. If we consider $N = 1000$, the Nash flow for the agents is to follow $s - v - w - d$. The system optimum flow is for half of the agents to take route $s - v - d$, and half of the agents $s - w - d$. This leads to a PoA of 4/3, which is the maximal possible PoA in networks with linear latencies [39][3]. This Braess network can be extended to higher order sizes with their corresponding PoAs, as shown in Figure 3 and described in [47]. The Braess networks are ideal testing grounds for our algorithm, since the PoAs are maximally suboptimal.

5.1.3 Sioux Falls. Thanks to their interesting properties, Braess and Pigou networks have been used as important benchmark problems for decades. They represent prominent examples of the inefficiency of selfish routing. Evaluating μ-MCT on these problems is therefore an important gauge of its properties. Additionally, there are larger benchmark networks that are modeled from real-world data, which tend to be more computationally intensive, and involve greater complexity in terms of nodes, edges and agents. However, these larger networks are often "easier", in the sense that the inefficiency caused by selfish routing is not maximal, but only to a lesser degree present in the network. To further assess μ-MCT's effectiveness, we apply it to the widely used Sioux Falls benchmark [17], a road network modeled after Sioux Falls, the most populous city in South Dakota (US), featuring multiple source-destination pairs and thousands of agents. In our experiments, we use the standard Bureau of Public Roads (BPR) latency functions, defined as:

$$l_e(f_e) = t\left(1 + a\left(\frac{f_e}{c}\right)^b\right), \tag{16}$$

where $a = 0.15$, and $b = 4$, which are representative for many real-world roads, and t and c are edge-specific parameters representing the free-flow latency and the edge capacity, respectively. The BPR function is widely used and representative of actual roads [19, 42], and offers a non-linear benchmark as well.

5.1.4 Preference distributions. We investigate the influence on the IPoA for several types of preference distributions. We sample from (i) uniform distributions, (ii) normal distributions $\mathcal{N}(0.5, \sigma^2)$, truncated between 0 and 1, with several standard deviations σ, and (iii) for the Sioux Falls benchmark, we sample from an inverse *Dagum distribution*. The Dagum distribution is the best fit to the personal

[3]To show the Braess' paradox, one can remove $v - w$, and the PoA becomes 1.

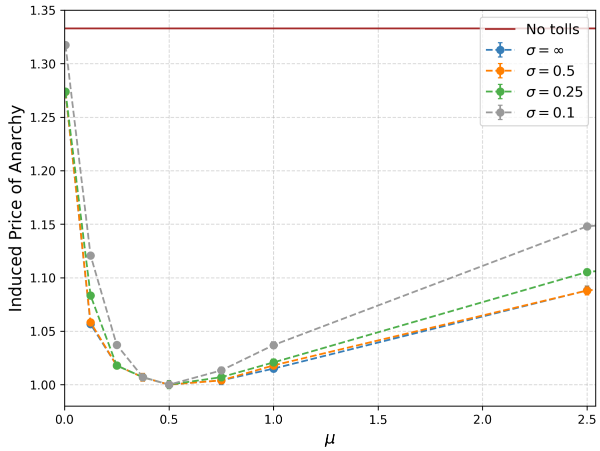

Figure 4: IPoA for the Pigou network. The preferences are drawn from different normal distributions $\mathcal{N}(0.5, \sigma)$. A larger σ (i.e., higher heterogeneity) generally leads to a higher IPoA. Observe that for a considerable range of values around $\mu = 0.5$, outcomes are close to the system optimum. Even for large values like $\mu = 2.5$ (meaning that tolls are only a fifth of the equivalent of traditional MCT tolls), μ-MCT still does considerably better than the no-toll setting. For even larger values (not shown here), all curves eventually align and converge to the no toll setting.

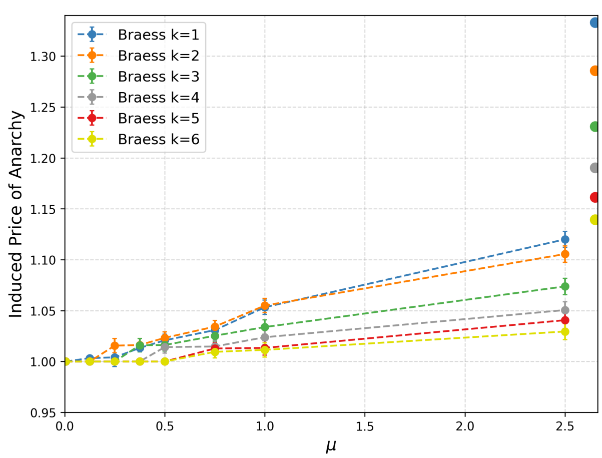

Figure 5: IPoA for a set of Braess networks tolled with μ-MCT. Preferences are sampled from a uniform distribution. The dots on the right represent the scenario of no tolling (or $\mu \to \infty$). We observe that even for relatively high values of μ (i.e., small tolls compared to standard MCT), μ-MCT still does considerably better than the no-toll setting. Note that in general, the best μ value is not 0.5, the average of the preference distribution.

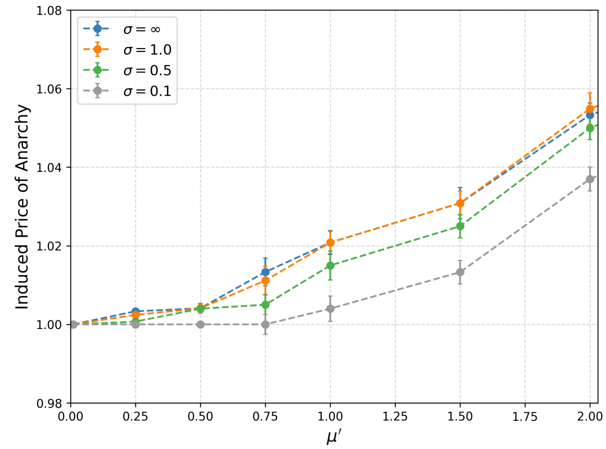

Figure 6: IPoA for the first order (classic) Braess network. The preferences are drawn from different normal distributions $\mathcal{N}(0.5, \sigma)$. A larger σ (greater heterogeneity), generally leads to a higher IPoA. For higher μ values (which are not shown here in order to focus on the most significant region), values for all σ gradually converge again.

income distribution for US citizens [22]. The probability density function of the Dagum distribution is given by:

$$f_D(x) = \frac{abc}{x^{b+1}(1 + ax^{-b})^{c+1}}$$

with the best fitting parameters for US incomes (2012) being $a = 22020.6$, $b = 2.7926$, and $c = 0.2977$ [22]. The reasoning behind sampling from the Dagum distribution is that the more income a user has, the lower its η_i will be, so we invert the sampled values to obtain η_i, making our model more representative of potential real-world heterogeneity.[4]

5.2 Results

In the following, we test μ-MCT for a range of μ values in the Braess and Pigou networks, and on the Sioux Falls benchmark. Nash flow convergence is obtained using the Method of Successive Averages [40], which we adjusted for our heterogeneous setting.

5.2.1 Pigou network. We consider a two-edge Pigou network (Figure 1) with $N = 1000$. The first edge has latency function $l_1(f_1) = 0.001f_1$, and the second edge $l_2(f_2) = 1$. The Nash flow without tolling consists of all agents choosing the first edge. The system optimum flow is again half-half over the two edges, leading to the maximal PoA of 4/3. In Figure 4, we show the results of applying μ-MCT to this setting for normal distributions with different standard deviations (the case of $\sigma = \infty$ is the uniform distribution case). The results show that the best μ is the mean of the distributions, i.e., 0.5. High heterogeneity (i.e., a high variance for the distribution from which the η_i values are sampled) leads generally to an increased

[4]In our experiments, due to the size of the Sioux Falls routing game, we group the sampled η_i values into 40 equally spaced bins, effectively creating a discrete version of the Dagum distribution. This makes it so that we only have 40 agent 'types', and we therefore only have to calculate shortest-path calculations for 40 types, instead of potentially N different types. This significantly speeds up numerical convergence.

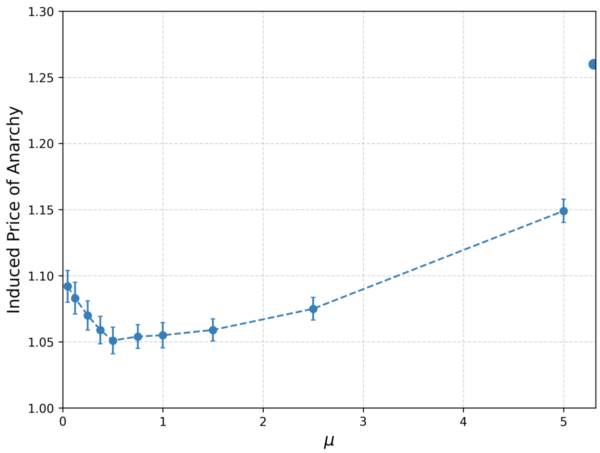

Figure 7: IPoA for the Sioux Falls benchmark when tolled with μ-MCT. The dot on the right represents no tolling. We observe that μ-MCT does surprisingly well for a broad range of μ values, with a minimum at 0.5, the average of the inverse Dagum distribution from which preferences were sampled.

IPoA. However, a broad range of values for μ significantly improves the IPoA compared to the no-toll setting (the blue dot), even in cases of high heterogeneity.

In testing μ-MCT on extensions of the two-edge Pigou network, we observed results fully consistent with those obtained for the original two-edge configuration. To avoid redundancy, we do not include these results here.

5.2.2 Braess networks. We consider k'th-order Braess networks with k ranging from 1 to 6 (Figure 3) and $N = 4200$ [47]. The results for the uniform sampling of $\eta_i \sim \mathcal{U}(0, 1)$ are presented in Figure 5. We notice that, again, for almost all values of μ, the IPoA is close to 1. Even for $\mu = 2.5$, meaning a five-fold reduction in tolling compared to the equivalent of classical marginal cost tolls (i.e., a normal distribution with mean 0.5 and variance 0), the maximal value for the IPoA is only 1.12 over all Braess networks. When compared to the untolled PoA (≈ 1.33) this is still a considerable improvement. This suggests that in general, even when small tolls are applied compared to classical marginal cost tolling, the PoA is still improved upon significantly, matching the results from the Pigou network.

Observe that for Braess networks, when μ gets smaller (meaning larger tolls), the IPoA nears or coincides with the system optimum. When μ is small, most or all agents have $\eta_i > \mu$, and experience a cost that pushes them towards routes that contain fewer edges with latency functions that scale with the flow. In Braess networks, these types of routes are exactly the routes used for the system optimal flow. This effect is therefore caused by the specific structure of the Braess networks and its symmetry, and cannot be generalized to all network types. However, this example already disproves the intuitive notion that the best possible μ is always the mean of the preference distribution, regardless of network type.

In Figure 6, we again investigate the influence of the variance for sampling the η_i values. The IPoA values here are shown only for the first order Braess network for preferences sampled from normal distributions $\mathcal{N}(0.5, \sigma)$ with different standard deviations σ (consistent results are observed for higher orders, so we omit them). We observe that for larger σ values (meaning greater heterogeneity), the IPoA increases for the same μ. However, μ-MCT (even with smaller tolls) performs again surprisingly well even in the strongly heterogeneous setting.

5.2.3 Sioux Falls. The results for the Sioux Falls network are shown in Figure 7. The optimal value for μ is 0.5, indicating that setting μ to the average of the preference distribution yields the best outcome. Nonetheless, both higher and lower values of μ significantly improve the IPoA compared to the no-toll scenario (represented by the blue dot).

Notably, the case with higher μ values, corresponding to small tolls, is again compelling. It demonstrates once more that even with relatively small tolls, substantial improvements over the no-tolling case can be achieved. To the best of our knowledge, this is the first work to identify this result.

6 OPTIMIZATION OF μ-MCT WITH LIMITED EVALUATIONS

Our results on benchmarks with high inefficiency and varying heterogeneity show that for a broad range of μ values, the PoA improves. This suggests that in practical settings, many μ values could improve the PoA. Especially large μ values – which correspond to small tolls – can be of interest, since they perform surprisingly well while reducing the amount of tolls.

However, simply setting a μ value, even a large one, can in certain adversarial cases still lead to an increased PoA, as shown by the upper bound on the IPoA (Section 4.1). If a toll designer can evaluate a broad range of μ values, the lack of an upper bound for the PoA is not as significant in practice. If all μ values decrease the PoA, the designer is free to choose which μ is best suited to the problem (being able to trade off minimizing latency and tolls). If only certain μ values increase the PoA, the designer can choose from the region that improves the PoA, or in the worst-case return to no tolling.

In contrast, when the toll designer has a limited budget for μ evaluations, we propose the following strategies to find a suitable μ. These strategies can also be combined.

6.1 Start with small tolls

A tolling agent can start with a large μ, corresponding to a safe setting of little tolling, and evaluate whether it decreases the PoA or not. Next, the agent can decrease μ gradually until satisfied (or the evaluation budget has been spent). Starting from a high μ is not only a safe option, but also toll-conserving.

6.2 Sample-efficient gradient-free optimization

Instead of arbitrarily decreasing μ from a high value, a learning approach can be used as well. The goal of the learner is to minimize the total induced latency $L(f^{NF}_{\eta,\mu})$ with respect to μ. It is not recommended to use gradient-based approaches, since there is no obvious analytical function mapping μ directly to $L(f^{NF}_{\eta,\mu})$. Therefore, we propose gradient-free optimization. Since finding the Nash flow for a given μ is the most expensive step in the optimization process (e.g., simulated or real-world convergence may require time and resources), we propose Bayesian optimization (BO), which is known to be sample-efficient (i.e., limiting the number of $L(f^{NF}_{\eta,\mu})$ evaluations). Bayesian optimization can either run until convergence (thereby guaranteeing to find an optimal μ), or work with a budget of function evaluations. An overview of Bayesian optimization is provided in [46].

An abstract implementation of Bayesian optimization with Gaussian Processes (GP) for heterogeneous routing games is provided in Algorithm 1. In line with the previous considerations, the initial sample set is recommended to contain at least one large μ. Several acquisition functions can be used, but we suggest the Expected Improvement function, as it balances exploration and exploitation well [18].

Finally, note that our discussion on sample-efficient gradient-free optimization is included solely to provide a constructive suggestion on how to use μ-MCT in practice, especially when evaluations are limited. However, we ran our benchmarks using Bayesian optimization (with a Gaussian process prior and the Expected Improvement acquisition function), which – given enough budget – provides similar curves as in Section 5. We therefore omit them here.

Algorithm 1 Bayesian Optimization for μ-MCT

Require: Method of evaluating $L(f^{NF}_{\eta,\mu})$ for given μ, prior GP model $\mathcal{GP}(\mu)$, acquisition function $\alpha(\mu)$, initial sample set $\{\mu_i, L_i\}_{i=1}^{n_0}$, budget B

Ensure: μ minimizing $L(f^{NF}_{\eta,\mu})$

Initialize Gaussian Process (GP) $\mathcal{GP}(\mu)$ using initial samples $\{\mu_i, L_i\}_{i=1}^{n_0}$

for $i = n_0$ to B **do**

  Select next query point:

$$\mu_i = \arg\max_{\mu} \alpha(\mu \mid \mathcal{GP})$$

  Evaluate $L_i = L(f^{NF}_{\eta,\mu})$

  Update GP model with new observation (μ_i, L_i)

end for

return $\mu = \arg\min_{\mu_i} L_i$

7 DISCUSSION

We introduced μ-**MCT**, a scalable tolling mechanism designed for routing games where users have unknown preferences. In contrast to a body of previous work that addresses heterogeneity by individually scaling marginal cost tolls per user, μ-MCT imposes a single toll per edge for all users and requires no knowledge of individual agent preferences, thus respecting privacy and addressing fairness concerns. μ-MCT is not suited for cases where only static tolls are allowed; here, methods like the ones presented in [12, 15] can offer an alternative.

Our experiments demonstrate that μ-MCT performs robustly across a variety of benchmark networks, improving the Price of Anarchy even when the tolls are kept relatively small and/or heterogeneity among users is high. Notably, μ-MCT's general effectiveness with small tolls underscores its potential for real-world applications, where excessive tolling may be impractical or undesirable.

When a toll designer has a limited budget to evaluate μ values, gradient-free methods like Bayesian optimization were proposed to find a suitable μ while considering sample efficiency.

Future work could explore the case of Dynamic Traffic Assignment [16], where demand and network conditions vary over time. Moreover, since our results indicate that surprisingly small tolls can lead to a significant improvement in total latency, formalizing the notion of working with a restricted toll budget and its implications for the Price of Anarchy might also be of interest in future work.

One could also investigate whether μ-MCT performs well in other domains with imperfect information in addition to traffic management. Of special interest are social dilemma games such as common-pool resource games [28], where inefficiencies caused by selfish decision making are prevalent. It would moreover be interesting to explore the interaction between the system designer and the users' information asymmetry, exploring information-constrained user equilibria [4, 33] and their relation to μ-MCT. In [21], a set of temporally and spatially extended social dilemma games are presented, in which agents deal with partial observability. Extending these games to include information-restricted mechanism design (e.g., extensions of μ-MCT) presents another interesting research direction.

Overall, μ-MCT presents a scalable tolling mechanism that addresses important challenges in heterogeneous routing games.

REFERENCES

[1] 2023. I-66 Commuter Information: Arlington Transportation. https://www.arlingtontransportationpartners.com/services/i-66-outreach [Online; accessed 12. Oct. 2024].

[2] 2024. How Dynamic Pricing Works - LBJ, NTE & NTE 35W TEXpress Lanes. https://www.texpresslanes.com/pricing/how-dynamic-pricing-works [Online; accessed 12. Oct. 2024].

[3] 2024. Transportstyrelsen: Congestion Taxes in Stockholm and Goteborg. https://www.transportstyrelsen.se/en/road/road-tolls/Congestion-taxes-in-Stockholm-and-Goteborg/congestion-tax-in-stockholm/hours-and-amounts-in-stockholm [Online; accessed 12. Oct. 2024].

[4] Daron Acemoglu, Ali Makhdoumi, Azarakhsh Malekian, and Asuman E. Ozdaglar. 2018. Informational Braess' Paradox: The Effect of Information on Traffic Congestion. *Oper. Res.* 66, 4 (2018), 893–917. https://doi.org/10.1287/OPRE.2017.1712

[5] Martin Beckmann, Charles B McGuire, and Christopher B Winsten. 1956. *Studies in the Economics of Transportation.* Technical Report.

[6] Dimitris Bertsimas and Amedeo Odoni. 1998. Airline Schedule Planning: Integrated Models and Algorithms for Schedule Design and Fleet Assignment. *Transportation Science* 32, 3 (1998), 239–255.

[7] Vincenzo Bonifaci, Mahyar Salek, and Guido Schäfer. 2011. Efficiency of restricted tolls in non-atomic network routing games. In *Algorithmic Game Theory: 4th International Symposium, SAGT 2011, Amalfi, Italy, October 17-19, 2011. Proceedings 4.* Springer, 302–313.

[8] Dietrich Braess. 1968. Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung* 12, 1 (1968), 258–268.

[9] Philip N Brown and Jason R Marden. 2020. Can taxes improve congestion on all networks? *IEEE Transactions on Control of Network Systems* 7, 4 (2020), 1643–1653.

[10] Geunyeong Byeon. 2016. *Equilibria for nonatomic routing games with heterogeneous players.* Ph.D. Dissertation. Seoul National University.

[11] George Christodoulou, Elias Koutsoupias, and Paul G Spirakis. 2011. On the performance of approximate equilibria in congestion games. *Algorithmica* 61 (2011), 116–140.

[12] Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. 2003. Pricing Network Edges for Heterogeneous Selfish Users. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing.* ACM, 521–530.

[13] Stella C Dafermos. 1973. Toll patterns for multiclass-user transportation networks. *Transportation science* 7, 3 (1973), 211–223.

[14] Michal Feldman, Nicole Immorlica, Brendan Lucier, Tim Roughgarden, and Vasilis Syrgkanis. 2016. The price of anarchy in large games. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing.* 963–976.

[15] Lisa Fleischer, Kamal Jain, and Mohammad Mahdian. 2004. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *45th Annual IEEE Symposium on Foundations of Computer Science.* IEEE,

277–285.
[16] Michael Florian and Donald Hearn. 1995. Network equilibrium models and algorithms. *Handbooks in operations research and management science* 8 (1995), 485–550.
[17] Transportation Networks for Research Core Team. 2016. Transportation Networks for Research. https://github.com/bstabler/TransportationNetworks [Online; accessed 12. Oct. 2024].
[18] Donald R. Jones, Matthias Schonlau, and William J. Welch. 1998. Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization* 13, 4 (1998), 455–492. https://doi.org/10.1023/A:1008306431147
[19] Lee Klieman, Wang Zhang, Vincent L Bernardin Jr, and Vladimir Livshits. 2011. *Estimation and comparison of volume delay functions for arterials and freeway hov and general purpose lanes.* Technical Report.
[20] Elias Koutsoupias and Christos Papadimitriou. 1999. Worst-case equilibria. In *Annual symposium on theoretical aspects of computer science.* Springer, 404–413.
[21] Joel Z. Leibo, Edgar Dué nez Guzmán, Alexander Sasha Vezhnevets, John P. Agapiou, Peter Sunehag, Raphael Koster, Jayd Matyas, Charles Beattie, Igor Mordatch, and Thore Graepel. 2021. Scalable Evaluation of Multi-Agent Reinforcement Learning with Melting Pot. *International Conference on Machine Learning.* https://doi.org/10.48550/arXiv.2107.06857
[22] P Łukasiewicz, Krzysztof Karpio, and Arkadiusz Orłowski. 2012. The models of personal incomes in USA. *Acta Physica Polonica A* 121, 2B (2012).
[23] Reshef Meir and David C Parkes. 2016. When are marginal congestion tolls optimal?. In *ATT@ IJCAI.*
[24] Igal Milchtaich. 1996. Congestion games with player-specific payoff functions. *Games and Economic Behavior* 13, 1 (1996), 111–124.
[25] Uri Nadav and Tim Roughgarden. 2010. The limits of smoothness: A primal-dual framework for price of anarchy bounds. In *International Workshop on Internet and Network Economics.* Springer, 319–326.
[26] John Nash. 1950. Equilibrium Points in n-Person Games. *Proceedings of the National Academy of Sciences* 36, 1 (1950), 48–49.
[27] Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani. 2007. *Algorithmic Game Theory.* Cambridge University Press, New York, NY, USA.
[28] Elinor Ostrom, Roy Gardner, and James Walker. 1994. *Rules, Games, and Common-Pool Resources.* University of Michigan Press. https://press.umich.edu/Books/R/Rules-Games-and-Common-Pool-Resources
[29] Dario Paccagnan, Rahul Chandan, Bryce L Ferguson, and Jason R Marden. 2021. Optimal taxes in atomic congestion games. *ACM Transactions on Economics and Computation (TEAC)* 9, 3 (2021), 1–33.
[30] A. C. Pigou. 1920. *The Economics of Welfare.* Macmillan and Co., London, UK.
[31] Gabriel de O Ramos, Bruno C Da Silva, Roxana Rădulescu, Ana LC Bazzan, and Ann Nowé. 2020. Toll-based reinforcement learning for efficient equilibria in route choice. *The Knowledge Engineering Review* 35 (2020), e8.
[32] Gabriel de O Ramos, Roxana Rădulescu, Ann Nowé, and Anderson R Tavares. 2020. Toll-based learning for minimising congestion under heterogeneous preferences. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems.* 1098–1106.
[33] Charlotte Roman and Paolo Turrini. 2019. Multi-Population Congestion Games With Incomplete Information. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, Sarit Kraus (Ed.). ijcai.org, 565–571. https://doi.org/10.24963/IJCAI.2019/80
[34] Charlotte Roman and Paolo Turrini. 2023. Fighting for Routes: Resource Allocation among Competing Planners in Transportation Networks. *Games* 14, 3 (2023), 37. https://doi.org/10.3390/G14030037
[35] Robert W Rosenthal. 1973. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2 (1973), 65–67.
[36] Tim Roughgarden. 2002. *Selfish Routing.* Ph.D. Dissertation. Cornell University.
[37] Tim Roughgarden. 2005. *Selfish Routing and the Price of Anarchy.* MIT Press, Cambridge, MA.
[38] Tim Roughgarden. 2015. Intrinsic robustness of the price of anarchy. *Journal of the ACM (JACM)* 62, 5 (2015), 1–42.
[39] Tim Roughgarden and Éva Tardos. 2002. How bad is selfish routing? *Journal of the ACM (JACM)* 49, 2 (2002), 236–259.
[40] Hayssam Sbayti, Chung-Cheng Lu, and Hani S Mahmassani. 2007. Efficient implementation of method of successive averages in simulation-based dynamic traffic assignment models for large-scale network applications. *Transportation Research Record* 2029, 1 (2007), 22–30.
[41] Guni Sharon, Josiah P Hanna, Tarun Rambha, Michael W Levin, Michael Albert, Stephen D Boyles, and Peter Stone. 2017. Real-time adaptive tolling scheme for optimized social welfare in traffic networks. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2017).* 828–836.
[42] Yosef Sheffi. 1985. *Urban Transportation Networks: Equilibrium Analysis With Mathematical Programming Methods.* Prentice-Hall.
[43] David Simchi-Levi, Philip Kaminsky, and Edith Simchi-Levi. 2014. *Designing and Managing the Supply Chain: Concepts, Strategies and Case Studies* (3 ed.). McGraw-Hill Education.
[44] National Library Board Singapore. 2024. Electronic Road Pricing system. *National Library Board* (April 2024). https://www.nlb.gov.sg/main/article-detail?cmsuuid=7b3bfea7-2173-4dd9-bd9d-626e8c621ede
[45] MJ Smith. 1979. The marginal cost taxation of a transportation network. *Transportation Research Part B: Methodological* 13, 3 (1979), 237–242.
[46] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. 2012. Practical Bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems.* 2951–2959.
[47] Fernando Stefanello and Ana LC Bazzan. 2016. *Traffic Assignment Problem-Extending Braess Paradox.* Technical Report. Universidade Federal do Rio Grande do Sul, Porto Alegre .... https://www.overleaf.com/read/xqkrzrsfxcyv
[48] J. G. Wardrop. 1952. Some theoretical aspects of road traffic research. In *Proceedings of the Institution of Civil Engineers, Part II*, Vol. 1. 325–378.