

Incentive-Compatible Robust Line Planning^{*}

Apostolos Bessas^{1,3}, Spyros Kontogiannis^{1,2}, and Christos Zaroliagis^{1,3}

¹ R.A. Computer Technology Institute, N. Kazantzaki Str., Patras University Campus, 26500 Patras, Greece

² Computer Science Department, University of Ioannina, 45110 Ioannina, Greece

³ Department of Computer Engineering and Informatics, University of Patras, 26500 Patras, Greece

mpessas@ceid.upatras.gr, kontog@cs.uoi.gr, zaro@ceid.upatras.gr

Abstract. The problem of *robust line planning* requests for a set of origin-destination paths (lines) along with their frequencies in an underlying railway network infrastructure, which are robust to fluctuations of real-time parameters of the solution. In this work, we investigate a variant of robust line planning stemming from recent regulations in the railway sector that introduce competition and free railway markets, and set up a new application scenario: there is a (potentially large) number of *line operators* that have their lines fixed and operate as competing entities issuing frequency requests, while the management of the infrastructure itself remains the responsibility of a single entity, the *network operator*. The line operators are typically unwilling to reveal their true incentives, while the network operator strives to ensure a fair (or socially optimal) usage of the infrastructure, e.g., by maximizing the (unknown to him) aggregate incentives of the line operators.

By investigating a resource allocation mechanism (originally developed in the context of communication networks), we show that a socially optimal solution can be accomplished in certain situations via an anonymous incentive-compatible pricing scheme for the usage of the shared resources that is *robust* against the unknown incentives and the changes in the demands of the entities. This brings up a new notion of robustness, which we call *incentive-compatible robustness*, that considers as robustness of the system its tolerance to the entities' unknown incentives and elasticity of demands, aiming at an eventual stabilization to an equilibrium point that is as close as possible to the social optimum.

1 Introduction

Problem Setting. An important phase in the strategic planning process of a railway (or any public transportation) company is to establish a suitable *line plan*, i.e., to determine the routes of trains that serve the customers. In the *line planning* problem, we are given a network $G = (V, L)$ (usually referred to as

^{*} This work was partially supported by the Future and Emerging Technologies Unit of EC, under contracts no. FP6-021235-2 (FP6 IST/FET Open/Project ARRIVAL), and no. ICT-215270 (FP7 ICT/FET Proactive/Project FRONTS).

the public transportation network), where the node set V represents the set of stations (including important junctions of railway tracks) and the edge set L represents the direct connections or links (of railway tracks) between elements of V . A line p is a path in G . The *frequency* of a line p is a rational number indicating how often service to customers is provided along p within the planning period considered. For an edge $\ell \in L$, the *edge frequency* f_ℓ is the sum of the frequencies of the lines containing ℓ and is upper bounded by the *capacity* c_ℓ of ℓ , i.e., a maximum edge frequency established for safety reasons (measured as the maximum number of trains per day). The goal of the line planning problem is to provide the final set of lines offered by the public transportation company, along with their frequencies (also known as the *line concept*). Typically, a *line pool* is also provided, i.e., a set of potential lines among which the final set of lines will be decided. In certain cases, there may be *multiple line pools* representing the availability of the network infrastructure at different time slots or zones. This is due to variations in customer traffic (e.g., rush-hour pool, late evening pool, night pool), maintenance (some part of the network at a specific time zone may be unavailable), dependencies between lines (e.g., the choice of a high-speed line may affect the choice of lines for other trains), etc.

The line planning problem has been mostly studied under two main approaches (see e.g., [7,10]). In the *cost-oriented* approach, the goal is to minimize the costs of the public transportation company, under the constraint that all customers can be transported. In the *customer-oriented* approach, the goal is to maximize the aggregate level of satisfaction for the customers (e.g., maximize the number of customers with direct connections, minimize the maximum number of intermediate changes of a single customer, or minimize the traveling time of the customers). A recent approach aims at minimizing the travel times over all customers including penalties for the transfers needed [18,20].

The aforementioned approaches do not take into account certain fluctuations of input parameters; for instance, due to disruptions to daily operations (e.g., delays), or due to fluctuating customer demands. This aspect introduces the so-called *robust line planning* problem: Provide a set of lines along with their frequencies, which are robust to fluctuations of input parameters. Very recently, a game theoretic approach for robust line planning was presented in [19]. In that model, the lines act as players and the strategies of the players correspond to line frequencies. Each player aims to minimize the expected delay of her own line. The delay depends on the traffic load and hence on the frequencies of all lines in the network. The objective is to provide lines *and* their frequencies, that are robust against delays. This is pursued by distributing the traffic load evenly over the network (respecting edge capacities) such that the probability of delays in the system is as small as possible.

In this work, we investigate a different perspective of robust line planning stemming from recent regulations in the railway sector (at least within Europe) that introduce competition and free railway markets, and set up a new application scenario: there is a (possibly large) number of *line operators* (LOPs in short) that should operate as commercial organizations, while the management

of the network remains the responsibility of a single (typically governmental) entity; we shall refer to the latter as the **network operator** (NOP in short). Under this framework, LOPs act as competing entities for the exploitation of shared goods and are (possibly) unwilling to reveal their actual level-of-satisfaction (or utility) functions that determine their true incentives. Nevertheless, the NOP would like to ensure the maximum possible level of satisfaction of these competing entities, e.g., by maximizing the (unknown due to privacy) aggregate levels of satisfaction. This would establish a notion of a socially optimal solution, which could also be seen as a fair solution, in the sense that the average level of satisfaction is maximized. Additionally, the NOP should ensure that the operational costs of the whole system are covered by a fair cost sharing scheme announced to the competing entities. This implies that a (possibly anonymous) pricing scheme for the usage of the shared resources should be adopted, that is also *robust* against changes in the demands of the LOPs. That is, we consider as *robustness* of the system its tolerance to the entities' unknown incentives and elasticity of demand requests, and the eventual stabilization at an equilibrium point that is as close as possible to the social optimum.

Contribution. In this paper, motivated by rate allocation in communication networks [13,14,21], we explore the aforementioned rationale by considering the case where the (selfishly motivated) LOPs request frequencies over a pool of already fixed line routes (one route per LOP). In particular, we investigate the resource allocation mechanism proposed in the pioneering work of Kelly [13]. Rather than requesting end-to-end frequencies, the LOPs offer bids, which they (dynamically) update for buying frequencies. Each LOP has a utility function determining her level of satisfaction that is *private*; i.e., she is not willing to reveal it to the NOP or her competitors, due to her competitive nature. The NOP announces an (anonymous) resource pricing scheme, which indirectly implies an allocation of frequencies to the LOPs, given their own bids.

Our first contribution is to show that for the case of a single line pool an adaptation of Kelly's approach [13] provides a distributed, dynamic, (LOP) bidding and (resource) price updating scheme, whose equilibrium point is the unknown social optimum – assuming strict concavity and monotonicity of the private utility functions. All dynamic updates of bids and prices can be done at the LOP and resource level respectively, based only on *local information* that concerns the particular LOP or resource. The key assumption is that the LOPs can control only a negligible amount of frequency along a single line, compared to its total frequency.

Our second contribution is a (non-trivial) extension of the approach for a single pool to the case of multiple line pools. By assuming that the NOP can periodically exploit a whole set of (disjointly operating) line pools and that each LOP may be interested in different lines from different pools, we show that there exists a globally convergent, dynamic, (LOP) bidding and (resource) price updating scheme, whose equilibrium point is the unknown social optimum. The NOP, similarly to the single pool case, uses a mechanism (a feasible frequency allocation rule and an anonymous resource pricing scheme) aiming to maximize

the aggregate level of satisfaction of LOPs. The NOP, contrary to the single pool case, decides on how to divide the whole infrastructure among the different pools so that the resource capacity constraints are preserved, aiming (again) to achieve the optimal welfare value.

Our third contribution is an experimental study on a discrete variant of the distributed, dynamic scheme developed for the single pool case on both synthetic and real-world data. We note that in both single and multiple line pool cases the proposed mechanisms assure *market clearance*, i.e., the entire network infrastructure (capacities) is eventually used by the LOPs and all the budget afforded by the LOPs is actually spent.

Our solution is robust against the imperfect knowledge imposed by the private (unknown) utility functions and the arbitrary (dynamically updated) bids, since the proposed protocol enforces convergence to an equilibrium which is the social optimum. Our approach introduces a new notion of robustness, which we call *incentive-compatible robustness*, that is complementary to the notion of *recoverable robustness* introduced in [2,16,17]. The latter appears to be more suitable in the context of railway optimization, as opposed to the classical notion of robustness within robust optimization; see [2,16,17] for a detailed discussion on the subject and for the limitations of the classical approach as suggested in [4].

Recoverable robustness is about computing solutions that are robust against a limited set of scenarios (that determine the imperfection of information) and which can be made feasible (recovered) by a limited effort. One starts from a feasible solution x of an optimization problem, which a particular scenario s , that introduces imperfect knowledge (i.e., by adding more constraints), may turn to infeasible. The goal is to have at our disposal a recovery algorithm A that takes x and turns it to a feasible solution under s (i.e., under the new set of constraints). In other words, in recoverable robustness there is uncertainty about the feasibility space: imperfect information generates infeasibility and one strives to (re-)establish feasibility.

Incentive-compatible robustness is about computing an incentive-compatible recovery scheme for achieving robustness (interpreted as convergence to optimality). By an incentive-compatible scheme, we mean that the players act (update their bids, in our application) in a selfish manner during the convergence sequence. In this context, the feasibility space is known and incomplete information refers to complete lack of information about the optimization problem, due to the unknown utility functions. The goal is to define an incentive-compatible (pricing) scheme so that the players converge (recover) to the system's optimum. In other words, in incentive-compatible robustness there is uncertainty about the objectives: feasibility is guaranteed, since imperfect knowledge does not introduce new constraints, and one strives to achieve optimality, exploiting the selfish nature of the players.

Note that incentive-compatible robustness is different from the concept of game-theoretic robustness as developed in [1]. The approach in [1] is a centralized, deterministic paradigm to uncertainty in strategic games, mainly in the

flavor of the Bertsimas and Sim approach [4] to robust LP optimization. We elaborate on the differences in Section 5.

Related Work and Approaches. Related to our work is that of Borndörfer et al. [5] that considers the allocation of slots in railway networks. That work considers the improvement of existing schedules of lines and frequencies, by reconsidering the allocation of (scarce) bundles of slots (i.e., lines with given frequencies in our own terminology) that have positive synergies with each other. The remaining schedule is assumed to remain intact, so that the resulting optimization problem is solvable. Initially, the involved users (LOPs) make some bids and consequently a centralized optimization problem is solved to determine the changes in the allocation of these slots so as to maximize the welfare of the whole system. This approach is different from ours in the following points: (i) It assumes no incentive-compatibility for the involved users and the eventual allocation is determined by a centralized scheduler. In our case, there is a simple pricing policy per resource (track), which is a priori known to all the players, and the winner is determined by the players' bids. The selfish behavior of the LOPs (in our case) is not only taken into account, but also exploited by the system in order to assure convergence to the social optimum of the whole network. (ii) The approach in [5] makes some local improvements *in hope* of improving the whole system, but does not exclude being trapped at some local optimum, which may be far away from the social optimum of the system. Our proposed scheme *provably* converges towards the social optimum, even if changes in the parameters of the game (e.g., in the players' secret utilities) change in the future. (iii) In [5], it is required that a centralized optimization problem is solved (considering the data regarding the whole network) and its solution is enforced in the current schedule. In our work (at least for the single-pool case) there is no need for global knowledge of the whole network. Each player dynamically adapts her bids according to her own (secret) utility and the aggregate cost she faces along her own path.

Another way to tackle the problem we consider here would be through the celebrated Vickrey-Clarke-Groves (VCG) class of mechanisms [6,12,22]. Such a mechanism would guarantee in our application scenario the existence of a dominant strategy equilibrium [11] in which the allocation of frequencies to the LOPs indeed maximizes the sum of their utilities, by encouraging LOPs to reveal their utility functions truthfully. Unfortunately, implementing VCG mechanisms is generally a very complex task, not only due to the huge size of the centralized optimization problem to be solved, but also due to the dynamic nature of an evolving market in which the parameters of the problem (railway infrastructure, number of participating LOPs, LOP utilities, etc) change over time. For instance, each LOP may vary her own utility function over time, due to changes in her own data, or her way of thinking. It may even be the case that some LOPs are unable to fully express their utility function, simply because they do not know it (for example, determining the parameters of such a function could be a hard optimization problem to solve by itself). On the other hand, it seems more plausible for a LOP to determine whether she would like to marginally increase

or decrease her budget for claiming usage of the network, given the current situation she faces in the system. Therefore, we opt to follow Kelly’s approach [13] by deploying a decentralized, dynamic updating scheme for the LOP budgets and the resource prices, whose updating rules are *simple* and are based on local information as much as possible that will monotonically converge to the socially optimal solution. Of course, the price to pay is some loss of efficiency w.r.t. how fast we converge to the optimum. Nevertheless, the self-stabilizing nature of our scheme, even to also dynamically changing optima (e.g., due to changes in the system infrastructure), is a very strong characteristic that compensates this efficiency drawback, compared to the adoption of a static, centralized VCG mechanism.

Structure. The rest of this paper is organized as follows. In Section 2, we provide the set up of our modeling, and present the adaptation of Kelly’s approach [13] to the case of a single line pool by showing that the social optimum can be found by a polynomial-time computable mechanism, and by providing a decentralized, dynamic scheme that globally converges to the social optimum. To adapt and cast Kelly’s approach to our problem setting, we recapitulate and re-prove certain results both for the sake of completeness and for providing the road-map for the extension to the multiple pools case. In Section 3, we provide our approach for the case of multiple line pools. We show that the social optimum can be found by a polynomial-time computable mechanism, and we provide a dynamic scheme (for implementing this mechanism) that globally converges to the social optimum. In Section 4, we present an experimental evaluation of our decentralized dynamic scheme for the single pool case, using synthetic and real-world data. In Section 5, we discuss incentive-compatible robustness and its comparison to other notions of robustness. We conclude in Section 6. A preliminary version of this work appeared in [15].

2 Single Line Pool: Modeling and Solution Approach

In this section, we present the modeling and the solution approach for the robust line planning problem we consider, for the case where a single line pool is provided. The development in this section is based on an adaptation of Kelly’s resource allocation mechanism [13] (originally proposed within the context of communication networks for allocating network capacity to potential users). To adapt and cast Kelly’s approach to our problem setting, we recapitulate and re-prove certain results for the sake of completeness and also to make this section the road-map for the development of our approach for the case of multiple line pools in Section 3.

Suppose that a set P of LOPs behave as competing service providers, willing to offer regular (train) line routes to the end users of a railway public transportation system. The NOP provides the (aforementioned) public transportation network $G = (V, L)$. The node set V represents train stations and junctions, while the edge set L (with each edge corresponding to a railway track establishing direct




Fig. 1. A simple network with two distinct lines, and the corresponding routing matrix

connection for some pair of nodes in G) is the set of shared *resources* of the network. These resources are assumed to be subject to (fixed) capacity constraints, described by the capacity vector $\mathbf{c} = (c_\ell)_{\ell \in L} > 0$; for an edge $\ell \in L$, c_ℓ represents the maximum number of trains passing through ℓ over a whole time period (e.g., a day).

There is a fixed pool of line routes (i.e., origin–destination paths) that the LOPs are willing to use, and we assume that there is one line route per LOP¹. This pool is represented by a **routing matrix** $\mathbf{R} \in \{0, 1\}^{|L| \times |P|}$, in which each row $\mathbf{R}_{\ell, \star}$ corresponds to a different edge $\ell \in L$, and each column $\mathbf{R}_{\star, p}$ corresponds (actually, is the characteristic vector of) the line route of a distinct LOP $p \in P$. Fig. 1 demonstrates a network with two distinct lines and the corresponding routing matrix. Each LOP $p \in P$ claims a *frequency* of trains that she wishes to route over her path, $\mathbf{R}_{\star, p}$, given that no edge capacity constraint is violated in the network. A utility function $U_p : \mathbb{R} \mapsto \mathbb{R}$ determines the level of satisfaction of $p \in P$ for committing an end-to-end frequency $x_p > 0$ along her route $\mathbf{R}_{\star, p}$, for the purposes of her clients. These utility functions are assumed to be strictly increasing, strictly concave, nonnegative real functions of the end-to-end frequency x_p allocated to $p \in P$. It is also assumed that these functions are *private*: each LOP is not willing to reveal it either to the NOP, or to her competitors, due to her competitive nature.

The NOP is only interested in having a socially optimal (fair) solution. This is usually interpreted as maximizing the aggregate satisfaction of the LOPs. Therefore, the social welfare objective is considered to be the maximization of the aggregate utilities of the LOPs, subject to the capacity constraints. That is, the NOP is interested in the solution of the following convex optimization² problem:

$$\boxed{\text{SOCIAL (SC)}} \quad \max \left\{ \sum_{p \in P} U_p(x_p) : \mathbf{R}\mathbf{x} \leq \mathbf{c}; \mathbf{x} \geq \mathbf{0} \right\}$$

¹ We can always enforce this assumption by considering a LOP with more than one routes as different LOPs distinguished by the specific route.

² We make the tacit assumption that convex optimization refers to minimizing a convex function f , which is equivalent to maximizing the concave function $-f$.

Since all (private) utility functions are strictly concave and the feasible space is also convex, **[SC]** has a *unique* optimal solution, which is called the **social optimum**. To solve **[SC]** directly, the NOP, apart from the inherent difficulty in centrally solving (even convex) optimization programs of the size of a railway network, faces the additional obstacle of *not knowing* the exact shape of the objective function. Moreover, there exist some operational costs that have to be split among the LOPs who use the infrastructure, and this has to be done also in a fair way: each LOP should only be charged for the usage of the resources standing on her own route. In addition, the per-unit cost for using a line should be independent of a LOP's identity (i.e., we would like to have an *anonymous* pricing scheme for using the resources). But of course, this cost depends on the aggregate frequency induced by all the LOPs in each of these edges, due to the congestion effect. Indeed, it would be desirable for the NOP to be able to exploit the announcement of a pricing scheme not only for covering these operational costs, but also in such a way that a fair solution for all the LOPs is induced, despite the fact that there is no global knowledge of the exact utility functions of the LOPs.

In this work, we explore the possibilities of having such a frequency allocation and resource pricing mechanism. We would like this mechanism to depend only on the information affecting either a specific LOP (e.g., the amount of money she is willing to afford) or a specific resource (e.g., the aggregate frequency induced by the LOPs' demands on this resource), but as we shall see this is not always possible.

As for the LOPs (the players), each of them is interested in selfishly utilizing her own payoff, which is determined by the difference of the private utility value minus the operational cost that the NOP charges her for claiming an amount of frequency along her own route. The strategy space of a LOP is to claim (via bidding) the value of the frequency she is willing to buy, subject to the global capacity constraints (for all the players). It is mentioned here that this linear combination of the private utility and the cost share is not a real restriction, as there is no restriction for the shape of the utility function, other than the strict concavity and the monotonicity, which are quite natural assumptions.

2.1 Social Optimum – Tractability

Our first goal is to demonstrate that, despite the hidden utilities of the LOPs, it is indeed possible for the NOP to induce the social optimum, i.e., the solution of **[SC]**, as the result of the LOPs' selfish behavior. In order to study the effect of the selfish behavior in this setting, we consider the following **Frequency Game** in Line Planning:

- Each player $p \in P$ is a LOP, whose strategy is to choose a line frequency over her (already fixed) route $\mathbf{R}_{*,p}$ connecting her own origin–destination pair (s_p, t_p) of stations/stops.

- The strategy space for all the players is the set of feasible flows from origin to destination nodes, so that the edge capacity constraints are preserved. That is, the strategy space of the game is the set of vectors $\{\mathbf{x} \in \mathbb{R}_{\geq 0}^{|P|} : \mathbf{R}\mathbf{x} \leq \mathbf{c}\}$.
- Each player's payoff is determined both by the value of the private utility function $U_p(x_p)$ (for having a frequency of x_p over her route) and the operational cost $C_p(\mathbf{x})$ she has to pay along her own route, due to the required frequency vector \mathbf{x} induced by all the players in the network. Hence, player p 's individual payoff is defined as: $IP_p(x_p, \mathbf{x}_{-p}) = U_p(x_p) - C_p(x_p, \mathbf{x}_{-p})$, where \mathbf{x}_{-p} is the frequency vector for all the players but for player p . Therefore, the sole goal of player $p \in P$ is to choose her frequency so as to maximize her individual payoff:

$$\boxed{\text{USER}} \quad \max \{IP(x_p, \mathbf{x}_{-p}) = U_p(x_p) - C_p(x_p, \mathbf{x}_{-p}) : x_p \geq 0\}$$

- We consider as shared resources the capacities of the available network edges, for which the LOPs compete with each other.

As we shall explain later, we actually view this game as a mechanism–design instance, in which the NOP is the game regulator that receives the players' bids (for buying frequencies) and consequently decides both a feasible allocation of frequencies to the players and the payments that they have to provide. In this setting, the players can only affect their own eventual choice (allocation of a frequency) *indirectly* via bidding, rather than freely setting her own frequency along her route. In order to receive a (hopefully) higher frequency, a player may only offer a higher bid.

Describing the Social Optimum. Due to our assumption on the convexity of $\boxed{\text{SC}}$, we know that a frequency vector $\hat{\mathbf{x}}$ is the social optimum if there exists a vector of Lagrange Multipliers $\hat{\boldsymbol{\lambda}} = (\hat{\lambda}_\ell)_{\ell \in L}$ satisfying the following Karush-Kuhn-Tucker (KKT) conditions (see e.g., [3, Chap. 3]):

KKT-SOCIAL (KKT-SC)

$$U'_p(\hat{x}_p) = \hat{\boldsymbol{\lambda}}^T \cdot \mathbf{R}_{\star,p}, \quad \forall p \in P, \quad (1)$$

$$\hat{\lambda}_\ell (c_\ell - \mathbf{R}_{\ell,\star} \cdot \hat{\mathbf{x}}) = 0, \quad \forall \ell \in L, \quad (2)$$

$$\mathbf{R}_{\ell,\star} \cdot \hat{\mathbf{x}} \leq c_\ell, \quad \forall \ell \in L, \quad (3)$$

$$\hat{\boldsymbol{\lambda}}, \hat{\mathbf{x}} \geq \mathbf{0} \quad (4)$$

Of course, the problem with the $\boxed{\text{KKT-SC}}$ system is that the utility functions (and hence their derivatives) are unknown to the system. The question is whether there exists a way for the network designer to enforce the optimal solution of $\boxed{\text{SC}}$, also described in $\boxed{\text{KKT-SC}}$, without demanding this knowledge. The answer to this is *partially affirmative*, and this is by exploiting the selfish nature of the LOPs as we shall see shortly.

Setting the Right Pricing Scheme for the Players. In order to allow usage of his resources (the capacities of the edges in the network), the NOP has to define a pricing scheme that will (at least) pay back the operational costs of the edges. This scheme should be *anonymous*, in the sense that all the LOPs willing to use a given edge, will have to pay the same per-unit-of-frequency price for using it. But these prices may vary for different edges, depending on the popularity and the availability of each edge.

For the moment let us assume that we already know the optimal Lagrange Multipliers, $(\hat{\lambda}_\ell)_{\ell \in L}$ of **KKT-SC**. Interpreting these values as the per-unit-of-frequency prices of the resources, we have a pricing scheme for the frequency induced by the LOPs to their own routes. Each LOP pays exactly for the marginal cost of her own frequency at the resources she uses in her route. That is,

$$\forall p \in P, C_p(x_p, \mathbf{x}_{-p}) = \hat{\mu}_p \cdot x_p$$

where $\hat{\mu}_p \equiv \sum_{\ell \in L: R_{\ell,p}=1} \hat{\lambda}_\ell = \hat{\boldsymbol{\lambda}}^T \mathbf{R}_{*,p}$ is the per-unit price for committing one unit of frequency along the route $\mathbf{R}_{*,p}$ of player $p \in P$.

One should mention here that there is indeed an indirect effect of the other players' congestion in the marginal cost of each player, despite the fact that this seems to be only linear in her own frequency. This is because the scalar $\hat{\mu}_p$ actually depends on the optimal primal-dual pair $(\hat{\mathbf{x}}, \hat{\boldsymbol{\lambda}})$.

We next assume that the players are actually controlling only *negligible* amounts of frequencies compared to the aggregate ones³. Then, their effect in the total congestion (and therefore in the values of the marginal prices) is also negligible. This implies that the players consider the per-unit-prices they face to be constant, even if this is actually affected by the frequency vector as well. In such a case we say that the players are **price takers**, i.e., they accept the prices without anticipating to have an effect on them by their own strategy. In such a case each player solves the following optimization problem:

$$\boxed{\text{USER-I}} \quad \max \{U_p(x_p) - \hat{\mu}_p x_p : x_p \geq 0\}$$

Due to the convexity of **USER-I**, $\tilde{x}_p \geq 0$ is an optimal solution if $U'_p(\tilde{x}_p) = \hat{\mu}_p$. That is, each player (selfishly) tries to satisfy her own part of equations (1) in **KKT-SC**. Of course, we still have to deal with the crucial problem that the optimal Lagrange Multipliers (that define the marginal prices for the users) cannot be directly computed, due to both the size of **SC** and the lack of knowledge of the private utility functions, in the framework of railway optimization.

To tackle this situation, we transform the Frequency Game to a mechanism design instance, in order to have a more active participation of the NOP, as the game regulator. In particular, we consider the following two-level scenario for dynamically setting per-unit prices of the edges and frequencies of the selfish players. Initially each LOP $p \in P$ announces a bid $w_p \geq 0$ concerning the total amount of money she is willing to pay for buying frequency along her own route.

³ For the considered application scenario, this is not unrealistic.

The exact amount of frequency that she will eventually buy, depends on the per-unit price that will be announced by the NOP, and is not yet known to her (nevertheless, it will be the case that, given the other players' bids, any LOP $p \in P$ may only increase her assigned frequency by raising, unilaterally, her own bid). Consequently, the NOP considers the following optimization problem, whose Lagrange Multipliers define the per-unit prices of the edges:

$$\boxed{\text{NETWORK (NET)}} \quad \max \left\{ \sum_{p \in P} w_p \cdot \log(x_p) : \mathbf{R}\mathbf{x} \leq \mathbf{c}; \mathbf{x} \geq \mathbf{0} \right\}$$

That is, the NOP considers that the private utility $U_p(x_p)$ is substituted by the (also strictly concave and increasing, for any given bid vector $\mathbf{w} \in \mathbb{R}_{\geq 0}^{|P|}$) function $w_p \log(x_p)$. The choice of this function, along with the selfishness of the LOPs, allows us to obtain a *convex program with linear inequalities*, whose KKT conditions are very similar (except for the first line) to those of $\boxed{\text{KKT-SC}}$:

$$\text{KKT-NETWORK (KKT-NET)}$$

$$\frac{w_p}{\bar{x}_p} = \bar{\lambda}^T \cdot \mathbf{R}_{\star, p}, \quad \forall p \in P, \quad (5)$$

$$\bar{\lambda}_\ell (c_\ell - \mathbf{R}_{\ell, \star} \cdot \bar{\mathbf{x}}) = 0, \quad \forall \ell \in L, \quad (6)$$

$$\mathbf{R}_{\ell, \star} \cdot \bar{\mathbf{x}} \leq c_\ell, \quad \forall \ell \in L, \quad (7)$$

$$\bar{\lambda}, \bar{\mathbf{x}} \geq \mathbf{0} \quad (8)$$

By $(\bar{\mathbf{x}}, \bar{\lambda})$, we denote the optimal primal-dual pair of $\boxed{\text{KKT-NET}}$. Observe that the only difference between $\boxed{\text{KKT-NET}}$ and $\boxed{\text{KKT-SC}}$ concerns the (left-hand side of) equations (5) and (1), respectively. But we shall demonstrate now that the selfish (and price taking) behavior of the LOPs is enough to make this difference vanish. Returning to the LOPs, we initially assumed that they announce some *fixed* bids, and consequently the NOP sets the per-unit prices of the resources. Given the bid vector and the resource prices, it is then easy to determine each LOP's assigned frequency. But the truth is that, since the pricing scheme changes over time, it is in the interest of each LOP to actually vary her own bid over time. Indeed, if the players are assumed to be price takers and act myopically (i.e., without anticipating to affect the prices via their own pricing policy), then they will try to solve the following system, which is parameterized by the instantaneous set of per-unit prices $\boldsymbol{\mu}(t) = (\mu_p(t))_{p \in P}$ (now seen by the LOPs as *constants*) they are charged at time $t \geq 0$:

$$\boxed{\text{USER-II}} \quad \max \left\{ U_p \left(\frac{w_p(t)}{\mu_p(t)} \right) - w_p(t) : w_p(t) \geq 0 \right\}$$

Due to convexity, the optimal solution $\tilde{w}_p(t)$ of the unconstrained optimization program $\boxed{\text{USER-II}}$, will be the bid chosen by player $p \in P$ at time $t \geq 0$, and is given by:

$$\begin{aligned}
\frac{1}{\mu_p(t)} \cdot U'_p \left(\frac{\tilde{w}_p(t)}{\mu_p(t)} \right) &= 1 \Leftrightarrow \\
U'_p(\tilde{x}_p(t)) = U'_p \left(\frac{\tilde{w}_p(t)}{\mu_p(t)} \right) &= \mu_p(t) \Leftrightarrow \\
\tilde{x}_p(t) U'_p(\tilde{x}_p(t)) &= \mu_p(t) \cdot \tilde{x}_p(t) = \tilde{w}_p(t)
\end{aligned}$$

That is, the price taking, myopic players have an incentive to set their bids properly so that $\forall t \geq 0, \forall p \in P, w_p(t) = x_p(t) U'_p(x_p(t))$. This will also hold at the optimal solution of **[NET]**, i.e., $\forall p \in P, \bar{w}_p = \bar{x}_p U'_p(\bar{x}_p)$. But when this is true, it also holds that **[KKT-NET]** and **[KKT-SC]** coincide. That is, the selfish-bidding behavior of the myopic, price taking players, under the pricing scheme $\bar{\lambda}$ determined by the Lagrange Multipliers of **[KKT-NET]**, leads to the optimal solution $(\bar{x}, \bar{\lambda}) = (\hat{x}, \hat{\lambda})$ of **[KKT-SC]**.

The discussion within this section establishes the following result.

Theorem 1. *Consider a transportation network $G = (V, L)$ and a set P of (selfish, price taking) LOPs with hidden utilities, whose lines are determined by a routing matrix $\mathbf{R} \in \{0, 1\}^{|L| \times |P|}$. There exists a polynomial-time computable mechanism (i.e., a pair of a frequency allocation rule and resource pricing scheme) which induces the optimal solution of **[SC]** as a result of the LOPs' selfish behavior.*

2.2 Social Optimum – Dynamic and Decentralized Computation

At this point, one could argue that, in order to solve the (partially determined) convex program **[SC]**, it suffices to determine the proper resource prices by the optimal solution of the (completely determined, and computationally tractable) convex program **[NET]**. The latter can be directly solved and provide the proper Lagrange Multipliers of **[SC]**. However, the huge scale of a railway network optimization instance makes this rationale rather unappealing.

Motivated by the pioneering work of Kelly et al. [13,14] and its excellent simplification and elaboration in [21], we shall try to compute an optimal solution of **[NET]** as the stable point of a system of differential equations that determines the updates of the resource prices, and (consequently) the LOPs' bids. The crucial observation at this point is that it suffices to enforce the resource prices to gradually converge to the optimal price vector $\bar{\lambda}$ provided by **[NET]**, and the “right bids” will follow.

We consider the following dynamic system of differential equations that actually constitutes our decentralized, dynamic algorithm for computing the social optimum.

1. Each resource (edge in the transportation network) is equipped with a dynamically updated charging mechanism, which is the same (per-unit) price for all the LOPs using it. This charging mechanism is updated according to the following system of differential equations:

$$\forall \ell \in L, \quad \dot{\lambda}_\ell(t) = \max\{y_\ell(t) - c_\ell, 0\} \cdot \mathbb{I}_{\{\lambda_\ell(t)=0\}} + (y_\ell(t) - c_\ell) \cdot \mathbb{I}_{\{\lambda_\ell(t)>0\}} \quad (9)$$

where $y_\ell(t) \equiv \sum_{p \in R: R_{\ell,p}=1} x_p(t) = \mathbf{R}_{\ell,\star} \cdot \mathbf{x}(t)$ is the aggregate frequency committed at edge $\ell \in L$ at time $t \geq 0$, and $\mathbb{I}_{\{\mathcal{E}\}}$ is the indicator variable of the truth of a logical expression \mathcal{E} .

2. Each LOP $p \in P$, at any time $t \geq 0$, is charged an instantaneous per-unit price $\mu_p(t) \equiv \sum_{\ell \in L: R_{\ell,p}=1} \lambda_\ell(t) = \boldsymbol{\lambda}(t)^T \cdot \mathbf{R}_{\star,p}$. It solves USER-II to determine $w_p(t)$, and consequently is allocated a frequency $x_p(t) = \frac{w_p(t)}{\mu_p(t)}$.

The system of differential equations (9) is obtained from the well-known approach (see e.g., [13,21]) that considers the Lagrange Multipliers of an optimization problem as the (per unit) prices of the resources corresponding to the constraints represented by each Lagrange Multiplier. Therefore, the above system has the following intuitive interpretation. For each resource ℓ that currently has a zero price, the tendency is to increase the price only if this resource is over-used (i.e., the aggregate frequency exceeds the capacity of the resource). When a resource has positive price, then the tendency is either to increase or reduce this price, depending on whether its current frequency exceeds or is below the capacity of the resource, respectively. Thus, the only stable situation is when a resource is either under-used and has zero price (since there is no interest in using the residual capacity), or its frequency has already reached its capacity. Observe that the equilibrium of this system of differential equations has $\forall \ell \in L, \bar{y}_\ell \equiv \mathbf{R}_{\ell,\star} \cdot \bar{\mathbf{x}} = c_\ell \quad \vee \quad \bar{\lambda}_\ell = 0$. That is, the complementarity conditions of both KKT-SC and KKT-NET (equations (2) and (6)) are satisfied.

Step 2 above implies that at equilibrium player p , given its commitment on spending w_p for buying frequency, is allocated a frequency of $\bar{x}_p = \frac{w_p}{\bar{\mu}_p}$. From this we deduce that at equilibrium also the equations (5) of KKT-NET are satisfied.

We are now ready to prove the following.

Theorem 2. *The above defined dynamic system of resource-pricing and LOP-bid-updating differential equations ensures monotonic convergence to the social optimum of NET from any initial point of resource prices and LOP bids.*

Proof. The above system of differential equations is a distributed algorithm, in which each LOP reacts to signals she gets about the aggregate frequency along her route. These signals are the per-unit prices $\mu_p(t)$ that the LOP gets from the NOP at any time.

The question is whether the above system converges at all. This is indeed true, if we assume that the routing matrix \mathbf{R} has full rank. This assures that given a set $\boldsymbol{\lambda}(t) = (\lambda_\ell(t))_{\ell \in L}$ of instantaneous per-unit prices at the resources, the set $\boldsymbol{\mu}(t) = (\mu_p(t))_{p \in P}$ of per-unit prices for the LOPs, that is computed as the solution of the system $\boldsymbol{\mu}(t) = \mathbf{R}^T \cdot \boldsymbol{\lambda}(t)$, is *unique*. Using a proper Lyapunov function argument, it can be shown (cf. [21, Chapter 3]) that this dynamic (and

distributively implemented) pricing scheme, for *fixed* player bids $(w_p)_{p \in P}$, is stable and converges to the optimal solution $(\bar{\mathbf{x}}, \bar{\boldsymbol{\lambda}})$ of NET.

In particular, consider the Lyapunov function $V(\boldsymbol{\lambda}(t)) = \frac{1}{2}(\boldsymbol{\lambda}(t) - \bar{\boldsymbol{\lambda}})^T(\boldsymbol{\lambda}(t) - \bar{\boldsymbol{\lambda}})$. To show stability of our scheme, it suffices to show that $dV(\boldsymbol{\lambda}(t))/dt \leq 0$. Then we have:

$$\begin{aligned}
& \frac{dV(\boldsymbol{\lambda}(t))}{dt} \\
&= \sum_{\ell \in L} (\lambda_\ell(t) - \bar{\lambda}_\ell) \cdot \dot{\lambda}_\ell(t) \\
&= \sum_{\ell \in L} (\lambda_\ell(t) - \bar{\lambda}_\ell) \cdot [\max\{y_\ell(t) - c_\ell, 0\} \cdot \mathbb{I}_{\{\lambda_\ell(t)=0\}} + (y_\ell(t) - c_\ell) \cdot \mathbb{I}_{\{\lambda_\ell(t)>0\}}] \\
&\leq \sum_{\ell \in L} (\lambda_\ell(t) - \bar{\lambda}_\ell) \cdot (y_\ell(t) - c_\ell) \\
&= \sum_{\ell \in L} (\lambda_\ell(t) - \bar{\lambda}_\ell) \cdot [(y_\ell(t) - \bar{y}_\ell) + (\bar{y}_\ell - c_\ell)] \\
&\leq \sum_{\ell \in L} (\lambda_\ell(t) - \bar{\lambda}_\ell) \cdot (y_\ell(t) - \bar{y}_\ell) \\
&= \sum_{\ell \in L} (\lambda_\ell(t) - \bar{\lambda}_\ell) \cdot \mathbf{R}_{\ell, \star} \cdot (\mathbf{x}(t) - \bar{\mathbf{x}}) \\
&= \sum_{p \in P} (\mu_p(t) - \bar{\mu}_p) \cdot (x_p(t) - \bar{x}_p) \\
&= \sum_{p \in P} \left(\frac{w_p}{x_p(t)} - \frac{w_p}{\bar{x}_p} \right) \cdot (x_p(t) - \bar{x}_p) = \sum_{p \in P} w_p \cdot \left(2 - \frac{x_p(t)}{\bar{x}_p} - \frac{\bar{x}_p}{x_p(t)} \right) \\
&\leq 0
\end{aligned}$$

The first inequality holds because: $\forall \ell \in L$, (i) if $\lambda_\ell(t) > 0$ then $\dot{\lambda}_\ell(t) = y_\ell - c_\ell$; (ii) if $\lambda_\ell(t) = 0$ then $\max\{y_\ell - c_\ell, 0\} \geq 0$ and $\lambda_\ell(t) - \bar{\lambda}_\ell = -\bar{\lambda}_\ell \leq 0$. Therefore, for $\lambda_\ell(t) = 0$ it holds that $(\lambda_\ell(t) - \bar{\lambda}_\ell) \max\{y_\ell(t) - c_\ell, 0\} = -\bar{\lambda}_\ell \max\{y_\ell(t) - c_\ell, 0\} \leq 0$. But so long as $\lambda(t) = 0$, it holds that the total frequency $y_\ell(t)$ is at most as large as the capacity c_ℓ (otherwise the price for this resource would have raised earlier). That is, $0 \leq -\bar{\lambda}_\ell(y_\ell(t) - c_\ell)$. The second inequality holds because at equilibrium no aggregate frequency \bar{y}_ℓ can exceed the capacity c_ℓ of the resource, and $\bar{\lambda}_\ell(\bar{y}_\ell - c_\ell) = 0$. The third inequality holds because $\forall z > 0, z + \frac{1}{z} \geq 2 \Rightarrow 2 - z - \frac{1}{z} \leq 0$. We have also exploited the facts that $\forall t \geq 0, \mathbf{y}(t) = \mathbf{R} \cdot \mathbf{x}(t)$ and $\boldsymbol{\mu}(t) = \boldsymbol{\lambda}(t)^T \cdot \mathbf{R}$. \square

3 Multiple Line Pools: Modeling and Solution Approach

In this section we extend the freedom of both the NOP and the LOPs. For the NOP we assume that he can now periodically exploit a whole set K of

(disjointly operating) line pools, rather than a single line pool, to serve the LOPs' connection requests. It is up to the NOP how to split a whole operational period of the railway infrastructure among the different pools, so that (in overall, for the whole period) the resource capacity constraints are not violated. A first assumption that we make at this point, is that the NOP divides the usage of the whole infrastructure (rather than each resource separately) among the pools. This is because we envision the line pools to be implemented, not concurrently, but in disjoint time intervals (e.g., via some sort of time division multiplexing), and also to concern different characteristics of the involved lines (e.g., high-speed pool, regular-speed pool, local-trains pool, rush-hour pool, night-shift pool, etc.). The capacity of each resource (as in the single pool case) refers to its usage (number of trains) over the whole time period we consider (e.g., a day), and if a particular pool consumes (say) 50% of the whole infrastructure, then this implies that for all the lines in this pool, each resource may exploit at most half of its capacity.

As for the LOPs, they can now even claim different lines from different pools. In accordance with the single pool case, each LOP may express interest in at most one line per pool. For simplicity we assume that each LOP is interested for *exactly* one line per pool, adding dummy origin-destination pairs connected with an edge of zero capacity, for every LOP that has no interest in some pool. Technically, our analysis would allow even the case where a LOP expresses interest for lines with different origin-destination pair (in different pools). Nevertheless, we assume that each LOP $p \in P$ has a *single* (strictly concave, as before) utility function $U_p : \mathbb{R}_{\geq 0} \mapsto \mathbb{R}_{\geq 0}$, which depends on the *aggregate frequency* x_p that she gets from all the pools in which she is involved. In order to be in compliance with this assumption, we consider the case where each LOP expresses interest for different lines (at most one per pool) over the same origin-destination pair. Of course, in reality, different ways of dividing the same aggregate frequency x_p among the various lines, could make a huge difference for the particular LOP, but we do not account for this effect in this paper.

In analogy with the single pool case, each pool $k \in K$ is represented by its own routing matrix $\mathbf{R}(k) \in \{0, 1\}^{|L| \times |P|}$. The frequency (number of trains over one time period) granted to the LOP $p \in P$ within the pool $k \in K$, is indicated by a nonnegative real variable $x_{p,k}$. The aggregate frequency that p gets is then $x_p = \sum_{k \in K} x_{p,k}$. The LOPs still try to have (indirect, via bidding) control over the aggregate end-to-end frequency x_p they get by the NOP along all their lines, from all the possible pools that may be of use by the NOP. It is up to the NOP to decide how to divide the whole railway infrastructure among the different pools, so that the resource capacity constraints are preserved, the goal being to achieve the optimal social welfare value. That is, the NOP now directly participates in the optimization problem via the variables $f_k : k \in K$ indicating the proportion of capacity that each pool consumes from every resource, over the whole time period we study. We will say that the NOP or the vector \mathbf{f} *completely divides* the infrastructure, if $\sum_{k \in K} f_k = 1$.

The NOP is now interested in solving the following optimization problem:

MULTI-SOCIAL (MSC)

$$\begin{aligned}
 & \text{maximize} \quad \sum_{p \in P} U_p(x_p) = \sum_{p \in P} U_p \left(\sum_{k \in K} x_{p,k} \right) \\
 & \text{s.t.} \quad \forall (\ell, k) \in L \times K, \quad \sum_{p \in P} R_{\ell,p}(k) \cdot x_{p,k} \leq c_\ell \cdot f_k \\
 & \quad \quad \quad \sum_{k \in K} f_k \leq 1 \\
 & \quad \quad \quad \mathbf{x}, \mathbf{f} \geq \mathbf{0}
 \end{aligned}$$

Once more, this is a strictly convex optimization problem (due to the strict concavity of the LOP utility functions, and the linearity of the feasible space), whose objective function is unknown to the NOP. We shall explain in this section how we can tackle this issue. The overall idea is that we can handle the multiple pools case as an expanded single pool case. We have $|K|$ replicas of the same railway infrastructure, and $|K|$ replicas $p_1, \dots, p_{|K|}$ of the same LOP $p \in P$, each being interested only in a single line (the one of interest to LOP p in the corresponding pool). Each LOP offers her bid w_p for buying aggregate frequency x_p . The NOP determines the proportions of railway infrastructure that are committed per pool. Exactly the same proportions are used (by the NOP) for splitting the LOPs' bids among the various pools.

3.1 Multi Social Optimum – Tractability

We start by an observation that exploits the economic interpretation of the Lagrange Multipliers of MSC.

Lemma 1. *Assuming that all the players adopt strictly increasing, concave utility functions, if the resource prices are determined by the vector $\hat{\mathbf{A}}$ of optimal Lagrange Multipliers of the resource constraints in MSC, then the following are true: (i) Each LOP is indifferent of the way her aggregate frequency is split among different pools. (ii) All the pools have the same (weighted) aggregate cost. (iii) The NOP completely divides the whole railway infrastructure among the different pools. Facts (i) and (iii) also hold even when the NOP fixes a priori the vector of proportions, for the corresponding optimal solution.*

Proof. Let \mathbf{A} be the vector of Lagrange Multipliers for the resource capacity constraints, and ζ the Lagrange Multiplier concerning the constraint for the capacity proportions per pool. The Lagrangian function is the following:

$$\begin{aligned}
 & L(\mathbf{x}, \mathbf{f}, \mathbf{A}, \zeta) \\
 &= \sum_{p \in P} U_p(x_p) - \sum_{\ell \in L} \sum_{k \in K} \Lambda_{\ell,k} \cdot \left[\sum_{p \in P} R_{\ell,p}(k) \cdot x_{p,k} - c_\ell \cdot f_k \right] - \zeta \left[\sum_{k \in K} f_k - 1 \right] \\
 &= \sum_{p \in P} \left[U_p(x_p) - \sum_{k \in K} x_{p,k} \left(\sum_{\ell \in L} \Lambda_{\ell,k} \cdot R_{\ell,p}(k) \right) \right] + \sum_{k \in K} f_k \cdot [\mathbf{c}^T \mathbf{A}_{\star,k} - \zeta] + \zeta \\
 &= \sum_{p \in P} \left[U_p(x_p) - \sum_{k \in K} x_{p,k} \cdot \mu_{p,k}(\mathbf{A}) \right] + \sum_{k \in K} f_k \cdot [\mathbf{c}^T \mathbf{A}_{\star,k} - \zeta] + \zeta
 \end{aligned}$$

where we set $\mu_{p,k}(\mathbf{A}) \equiv \sum_{\ell \in L} \Lambda_{\ell,k} \cdot R_{\ell,p}(k)$. If we consider $\Lambda_{\ell,k}$ as the per-unit-of-frequency price of resource ℓ with respect to the pool k , then $\mu_{p,k}(\mathbf{A})$ is again the end-to-end per-unit cost that p has to pay in pool k . The strict concavity of the utility functions, along with the linearity of the feasible space, assure that we indeed have to solve a strictly convex optimization problem, which has a *unique* optimal solution, $(\hat{\mathbf{x}}, \hat{\mathbf{f}})$. The system of KKT conditions of MSC describing this solution, is the following:

KKT-MULTI-SOCIAL (KKT-MSC)

$$\begin{aligned}
 U'_p(\hat{x}_p) &= \hat{\mu}_{p,k} \equiv \mu_{p,k}(\hat{\mathbf{A}}), \quad \forall (p, k) \in P \times K & (10) \\
 \mathbf{c}^T \cdot \hat{\mathbf{A}}_{\star,k} \equiv \sum_{\ell \in L} \hat{\Lambda}_{\ell,k} \cdot c_\ell &= \hat{\zeta}, \quad \forall k \in K & (11) \\
 \hat{\Lambda}_{\ell,k} \left[\sum_{p \in P} R_{\ell,p}(k) \hat{x}_{p,k} - c_\ell \hat{f}_k \right] &= 0, \quad \forall (\ell, k) \in L \times K & (12) \\
 \hat{\zeta} \cdot \left(\sum_{k \in K} \hat{f}_k - 1 \right) &= 0 & (13) \\
 \sum_{p \in P} R_{\ell,p}(k) \cdot \hat{x}_{p,k} &\leq c_\ell \cdot \hat{f}_k, \quad \forall (\ell, k) \in L \times K & (14) \\
 \sum_{k \in K} \hat{f}_k &\leq 1 & (15) \\
 \hat{\mathbf{x}} \geq \mathbf{0}, \quad \hat{\mathbf{f}} \geq \mathbf{0}, \quad \hat{\mathbf{A}} \geq \mathbf{0}, \quad \hat{\zeta} \geq 0 & & (16)
 \end{aligned}$$

Observe that from KKT-MSC we can easily deduce the following facts with respect to the optimal solution:

- (i) By equation (10), each LOP faces exactly the same end-to-end per-unit-of-frequency cost $\hat{\mu}_p = \hat{\mu}_{p,k} = U'_p(\hat{x}_p)$, $\forall k \in K$, along any line of interest to her. This justifies the fact that p is not really concerned about how the aggregate frequency $\hat{x}_p = \sum_{k \in K} \hat{x}_{p,k}$ is distributed among the different lines of interest to her. The pricing scheme induced by $\hat{\mathbf{A}}$ makes all these lines look of equal importance.

- (ii) By equation (11), in the optimal solution all the pools have the same (weighted) aggregate per-unit-of-frequency cost, equal to $\hat{\zeta}$, if we interpret the resource capacities as their weights.
- (iii) Due to equation (13), unless this optimal (identical for all pools) aggregate per-unit-of-frequency cost is zero, it holds that the resource capacities are totally distributed among the distinct pools: **if $\hat{\zeta} > 0$ then $\sum_{k \in K} \hat{f}_k = 1$** . But if we consider the non-trivial case in which the network has positive resource capacities, then clearly (due to strict concavity of the utilities) some of the resource prices will have to be positive. This directly implies the positivity of $\hat{\zeta}$.

Observe finally that facts (i) and (iii) still hold for the unique optimal primal-dual solution $(\bar{\mathbf{x}}, \bar{\mathbf{A}})$, in the case that the NOP fixes a particular vector of proportions $\bar{\mathbf{f}}$ (that completely divides the infrastructure among the pools), which is then considered to be constant both in **MSC** and in **KKT-MSC**. Of course, this time we cannot assure the same aggregate cost per pool. \square

To tackle the issue of limited information, we consider again (as in the single pool case) a mechanism in which the LOPs are initially required to propose their own bids for buying frequencies, and consequently the NOP somehow determines the resource prices and the frequencies granted to the LOPs (per pool) according to this pricing scheme and their bids. In particular, we construct a new (strictly convex) program, by substituting the (unknown) LOP utility functions with the pseudo-utilities $w_p \log(x_p)$, where $w_p \geq 0$ is the (fixed) amount of money that $p \in P$ is willing to spend for buying frequency (across all pools). This program is the following:

MULTI-NETWORK (MNET)

$$\begin{aligned}
 & \text{maximize} \quad \sum_{p \in P} w_p \log(x_p) = \sum_{p \in P} w_p \log \left(\sum_{k \in K} x_{p,k} \right) \\
 & \text{s.t. } \forall (\ell, k) \in L \times K, \quad \sum_{p \in P} R_{\ell,p}(k) \cdot x_{p,k} \leq c_\ell \cdot f_k \\
 & \quad \sum_{k \in K} f_k \leq 1 \\
 & \quad \mathbf{x}, \mathbf{f} \geq \mathbf{0}
 \end{aligned}$$

Therefore, for any (fixed) vector $\mathbf{w} = (w_p)_{p \in P}$ of LOP bids, the NOP computes (in polynomial time) the optimal solution of **MNET**, considering as resource prices the optimal Lagrange Multipliers of the resource constraints in **KKT-MNET**, which is almost identical to **KKT-MSC**, except for the equations (10), which are substituted by the following:

$$\frac{w_p}{\bar{x}_p} = \bar{\mu}_{p,k} \equiv \mu_{p,k}(\bar{\mathbf{A}}), \quad \forall (p, k) \in P \times K \quad (17)$$

The properties of Lemma 1 for the optimal solution of **MSC** also hold for the optimal solution (for any fixed bid vector) of **MNET**, even when the NOP decides to fix a particular vector of proportions $\bar{\mathbf{f}}$. In particular, for $(\bar{\mathbf{x}}, \bar{\mathbf{f}}, \bar{\mathbf{A}}, \bar{\zeta})$ it holds that each LOP faces exactly the same cost $\bar{\mu}_p = \bar{\mu}_{p,k}$ in every pool $k \in K$, this time equal to $\frac{w_p}{\bar{x}_p}$ rather than $U'_p(\bar{x}_p)$. Moreover, if the capacity proportions are also variables (rather than constants), then all the pools have the same (weighted) aggregate cost $\bar{\zeta}$.

Consequently, the NOP announces all the optimal frequencies $\bar{x}_{p,k}$ for each LOP $p \in P$ and each pool $k \in K$, for which we know that $\bar{x}_p = \sum_{k \in K} x_{p,k} = \frac{w_p}{\bar{\mu}_p}$. Based once more on our assumption that the LOPs are price taking selfish entities, as in the single pool case, we exploit the fact that each LOP will choose her own bid \bar{w}_p as the optimal solution of **USER-II**, which assures then that $U'_p(\bar{x}_p) = \bar{\mu}_p = \bar{\mu}_{p,k}$, $\forall k \in K$, exactly as required in **KKT-MSC**. This holds for any vector of resource prices that assures for every LOP exactly the same per-unit cost in all the pools, and in particular, for the optimal Lagrange Multipliers vector $\bar{\mathbf{A}}$ of **KKT-MNET**. Therefore, we again conclude that at equilibrium the LOPs will choose their bids in such a way that the optimal solution of **KKT-MSC** coincides with the optimal solution of **KKT-MNET**, for any fixed vector of capacity proportions, $\bar{\mathbf{f}}$. The above discussion thus leads to the following conclusion.

Theorem 3. *Consider a transportation network $G = (V, L)$ and a set P of (selfish, price taking) LOPs with private utility functions of the aggregate frequency assigned to them. Each LOP expresses interest for at most one line in each pool from a set K of pools. There exists a polynomial-time computable mechanism (i.e., a pair of a frequency allocation rule and resource pricing scheme) which induces (as the only equilibrium point) the optimal solution with respect to the aggregate utility value, as a result of the LOPs' selfish behavior.*

Once more, this tractable mechanism, which is based on the solvability of **MNET**, is totally centralized and rather inconvenient for a dynamically changing (over time), large-scale railway system. Therefore, in the next subsection we shall devise an almost-localized analogue to the single pool case that is based on a system of updating rules for the resource prices (determined by each resource), the LOP bids, and the vector of proportions (determined by the NOP), which converges to this optimal solution of **MSC**.

3.2 Multi Social Optimum – Dynamic and Decentralized Computation

Our first argument has to do with the independence of the adopted pricing scheme from the way that the NOP chooses to split the railway infrastructure among the different pools. In particular, as we shall shortly explain (Lemma 2), for any *fixed* vector of capacity proportions $\bar{\mathbf{f}}$ that the NOP chooses, the optimal value of the corresponding dual program of **MSC** exclusively depends on the choice of the vector \mathbf{A} of resource prices. The dynamic updating system that we

shall later propose will exploit exactly this fact and let (in a continuous fashion) the resource prices gradually converge to the optimal price vector (which then forces the LOP bids and the corresponding frequencies to get the right values), for the currently adopted vector of capacity proportions. This vector of proportions will be updated *periodically* by the NOP, only after the system has stabilized to that optimal point (of optimal prices and bids).

Lemma 2. *For any (fixed) vector \mathbf{f} of capacity proportions that completely divides the network infrastructure among the pools, the optimal value of MSC exclusively depends on the optimal vector $\mathbf{\Lambda}$ of per-unit prices for the resources.*

Proof. Using the Lagrange function previously defined, the dual problem of MSC is the following:

$$\text{DUAL-MSC} \quad \max \{D(\mathbf{\Lambda}, \zeta) : \forall \ell \in L, \forall k \in K, \Lambda_{\ell,k} \geq 0; \zeta \geq 0\}$$

where:

$$\begin{aligned} D(\mathbf{\Lambda}, \zeta) &= \max_{\mathbf{x}, \mathbf{f} \geq \mathbf{0}} \{L(\mathbf{x}, \mathbf{f}, \mathbf{\Lambda}, \zeta) : \mathbf{x}, \mathbf{f} \geq \mathbf{0}\} \\ &= \max_{\mathbf{x}, \mathbf{f} \geq \mathbf{0}} \left\{ \sum_{p \in P} \left[U_p(x_p) - \sum_{k \in K} x_{p,k} \sum_{\ell \in L} \Lambda_{\ell,k} R_{\ell,p}(k) \right] + \sum_{k \in K} f_k \left[\sum_{\ell \in L} \Lambda_{\ell,k} c_{\ell} - \zeta \right] + \zeta \right\} \\ &= \max_{\mathbf{x} \geq \mathbf{0}} \left\{ \sum_{p \in P} \left[U_p(x_p) - \sum_{k \in K} x_{p,k} \mu_{p,k}(\mathbf{\Lambda}) \right] \right\} + \max_{\mathbf{f} \geq \mathbf{0}} \left\{ \sum_{k \in K} f_k \left[\sum_{\ell \in L} \Lambda_{\ell,k} c_{\ell} - \zeta \right] \right\} + \zeta \end{aligned}$$

Observe that the dual objective $D(\mathbf{\Lambda}, \zeta)$ can be split in two parts. The first part:

$$F(\mathbf{\Lambda}) = \max_{\mathbf{x} \geq \mathbf{0}} \left\{ \sum_{p \in P} \left[U_p(x_p) - \sum_{k \in K} x_{p,k} \cdot \mu_{p,k}(\mathbf{\Lambda}) \right] \right\}$$

is a maximization problem similar to the one already dealt with in the single pool case (i.e., for $|K| = 1$) of the previous single-pool case. Its value is a function of the resource prices, and the vector of proportions has no involvement at this point. The only difference from the single pool case, is that we now have distinct LOP frequencies, as well as LOP end-to-end costs, per pool. But this technical issue can be tackled by a proper choice of the dynamic updating system, as we shall see later. The second part of $D(\mathbf{\Lambda}, \zeta)$ is the following:

$$\begin{aligned} G(\mathbf{\Lambda}, \zeta) &= \max_{\mathbf{f} \geq \mathbf{0}} \left\{ \sum_{k \in K} f_k \cdot \left[\sum_{\ell \in L} \Lambda_{\ell,k} \cdot c_{\ell} - \zeta \right] \right\} + \zeta \\ &= \max_{\mathbf{f} \geq \mathbf{0}} \left\{ \sum_{k \in K} f_k \cdot (c^T \mathbf{\Lambda}_{\star,k}) + \zeta \cdot \left(1 - \sum_{k \in K} f_k \right) \right\} \end{aligned}$$

Recall that at global optimality (when we consider the capacity proportions as variables), the term $\zeta \cdot (1 - \sum_{k \in K} f_k)$ has zero contribution in $G(\mathbf{\Lambda}, \zeta)$. But this

also holds for any solution in which the NOP chooses some vector \mathbf{f} of capacity proportions that sums up to 1. Additionally, we have already seen that this is indeed the case for the optimal vector of capacity proportions as well, as was explained in Lemma 1, fact (iii). Therefore, the optimal choice $\hat{\mathbf{f}}$ of capacity proportions can be seen as a *probability distribution* that assigns positive mass only to pools of maximum aggregate price (according to $\mathbf{\Lambda}$). We demand this restriction explicitly from $G(\mathbf{\Lambda}, \zeta)$:

$$\begin{aligned} G(\mathbf{\Lambda}, \zeta) &= \max_{\mathbf{1}^T \mathbf{f} = 1; \mathbf{f} \geq \mathbf{0}} \left\{ \sum_{k \in K} f_k \cdot (\mathbf{c}^T \cdot \mathbf{\Lambda}_{\star, k}) \right\} = \max_{k \in K} \{ \mathbf{c}^T \cdot \mathbf{\Lambda}_{\star, k} \} \\ &= \min \{ z : z \cdot \mathbf{1}^T \geq \mathbf{c}^T \cdot \mathbf{\Lambda} \} \end{aligned}$$

That is, $G(\mathbf{\Lambda}, \zeta)$ simply calculates the maximum (rather than the average, indicated by ζ) aggregate (per-unit) cost among the pools, which only depends on the given resource pricing vector $\mathbf{\Lambda}$. \square

Lemma 2 is crucial in deriving a dynamic algorithm that computes the social optimum, in analogy with the one derived for the single pool case. In particular, Lemma 2 and the framework of the single pool case suggest the following dynamic algorithm, whose high-level description is as follows.

1. Each resource continuously updates its own (anonymous) per-unit-of-frequency price.
2. Each LOP updates her offer (bid) for claiming frequency, only when the resource prices (and thus her own per-unit costs in the pools) have stabilized.
3. The NOP updates *periodically* the vector of capacity proportions of the railway infrastructure given to the different line pools, only when both the resource prices and the LOP bids have stabilized.

In particular, assume that at some time $t \geq 0$ we have the following situation:

- $\mathbf{\Lambda}(t)$ is the vector of current resource prices. $\forall p \in P, \forall k \in K, \mu_{p,k}(t) = \sum_{\ell \in L} R_{\ell,p}(k) \cdot \Lambda_{\ell,k}$ is the per-unit cost of player p at pool k , while $\mu_p(t) = \frac{1}{|K|} \sum_{k \in K} \mu_{p,k}(t)$ is the *average* per-unit cost of p over all the pools.
- $\mathbf{w}(t) = (w_p(t))_{p \in P}$ is the vector of the LOPs' current bids.
- $\mathbf{f}(t) = (f_k(t))_{k \in K}$ is the current vector of proportions of resource capacities of the railway infrastructure to each of the pools (as determined by NOP). We always assure the **invariant** that the entire railway infrastructure is provided to the pools: $\mathbf{1}^T \cdot \mathbf{f}(t) = 1$.
- We calculate the frequencies that each LOP gets per pool as follows. We split each LOP's bid $w_p(t)$ among the different pools according to the vector of capacity proportions. Then each LOP buys the corresponding frequency, given her bid and the per-unit cost for this LOP at each particular pool: $\forall (p, k) \in P \times K, x_{p,k}(t) = \frac{f_k(t) \cdot w_p(t)}{\mu_{p,k}(t)}$. The aggregate frequency of the LOP $p \in P$ is obviously $x_p(t) = \sum_{k \in K} x_{p,k}(t)$.

- The resource frequencies at time t are then calculated as follows: $\forall(\ell, k) \in L \times K$, $y_{\ell,k}(t) = \sum_{p \in P} R_{\ell,p}(k) \cdot x_{p,k}(t)$ and $y_{\ell}(t) = \sum_{k \in K} y_{\ell,k}(t)$.

We assume that the resource price updating scheme operates continuously, the LOP bidding updating scheme applies only when the resource prices have stabilized, and finally the updating of the capacity proportions (conducted by the NOP) is carried out only when both the resource prices and the LOP bids have stabilized. This is explained as follows. Each resource continuously updates its price as a function of the aggregate frequency over it (in each pool), and this is instantly known local information to the resource. As for the LOPs, they would like to update their bids only when there is a clear picture of what should be paid in each pool. This can only happen when the resource prices have stabilized. Additionally, each LOP has to gather the pricing information along the lines she uses, which is somehow local information (only refers to resources actually used by the LOP) but not instantly available. Finally the NOP wishes to: (i) let the whole situation with the resource prices and LOP bids stabilize before it intervenes to determine the new capacity proportions of infrastructure, and (ii) avoid too frequent changes in the capacity proportions, since this updating scheme does not depend only on local information (either on each LOP, or on each resource) but on the aggregate costs of all the pools, as we shall see shortly. Therefore, the NOP prefers this update to happen only occasionally, in order to be able to amortize its heavy cost over a large period of time.

Let's now see the exact shape of the dynamic protocol at time $t \geq 0$.

Resource Price Updating. $\forall t \geq 0$, the resource per-unit prices are updated according to the following differential equation: $\forall \ell \in L, \forall k \in K$,

$$\dot{A}_{\ell,k}(t) = \max\{0, y_{\ell,k}(t) - c_{\ell}f_k\} \cdot \mathbb{I}_{\{\Lambda_{\ell,k}(t)=0\}} + (y_{\ell,k}(t) - c_{\ell}f_k) \cdot \mathbb{I}_{\{\Lambda_{\ell,k}(t)>0\}}$$

LOP Bid Updating. Assuming now that the LOPs are selfish entities, their (instantaneous) bids are chosen as the solutions of the analogue of USER-II (per LOP), which is the following:

$$\text{MUSER-II} \quad \text{maximize} \left\{ U_p \left(\sum_{k \in K} \frac{f_k w_p}{\bar{\mu}_{p,k}} \right) - w_p : w_p \geq 0 \right\}$$

where, $\forall k \in K$, $\bar{\mu}_{p,k} = \bar{\mu}_p$ is the common per-unit cost that the LOP $p \in P$ faces in each pool, as soon as the resource prices stabilize. The optimality condition of MUSER-II is now that

$$\begin{aligned} U'_p \left(\sum_{k \in K} \frac{f_k w_p}{\bar{\mu}_{p,k}} \right) \cdot \sum_{k \in K} \frac{f_k}{\bar{\mu}_{p,k}} &= 1 \\ \Leftrightarrow U'_p(x_p) = U'_p \left(\sum_{k \in K} \frac{f_k w_p}{\bar{\mu}_{p,k}} \right) &= \left(\sum_{k \in K} \frac{f_k}{\bar{\mu}_{p,k}} \right)^{-1} = \bar{\mu}_p \end{aligned}$$

Capacity Proportions Updating. After the LOPs have stabilized the bids $(w_p(t))_{p \in P}$ and the resources have updated their per-unit prices in each pool $(\Lambda_{\ell,k}(t))_{\ell \in L, k \in K}$, the NOP sets $\zeta(t)$ to the *average* price of a pool:

$$\zeta(t) = \frac{1}{|K|} \sum_{k \in K} \mathbf{c}^T \cdot \mathbf{A}_{\star,k}(t) \quad (18)$$

Then the NOP updates the proportions of the railway infrastructure granted to each of the pools, so that pools exceeding the current average cost $\zeta(t)$ increase their proportion (in hope of decreasing their weighted cost), while pools that are cheaper than the average price slightly decrease their proportion (recall that in the optimal solution all the pools have exactly the same weighted aggregate cost). The proportions are updated according to the following system of differential equations:

$$\forall k \in K, \dot{f}_k(t) = \max \{0, \mathbf{c}^T \cdot \mathbf{A}_{\star,k}(t) - \zeta(t)\} \quad (19)$$

It should be noted here that, in order for the vector $\mathbf{f}(t+1)$ of capacity proportions to sum up to 1, we must divide the resulting vector of new proportions by a proper scaling factor $\phi(t) > 1$ (since the expensive pools increased their proportions, while the cheap pools kept their old proportion, according to the proposed derivative in equation (19)).

The resource updating in this differential system assures the validity of equations (12) at equilibrium, for any fixed vector \mathbf{f} of capacity proportions provided by the NOP, and any fixed bid vector \mathbf{w} provided by the LOPs. Moreover, if we assume that the LOPs are price taking and myopic entities, the LOP bid updating again leads us to the validity of equations (10). We shall now prove the convergence to the optimal resource prices, with respect to any given vector of capacity proportions, and any given vector of LOP bids.

Lemma 3. *For any choice of fixed bid vector $\bar{\mathbf{w}} = (w_p)_{p \in P}$ offered by the LOPs, and any fixed vector of proportions $\bar{\mathbf{f}} = (f_k)_{k \in K}$ determined by the NOP, the resource price updating scheme makes the resource prices converge to the corresponding optimal vector $\bar{\mathbf{\Lambda}}$ (for these particular given bids and proportions).*

Proof. We use again the Lyapunov function $V(\mathbf{\Lambda}(t)) = \frac{1}{2} \cdot (\mathbf{\Lambda}(t) - \bar{\mathbf{\Lambda}})^T \cdot (\mathbf{\Lambda}(t) - \bar{\mathbf{\Lambda}})$, we can once more prove convergence to the optimal resource prices, $\bar{\mathbf{\Lambda}}$, for any fixed vector of LOP bids, $\bar{\mathbf{w}}$ and any vector of pool proportions, $\bar{\mathbf{f}}$ (determined by the NOP):

$$\begin{aligned} \frac{dV(\mathbf{\Lambda}(t))}{dt} &= \sum_{\ell \in L} \sum_{k \in K} (\Lambda_{\ell,k}(t) - \bar{\Lambda}_{\ell,k}) \cdot \dot{\Lambda}_{\ell,k}(t) \\ &= \sum_{\ell \in L} \sum_{k \in K} (\Lambda_{\ell,k}(t) - \bar{\Lambda}_{\ell,k}) \cdot [\max \{0, y_{\ell,k}(t) - c_{\ell} \bar{f}_k\} \cdot \mathbb{I}_{\{\Lambda_{\ell,k}(t)=0\}} \\ &\quad + (y_{\ell,k}(t) - c_{\ell} \bar{f}_k) \cdot \mathbb{I}_{\{\Lambda_{\ell,k}(t)>0\}}] \end{aligned}$$

$$\begin{aligned}
&\leq \sum_{\ell \in L} \sum_{k \in K} (\Lambda_{\ell,k}(t) - \bar{\Lambda}_{\ell,k}) \cdot [y_{\ell,k}(t) - c_{\ell} \bar{f}_k] \\
&= \sum_{\ell \in L} \sum_{k \in K} (\Lambda_{\ell,k}(t) - \bar{\Lambda}_{\ell,k}) \cdot [y_{\ell,k}(t) - \bar{y}_{\ell,k} + \bar{y}_{\ell,k} - c_{\ell} \bar{f}_k] \\
&\leq \sum_{\ell \in L} \sum_{k \in K} (\Lambda_{\ell,k}(t) - \bar{\Lambda}_{\ell,k}) \cdot [y_{\ell,k}(t) - \bar{y}_{\ell,k}] \\
&= \sum_{\ell \in L} \sum_{k \in K} (\Lambda_{\ell,k}(t) - \bar{\Lambda}_{\ell,k}) \cdot \sum_{p \in P} R_{\ell,p}(k) \cdot [x_{p,k}(t) - \bar{x}_{p,k}] \\
&= \sum_{p \in P} \sum_{k \in K} [x_{p,k}(t) - \bar{x}_{p,k}] \cdot (\mu_{p,k}(t) - \bar{\mu}_{p,k}) \\
&= \sum_{p \in P} \sum_{k \in K} [x_{p,k}(t) - \bar{x}_{p,k}] \cdot \left(\frac{\bar{f}_k \bar{w}_p}{x_{p,k}(t)} - \frac{\bar{f}_k \bar{w}_p}{\bar{x}_{p,k}} \right) \\
&= \sum_{p \in P} \sum_{k \in K} \bar{f}_k \bar{w}_p \cdot \left[2 - \frac{x_{p,k}(t)}{\bar{x}_{p,k}} - \frac{\bar{x}_{p,k}}{x_{p,k}(t)} \right] \\
&\leq 0
\end{aligned}$$

The first inequality holds trivially for each $(\ell, k) : \Lambda_{\ell,k}(t) > 0$, but also holds when for $(\ell, k) : \Lambda_{\ell,k}(t) = 0$, because then either $y_{\ell,k}(t) - c_{\ell} \bar{f}_k \geq 0$ and

$$(\Lambda_{\ell,k}(t) - \bar{\Lambda}_{\ell,k}) \cdot \max \{0, y_{\ell,k}(t) - c_{\ell} \bar{f}_k\} = -\bar{\Lambda}_{\ell,k} \cdot [y_{\ell,k}(t) - c_{\ell} \bar{f}_k]$$

or $y_{\ell,k}(t) - c_{\ell} \bar{f}_k < 0$ and then:

$$(\Lambda_{\ell,k}(t) - \bar{\Lambda}_{\ell,k}) \cdot \max \{0, y_{\ell,k}(t) - c_{\ell} \bar{f}_k\} = -\bar{\Lambda}_{\ell,k} \cdot 0 < -\bar{\Lambda}_{\ell,k} \cdot [y_{\ell,k}(t) - c_{\ell} \bar{f}_k]$$

The second inequality holds because for the optimal vector $\bar{\Lambda}$ (for the given vectors $\bar{\mathbf{w}}$ and $\bar{\mathbf{f}}$) it holds that $\sum_{\ell \in L} \sum_{k \in K} \bar{\Lambda}_{\ell,k} (\bar{y}_{\ell,k} - c_{\ell} \bar{f}_k) = 0$ (cf. equation (12), which is also a KKT condition for **MNET**). The third inequality holds again because $\forall z > 0, 2 - z - \frac{1}{z} \leq 0$. \square

Of course, when the resource prices and LOP bids have stabilized, we still cannot be sure that we have reached the optimal solution of **MNET**, because we cannot guarantee for the time being that all the pools have the same (weighted) aggregate cost, as required by equation (11). Due to the strict concavity of **MSC**, we know that the current optimal value (for the given proportions) of its dual is strictly less than the globally optimal value (with respect to the optimal proportions). This is because we obviously have not chosen yet the optimal vector of proportions. But, as it was shown in Lemma 2, the optimal value of **DUAL-MSC** exclusively depends on the vector of resource prices. Therefore we also know that we do not have the optimal resource prices as well. At this point exactly, the NOP intervenes with the capacity proportions updating procedure, which increases (in a continuous fashion) the proportions of pools that are more expensive than the current value of the average cost $\zeta(t)$. That is, the NOP chooses to increase the pool-capacity proportions to already expensive pools

(therefore allowing, at the next optimal point, lower aggregate costs for them) and decreases the proportion of infrastructure for cheap pools (which can afford slightly larger aggregate costs). This way we get closer to the optimal vector of capacity proportions in $\boxed{\text{MSC}}$, since the vector of aggregate pool costs will now become smoother. Consequently, the new optimal value of $\boxed{\text{DUAL-MS}}$ (with respect to the new capacity proportions) will strictly increase due to the intervention of the NOP, because the dominant parameter for it is the vector of resource prices.

Eventually, by Lemmata 2 and 3, we shall converge to an equilibrium point of the whole system in which equation (10) is guaranteed by the selfish, price taking behavior of the LOPs, equation (11) is assured by the NOP, equation (12) is assured by the resource price updating scheme, and equation (13) is assured by our invariant on the vector of capacity proportions. This is exactly the optimal solution of both $\boxed{\text{KKT-MS}}$ and $\boxed{\text{MNET}}$, as required. The following theorem summarizes the previous discussion.

Theorem 4. *The aforementioned dynamic system of resource-pricing, LOP-bid-updating and capacity-proportions-updating differential equations ensures monotonic convergence to the social optimum of $\boxed{\text{MNET}}$ from any initial point of resource prices, LOP bids and proportions of capacities for the pools.*

4 Implementation and Experimental Evaluation

In this section, we present the implementation of a discrete version of our decentralized algorithm for the single pool case and its experimental evaluation on synthetic and real-world data.

4.1 The Algorithm

We have implemented a discrete version of our distributed algorithm given in Section 2.2, and which is provided below. Parameter b determines the desired accuracy at equilibrium, B_ℓ is an upper bound on the value of $\lambda_\ell(t)$, ε_ℓ represents the interval upon which $\dot{\lambda}_\ell(t)$ is defined and which gradually reduces via the UPDATE routine, and the boolean variable S_ℓ is used to determine the termination condition of the repeat-until loop. The algorithm is as follows.

1. INITIALIZATION ($t = 0$).
 - (a) For all $p \in P$: $\{ w_p(0) = 1; x_p(0) = \min_{\ell \in p} \{c_\ell\}; \}$
 - (b) For all $\ell \in L$: $\{ \lambda_\ell(0) = 0; \varepsilon_\ell = 1; \delta_\ell = 10^{-b}; B_\ell = +\infty; S_\ell = \text{FALSE}; \}$
2. REPEAT FOR $t > 0$
 - (a) For all $\ell \in L$:
 - i. $y_\ell(t) = \sum_{\ell \in p} x_p(t-1)$;
 - ii. $\alpha_\ell(t) = y_\ell(t) - c_\ell$;
 - iii. $\dot{\lambda}_\ell(t) = \max\{0, \alpha_\ell(t)\} - \min\{\frac{\lambda_\ell(t)}{2}, \max\{0, -\alpha_\ell(t)\}\}$;
 - iv. **if** $\lambda_\ell(t-1) = 0 \wedge \max\{0, \alpha_\ell(t)\} < \delta_\ell$ **then** $S_\ell = \text{TRUE}$;

```

v. if  $\lambda_\ell(t-1) > 0 \wedge |\alpha_\ell(t)| < \delta_\ell$  then  $S_\ell = \text{TRUE}$ ;
vi.  $\text{UPDATE}(\varepsilon_\ell)$ ;
vii.  $\lambda_\ell(t) = \lambda_\ell(t-1) + \varepsilon_\ell \dot{\lambda}_\ell(t)$ ;
(b) if  $\bigcap_{\ell \in L} S_\ell = \text{TRUE}$  then BREAK;
(c) For all  $p \in P$ :
    i.  $\mu_p(t) = \sum_{\ell \in p} \lambda_\ell(t)$ ;
    ii. Solve USER-II to determine  $w_p(t)$ ;
    iii.  $x_p(t) = \frac{w_p(t)}{\mu_p(t)}$ ;
(d)  $t = t + 1$ ;
UNTIL TRUE

```

The routine for updating ε_ℓ is as follows.

$\text{UPDATE}(\varepsilon_\ell)$

1. **if** $\dot{\lambda}_\ell(t) < 0 \wedge \dot{\lambda}_\ell(t-1) > 0$ **then**
 if $\lambda_\ell(t-1) < B_\ell$ **then** $B_\ell = \lambda_\ell(t-1)$ **else** $\varepsilon_\ell = \varepsilon_\ell/2$;
2. **RETURN** ε_ℓ ;

4.2 Experimental Setup

The algorithm was implemented in C++ using the GNU g++ compiler (version 4.3.2) with -O2 optimization level, and the LEDA C++ library (version 6.2). Our experiments were performed on a computer having an Intel Core 2 Duo Processor clocked at 2.00GHz (T7300 model) and a total of 2GB RAM.

Synthetic data consisted of grid graphs having a vertical dimension of 3 and a horizontal dimension ranging from 120 to 36000 (i.e., the graph sizes range from 120×3 to 36000×3). The edge capacity was set to 10. In these graphs, we define three paths in a way that they have a fair amount of edges in common. By considering the graph nodes as points in the plane, we define three directions:

UP: The next edge of the path is headed upwards.

RIGHT: The next edge of the path is headed to the right.

DOWN: The next edge of the path is headed downwards.

We consider two families of three paths. The first family is defined as follows; see Figure 2. All paths start at node $(0, 1)$. The first path follows the **RIGHT** direction until it can no longer proceed. We will call this the *middle path* (yellow-grey path in Fig. 2). The second path (red path in Fig. 2) first goes **RIGHT**, then **UP**, then **RIGHT**, then **DOWN**, and then continues the same pattern until it can no longer proceed. The third path (green path in Fig. 2) first goes **RIGHT**, then **DOWN**, then **RIGHT**, then **UP**, and then continues the same pattern until it can no longer proceed. Observe that all paths share the odd edges of the middle path (i.e., the edges $(2i, 2i+1)$, $i = 0, 1, 2, \dots$).




Fig. 2. Deterministic paths on grids

The second family of paths is defined as follows. All paths start at node $(0, 1)$ and go RIGHT. Then, at node $(2i + 1, 1)$, $i = 0, 1, 2, \dots$, each path makes a (uniformly) random decision on whether it will go UP, RIGHT, or DOWN. If the random choice is to go UP, then it follows the pattern RIGHT, DOWN, RIGHT, reaching the next node where it will make a new random decision. If the random choice is to go DOWN, then it follows the pattern RIGHT, UP, RIGHT, reaching the next node where it will make a new random decision. If the random choice is to go RIGHT, then it follows the pattern RIGHT, RIGHT, reaching the next node where it will make a new random decision. These choices ensure that all three paths share the edges $(2i, 2i + 1)$, $i = 0, 1, 2, \dots$

Real-world data concern parts of the German railway network (concerning mainly intercity train connections). We have considered three instances with 280 (354), 296 (393), and 319 (452) nodes (edges), and a single line pool of varying size. The capacity of the edges varied from 8 up to 16.

For both synthetic and real-world data, we used the function $U_p(x) = a\sqrt{x}$ as utility function of all LOPs $p \in P$, where a is a constant ($a \geq 10^4$).

4.3 Experimental Results

We start with the experimental results on our synthetic data. Figure 3 shows the number of iterations required by our distributed algorithm to converge to the social optimum in the grid graphs of sizes 12000×3 to 36000×3 . The top diagram does this for the first family of paths (deterministically defined paths), while the bottom diagram does it for the second family of paths (that include random choices at certain nodes). We observe that despite the graph size, the algorithm converges quite fast to the optimal solution. It is worth mentioning that the real execution time never exceeded 1.5 minutes.

We now turn to the real-world graphs. Figures 4 and 5 show the number of iterations required by our distributed algorithm to converge to the social optimum with respect to the size of the line pool, which varies from 100 to 2000 lines (in Fig. 5). We observe again that the algorithm converges fast to the optimal solution; the maximum execution time never exceeded 2 minutes.




Fig. 3. Grid graphs. Top: deterministic paths. Bottom: random paths.

In both synthetic and real-world experiments, we observe that the convergence rate, determined by the number of iterations required to reach the optimum, varies not only between graph classes but also within the same graph class. This can be explained as follows.

It is clear from the description of the algorithm that the number of iterations depends on how fast $\lambda_\ell(t)$ reach their optimal values, which also depends on $\dot{\lambda}_\ell(t)$. The latter depends on $\alpha_\ell(t)$, which in turn depends on $y_\ell(t)$.

The quantity $y_\ell(t)$ depends on the number of paths that use edge ℓ . Fast convergence implies an as small as possible value for $\alpha_\ell(t)$, which implies a value for $y_\ell(t)$ that is as close as possible to c_ℓ . At a first place observe that initially $y_\ell(t)$ can be much larger than c_ℓ , especially in the case where many paths use edge ℓ . At a second place observe that the initial value of $w_p(t)$ can be quite




Fig. 4. Real-world graphs. Top: 280 nodes and 354 edges. Bottom: 296 nodes and 393edges.

large since its value depends on the utility function. This in turn implies that the initial frequency value $x_p(t)$ can be very large, leading to large values of $y_\ell(t)$ in subsequent iterations. For these two reasons, the value of $\alpha_\ell(t)$ (that depends on $y_\ell(t)$) can start, for a specific input instance, from a rather high value and therefore it may take more iterations to reach its proper value, demonstrating a slower rate of convergence.

In conclusion, the convergence rate depends on a combination of input-specific factors that can vary considerably even between instances of the same graph class. These factors include the edge capacity values, the specific form of the utility function, the number of edges in a line route, and the number of line routes that share edges.




Fig. 5. Real-world graph with 319 nodes and 452 edges

5 Incentive-Compatible Robustness and Railway Optimization

The approach pursued in the preceding sections can be generalized to deal with robustness issues in the broader context of railway optimization (i.e., not only within line planning). In this section, we will argue on this matter and also compare incentive-compatible robustness with other known notions of robustness.

Railway optimization deals with large-scale planning and scheduling problems over several time horizons. Due to their complexity and sheer size, quite often such problems are provided with incomplete or uncertain data. For instance, some data may in advance be completely unknown or of low precision, while other data are subject to changes during the operational phase (e.g., due to delays). As there are several different types of imperfect information, there are also several different concepts for optimizing with respect to imperfect information. Each of the known concepts, such as multi-stage stochastic programming, chance constraint programming, robust optimization or online optimization, has its own strengths and weaknesses, making the various concepts the models of choice for different practical problems.

Two main approaches have been pursued in the literature to handle uncertainty: (i) Stochastic programming models, offering great flexibility but often too large in size to be handled efficiently, and also requiring that probability distributions are given a priori and can be handled in the solution procedure. (ii) Robust optimization models, which are easier to solve but sometimes leading to very conservative solutions of little practical use. Under this classical concept, feasibility is guaranteed if the number of constraints affected by data changes is bounded; see, for instance, the seminal work of Betsimas and Sim [4] on this subject.

A third way to model uncertainty, leading to a modeling framework called *light robustness*, was recently proposed in [8,9]. It couples robust optimization with a simplified two-stage stochastic programming approach, and constitutes a flexible counterpart of (classical) robust models that turned out to be quite promising within railway optimization. It is also worth mentioning that a variant of the Bertsimas and Sim method [4] for robustness was also applied to a game-theoretic scenario, in which competing entities act selfishly so as to optimize their own utility functions. In this case, uncertainty can involve both the rules of the game and the players' utility functions. It is showed that the robust counterpart of a game (under bounded, polyhedral uncertainty) is approximately as hard as the nominal game, at least in certain cases of finite games [1].

The classical robustness framework neglects the realistic possibility of a recovery phase; i.e., that small infeasibilities may be corrected. If such a phase is excluded, then solutions should be either unacceptably expensive or too conservative in order to be feasible in all scenarios that determine the imperfection of information.

Recoverable robustness is a new notion introduced in [2,16,17] that appears more suitable to deal with data uncertainty and recovery in the context of railway optimization. Recoverable robustness is about computing solutions that are robust against a limited set of scenarios and which can be made feasible (recovered) by a limited effort. One starts from a feasible solution x of an optimization problem which a particular scenario s , that introduces imperfect knowledge (i.e., by adding more constraints), may turn to infeasible. The goal is to have handy a recovery algorithm A that takes x and turns it to a feasible solution under s (i.e., under the new set of constraints). In other words, in recoverable robustness there is uncertainty about the feasibility space: imperfect information generates infeasibility and one strives to (re-)achieve feasibility.

The aforementioned approaches provide a quite powerful set of methods to deal with some kind of *predictable* and statically described level of uncertainty mainly in the constraints⁴. But what happens when the exact shape of the global objective function is unknown to the system? This may happen in application scenarios where many entities compete for common resources and each one acts selfishly. For example, in the line planning setting which we considered in the previous sections, where commercial line operators (competing entities) with fixed choices of lines compete for the utilization of these lines (common resources) via frequency negotiation with the (possibly public-sector or governmental) network operator. In such settings, for obvious reasons, each competing entity is not willing to reveal her own (private) utility function; that is, to reveal her level of satisfaction for acquiring a specific usage of resources. Nevertheless, the goal of the dynamic market designer (network operator in the line planning setting) – corresponding to the socially optimal solution – is to guarantee a *fair and feasible* solution, that is, a certain usage of subsets of resources to the competing entities in such a way that constraints regarding the usage of the resources

⁴ Uncertainty can also be transferred to the objective function as well, by incorporating it into the constraints.

are not violated and at the same time the average satisfaction of the competing entities (players) is maximized.

All aforementioned approaches seem to be inadequate to deal with such an application scenario, because the nature of the uncertainty itself is not quantified in any sense, and indeed may vary with time. Additionally, this situation should not be dealt with as a static problem to be centrally solved, but rather as a dynamic decentralized scheme, that continuously adapts the usage of resources to the players, in order to always keep them as close as possible to the socially optimal solution, as the utility functions of the players may also evolve with time.

In this work, we propose a new notion of robustness along with a corresponding solution framework, which we call *incentive-compatible robustness* and which is complementary to the notion of recoverable robustness. It is concerned with the computation of an *incentive-compatible recovery scheme* that achieves robustness by enforcing the system to converge to its optimal solution. By an incentive-compatible recovery scheme we mean a decentralized price-updating and resource-usage allocation method, that exploits the selfish nature of the competing entities, in order to lead them back to the socially optimal solution, even if the social optimum itself varies with time. Each resource gets a dynamic pricing scheme, and each competing entity is allowed to continuously change her bidding for getting (in the near future) usage of resources. In this context, the feasibility space is known and incomplete information refers to lack of information about the optimization problem, due to the unknown utility functions. In incentive-compatible robustness, there is uncertainty about the objectives: feasibility is guaranteed, since imperfect knowledge does not introduce new constraints, and one strives to achieve optimality, exploiting the selfish nature of the players.

Note that incentive-compatible robustness is different from the concept of game-theoretic robustness as developed in [1]. The approach in [1] is a centralized, deterministic paradigm to uncertainty in strategic games. Our approach differs from that in the following: (i) It is decentralized to a large extent, based only on local information that the participating entities (line operators and resources) have at any time; (ii) we impose no restriction on the kind of the utility functions of the players other than their strict concavity, whereas the approach in [1] has to somehow quantify the “magnitude” of uncertainty of the constraints and/or the payoffs, in order to keep the solvability of the problem comparable to that of the nominal counterpart; (iii) the solvability of the robust counterpart in [1] is largely based on the solvability of the nominal counterpart (which is strongly questionable for the general game-theoretic framework).

To summarize, incentive-compatible robust optimization suggests a generic approach to deal with robustness issues in railway optimization applications that require setting up a dynamic market for negotiating usage of resources, over subsets of resources, by selfish entities that do not reveal their incentives and having non-fixed (elastic) demands.

6 Conclusions and Open Issues

We investigated a new application scenario in line planning that achieves *incentive-compatible robust* solutions by exploiting a resource allocation mechanism introduced by Kelly [13] in the context of communication networks. For the case of a single line pool, an adaptation of Kelly's approach can provide (under certain assumptions) a decentralized algorithm that provably converges to the socially optimal solution. For the case of multiple line pools, an extension and further elaboration of Kelly's approach is required in order to derive such an algorithm. We also conducted experiments on a discrete variant of the pricing scheme for the single-pool case over synthetic and real-world data. Our algorithms allow LOPs to negotiate line frequencies over fixed lines in a dynamic fashion. In a broader context, our approach comprises a generic technique to set up a dynamic market for (re-)negotiating usage of resources over subsets of resources. Consequently, it could be applied to set up a dynamic frequency market over other transportation settings (e.g., in the airline domain).

A crucial question would be to devise protocols that demonstrate faster convergence to the equilibrium point, even approximately. Additionally, it would be interesting to find ways to tackle the assumption on price taking and myopic behavior of the users. It would be nice to do this even at the cost of suboptimal equilibrium points. It is noted that when the LOPs are not price takers and myopic (called *price anticipators* in the congestion control jargon), then the above scheme does not lead to socially optimal solutions, even for the case where there is only a single resource to share. Nevertheless, it would be quite interesting to know how far one can be from the social optimum, given that a decentralized updating scheme is adopted for the user requests and the prices of the resources.

Further open issues concern: (i) a theoretical analysis of the discrete variants of our algorithms; (ii) an extension of our approach to introduce proportions per resource (rather than per line pool); (iii) the investigation of other types of LOP's utility functions, or the case for a LOP to pursue a different utility function per line pool; (iv) looking for other parameters of robustness and recoverability (e.g., introduction of delays).

Acknowledgements. We would like to thank the referees for their valuable comments and suggestions that improved the presentation. We are also indebted to Rolf Möhring, Christian Liebchen, and Sebastian Stiller for their comments on an earlier draft of this work, and to Kostas Tsihlas for many fruitful discussions.

References

1. Aghassi, M., Bertsimas, D.: Robust Game Theory. Mathematical Programming Ser. B 107, 231–273 (2006)
2. ARRIVAL Deliverable D1.2, New Theoretical Notion of the Prices of Robustness and Recoverability, ARRIVAL Project, Version 2 (July 2007)
3. Bertsekas, D.: Nonlinear Programming, 2nd edn. Athena Scientific (1999)

4. Bertsimas, D., Sim, M.: The Price of Robustness. *Operations Research* 52(1), 35–53 (2004)
5. Borndörfer, R., Grötschel, M., Lukac, S., Mitusch, M., Schlechte, T., Schultz, S., Tanner, A.: An Auctioning Approach to Railway Slot Allocation. *Competition and Regulation in Network Industries* 1(2), 163–196 (2006)
6. Clarke, E.H.: Multipart pricing of public goods. *Public Choice* 11, 19–33 (1971)
7. Dienst, H.: Linienplanung im spurgeführten Personenverkehr mit Hilfe eines heuristischen Verfahrens, PhD thesis, Technische Universität Braunschweig (1978)
8. Fischetti, M., Monaci, M.: Light Robustness. Technical Report ARRIVAL-TR-0119, ARRIVAL Project (January 2008)
9. Fischetti, M., Monaci, M.: Light Robustness. In: Ahuja, R.K., Möhring, R.H., Zaroliagis, C.D. (eds.) *Robust and Online Large-Scale Optimization*. LNCS, vol. 5868, pp. 61–84. Springer, Heidelberg (2009)
10. Goossens, J., van Hoesel, C., Kroon, L.: A branch and cut approach for solving line planning problems. *Transportation Science* 38, 379–393 (2004)
11. Green, J.R., Laffont, J.J.: *Incentives in Public Decision-Making*. North-Holland Publishing Company, Amsterdam (1979)
12. Groves, T.: Incentives in teams. *Econometrica* 41(4), 617–631 (1973)
13. Kelly, F.: Charging and rate control for elastic traffic. *European Transactions on Telecommunications* 8, 33–37 (1997)
14. Kelly, F., Maulloo, A., Tan, D.: Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society* 49, 237–252 (1998)
15. Kontogiannis, S., Zaroliagis, C.: Robust Line Planning under Unknown Incentives and Elasticity of Frequencies. In: *Proc. 8th Workshop on Algorithmic Approaches for Transportation Modeling, Optimization, and Systems – ATMOS 2008* (2008)
16. Liebchen, C., Lübbecke, M., Möhring, R., Stiller, S.: Recoverable Robustness. Technical Report ARRIVAL-TR-0066, ARRIVAL Project (August 2007)
17. Liebchen, C., Lübbecke, M., Möhring, R., Stiller, S.: The Concept of Recoverable Robustness, Linear Programming Recovery, and Railway Applications. In: Ahuja, R.K., Möhring, R.H., Zaroliagis, C.D. (eds.) *Robust and Online Large-Scale Optimization*. LNCS, vol. 5868, pp. 1–27. Springer, Heidelberg (2009)
18. Schöbel, A., Scholl, S.: Line Planning with Minimal Traveling Time. In: *Proc. 5th Workshop on Algorithmic Methods and Models for Optimization of Railways – ATMOS* (2005)
19. Schöbel, A., Schwarze, S.: A Game-Theoretic Approach to Line Planning. In: *Proc. 6th Workshop on Algorithmic Methods and Models for Optimization of Railways – ATMOS* (2006)
20. Scholl, S.: Customer-oriented line planning. PhD thesis, Technische Universität Kaiserslautern (2005)
21. Srikant, R.: *The Mathematics of Internet Congestion Control*. Birkhäuser (2004)
22. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tender. *Journal of Finance* 16(1), 8–37 (1961)