Robust Line Planning under Unknown Incentives and Elasticity of Frequencies *

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, Ioannina, Greece
[3] Department of Computer Engineering and Informatics, University of Patras, 26500 Patras, Greece
Email: 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 traffic rates (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 struggling to exploit the underlying network infrastructure via frequency requests, while the management of the infrastructure itself remains the responsibility of a single (typically governmental) entity, the *network operator*.

The line operators are typically unwilling to reveal their true incentives. Nevertheless, the network operator would like 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. We show that this can be accomplished in certain situations via a (possibly 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

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

* This work was partially supported by the Future and Emerging Technologies Unit of EC (IST priority – 6th FP), under contract no. FP6-021235-2 (project ARRIVAL).

ATMOS 2008
8th Workshop on Algorithmic Approaches for Transportation Modeling, Optimization, and Systems
http://drops.dagstuhl.de/opus/volltexte/2008/1581

are given a network $G = (V, L)$ (usually referred to as 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 is a path in G. Typically, a *line pool* is also provided, i.e., a set of potential lines among which the final set of lines will be decided. The *frequency* of a line l is a rational number indicating how often service to customers is provided along l 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. 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*).

The line planning problem has mainly been studied under two main approaches (see e.g., [6, 7]). 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 number of customers with direct connections (under a similar constraint), or at least 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 [9, 11].

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 [10]. 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 lines. The delay depends on the traffic load and hence on the frequencies of all lines in the network. The objective is to provide lines 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* 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*. Under this framework, line operators act as competing entities for the exploitation of shared goods and are (possibly) unwilling to reveal their actual level-of-satisfaction functions that determine their true incentives. Nevertheless, the network operator 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 network operator 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 *robust* against changes in the demands of the entities (line operators). 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.

In this paper, we explore this rationale by considering the case where the (selfishly motivated) line operators request frequencies (traffic demands) over a pool of already fixed line routes (one per line operator). Rather than requesting end-to-end frequencies, the line operators offer bids, which they (dynamically) update, for buying frequencies. Each line operator has a utility function determining her level of satisfaction that is *private*; i.e., she is not willing to reveal it to the network operator or her competitors, due to her competitive nature. The network operator announces an (anonymous) resource pricing scheme, which indirectly implies an allocation of frequencies to the line operators, given their own bids. By applying techniques from the network congestion control literature, we show that for the case of a single pool of routes, there exists a distributed, dynamic, (user) bidding and (resource) price updating protocol, whose equilibrium point is the unknown social optimum. We first study the single pool case, assuming strict concavity and monotonicity of the private utility functions. All dynamic updates of bids or prices may be done at the line operator or resource level, based only on local information, that concerns the particular line operator or resource. The key assumption is that the line operators can control only a negligible amount of frequency along a single line compared to its total frequency. We extend our technique to the case of multiple line pools, whose mix is determined by the network operator for the sake of social optimality, and prove similar results.

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]. 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] for a detailed discussion on the subject as well as 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 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.

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. 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).

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 (line operators) 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 line operators (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.

The rest of this paper is organized as follows. Section 2 provides the set up of our modeling. The decentralized pricing mechanism both for the single and the multiple line pool case is given in Section 3. We conclude in Section 4.

2 The Model

Suppose that a set P of line operators behave as competing service providers, willing to offer regular train routes to the end users of a railway public transportation system. The railway network operator provides the (aforementioned) public transportation network $G = (V, L)$, with the set L of edges (railway tracks connecting directly two nodes of G) being the *resources* of the network. These resources are assumed to be subject to (fixed) capacity constraints, described by the capacity vector $\mathbf{c} = (c_\ell)_{\ell \in L} > 0$. The capacity of each edge is considered as a shared resource provided by the network operator.

There is a fixed pool of routes (i.e., origin–destination paths), one per line operator, that the line operators are willing to use. This pool is represented by a **routing matrix** $R \in \{0,1\}^{|L| \times |P|}$, in which each row $R_{\ell,\star}$ corresponds to a different edge $\ell \in L$, and each column $R_{\star,p}$ corresponds (actually, is the characteristic vector of) the route of a distinct line operator $p \in P$. Each line operator $p \in P$ has complete control over the *frequency* or *traffic rate* (of trains) she decides to route over her path, $R_{\star,p}$, 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 the line operator $p \in P$ for committing an end-to-end traffic rate $x_p > 0$ along her route $R_{\star,p}$, 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 traffic rate x_p allocated to the line operator $p \in P$. It is also assumed that these functions are *private*: Each line operator is not willing to reveal it to the network operator or her competitors, due to her competitive nature.

The railway network operator is only interested in having a socially optimal (fair) solution. This is usually interpreted as maximizing the aggregate satisfaction of the line operators. Therefore, the social welfare objective is considered to be the maximization of the aggregate utilities of the line operators, subject to the capacity constraints. That is, the network operator is interested in the solution of the following convex optimization[1] problem:

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

Since all utility functions are strictly concave, then $\boxed{\text{SOCIAL}}$ has a unique optimal solution, which is the social optimum. To solve $\boxed{\text{SOCIAL}}$ directly, the network operator, 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 line operators who use the infrastructure, and this has to be done also in a fair way: Each line operator 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 the line operator'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 congestion induced by all the line operators in these edges, due to the congestion effect. Indeed, it would be desirable for the network operator 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 line operators is induced, despite the fact that there is no global knowledge of the exact utility functions of the line operators.

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

As for the line operators (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 network operator charges her for claiming an amount of traffic rate along her own route. The strategy space of a line operator is to claim (via bidding) the value of the traffic rate 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.

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

3 A Decentralized Pricing Scheme

The selfish perspective of the competing line operators (the players) implies a strategic game among them, in which the network operator intervenes only implicitly (as the game designer), by setting the resource usage (per–unit) costs. 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 line operator, whose strategy is to choose a line frequency (traffic rate) over her (already fixed) route $R_{\star,p}$ 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 $\left\{\mathbf{x} \in \mathbb{R}_{\geq 0}^{|P|} : R\mathbf{x} \leq \mathbf{c}\right\}$.
- Each player's payoff is determined both by the value of the private utility function $U_p(x_p)$ (for having a traffic rate 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 traffic rate 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 traffic rate vector for all the players but for player p. Therefore, the sole goal of player $p \in P$ is to choose her traffic rate so as to maximize her individual payoff:

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

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

As we shall see shortly, we will allow the players to affect their own choices (traffic rates) only indirectly, via bidding. That is, each player is not assumed to freely choose her own traffic rate along her route, but rather offer a larger bid for (hopefully) getting higher traffic rate.

3.1 Describing the Social Optimum

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

KKT-SOCIAL

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

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

$$R_{\ell,\star} \cdot \hat{\mathbf{x}} \leq c_\ell, \quad \forall \ell \in L, \tag{3}$$

$$\hat{\lambda}, \hat{\mathbf{x}} \geq \mathbf{0} \tag{4}$$

Of course, the problem with the KKT-SOCIAL 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 SOCIAL, also described in KKT-SOCIAL, without demanding this knowledge. The answer to this is *partially affirmative*, and this is by exploiting the selfish nature of the players (line operators) as we shall see shortly.

3.2 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 network operator 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 line operators 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's assume that we already know the optimal Lagrange Multipliers, $(\hat{\lambda}_\ell)_{\ell \in L}$ of KKT-SOCIAL. Interpreting these values as the per–unit–of–frequency prices of the resources, we have a pricing scheme for the traffic induced by the line operators to their own routes: Each line operator pays exactly for the marginal cost of her own traffic rate 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 $\hat{\mu}_p \equiv \sum_{\ell \in L: R_{\ell,p}=1} \hat{\lambda}_\ell = \hat{\lambda}^T R_{\star,p}$ is the per–unit price for committing a unit of traffic along the route $R_{\star,p}$ of player $p \in P$.

One should mention here that indeed there is 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 traffic rate. This is because the scalar $\hat{\mu}_p$ actually depends on the optimal primal–dual pair $(\hat{\mathbf{x}}, \hat{\lambda})$.

We next assume that the players are actually controlling only *negligible* amounts of traffic rates compared to the aggregate ones[2]. 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 traffic rate 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:

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

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 the first set of equalities in KKT-SOCIAL. Of course, we still have to deal with the crucial

[2] For the considered application scenario, this is not unrealistic.

problem that the optimal Lagrange Multipliers (that define the marginal prices for the users) cannot be directly computed, due to both the size of $\boxed{\text{SOCIAL}}$ and the lack of knowledge of the private utility functions, in the framework of railway optimization.

In order to handle this situation, we consider the following two-level scenario for dynamically setting per–unit prices of the edges and frequencies of the selfish players: Initially each player $p \in P$ announces a bid $w_p \geq 0$ concerning the total amount of money she is willing to pay for buying traffic rate along her own route. The exact amount of traffic rate that she will eventually buy, depends on the per-unit price that will be announced by the network operator, and is not yet known to her. Consequently, the network designer considers the following optimization problem, whose Lagrange Multipliers define the per–unit prices of the edges:

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

That is, the network operator considers that the private utility $U_p(x_p)$ is substituted by the (also strictly concave and increasing) function $w_p \log(x_p)$. The choice of this function along with the selfishness of the players allows us to obtain a convex program with linear inequalities, whose KKT system is very similar (except for the first line) to $\boxed{\text{KKT-SOCIAL}}$:

KKT-NETWORK

$$\frac{w_p}{\bar{x}_p} = \bar{\lambda}^T \cdot R_{\star,p}, \quad \forall p \in P, \tag{5}$$

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

$$R_{\ell,\star} \cdot \bar{\mathbf{x}} \leq c_\ell, \quad \forall \ell \in L, \tag{7}$$

$$\bar{\lambda}, \bar{\mathbf{x}} \geq \mathbf{0} \tag{8}$$

At this point, one could argue that the convex program $\boxed{\text{NETWORK}}$ could be directly solved, and compute (along with $\boxed{\text{KKT-NETWORK}}$) the requested Lagrange Multipliers. The huge scale of a railway network optimization instance makes this approach rather unappealing. Therefore, we shall compute an optimal solution of $\boxed{\text{NETWORK}}$ in a distributed fashion, as follows:

– Each edge is equipped with a dynamically updated charging mechanism, which is the same (per–unit) price for all the line operators using it. This charging mechanism is updated according to the following system of differential equations:

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

where $y_\ell(t) \equiv \sum_{p\in R:R_{\ell,p}=1} x_p(t) = R_{\ell,\star} \cdot \mathbf{x(t)}$ is the cumulative traffic rate 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 $\mathcal{E}$. The system of differential equations (9) is obtained from the well-known approach (see e.g., [12]) 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 traffic rate 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 traffic rate is below or exceeds the capacity of the resource, respectively. Thus, the only stable situation is only when a resource is either under-used and has zero price (since there is no interest in using the residual capacity), or its traffic has already reached its capacity. Observe that the equilibrium of this system of differential equations has $\forall \ell \in L, \bar{y}_\ell \equiv R_{\ell,\star} \cdot \bar{\mathbf{x}} = c_\ell \quad \vee \quad \bar{\lambda}_\ell = 0$. That is, the complementarity conditions of both KKT-SOCIAL and KKT-NETWORK (equations (2) and (6)) are satisfied.

– Each line operator $p \in P$ is charged an instantaneous per-unit price $\mu_p(t) \equiv \sum_{\ell \in L: R_{\ell,p}=1} \lambda_\ell(t) = \lambda(\mathbf{t})^T \cdot R_{\star,p}$, at any time $t \geq 0$. Therefore, given their commitment on spending w_p for buying traffic rate, at equilibrium player p is allocated a traffic rate $\bar{x}_p = \frac{w_p}{\bar{\mu}_p}$. From this we deduce that at equilibrium also the equations (5) of KKT-NETWORK are satisfied.

The above distributed scheme is a congestion control algorithm, in which each player (line operator) reacts to signals she gets about the congestion along her route. These signals are the per-unit prices $\mu_p(t)$ that the line operator gets from the network operator at any time.

The question is whether the above system converges at all. This is indeed true, if we assume that the routing matrix R has full rank. This assures that given a set $\lambda(t) = (\lambda_\ell(t))_{\ell \in L}$ of instantaneous per-unit prices at the resources, the set $\mu(t) = (\mu_p(t))_{p \in P}$ of per-unit prices for the line operators, that is computed as the solution of the system $\mu(t) = R^T \cdot \lambda(t)$, is *unique*. Using a proper Lyapunov function argument, it can be shown (cf. [8, Chapter 22]) that this dynamic (and distributedly implemented) pricing scheme, for *fixed* player bids $(w_p)_{p \in P}$, is stable and converges to the optimal solution $(\bar{\mathbf{x}}, \bar{\lambda})$ of NETWORK.

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

Then we have:

$$\begin{aligned}
\frac{dV(\lambda(t))}{dt} &= \sum_{\ell\in L}(\lambda_\ell(t)-\bar{\lambda}_\ell)\cdot\dot{\lambda}(t)\\
&= \sum_{\ell\in L}(\lambda_\ell(t)-\bar{\lambda}_\ell)\cdot[\max\{y_\ell(t)-c_\ell,0\}\cdot\mathbb{I}_{\{\lambda_\ell(t)=0\}}+(y_\ell(t)-c_\ell)\cdot\mathbb{I}_{\{\lambda_\ell(t)>0\}}]\\
&\le \sum_{\ell\in L}(\lambda_\ell(t)-\bar{\lambda}_\ell)\cdot(y_\ell(t)-c_\ell)\\
&= \sum_{\ell\in L}(\lambda_\ell(t)-\bar{\lambda}_\ell)\cdot[(y_\ell(t)-\bar{y}_\ell)+(\bar{y}_\ell-c_\ell)]\\
&\le \sum_{\ell\in L}(\lambda_\ell(t)-\bar{\lambda}_\ell)\cdot(y_\ell(t)-\bar{y}_\ell)\\
&= \sum_{\ell\in L}(\lambda_\ell(t)-\bar{\lambda}_\ell)\cdot R_{\ell,\star}\cdot(\mathbf{x}(t)-\bar{\mathbf{x}})
\end{aligned}$$

$$\begin{aligned}
&= \sum_{p\in P}(\mu_p(t)-\bar{\mu}_p)\cdot(x_p(t)-\bar{x}_p)\\
&= \sum_{p\in P}\left(\frac{w_p}{x_p(t)}-\frac{w_p}{\bar{x}_p}\right)\cdot(x_p(t)-\bar{x}_p)=\sum_{p\in P}w_p\cdot\left(2-\frac{x_p(t)}{\bar{x}}-\frac{\bar{x}_p}{x_p(t)}\right)\\
&\le 0
\end{aligned}$$

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

Let's now return to the line operators. We have assumed that these players announce some fixed bids, but the truth is that since the pricing scheme changes over time, it is in the interest of each of them to vary her own bid. 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 $\mu(\mathbf{t})=(\mu_p(t))_{p\in P}$ (seen as constants) they are charged at time $t\ge 0$:

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

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

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

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 equilibrium, 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-NETWORK and KKT-SOCIAL coincide. Therefore, the selfish behavior of the myopic, price taking players, under the dynamic price setting mechanism and bidding scheme, converges to the optimal solution $(\hat{\mathbf{x}}, \hat{\lambda})$ of SOCIAL.

3.3 Extension to Multiple Pools of Routes

In this subsection we extend the freedom of the railway network operator, assuming that he can periodically use different pools of routes for the players, from a set K of pools. The set of different pools is motivated by the fact that usually there are dependencies between lines; for instance, the choice of a high-speed line affects the choice of lines for other trains. These dependencies split naturally the set of all lines into a small number of subsets determined by the network operator, resulting in different line pools.

Each pool $k \in K$ is represented by its routing matrix $R(k) \in \{0,1\}^{|L|\times|P|}$, as before. The line operators still try to have (indirect) control only over the end-to-end traffic rates they get by the network operator. We assume that player $p \in P$ gets a unique traffic rate x_p for the whole period of time considered. It is up to the network operator to decide how to multiplex the distinct pools of routes, in order to achieve the optimal social welfare value. That is, the network operator now directly participates in the optimization problem, via the variables $f_k : k \in K$ indicating the portion of time each pool consumes over the whole time period we study. This social welfare optimization problem is the following:

MULTI-SOCIAL

$$\max\left\{\sum_{p\in P} U_p(x_p) : \forall \ell \in L, \forall k \in K, (R(k))_{\ell,\star} \cdot \mathbf{x} \leq c_\ell \cdot f_k;\ \sum_{k\in K} f_k \leq 1;\ \mathbf{x}, \mathbf{f} \geq \mathbf{0}\right\}$$

The Lagrangian function is the following:

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

If we set $\mu_p(\Lambda) = \sum_{\ell\in L}\sum_{k\in K}\Lambda_{\ell,k}\cdot(R(k))_{\ell,p} = \sum_{k\in K}{\Lambda_{\star,k}}^T(R(k))_{\star,p}$, the system of KKT conditions of MULTI-SOCIAL is written as follows:

KKT-MULTI-SOCIAL

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

Observe that, by equation (11), in the optimal solution all the pools have the same weighted aggregate price, equal to $\hat{\zeta}$, if we use the edge capacities as weights. Moreover (due to equation (13)), unless this optimal aggregate price is zero, it holds that the edge capacities are totally distributed among the distinct pools: **if** $\hat{\zeta} > 0$ **then** $\sum_{k\in K}\hat{f}_k = 1$.

We can once more set the instantaneous per-unit prices for the players as a (linear) function of the Lagrange Multipliers, as follows:

$$\forall t\geq 0,\ \forall p\in P,\ \mu_p(t) = \sum_{\ell\in L}\sum_{k\in K}\Lambda_{\ell,k}(t)\cdot(R(k))_{\ell,p} = \sum_{k\in K}(\Lambda_k(t))^T\cdot(R(k))_p$$

The congestion at edge $\ell\in L$ due to route $R(k)$ at time $t\geq 0$, is given by $y_{\ell,k}(t) = (R(k))_{\ell,\star}\cdot\mathbf{x}(t)$. A dynamic (by the edges) pricing scheme, is described by the following system of differential equations:

$$\begin{aligned}
&\forall\ell\in L,\ \forall k\in K,\\
&\dot{\Lambda}_{\ell,k}(t) = \max\{y_{\ell,k}(t) - c_\ell f_k, 0\}\cdot\mathbb{I}_{\{\Lambda_{\ell,k}(t)=0\}} + (y_{\ell,k}(t) - c_\ell f_k)\cdot\mathbb{I}_{\{\Lambda_{\ell,k}(t)>0\}} \qquad (17)
\end{aligned}$$

This differential system would then assure the validity of equations (12), at equilibrium, if the pool frequencies provided by the network operator were fixed. Assuming again that the players announce (instantaneous) bids, and providing them with a traffic rate $x_p(t) = w_p(t)/\mu_p(t)$ at any time, along with their price taking and myopic behavior, would also assure (at equilibrium) the validity of equations (10).

But we also have to assure the validity of equations (11) and (13). In order to achieve this, we study the dual of $\boxed{\text{MULTI-SOCIAL}}$: The dual problem of $\boxed{\text{MULTI-SOCIAL}}$ is the following:

$$\boxed{\text{DUAL-MULTI-SOCIAL}}\quad \max\left\{D(\Lambda,\zeta) : \forall \ell \in L, \forall k \in K, \Lambda_{\ell,k} \geq 0; \zeta \geq 0\right\}$$

where:

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

Observe that the dual objective $D(\Lambda,\zeta)$ is split in two parts. The first part

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

is a maximization problem similar to the one already dealt with in the single pool case of the previous section. The second part

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

It is now clear that so long as there exists a pool with weighted aggregate price (using the capacities as weights) strictly larger than the value of ζ, the optimal choice of $\mathbf{f}$ is to be a probability distribution that assigns positive mass only to pools of maximum aggregate price (according to Λ). Therefore:

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

We augment our price updating scheme described in the system (17) as follows:

– After the players having announced the current bids $(w_p(t))_{p\in P}$ and (consequently) the edges having updated their current prices in each pool $(\Lambda_{\ell,k}(t))_{\ell\in L,k\in K}$, $\zeta(t)$ is set to the *average* price of a pool:

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

– The network operator then updates the portions of time granted to each of the pools, so that pools exceeding the current average price tend to increase their portion of time (in hope of decreasing their weighted cost), while pools that are cheaper than the average price slightly decrease their portion of time. That is:

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

where, $\phi > 0$ is a scaling factor ensuring that the resulting vector $\mathbf{f}$ is again a probability distribution over the pools.

Observe that, at equilibrium $(\hat{\mathbf{x}}, \hat{\mathbf{f}}, \hat{\Lambda}, \hat{\zeta})$, our augmented pricing scheme has all pools with the same aggregate price: $\forall k \in K, \hat{\mathbf{c}}^T\cdot\hat{\Lambda}_{\star,k} = \hat{\zeta}$. It is only then that the time portions of the pools stabilize (cf. equation (19)). This assures the validity of equations (11). Moreover, by brute force we assure the validity of equation (13).

It is mentioned at this point that the two updating schemes concerning the average pool price and the vector of time portions of the pools, have to be centrally computed by the network operator, since it is his decision how to change them and these changes take into account the state of the whole network. Unfortunately we cannot (at least in this approach) avoid this bottleneck, since we involve the network operator in the competing environment.

4 Conclusions and Open Issues

We presented a distributed protocol for a new application scenario in line planning that achieves *incentive-compatible robust* solutions. Our protocol allows line operators 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 *fast* 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 players are not price takers and myopic (they are called then *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 distributed (and localized) updating scheme is adopted for the user requests and the prices of the resources.

Acknowledgements. We are 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. M. Aghassi and D. Bertsimas, "Robust Game Theory", *Mathematical Programming Ser. B*, 107 (2006), pp. 231-273.
2. ARRIVAL Deliverable D1.2, "New Theoretical Notion of the Prices of Robustness and Recoverability", ARRIVAL Project, Version 2, July 2007.
3. D. Bertsekas, *Nonlinear Programming*, Athena Scientific, 2nd Edition, 1999.
4. D. Bertsimas and M. Sim, "The Price of Robustness", *Operations Research*, 52:1 (2004), pp. 35-53.
5. R. Borndörfer, M. Grötschel, S. Lukac, M. Mitusch, T. Schlechte, S. Schultz, A. Tanner, "An Auctioning Approach to Railway Slot Allocation", *Competition and Regulation in Network Industries*, 1:2 (2006), pp. 163–196.
6. H. Dienst, "Linienplanung im spurgeführten Personenverkehr mit Hilfe eines heuristischen Verfahrens", PhD thesis, Technische Universität Braunschweig, 1978.
7. J. Goossens, C. van Hoesel, and L. Kroon, "A branch and cut approach for solving line planning problems", *Transportation Science* 38 (2004), pp. 379393.
8. A. Ozdaglar and R. Srikant, "Incentives and Pricing in Communications Networks", Chapter 22 in *Algorithmic Game Theory*, ed. N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, Cambridge University Press, 2007.
9. A. Schöbel and S. Scholl, "Line Planning with Minimal Traveling Time", In *Proc. 5th Workshop on Algorithmic Methods and Models for Optimization of Railways* – ATMOS 2005.
10. A. Schöbel and S. Schwarze, "A Game-Theoretic Approach to Line Planning", in *Proc. 6th Workshop on Algorithmic Methods and Models for Optimization of Railways* – ATMOS 2006.
11. S. Scholl, "Customer-oriented line planning", PhD thesis, Technische Universität Kaiserslautern, 2005.
12. R. Srikant, *The Mathematics of Internet Congestion Control*, Birkauser, 2004.