Railway capacity auctions with dual prices[‡]

Thomas Schlechte[**] Andreas Tanner[*]

Abstract

Railway scheduling is based on the principle of the construction of a conflict-free timetable. This leads to a strict definition of capacity: in contrast with road transportation, it can be said in advance whether a given railway infrastructure can accommodate – at least in theory – a certain set of train requests. Consequently, auctions for railway capacity are modeled as auctions of discrete goods – the train slots. We present estimates for the efficiency gain that may be generated by slot auctioning in comparison with list price allocation. We introduce a new class of allocation and auction problems, the *feasible assignment problem*, that is a proper generalization of the well-known combinatorial auction problem. The feasible assignment class was designed to cover the needs for an auction mechanism for railway slot auctions, but is of interest in its own right. As a practical instance to state and solve the railway slot allocation problem, we present an integer programming formulation, briefly the ACP , which turns out to be an instance of the feasible assignment problem and whose dual problem yields prices that can be applied to define a useful activity rule for the linearized version of the Ausubel Milgrom Proxy auction. We perform a simulation aiming to measure the impact on efficiency and convergence rate.

1 Introduction

Railway scheduling is based on the principle of the construction of a conflict-free timetable. This leads to a strict definition of capacity: in contrast with road transportation, it can be said in advance whether a given railway infrastructure can accommodate – at least in theory – a certain set of train requests. Consequently, auctions for railway capacity are modeled as auctions of discrete goods – the train slots. These slots can be understood as the right to use a certain set of infrastructure segments at a given time. Brewer and Plott [6] suggest a model where feasibility of a schedule is based on the *binary exclusion* property, which says that a schedule of trains is feasible if any two trains are conflict-free.

[‡]This work was funded by the Federal Ministry of Economics and Technology (BMWi), project *Trassenbörse*, grant 19M7015B.

[*]WIP TU Berlin, Workgroup for Infrastructure Policy, Strasse des 17. Juni 135, 10623 Berlin, Germany, Email `ant@wip.tu-berlin.de`

[**]Zuse-Institute Berlin (ZIB), Takustrasse 7, 14195 Berlin-Dahlem, Germany, Email `schlechte@zib.de`

Parkes and Ungar [11] present an auction-based track allocation mechanism for the case that single-track, double-track, and yard segments have to be concatenated to form a single line. They suggest a hybrid mechanism that combines elements of the simultaneous and the combinatorial auction formats.

Our approach is more general and is applicable to arbitrary infrastructure networks where a discrete allocation has to be found. Airport landing slots allocation, for instance, has long been thought of as a possible allocation problem that could be solved by auctioning [2]. In fact, we take an abstract look on scheduling conflicts by defining a general *feasibility predicate* that is thought to express the absence of scheduling conflicts for arbitrary sets of *train slots*.

It is obvious that the valuations of slots are not independent from each other. Both complementary and substitutional relationships are seen in practice and are modeled in some detail in [4]. However, in this paper, we consider *linear valuation functions*, that is, the network value is modeled as the sum of the values of all implemented slots. We formulate the feasible assignment problem in Section 2 and describe its relationship to the combinatorial auction problem.

Borndörfer and Schlechte [5] introduce the ACP , an integer programming formulation of the train timetabling problem with nice computational properties for large-scale problem instances. In Section 2.2, we present the ACP as an example of the feasible assignment problem. In Section 4, we present an iterative auction format derived from the Ausubel Milgrom ascending auction [1] that uses linear prices. We then show how the dual prices of the ACP can be used to define an activity rule that tolerates late introduction of new bids and still has acceptable convergence properties.

2 The feasible assignment problem

In the most general sense, feasibility is just a predicate applied to a schedule, or set of train slots with a very natural monotonicity property:

Definition 1. *Let $\mathcal{S}$ be a set of slots. A* feasibility predicate *is a predicate on $\wp(\mathcal{S})$ which is* monotonous*, that is,*

$$\mathbf{feas}(S) \text{ and } S' \subseteq S \implies \mathbf{feas}(S') \tag{1}$$

Following auction-theoretic convention, the traffic operating companies are represented by agents $l \in L$ with valuations $v_l(S)$ for bundles $S \subseteq \mathcal{S}$ of slot requests. An *assignment* is a vector $\chi = (\chi_l : l \in L)$ of mappings $\chi_l : S \mapsto \{0,1\}$, where $\chi_l(s) = 1$ means that agent l receives item s. We also write $\chi(l) = \{s \in S : \chi_l(s) = 1\}$.

Definition 2. *The* feasible assignment problem *then is defined to be the following optimization problem:*

$$\max_{\chi:\chi_l:S\mapsto\{0,1\}} \sum_{l\in L} v_l(\chi(l)) \tag{2}$$

such that

$$(\forall s \in S) \sum_{l \in L} \chi_l(s) \leq 1 \tag{3}$$

$$\mathbf{feas}\left\{ s \in S : \sum_{l \in L} \chi_l(s) = 1 \right\} \tag{4}$$

The objective function (2) expresses the aim of maximizing total value. Condition (3) states that every slot can be assigned only once. The assumption that all slot requests are unique can be done without loss of generality – if there are two identical slot requests, we could simply tag them with different IDs in order to make them different.

Condition (4) ascertains that the set of assigned slots is feasible. The feasible assignment problem searches for those assignment of slots that maximizes total revenue while maintaining that the set of all assigned slots is feasible.

2.1 Relationship to the Combinatorial Allocation Problem

We use the terminology of an *allocation problem* as the optimization problem of finding an optimal with respect to some target function allocation subject to a set of constraints, and the corresponding *auction problem* of designing an as efficient as possible auction scheme for bidders with private valuations who may behave strategically and possibly collude.

The following example shows that the feasible assignment problem is a generalization of the combinatorial allocation problem, the allocation problem that is underlying the well-known combinatorial auction problem.

Example 1. *Let G be a set of goods, and L a set of agents.*

Define $\mathcal{S} = \wp(G)$, and the predicate **feas_ca** *on $\wp(\mathcal{S})$ by*

$$\mathbf{feas\_ca}(S) \Leftrightarrow ((\forall s, s' \in S)\, s \cap s' = \emptyset) \tag{5}$$

Then the combinatorial allocation problem for G and the feasible assignment problem for **feas_ca** *are equivalent in the following sense:*

If $\chi : L \mapsto \wp(G)$, then χ satisfies constraints (3) and (4) if and only if χ is a valid allocation for G.

Now we show that every assignment problem can be *embedded* into some combinatorial allocation problem.

Definition 3. *A* token system *for a set $\mathcal{S}$ of slots is a multiset*[1] *$\mathcal{T}$ of* tokens *$t \in \mathcal{T}$ together with a map*

$$\phi : \mathcal{S} \ni s \mapsto \phi(s) \subseteq \mathcal{T} \tag{6}$$

[1] The notion of *multiset* generalizes the notion of a set by allowing and accounting for multiple copies of identical elements. The usual set-theoretical operations and relations (union, intersection, subset-relation) all have their natural generalization to multisets.

We say that for $s \in \mathcal{S}$, $\phi(s)$ is the set of tokens required *for s.*

We say that a token system $(\mathcal{T}, \phi)$ implements *the feasibility predicate* **feas** *if*

$$\mathbf{feas}(S) \iff \bigcup_{s \in S}^{*} \phi(s) \subseteq^{*} \mathcal{T} \tag{7}$$

where $\bigcup^*$ and $\subseteq^*$ are operations and relations of multisets.

Proposition 1. *Let* **feas** *be a feasibility predicate for a set $\mathcal{S}$ of slots. Then there is a token system $(\mathcal{T}, \phi)$ that implements* **feas***.*

Proposition 1 implies that to every monotonous slot assignment problem $(\mathcal{S}, \mathbf{feas})$, there is a combinatorial auction problem $(\mathcal{T}, \phi)$ that $(\mathcal{S}, \mathbf{feas})$ can be *embedded into*. In particular, there is an integer program (IP) implementing the problem.

Corollary 2. *To any feasible assignment problem, there is an IP implementing the problem.*

Proof. Let $(\mathcal{T}, \phi)$ be a token system implementing the feasible assignment predicate **feas**.

Define

$$\phi_s^*(t) = \begin{cases} 1 & \text{if } t \in \phi(s) \\ 0 & \text{otherwise.} \end{cases} \tag{8}$$

The feasible assignment problem then is equivalent to the following IP:

$$
\begin{array}{llllr}
(\mathrm{IP}(\mathcal{T}, \phi)) & & & & \\
\max & \sum_{l \in L} \sum_{S \subseteq \mathcal{S}} v_l(S) \chi_l(S) & & & \text{(i)} \\
\text{s.t.} & \sum_{S} \chi_l(S) & \leq 1, & \forall l \in L & \text{(ii)} \\
 & \sum_{S \subseteq \mathcal{S}} \sum_{s \in S} \chi_l(S) \phi_s^*(t) - y_l(t) & \leq 0, & \forall l \in L, t \in \mathcal{T} & \text{(iii)} \\
 & \sum_{l \in L} y_l(t) & \leq 1, & \forall t \in \mathcal{T} & \text{(iv)} \\
 & \chi_l(S), y_l(t) & \in \{0, 1\} & \forall l \in L, t \in \mathcal{T}, S \in \mathcal{S} & \text{(v)}
\end{array}
$$

The target function in (i) maximizes the total valuation realized by the allocated bundles. Constraint (ii) expresses that to every agent, at most one bundle is allocated. The binary decision variables $y_l(t)$ express if a token t is acquired by agent l, and constraint (iii) expresses that if a bundle is awarded to l, then l has to acquire all necessary tokens. Constraint (iv) yields that every token can be acquired only once, and constraint (v) are the integrality constraints which guarantee that neither tokens nor bundles are split. □

2.2 The ACP model for the train scheduling problem

In this section we recapitulate the track allocation problem, a network formulation, and a track configuration based integer programming model, for details we refer to [5] and a recent survey [10]. It is interesting that this formulation yields dual information, which can nicely be interpreted as shadow prices for using resources. This is something standard packing models for the train timetabling problem can not provide that easily as introduced in [8].

The problem to construct a feasible railway track allocation can be formulated in terms of an acyclic digraph $D = (V, A)$, the train scheduling network. Its nodes represent events, such as arrival, departure or passing of trains, at a set S of stations at discrete times $T \subseteq \mathbb{Z}$, its arcs model activities, primarily runs of trains between stations. By δ_{in} we denote the set of incoming arcs $a \in A$ for $v \in V$, by $\delta_{out}(v)$ the set of outgoing arcs, respectively. Denote by $s(v) \in S$ the station associated with event $v \in V$, and by $t(v) \in T$ the time of this event; we can assume $t(u) < t(v)$ for each arc $uv \in A$ such that D is acyclic. Let $J = \{s(u)s(v) : uv \in A, s(u) \neq s(v)\}$ be the set of all railway tracks.[2] We are further given a set I of requests to route trains through D. More precisely, train $i \in I$ can be routed on a path through some induced subdigraph $D_i = (V_i, A_i) \subseteq D$ from an artificial source node $s_i \in V_i$ to a sink node $t_i \in V_i$. Given arc weights w_a, $a \in A$ and considering each train individually, the problem of finding a best allocation for request i is to construct a longest (s_i, t_i)-path in D_i w.r.t. w. But due to the fact that trains are potentially using the same infrastructure resources, we have to guarantee that the set of chosen train routes is conflict-free.

We distinguish between two types of conflicts. The first cause are invalid occupation times on tracks. In the ACP integer programming formulation, conflicts on the same track $j \in J$ are handled by another subgraph structures $D_j = (V_j, A_j)$ of D. These are induced by all running arcs in D_i and coupled in the integer programming formulation. In [5] a construction was introduced such that each feasible flow from an artificial source node $s_j \in V_j$ to the sink node $t_j \in V_j$ can be mapped unique to an allocation on track j that is feasible w.r.t. the safety block occupation or headway times[3]. Let $A_c^j = \{uv \in A : s(u)s(v) = j\}$ be the set of all running arcs associated with some track $j \in J$. By $A_c = \bigcup_{j \in J} A_c^j$, we denote the set of all "coupling" arcs. For each arc $a \in A_c$ the subset $A_c^a \subseteq A_c$ contains all conflicting coupling arcs w.r.t. the headway times on that track.

Second type of conflicts is station capacity, which cannot be exceeded at any discrete time. Let $\kappa_{s,t}$ be the upper bound of trains that can be handled in station $s \in S$ at time $t \in \{0, T\}$. By $A_{s,t} = \{uv \in A : s(u) \neq s, s(v) = s, t(v) = t\} \cup \{uv \in A : s(u) = s(v) = s, t(u) \leq t \leq t(v)\}$ we denote the set of all arcs, which correspond to active train events in station s at time t, that are arriving

[2] For notational reasons, we assume that no parallel tracks exists in J, otherwise the projection $s(u)s(v)$ would not be unique.

[3] The same construction works if and only if the triangle inequality holds for the headway matrix.

trains or already waiting trains at time t.

A *track allocation* or timetable is a set of conflict-free routes, satisfying each request at most once. The *track allocation problem* or train timetabling problem is to find a track allocation of maximum weight.

Variables x_a, $a \in A_i$, $i \in I$ control the use of trip a in D_i and y_a, $a \in A_j$, $j \in J$ in D_j, respectively. Note that each arc in set A_c corresponds to two variables x_a and y_a.

$$
\begin{array}{llrll}
\text{(ACP)} & & & & \\
\max & \sum_{a\in A} w_a x_a & & & \text{(i)} \\
\text{s.t.} & \sum_{a\in\delta_{out}(v)} x_a - \sum_{a\in\delta_{in}(v)} x_a & = 0, & \forall i \in I, v \in V_i \backslash \{s_i, t_i\} & \text{(ii)} \\
& \sum_{a\in\delta_{out}(s_i)} x_a & \le 1, & \forall i \in I, s_i \in V_i & \text{(iii)} \\
& \sum_{a\in\delta_{out}(v)} y_a - \sum_{a\in\delta_{in}(v)} y_a & = 0, & \forall j \in J, v \in V_j \backslash \{s_j, t_j\} & \text{(iv)} \\
& \sum_{a\in\delta_{out}(s_j)} y_a & \le 1, & \forall j \in J, s_j \in V_j & \text{(v)} \\
& x_a - y_a & \le 0, & \forall a \in A_c & \text{(vi)} \\
& \sum_{a\in A_{s,t}} x_a & \le \kappa_{s,t} & \forall s \in S, t \in \{0, T\} & \text{(vii)} \\
& x_a, y_a & \ge 0, & \forall a \in A & \text{(viii)} \\
& x_a, y_a & \in \{0,1\}, & \forall a \in A & \text{(ix)}
\end{array}
$$

The objective, denoted in (ACP) (i), is to maximize the weight of the track allocation. Equalities (ii) and (iv) are well-known *flow conservation constraints* at intermediate nodes for all trains flows $i \in I$ and for all flows on tracks $j \in J$, (iii) and (v) state that at most one flow unit is realized. Inequalities (vi) link arcs used by train routes and track configurations to ensure a conflict-free allocation on each track. Constraints (vi) ensure that at each time no station capacity is exceeded. Finally (viii) and (ix) are the non-negativity and the integrality constraints. Note that, all variables are implicitly 0/1.

Proposition 3. *Let $\mathcal{S} = I$ and suppose that agents $l \in L$ have additive valuations $v_l(S) = \sum_{s\in S} v_l(s)$ for subsets $S \subseteq \mathcal{S}$. Define*

$$
w_a = \begin{cases} v_l(i) & \text{for } a = (s_i, v) \in A_i \\ 0 & \text{otherwise.} \end{cases} \tag{9}
$$

Define the feasibility predicate

$$
\begin{aligned}
\mathbf{feas_{ACP}}(S) \text{ if and only if} & \qquad (10) \\
& \textit{There is a mapping } a \mapsto x_a \textit{ such that} \\
& \textit{constraints (ii) to (ix) of (ACP) hold.}
\end{aligned}
$$

Then (ACP) is equivalent to the feasible assignment problem with predicate $\mathbf{feas_{ACP}}$.

Consider the dual DLP of the LP-relaxation of ACP, i.e. neglecting constraints ACP (ix):

$$
\begin{array}{llllr}
\text{(DLP)} & & & & \\
\min & \sum_{s\in S}\sum_{t\in\{0,T\}} \kappa_{s,t}\mu_{s,t} & +\sum_{i\in I}\alpha_{s_i} & +\sum_{j\in J}\beta_{s_j} & \text{(i)} \\
\text{s.t.} & \alpha_u-\alpha_v+\lambda_{uv}+\sum_{\substack{s\in S\\ t\in\{0,T\}\\ a\in A_{s,t}}}\mu_{s,t} & \geq w_{uv} & \forall uv\in A_i\cap A_c,\ i\in I & \text{(ii}_\text{a}\text{)} \\
& \alpha_u-\alpha_v+\sum_{\substack{s\in S\\ t\in\{0,T\}\\ a\in A_{s,t}}}\mu_{s,t} & \geq w_{uv} & \forall uv\in A_i\backslash A_c,\ i\in I & \text{(ii}_\text{b}\text{)} \\
& \beta_u-\beta_v-\lambda_{uv} & \geq 0 & \forall uv\in A_j\cap A_c,\ j\in J & \text{(iii}_\text{b}\text{)} \\
& \beta_u-\beta_v & \geq 0 & \forall uv\in A_j\backslash A_c,\ j\in J & \text{(iii}_\text{b}\text{)} \\
& \alpha_{s_i} & \geq 0 & \forall i\in I & \text{(iv)} \\
& \alpha_{t_i} & =0 & \forall i\in I & \text{(v)} \\
& \beta_{s_j} & \geq 0 & \forall j\in J & \text{(vi)} \\
& \beta_{t_j} & =0 & \forall j\in J & \text{(vii)} \\
& \lambda_a & \geq 0 & \forall a\in A_c & \text{(viii)} \\
& \mu_{s,t} & \geq 0 & \forall s\in,t\in\{0,T\} & \text{(ix)}
\end{array}
$$

Here, $\alpha_v, v\in V$, β_w, $w\in W=\cup_{j\in J}V_j$, λ_a, $a\in A_c$, and $\mu_{s,t}, s\in S, t\in\{0,T\}$ are the dual variables associated with constraints ACP (ii) to (vii), respectively. Note that we introduce artificial variables α_{t_j} to get a compact notation of constraints (DLP)(ii_b) and (DLP)(iii_b), respectively.

3 The surplus of slot auctioning

The aim of this section is to give an estimation of the surplus that an auction generates, compared to the value generated by an allocation based on list prices. Let slot requests be triples $(s, u(s), l(s))$, where s is a slot from a set of possible slots $\mathcal{S}$ with some consistency predicate **feas**, and $u(s)>0, l(s)>0$ are the private valuation and the list price associated with s. Furthermore we assume that there is a cost-function c that associates with every slot request s, a cost $c(s)\geq 0$. This cost is to be understood as *marginal cost to the infrastructure provider* and not as the cost of operation to the train operating company. The latter cost are assumed to be internalized into the private value function u.

For a schedule $S\subset\mathcal{S}$, the *value* of S is defined to be

$$v(S)=\sum_{s\in S}u(s)-c(s) \tag{11}$$

Note that the value is always defined with respect to the private utility function u. The *auction value* is the value of the schedule maximizing $u - c$ subject to feasibility:

$$v^{\text{auction}} = v\left(\arg\max\{\sum_{s\in T} u(s) - c(s) : \mathbf{feas}(T)\}\right) \tag{12}$$

$$= \max\{\sum_{s\in T} u(s) - c(s) : \mathbf{feas}(T)\} \tag{13}$$

The *list value* is the value (with u used as valuation function!) of the schedule maximizing l subject to feasibility, and assuming that the list price also serves as minimum price (that is, if $u(s) < l(s)$, then the slot request is not submitted): With

$$I_{u(s)\geq l(s)}(s) = \begin{cases} 1 & \text{if } u(s) \geq l(s) \\ 0 & \text{if } u(s) < l(s) \end{cases} \tag{14}$$

we define

$$v^{\text{list}} = v\left(\arg\max\{\sum_{s\in T} l(s) \cdot I_{u(s)\geq l(s)}(s) : \mathbf{feas}(T)\}\right) \tag{15}$$

Note that we always have $v^{\text{list}} \leq v^{\text{auction}}$ (since the schedule maximizing the list price is also feasible when maximizing private value, and the value of a schedule is always computed based on the private value function u). Obviously for all non-trivial cases we can assume $v^{\text{list}} > 0$ and can interpret the fraction

$$\theta = \frac{v^{\text{auction}} - v^{\text{list}}}{v^{\text{list}}} \tag{16}$$

as relative surplus of a perfect auction versus a list price allocation.

3.1 A simple one-track model

We now derive an estimation of ρ for a simple model, giving an upper bound on the potential auctioning surplus. We elaborate two cases which differ in terms of the marginal costs of infrastructure provision. We first consider the case where marginal costs vanish – this is justified when the train operator carries all costs associated with the train run (energy, charge for wear and tear). We then consider the case of $c = l$, that is, that the list price reflects exactly the marginal cost. This is the approach officially taken by EU legislative, without specifying the exact price components to be included into calculation.

Let there be a sample of slot requests $S = \{(s_1, u_1, l_1), \ldots, (s_n, u_n, l_n)\}$. Furthermore, we assume

$$\mathbf{feas}(S) \text{ if and only if } |S| \leq m \tag{17}$$

This means that any m of the n slot requests may be scheduled.

3.1.1 Zero marginal costs

Now we assume that $c(s) = 0$ and $l_i = l$ for all i and some constants $l > 0$ and $m > 0$. This could be the case, for instance, if the slot requests all compete over usage of a single track and the list price does not differ in time. For the private valuations, we assume that they are independently and normally distributed with parameters (μ, σ). For a given sample, the auction value then is just the sum of the m largest values of the u_i: Let $(i_1, \ldots, i_n)$ be indexed such that $u_{i_1} \leq \ldots \leq u_{i_n}$. Then

$$v^{\text{auction}}(S) = \sum_{j=0}^{m-1} u_{i_{n-j}} \tag{18}$$

The distribution of the random variable $v^{\text{auction}}(S)$ is given explicitly in [13]; however, the term derived there becomes quite complex if $n - m$ grows, and we rely on expected values from a numerical simulation. To derive the list price allocation, note that any allocation that uses the capacity m to the most possible extent maximizes the generated revenue und thus maximizes the target function in equation (15). Let $S' = \{(s, u, l) \in S : u \geq l\}$. Since the list prices of all slot requests are equal and the feasibility constraint in our model just implements a cardinality restriction, the list price allocation randomly chooses a subset $T \subseteq S'$ of cardinality $\max\{m, |S'|\}$. Thus, the list value can be represented by the random variable

$$v^{\text{list}}(T) = \sum_{(s,u,l) \in T} u \tag{19}$$

Let

$$v^{\text{auction}} = E\left(v^{\text{auction}}(S)\right) \tag{20}$$

$$v^{\text{list}} = E\left(v^{\text{list}}(S)\right) \tag{21}$$

For the list value mean, we derive with $\phi^{(\mu,\sigma)}$ the distribution function of the normal distribution with parameters μ and σ, and $\mu^* = \mu^*(l) = \frac{\mu}{1-\phi^{(\mu,\sigma)}(l)}$

$$v^{\text{list}} = \mu^* \cdot \left(\sum_{k=0}^{m} k \binom{n}{k} (1 - \phi^{(\mu,\sigma)}(l))^k \phi^{(\mu,\sigma)}(l)^{n-k} \right. \tag{22}$$

$$\left. + \sum_{k=m+1}^{n} m \binom{n}{k} (1 - \phi^{(\mu,\sigma)}(l))^k \phi^{(\mu,\sigma)}(l)^{n-k} \right) \tag{23}$$

The term

$$\rho(l, \mu, \sigma, n, m) = \frac{1}{m} \sum_{k=0}^{n} k \binom{n}{k} (1 - \phi^{(\mu,\sigma)}(l))^k \phi^{(\mu,\sigma)}(l)^{n-k} \tag{24}$$

$$= \frac{n(1 - \phi^{(\mu,\sigma)}(l))}{m} \tag{25}$$

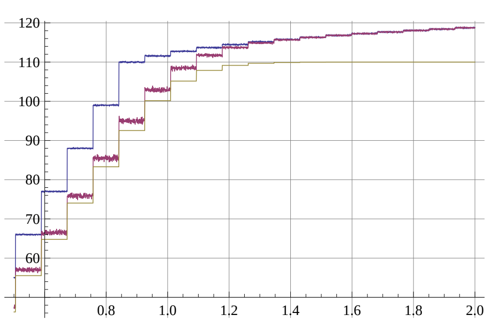

(a) $\mu = 11, \sigma = 1$

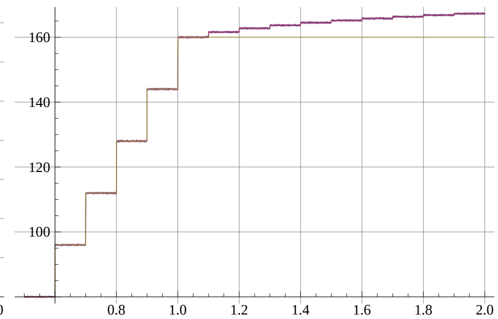

(b) $\mu = 16, \sigma = 0.9$, note that upper and middle curves coincide

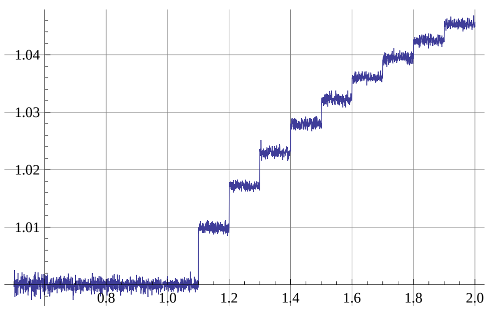

(c) Relative auction surplus θ for $\mu = 16, \sigma = 0.9$

Figure 1: v^{auction} (upper curve), $v^{\text{auction}}_{\text{min}=10}$ (middle), and v^{list} (lower curve) for $l = 10, m = 10$ and $0.5 \leq \rho \leq 2$

is called the *overdemand ratio*. Assuming that slot requests with a willingness to pay below the list price are not submitted, the overdemand ratio expresses by how much, on average, the demand on the given track supersedes capacity.

Figure (1a) (upper and lower curve) shows a comparison of auction and list value for a track with capacity $m = 10$ and overdemand ratio between 0.5 and 2. The ascend of both curves approximately equals the mean private value $\mu = 11$ until a critical point that can be seen as "saturation point". For the auction value, this is the point for which total demand equals capacity, that is, where $10 = n = \frac{m\rho}{1-\phi(\mu,\sigma)}$ or $\rho = 0.84$. For the list value, the critical point is at $\rho = 1$ or $n = 11.8$. As long as capacity saturation is not reached, the superiority of the auction is due to the fact that all slot requests are honored while the list price allocation discards the slot requests with values below list price. Once capacity is reached, the value of the list price allocation converges to $l \cdot m$, while the auction allocation benefits from selection of the slot requests with the highest private value.

3.1.2 List prices as marginal costs

Let us now consider the case of $c(s) = l(s) = l$ for all s. The list price allocation is unaffected by this change. Formula (18) transforms to

$$v^{\text{auction}}_{\text{min}=l}(S) = \sum_{j=0}^{m-1} u_{i_{n-j}} I_{u_{i_{n-j}} \geq l} \tag{26}$$

The middle curve in figure (1) shows the value of auctioning with minimum price equal to the list price. We observe that while spare capacity is available ($\rho \leq 1$), performance of auctioning and list price coincide. For $\rho > 1$, we have overdemand, the minimum price looses impact and performance of auctioning with and without minimum price converge.

3.1.3 Estimating the private valuations and relative auction surplus

The auction surplus depends on the relationship between the list price and the private valuations. If they coincide, or demand is below capacity, both allocation schemes perform similar. If there is variance in the private valuations which is not reflected in the list price, auctioning leads to more efficient allocation.

Data on private valuations is hard to come by as it generally considered to be sensitive information. We have performed some interviews with representatives from different freight-train operating companies and have received cautious word that the margin for a freight service could be somewhere between 2 and 6 percent. Furthermore, a rough estimate on costs results in the share of the infrastructure charge (at the list price system currently in place in Germany) amounts to about 25% of the total cost (infrastructure charges, energy, locomotive maintenance, wages). This leads as to estimate the additional willingness to pay over the list price at between 8 and 24%. In order to estimate the variance, we assume that the given bounds – both upper and lower ones – apply in 80% of the cases. Reading from the CDF of the normal distribution, we get

$$u \sim \text{NormalDistribution}(\mu, \sigma) \tag{27}$$

with

$$\mu = 0.16l \tag{28}$$

$$\sigma = 0.09l \tag{29}$$

Figure (1b) shows auction values with and without minimum price, and list price value, for the parameters for μ and σ given above. Since valuations are well above list price (the probability that a private valuation is below the list price is $P(u < 10) \sim 1.310^{-11}$), $v^{\text{auction}}_{\text{min}=l}(S)$ and $v^{\text{auction}}(S)$ coincide. From Figure (1c) one can read that the potential gain from auctioning is between 1 and 4%.

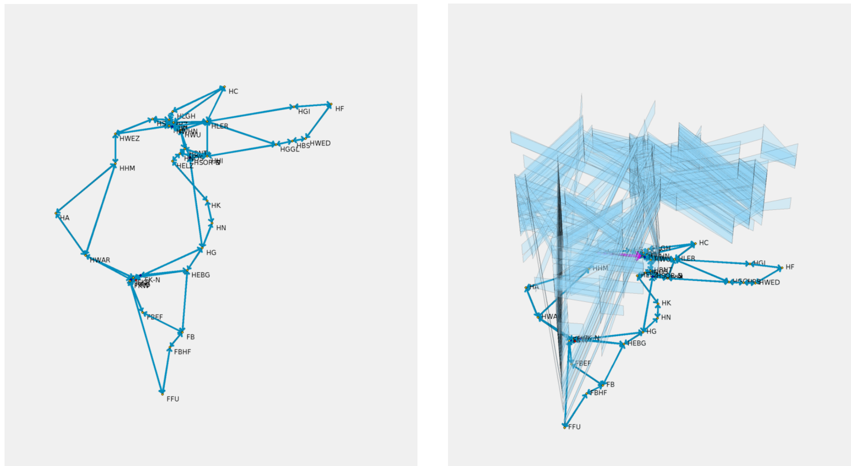

(a) Hanover-Kassel-Fulda sub-net

(b) Request set with 370 slot requests in space and time

Figure 2: Visualization of scenario with TraVis

3.2 An auction surplus estimation based on empirical demand data

In the previous section we derived an estimation of the auction surplus in a model where the graph structure of the railway network was omitted and, besides the approximation of the willingness to pay in relation to the list price, no empirical data on slot demand was used. Now we present an alternative way that uses both these information. Again we compare the two allocation schemes – list price and auction. For both allocation schemes, we base the schedule construction on combinatorial optimization so that the maximizations in equations (12) and (15) are really computed. We are aware of the fact that currently, a "fully automatic" construction of timetables based on combinatorial optimization is not in fact implemented anywhere currently. There is no doubt that if this was possible for real railways, this would be an very welcome innovation and enhance the efficiency of railway infrastructure usage. For our study, however, we wish to separate the effect of mathematical optimization from the effect of auctioning and therefore, choose the approach to estimate the auction surplus as pointed out above.

3.2.1 Data source and model construction

We work with a sub-net of the German railway network that covers the area between Hanover, Kassel, and Fulda, see Figure (2a). This is a central segment of the German network, and even as of 2008, there is reportedly a scarcity of slots.

The Federal office for statistical data (*Statistisches Bundesamt*) publishes data [7] concerning goods transported by railway. We used their tables that give yearly (we used data for 2008) quantities of transported goods, for national and international transport, as origin-destination pairs of *traffic districts* (*Verkehrsbezirke*), grouped according to the category of goods (*Gueterhauptgruppen*). The federal state of Hessia, for instance, is divided into 6 traffic districts of between, and there is a total of 101 traffic districts in Germany. Concerning international transportation, of the countries relevant for traffic through our sub-net, Austria, Czech republic, Denmark and Luxembourg are represented by one traffic district, while Italy and France are subdivided into 14 resp. 15 traffic districts. For goods being carried in transit through our subnet, we relied on expert knowledge[4] regarding the nodes where traffic on a given OD pair is entering and leaving the sub-net. Internal traffic and traffic originating from or with destination to the sub-net was assumed to be distributed among freight terminals according to our own estimations. Thus, a three-dimensional matrix $yt(or, dest, gtype)$ was produced that for pairs of nodes of the sub-net and a freight category, gives the quantity of freight of this category that is being transported yearly on paths between these nodes.

Next, we estimate the yearly number of trains used for implementing this matrix. Unfortunately, the classification of goods that the statistics is built upon leaves much to be desired. About 20% of the goods quantity – the largest among all shares – is being categorized as "other goods". For this position, we assumed that these goods were container freight. For the remaining categories, expert knowledge[5] was used to come up with an estimation of the weight $w(gtype)$ of a train homogeneously loaded with goods of this category could carry at maximum.

We expressed the train density, for any OD-pair and goods category, in *trains per hour* and divided the 24h hour period into segments of 1hour with constant train density. To get the train density for some hourly segment, we used the formula

$$tph(or, dest, gtype, k) = \frac{yt(or, dest, gtype)}{w(gtype) \cdot lf(gtype) \cdot d \cdot w(k)} \tag{30}$$

[4]Thanks to Martin Balser for providing this information.

[5]Thanks again to Martin Balser.

where

$$
\begin{aligned}
or &\text{ is the node of origin,}\\
dest &\text{ is the node of destination,}\\
gtype &\text{ is the type of good,}\\
k &\text{ is the hour under consideration}(0 \leq k \leq 23),\\
yt(or, dest, gtype) &\text{ is the tons } gtype \text{ transported from } or \text{ to } dest \text{ per year,}\\
w(gtype) &\text{ is the weight of a train fully loaded with } gtype,\\
lf(qtype) &\text{ is the load factor for trains carrying type } gtype,\\
d &\text{ is the number of working days per year}(d = 250),\\
w(k) &\text{ is the relative load share of the hour } k,\\
&\text{ with } w(k) \geq 0 \text{ and } \sum w(k) = 1.
\end{aligned}
$$

Our demand model has a relatively fine-granulated structure and thus the single entries in the table $(tph(or, dest, gtype, k) : or, dest$ traffic regions, $gtype, k)$ are quite small: between 10^{-6} and 0.1. This reflects the probability that at any given hour a freight train with specific origin, destination and load type enters the sub-net is small. In sum, nevertheless, the model yields around 600 daily freight trains for that area, grouped into departure time segments of 1 hour each. The precise requested departure times are generated by a Poisson process with fitting parameter. List-prices approximate the system of slot prices currently in place at the Deutsche Bahn, and private valuations are generated stochastically with normal distribution as pointed out in the previous section. For the passenger traffic we relied on published time-table data for a three-hour interval in the morning. Passenger trains do not participate in the auction, instead they receive slots (which may be elastic in time) before slots for the freight trains are auctioned.[6]

Extrapolated to a whole 24h day, we get around 1000 trains traveling in total through the sub-net. This seems to be a realistic number. We are interested in different economic scenarios that reflect different prognosis of freight demand and thus include sets of requests with the number of trains varying between 800 and 2000. For computational purposes, we trim these request sets to time windows of size 1 hour. Figure (2b) shows the time extended slot request for a request set with 118 slot request.

3.2.2 Simulation

Now we use the demand model in order to give an alternative estimation of auction surplus. As in the previous section we compare the economical value

[6] The reason for this is that passenger service in Germany is subsidized by the state (at least for local transit), and the Deutsche Bahn, and thus at least indirectly the railway infrastructure, is state-owned at the same time. Lastly, passenger slots are highly interdependent due to connection and regular service requests, and these interdependencies are not considered in our auction design.

of the allocation generated by (perfect) auctions with the list price allocation. All allocations were computed by solving model ACP on a standard PC using `TS-OPT` and `CPLEX 10.0`, see [3, 9]. In order to do that, we implemented a simplified version of the Deutsche Bahn slot pricing system[7]. We then computed list price and auction allocation for a range of request sets with varying factors of demand scaling, each covering a 24h period of traffic in the Hanover Kassel Fulda sub-net. For every demand scaling factor we generated 100 slot requests. To accommodate computational challenge we trimmed these slot requests to an excerpt of 1 hour and a subset of our sub-net extending from Fulda to Kassel. This sub-sub-net comprises 6 station nodes and 12 tracks. After trimming, the number of slots requested per set is between 15 and 150.

Table 1 gives the results. One can see that the surplus is between 0 and 10% and that it tends to be increasing with the demand scaling factor. For a scaling factor of 1.5, which seems realistic, the surplus is 4%. However there is a large statistical uncertainty that reflects the dependency of scheduling on the combinatorial structure of the time-extended network-request-graph. Nevertheless are the results consistent with those from the one track model.

scaling factor	slots requested (min-av-max)	slots scheduled (min-av-max)	ratio $\frac{\text{auction surplus}}{\text{list price surplus}}$ (min-av-max)
0.8	14-19.2-26	9-12.2-14	1-1.009-1.047
1	17-24.3-29	13-15.5-17	1-1.010-1.033
1.5	28-37.6-45	18-20.5-28	1-1.041-1.096
2	34-48.3-63	18-24.8-34	1.002-1.031-1.072
2.5	44-60-72	31-30.3-44	1.000-1.036-1.189
4	68-86.6-105	38-38.7-68	1.031-1.071-1.105
6	121-138.3-152	45-48.6-121	1.053-1.098-1.140

Table 1: Demand characteristics and auction surplus for the Hanover Kassel Fulda demand model

3.3 Interpretation

The authors were surprised – and admittedly somewhat disappointed – by the relative modest efficiency gains that can be contributed to slot auctioning. While a 4% efficiency gain seems a respectable result, one has to caution the reader that this is for the case that we are able to implement a perfect auction, that is, one that incites bidders to truthful bidding such that the most efficient allocation can be computed bases on the true valuations. We will see in the next section that the auctioning problem, from an auction-theoretical point of view, is not perfectly solvable if bidders behave strategically and may collude, and one has to relay on suboptimal auction schemes. Thus, we have to expect *bid shading*,

[7] omitting such peculiarities as the *Regionalfaktoren*

that is, bidders bid less than there true valuations, and the efficiency gains will shrink.

On the other hand our estimation uses conservatively chosen parameters. Most notably, the variance σ of the distribution of the private valuations has a large impact on the auction surplus and our choice is solely based on the assumption that the margin for a train operating company is between 2 and 6%. While managers naturally tend to understate the profitability of their business, the auction surplus increases with growing σ.

Even more important is the following argument: we assume that the effect of slot auctioning is just a re-allocation of slot requests that are present already now. But the auction allocation scheme has the potential to create new business opportunities for train operating companies (or enterprises from the logistic sector, or whoever wants to take the challenge) to offer high-quality services that may not be feasible presently. Such back-coupling effects are not present in our model and need more empirical research to be quantified.

4 The Linearized proxy auction with dual prices

We define an iterative auction with linear prices that generalizes the Ausubel Milgrom Proxy Auction. The result lies in the core. The generalization of the assignment problem leads to the possibility of prices lying above the bidder-optimal core frontier, in contrast to the Ausubel Milgrom Proxy auction. We give some examples for that.

We then turn to the use of dual prices to enforce *activity* in the iterative auction. We propose an activity rule that is based on dual prices for minimum bids for new bids introduced during the auction. We give a numerical simulation that measures efficiency, and compares the number of required rounds until auction termination, with and without the new activity rule.

4.1 Ausubel Milgrom Proxy Auction

Ausubel and Milgrom [1] propose an iterative auction (AMP) with results that lie in the core of the assignment problem game, thus being immune to collusion.

Ausubel and Milgrom formulate their proxy auction for agents bidding for contracts which are partially ordered by a preference-relation. We can apply their results immediately to the feasible assignment simultaneous proxy auction:

The auction is performed in a sequence of rounds. At every round i , a price vector $(p_{i,l}(S) : S \subseteq \mathcal{S})$ is maintained, and bidder l submits a set $D_{i,l} \subseteq \wp(\mathcal{S})$ of *most preferred bundles at given prices* (to be understood as XOR-connected bid).

The provisional allocation at round i is then computed as

$$\chi_i \in \arg \max_{\{\chi:\mathbf{feas}(\cup_l \chi(l))\}} \sum_l p_{i,l}(\chi(l)) \tag{31}$$

At round 0, the auction starts with price vector $p_{0,l} \equiv 0$. At round $i+1$, prices $p_{i+1,l}(S)$ are computed by

$$p_{i+1,l}(S) = \begin{cases} p_{i,l}(S) & \text{if } \chi_i(l) \in D_{i,l} \\ p_{i,l}(S) + \epsilon & \text{otherwise.} \end{cases} \tag{32}$$

The auction terminates at round i if for all bidders l, $\chi_i(l) \in D_{i,l}$. The outcome is the allocation χ_i and payments $p_{l,i}(\chi(l))$ of bidder l to the seller.

The AMP auction converges to a core outcome if agents bid truthfully at every round. In fact, the outcome lies within distance ϵ from the bidder-optimal frontier of the core. For substitute valuations, resulting prices are ϵ-minimal core prices.

4.2 An auction format for Railway slot auctioning: The linearized proxy auction

The AMP relies on the XOR bidding language. In particular, the indifference sets submitted by the bidders are dependent on *package prices*. It is trivial that in the case of a combinatorial auction, if bidder's valuations are linear, the usage of nonlinear prices is unnecessary.

Bidders for the assignment problem are assumed to have linear valuations. It is therefore natural to consider a restricted bidding language in which only linear valuations can be expressed, by interpreting replacing the set $D_{i,l}$ of most preferred bundles by set $d_{i,l}$ of most preferred *items*.

The provisional allocation rule (31) remains unchanged with

$$p_{i,l}(S) = \sum_{s \in S} p_{i,l}(s), \tag{33}$$

while the price update rule (32) changes to

$$p_{i+1,l}(s) = \begin{cases} p_{i,l}(s) & \text{if } s \in d_{l,i} \text{ but } s \notin \chi_i(l) \\ p_{i,l}(s) + \epsilon & \text{otherwise.} \end{cases} \tag{34}$$

Pricing rule (34) can be understood of generalizing CC pricing [12] to feasible assignment auctions: a slot is in overdemand if it is in the set of most preferred slots of some bidder but not allocated.

Definition 4. *The auction mechanism with the modified price update rule (34) is called* Linearized Proxy Auction *(LPA).*

It can be shown that the LPA terminates with a core outcome. However, even with linear valuations, prices need not be minimal in the core.

4.3 Using dual prices for late introduction of bids

AMP maintains bundle prices only on bundles that bidders have bidden for. A new bundle starts always at price zero. For combinatorial auctions, monotonicity of the bidder's valuation functions ensures that the sequence $(D_{i,l} : i = 1, 2, \ldots)$ of most-preferred bundle sets contains no increasing sequence of bundles.

In contrast, there is no obvious such monotonicity for train slot auctions. If at any round, a bidder may introduce a bid to a new slot starting at price zero, the auction will still converge as long as the set of train slots is finite, but only after exhausting the complete solution space. For example, a bidder could bid again and again on some slight variation of a slot request that was already rejected earlier, but the variation would still not have an inherent ask price.

If the computation of the provisional allocation is based on corollary 2, prices for token bundles can be used to compute ask prices for arbitrary slots. Because of the computational advantages of the ACP formulation, we used the prices from DLP to define an activity rule.

ACP-based minimum price rule for LPA. Let λ_a and $\mu_{s,t}$ be part of the solution to the DLP of the provisional allocation produced by ACP for some auction round n. λ_a can be understood as a price for usage of a certain arc a in the time-extended request graph, and $\mu_{s,t}$ as price for using a station s at time t. Now let $p = \{x_1, x_2, \ldots x_m\}$ be some implementation path of some request P_i. There may not be dual price information for the exact arcs that p makes usage of, but there may be prices for arcs that are relevant for p by being in conflict with p. Our dual price for path p then sums up, for every track and station, the average the dual prices of all arcs that are in conflict with path p.

$$w(p) = \sum_{x \in p} \left(\frac{1}{|A_c^a|} \sum_{a \in A_c, x \in A_c^a} \lambda_a + \frac{1}{|A_{s,t}|} \sum_{s \in S, t \in \{0,T\}, x \in A_{s,t}} \mu_{s,t} \right) \tag{35}$$

For any slot request i, which can be interpreted as set of implementing paths P_i, the path with the lowest cost provides a minimal price:

$$w(i) = \min_{p \in P_i} w(p) \tag{36}$$

4.3.1 Example for dual prices

Consider a path with just 4 nodes A, B, C and D, connected with 3 tracks AB,BC,CD, respectively. Let the running time be 1 for every track and suppose that track AB has a headway time of 10 and track BC and BD a headway time of 6, as shown in Figure 3. Now assume that there are 3 bidders and slot requests as given in the following table 2.

Note that bid 2_2 is the only one requesting track CD and therefore there is no scarcity on that track. The possible inclusion-maximal allocations for bids

Table 2: Definition of slot requests

			departure			
bidder_bidnr	source	destination	AB	BC	CD	private value
0_1	B	C	-	1	-	4
1_1	A	C	9	10	-	10
1_2	B	C	-	1	-	5
2_1	A	C	0	1	-	9
(2_2	B	D	-	4	5	*)

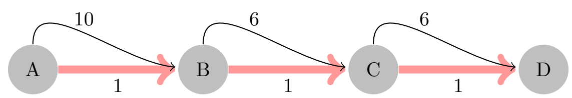

Figure 3: Example network

0_1 . . . 2_1 are

$$X = \{\{1\_1, 0\_1\}, \{1\_1, 1\_2\}, \{1\_1, 2\_2\}, \{2\_1, 0\_1\}, \{2\_1, 1\_2\}, \{2\_1, 2\_2\}\}$$

and the excluded bid 2_2 is in conflict with 0_1, 1_2 and 2_1.

Now assume that bidder 2 discovers the option of bidding on 2_2 only after his potential profit on 2_1 has shrieked below a certain threshold. The auction starts without that bid and after some rounds, the following bids could be submitted:

			departure			
bidder_bidnr	source	destination	AB	BC	CD	bid amount
0_1	B	C	-	1	-	4
1_1	A	C	9	10	-	5
1_2	B	C	-	1	-	5
2_1	A	C	0	1	-	4

Bidder 1 would win both his bids, and without the option of bid 2_2, the auction would stop now. Bid 2_2 is a new bid (since no bid to destination D was submitted so far), so there is no canonical minimum price for it so far. However, using DLP, the following dual arc prices can be computed:

track	departure	shadow price
AB	9	5
BC	1	5

Bid 2_2 conflicts with the arc on track BC, so the computed minimum price

according to equation (35) would be 5. Indeed, bid 2_2 has impact on the auction only if the bid amount exceeds 5, and the minimum bid saves the auction rounds that would otherwise be necessary to discover this.

4.3.2 Performance measurement simulation

We now report on a simulation that measures the performance of the dual prices generated by the ACP formulation. On the one hand, it is hoped that using these prices as minimal prices for newly introduced bids, the auction is sped up. On the other hand, there is a risk that these prices may be too high and may prevent bidders from submitting bids that would contribute to welfare.

Starting with demand sets from Section 3.2 – 10 for each scaling factor as shown in Table 3 – we run two versions of the LPA auction (with minimum increment 50). We assume that bidder bid myopically except that after any 10 rounds, every bid is randomly perturbed with probability 0.25, by translating the slot request's starting time by a random between -10 and +10.

The first version of the LPA does not know any minimum price rule for newly introduced slots, so bidders start bidding for perturbed slots from price 0. The second version of LPA uses the ACP-based minimum price rule. We compare the results in efficiency and convergence rate. The first columns of Table 3 show the efficiency for both LPA versions: one can see that the minimum price rule does not essentially affect efficiency. However, the last columns demonstrate that the number of rounds is significantly lower with the ACP-dual minimum price rule. We conclude that using dual prices as minimum prices may speedup the auction while the efficiency loss is moderate.

Scaling factor	incr. profit	dual profit	dual efficiency	incr. auction rounds	dual auction rounds	speedup
0.8	2983	2932	0.983	17.65	13.61	25%
1	3658	3597	0.984	19.43	14.11	27%
1.5	4941	4843	0.980	20.06	15.4	23%
2	6144	5967	0.971	21.53	17.2	20%
2.5	7272	7065	0.972	21.77	18.23	16%
4	9720	9374	0.964	22.96	19.84	14%
6	12233	11879	0.971	23.12	19.59	15%

Table 3: Incremental auction with and without dual prices: profit and number of rounds until termination

4.4 Conclusion and outlook

In this paper we provided an estimate of the potential social gain the could be generated by auctioning railway slots instead of allocating them by a list price schema. In both our models, a simple one-track model and an empirical real-world-data model, the efficiency gains is moderate, until 4% depending on the magnitude of overdemand and the spread of the private valuations of the train operating companies. Further research should illuminate the dependencies in more detail in order to produce well-founded prognosis for specific auctioning scenarios. Our model does not explicitly address the question of innovative use of the railway network that could be triggered by the auctioning, and we think that it would be interesting to investigate this point.

We then presented a possible auction design that built in the Ausubel Milgrom proxy auction but uses linear prices for slots and allows late introduction of bids. We demonstrated that using minimal prices derived from the dual of the ACP formulation of the scheduling problem may speed up the auction while suffering only moderate efficiency losses. It would be interesting to compare the dual prices used here with other linear pricing schemes for combinatorial auctions that are discussed in the literature. We leave this for future work.

References

[1] Lawrence M. Ausubel and Paul R. Milgrom. Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1(1):1–42, 2002.

[2] Michael Ball, George L. Donohue, and Karla Hoffman. Auctions for the safe, efficient, and equitable allocation of airspace system resources. In *Combinatorial auctions*, chapter 20. MIT press, 2008.

[3] Ralf Borndörfer, Berkan Erol, and Thomas Schlechte. Optimization of macroscopic train schedules via TS-OPT. In I.A. Hansen, E. Wendler, U. Weidmann, M. Lüthi, J. Rodriguez, S. Ricci, and L. Kroon, editors, *Proceedings of the 3rd International Seminar on Railway Operations Modelling and Analysis - Engineering and Optimisation Approaches*, Zürich, Switzerland, 2009.

[4] Ralf Borndörfer, Martin Grötschel, Sascha Lukac, Kay Mitusch, Thomas Schlechte, Sören Schultz, and Andreas Tanner. An auctioning approach to railway slot allocation. *Competition and Regulation in Network Industries*, 1(2), 2006.

[5] Ralf Borndörfer and Thomas Schlechte. Models for railway track allocation. In Christian Liebchen, Ravindra K. Ahuja, and Juan A. Mesa, editors, *ATMOS 2007 - 7th Workshop on Algorithmic Approaches for Transportation Modeling, Optimization, and Systems*, Dagstuhl, Germany, 2007. Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany.

[6] P. J. Brewer and C. R. Plott. A binary conflict ascending price (BICAP) mechanism for the decentralized allocation of the right to use railroad

tracks. *International Journal of Industrial Organization*, 14(6):857–886, 1996.

[7] Statistisches Bundesamt. Genesis database. https://www-genesis.destatis.de/genesis/online.

[8] Alberto Caprara, Matteo Fischetti, and Paolo Toth. Modeling and solving the train timetabling problem. *Operations Research*, 50(5):851–861, 2002.

[9] IBM ILOG CPLEX. *User-Manual CPLEX 10.0*, 2007.

[10] Richard Lusby, Jesper Larsen, Matthias Ehrgott, and David Ryan. Railway track allocation: models and methods. *OR Spectrum*, December 2009.

[11] David C. Parkes and Lyle H Ungar. An auction-based method for decentralized train scheduling. In *Proc. 5th International Conference on Autonomous Agents (AGENTS-01)*, pages 43–50, 2001.

[12] David Porter, Stephen Rassenti, Anil Roopnarine, and Vernon Smith. Combinatorial auction design. *Proc. Natl. Acad. Sci. USA*, 100(19):11153–11157, 2003.

[13] Douglas P. Wiens, Norman C. Beaulieu, and Pavel Loskot. On the exact distribution of the sum of the largest n-k out of n normal random variables with differing mean values. *Statistics*, 40:165–172, 2006.