Online Cooperative Cost Sharing*

Janina Brenner[1] and Guido Schäfer[2]

[1] Institute of Mathematics, Technical University Berlin, Germany
brenner@math.tu-berlin.de
[2] Algorithms, Combinatorics and Optimization, CWI Amsterdam, The Netherlands
g.schaefer@cwi.nl

Abstract. The problem of sharing the cost of a common infrastructure among a set of strategic and cooperating players has been the subject of intensive research in recent years. However, most of these studies consider cooperative cost sharing games in an *offline* setting, i.e., the mechanism knows all players and their respective input data in advance. In this paper, we consider cooperative cost sharing games in an *online* setting: Upon the arrival of a new player, the mechanism has to take instantaneous and irreversible decisions without any knowledge about players that arrive in the future. We propose an online model for general demand cost sharing games and give a complete characterization of both weakly group-strategyproof and group-strategyproof online cost sharing mechanisms for this model. Moreover, we present a simple method to derive incremental online cost sharing mechanisms from online algorithms such that the competitive ratio is preserved. Based on our general results, we develop online cost sharing mechanisms for several binary demand and general demand cost sharing games.

1 Introduction

The pivotal point in mechanism design is to achieve a global objective even though part of the input information is owned by selfish players. In cost sharing, the aim is to share the cost of a common service in a fair manner while the players' valuations for the service are private information. Based on the declared bids of the players, a cost sharing mechanism determines a service allocation and distributes the incurred cost among the served players. In many cost sharing games, the common service is represented by a combinatorial optimization problem like minimum Steiner tree, machine scheduling, etc., which defines a cost for every possible service allocation. We consider *cooperative* cost sharing games, i.e., players may form coalitions to coordinate their bidding strategies.

During the last decade, there has been substantial research on *binary demand* cost sharing games, where a service allocation defines simply whether or not a player is served. In this paper, we consider the *general demand* setting, in which players require not only one but several *levels* of service, and the mechanism

* This work was supported by the DFG Research Center Matheon "Mathematics for key technologies".

determines which service level is granted to each player and at what price. We assume that players are concerned only about the *quantity* of service levels they obtain, e.g., the number of distinct connections to a source, executions of their job, etc. Moreover, once a player's request for a certain service level was refused, she will not be granted a higher level. This general demand cost sharing model has recently been investigated quite intensively; see [1, 4, 11, 12].

To the best of our knowledge, all previous works on cooperative cost sharing consider *offline* settings, where the entire instance is known in advance. Hence, when determining the allocation and payment scheme, the mechanism can take into account all input data associated with every player (bids for different service levels and other relevant player characteristics). However, many natural cost sharing games inherently bear an *online* characteristic in the sense that players arrive over time and reveal their input data only at their arrival. In such settings, the mechanism needs to take instantaneous and irreversible decisions with respect to the assigned service level and payment of the player without any knowledge about players that arrive in the future. Problems in which the input data is revealed gradually and irreversible decisions have to be taken without knowledge of future requests are the subject of *online computation* [2]. The standard yardstick to assess the quality of an online algorithm is by means of its *competitive ratio*, i.e., the worst case ratio of the cost of the solution produced by the online algorithm compared to the cost of an optimal offline algorithm that knows the entire input data in advance.

Our Contributions. The main contributions of this paper are as follows:

1. We propose the first online model for general demand cost sharing games: In its most general form, every player arrives several times to request an additional service level. Upon the arrival of a player, the online mechanism immediately determines a price for her new request. We require that at each point of time, the sum of the collected payments approximates the cost of the (optimal offline) solution for the current allocation.
2. We completely characterize weakly group-strategyproof and group-strategyproof (formal definitions in Sec. 2) online mechanisms for general demand cost sharing games: We show that online cost sharing mechanisms are automatically weakly group-strategyproof for binary demand games. In the general demand case, this is true if the marginal costs of the underlying cost function are increasing. Moreover, we prove necessary and sufficient conditions for group-strategyproofness of online cost sharing mechanisms.
3. We present a simple yet effective method to derive online cost sharing mechanisms from competitive online algorithms: Given a ρ-competitive algorithm for the underlying problem, we show that the induced *incremental* online mechanism is ρ-budget balanced at all times. Using the above characterization, this enables us to derive incentive compatible online mechanisms for several binary demand and general demand cost sharing games for network design and scheduling problems.

Related Work. Immorlica et al. [9] partially characterized group-strategyproof cost sharing mechanisms in the offline case. They state that *upper-continuous* group-strategyproof β-budget balanced binary demand cost sharing mechanisms correspond to *cross-monotonic* cost sharing schemes. Juarez [10] very recently gave a similar characterization for mechanisms fulfilling the MAX property, meaning that indifferent players are always accepted.[3] He also showed that group-strategyproof cost sharing mechanisms correspond to feasible *sequential mechanisms* if indifferent players are always rejected. A sequential mechanism offers players service one after another according to an order that may change with previous decisions (see [10] for precise definitions).

Moulin [12] introduced *incremental cost sharing mechanisms* in the offline setting. An incremental mechanism is a sequential mechanism in which the payment offered to a player is equal to her *incremental cost*, i.e. the increase in cost caused by adding her to the set of previously selected players. He claimed that for supermodular cost functions, incremental mechanisms are group-strategyproof and budget balanced. However, this statement is flawed (as indicated in [10]) and holds only under the assumption that players are never indifferent.

We extend the characterizations for group-strategyproof mechanisms to the general demand online setting. The mechanisms in our online model always accept indifferent players and thus fulfill Juarez' MAX property. This allows us to guarantee group-strategyproofness for all incremental mechanisms derived from *submodular* cost functions. Moreover, we achieve weak group-strategyproofness for the whole class of games with increasing marginal cost functions.

2 Online General Demand Cost Sharing Games

We first review *offline general demand cost sharing games* as studied in [1, 4, 11, 12]. Let U be a set of players that are interested in a common service. In a *general demand* cost sharing game, every player has valuations for a finite number of service levels, i.e. the maximum service level requested is bounded by a constant $L \in \mathbb{N}$. Let (i, l) denote player i's request for service level l. Each player $i \in U$ has a *valuation* vector $v_i \in \mathbb{R}_+^L$, where $v_{i,l}$ denotes how much more (additive) player i likes service level l compared to service level $l-1$. The valuation vectors are private information, i.e. v_i is known to i only. Additionally, each player i announces a *bid* vector $b_i \in \mathbb{R}_+^L$. $b_{i,l}$ represents the maximum price player i is willing to pay for service level l (in addition to service level $l-1$).

An *allocation* of goods or service to the set of players U is denoted by a vector $\mathbf{x} \in \mathbb{N}_0^U$, where $x_i \in \mathbb{N}_0$ indicates the level of service that player i obtains; here $x_i = 0$ represents that i does not receive any good or service. Note that as a characteristic of this model, only subsequent service levels can be allocated to a player (i.e. if a player obtains service level l, then she also obtains service levels $1, \ldots, l-1$). We denote by $\mathbf{e}_i \in \mathbb{N}_0^U$ the ith unit vector.

The *servicing cost* of an allocation $\mathbf{x} \in \mathbb{N}_0^U$ is given by a cost function $C : \mathbb{N}_0^U \to \mathbb{R}_+$. We assume that C is non-decreasing in every component and

[3] A player is said to be *indifferent* if her bid is equal to her requested payment.

$C(\mathbf{0}) = 0$ for the all-zero allocation $\mathbf{0}$. In the examples we study, the common service is represented by a combinatorial optimization problem like e.g. Steiner tree, machine scheduling, etc. In these cases, we define $C(\mathbf{x})$ as the cost of an offline optimal solution to the underlying optimization problem.

A general demand *cost sharing mechanism* solicits the bid vectors b_i from all players $i \in U$, and computes a service allocation $\mathbf{x} \in \mathbb{N}_0^U$ and a payment $\phi_{i,l} \in \mathbb{R}$ for every player $i \in U$ and service level $l \leq L$. We sometimes write $\mathbf{x}(\mathbf{b})$ and $\phi(\mathbf{b})$ to refer to the outcome resulting from bid vector $\mathbf{b}$. We assume that the mechanism complies with the following standard assumptions:

1. *Individual rationality*: A player is charged only for service levels that she receives, and for any service level, her payment is at most her bid, i.e. for all i, l: $\phi_{i,l} = 0$ if $x_i < l$ and $\phi_{i,l} \leq b_{i,l}$ if $x_i \geq l$.
2. *No positive transfer*: A player is not paid for receiving service, i.e. $\phi_{i,l} \geq 0$ for all i, l.
3. *Consumer sovereignty*: A player is guaranteed to receive an additional service level if she bids high enough, i.e. there exists a threshold value $b^*_{i,l}$ for each player i and service level l such that $x_i \geq l$ if $b_{i,l} \geq b^*_{i,l}$ and $x_i \geq l - 1$.

For notational convenience, we define $v_{i,0} = \phi_{i,0} = 0$ for all players $i \in U$.

Let $\bar{C}(\mathbf{x})$ denote the cost of the actually computed solution for allocation $\mathbf{x}$. A cost sharing mechanism is *β-budget balanced* if the total payment obtained from all players deviates by at most a factor $\beta \geq 1$ from the total cost, i.e. $\bar{C}(\mathbf{x}) \leq \sum_{i \in U} \sum_{l=1}^{L} \phi_{i,l} \leq \beta \cdot C(\mathbf{x})$. If $\beta = 1$, we simply call the cost sharing mechanism budget balanced.

We assume that players act strategically and each player's goal is to maximize her own utility. The *utility* of player i is defined as $u_i(\mathbf{x}, \phi) := \sum_{l=1}^{x_i} (v_{i,l} - \phi_{i,l})$. Since the outcome $(\mathbf{x}, \phi)$ computed by the mechanism depends on the bids $\mathbf{b}$ of the players (and not on their true valuations), a player may have an incentive to declare a bid vector b_i that differs from her true valuation vector v_i. We say that a mechanism is *strategyproof* if bidding truthfully is a dominant strategy for every player. That is, for every player $i \in U$ and every two bid vectors $\mathbf{b}, \mathbf{b}'$ with $b_i = v_i$ and $b_j = b'_j$ for all $j \neq i$, we have $u_i(\mathbf{x}, \phi) \geq u_i(\mathbf{x}', \phi')$, where $(\mathbf{x}, \phi)$ and $(\mathbf{x}', \phi')$ are the solutions output by the mechanism for bid vectors $\mathbf{b}$ and $\mathbf{b}'$, respectively. Note that in our model, a player cannot lie about the characteristics or arrival times of her requests.

In *cooperative* mechanism design, it is assumed that players can form coalitions in order to coordinate their bids. A mechanism is *group-strategyproof* if no coordinated bidding of a coalition $S \subseteq U$ can strictly increase the utility of some player in S without strictly decreasing the utility of another player in S. Formally, for every coalition $S \subseteq U$ and every two bid vectors $\mathbf{b}, \mathbf{b}'$ with $b_i = v_i$ for all $i \in S$ and $b_i = b'_i$ for all $i \notin S$, if there is some $i \in S$ with $u_i(\mathbf{x}', \phi') > u_i(\mathbf{x}, \phi)$ then there is some $j \in S$ with $u_i(\mathbf{x}', \phi') < u_i(\mathbf{x}, \phi)$. A mechanism is *weakly group-strategyproof* if no coordinated bidding can strictly increase the utility of *every* player in a coalition. That is, for every coalition $S \subseteq U$ and every two bid vectors $\mathbf{b}, \mathbf{b}'$ with $b_i = v_i$ for all $i \in S$ and $b_i = b'_i$ for all $i \notin S$, there is some $i \in S$

with $u_i(\mathbf{x}', \phi') \leq u_i(\mathbf{x}, \phi)$. Intuitively, weak group-strategyproofness suffices if we assume that players adopt a slightly more conservative attitude with respect to their willingness of joining a coalition: Group-strategyproofness is needed if a player will participate in a coalition even if her utility is not affected, while weak group-strategyproofness suffices if she only joins if she is strictly better off.

Online Model. Many cost sharing games studied in the literature are derived from combinatorial optimization problems. This motivates us to define *online cost sharing games* very generally depending on the varying online characteristics inherited from different *online* optimization problems [2].

The most important characteristic of our model is that an online mechanism must immediately fix the payment for a requested service at the point of time when it is revealed, without any knowledge about future requests. As in the offline setting, we assume that an online mechanism never accepts further requests from players that have previously been rejected. For cost sharing games that are derived from combinatorial optimization problems, the mechanism has to maintain a (possibly suboptimal) feasible solution for the current service allocation. The feasible modifications of this current solution are inherited from the underlying online optimization problem.

We use the *online list* model by Borodin et al. [2] to describe the proceeding of an online mechanism: Service requests (i, l) arrive according to an online order. (For certain problems like online scheduling, jobs may have release dates which are then treated as arrival times of the respective requests.) Upon arrival, the player reveals the characteristics of her new request (i.e. the input information for the underlying combinatorial optimization problem) and her bid $b_{i,l}$. The mechanism immediately offers her the additional service level at a price p that may depend on previous inputs and decisions only. If her bid $b_{i,l}$ is larger or equal to this price, the request is accepted and added to the current allocation. Otherwise, the request is rejected and all further appearances of player i are removed from the online list (formally, we set $p = \infty$ for all subsequent requests of player i). A more formal description is given in Algorithm 1.

Algorithm 1: Online general demand cost sharing mechanism.

Input: online cost sharing game
Output: allocation vector $\mathbf{x} = (x_i)_{i \in U}$, payment vector $\phi = (\phi_{i,l})_{i \in U, l \leq L}$

Initialize $\mathbf{x}^0 = \mathbf{0}$
forall *requests* $t \in T$ **do**
    Read out input data and bid $b_{i,l}$ of newly arrived request $t =: (i, l)$.
    Determine payment p for new request.
    if $b_{i,l} \geq p$ **then** set $\mathbf{x}^t = \mathbf{x}^{t-1} + \mathbf{e}_i$ and $\phi_{i,l} = p$
    else set $\mathbf{x}^t = \mathbf{x}^{t-1}$ and $\phi_{i,l} = 0$; delete all further appearances of player i.
end
Output allocation vector $\mathbf{x}$ and payments ϕ

Let $\mathbf{x}^t$ denote the current allocation after processing request $t \in T = \{1, 2, \ldots\}$. Let $\bar{C}(\mathbf{x}^t)$ denote the cost of the actually computed solution for $\mathbf{x}^t$. We call an online cost-sharing mechanism *β-budget balanced at all times* for some $\beta \geq 1$ if for all requests $t \in T$:

$$\bar{C}(\mathbf{x}^t) \leq \sum_{i \in U} \sum_{l=1}^{x_i^t} \phi_{i,l} \leq \beta \cdot C(\mathbf{x}^t).$$

The conditions of individual rationality and no positive transfer as well as the different forms of incentive compatibility transfer in a straightforward way.

3 Incentive Compatibility

The following characterizations hold for all online mechanisms in our framework. Note that the requirements for group-strategyproofness highly depend on the fact that requests are accepted if the announced bid is equal to the offered price.

Strategyproofness. To guarantee strategyproofness, we must bound the increase in marginal valuations of individual players. This is essential to prevent players from overbidding for some level to obtain positive utility for higher levels. In previous works on general demand cost sharing [1, 11], players' valuations were assumed to be non-increasing. However, we can slightly relax this condition by introducing a positive factor λ:

Definition 1. *A valuation vector* $v_i \in \mathbb{R}^L$ *is* λ-decreasing *if for all* $1 < l \leq L$,

$$v_{i,l} \leq \lambda \cdot v_{i,l-1}.$$

Given λ-decreasing valuations for all players, an online mechanism is guaranteed to be strategyproof if and only if its cost shares grow faster than the valuations.

Definition 2. *A cost sharing mechanism has* λ-increasing prices *if for every bid vector* $\mathbf{b}$ *and player* $i \in U$*, the price for any service level* $1 < l \leq x_i(\mathbf{b})$ *is at least* λ *times the price for the previous service level, i.e.* $\phi_{i,l}(\mathbf{b}) \geq \lambda \cdot \phi_{i,l-1}(\mathbf{b})$.

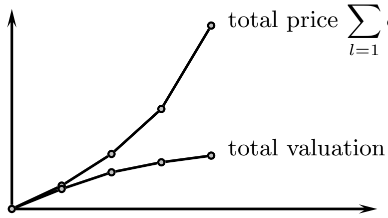

Fig. 1. Illustration of λ-decreasing valuations and λ-increasing prices for $\lambda = 1$

The above conditions can be further generalized by letting λ vary for every player (and/or level) or by adding constant terms to the right hand sides. However, the following fact emphasizes that a similar set of conditions are necessary to achieve strategyproofness. We omit the proof due to space restrictions.

Fact 1 *A general demand online mechanism is not strategyproof if cost shares do not increase by more than valuations per service level.*

Weak Group-Strategyproofness. We prove now that under the above conditions, an online mechanism is in fact weakly group-strategyproof.

Theorem 1. *If valuations are λ-decreasing, a general demand online cost sharing mechanism with λ-increasing prices is weakly group-strategyproof.*

Proof. Fix a coalition $S \subseteq U$ and a bid vector $\mathbf{b}$ with $b_i = v_i$ for all $i \in S$. Assume for contradiction that all members of the coalition can strictly increase their utilities by changing their bids to $\mathbf{b}'$ (while $b_i = b_i'$ for all $i \notin S$). Let (i, l) be the first request for which the mechanism makes different decisions in the runs on $\mathbf{b}$ and $\mathbf{b}'$. By the online character of the mechanism, the price offered for request (i, l) only depends on previous decisions and is thus equal in both runs. Let p denote this offer price. There are two possible cases:

1. $v_{i,l} < p \leq b'_{i,l}$. Then, $\phi_{i,l}(\mathbf{b}') = p$, and λ-decreasing valuations and λ-increasing prices yield $\ldots \leq \lambda^{-2} v_{i,l+2} \leq \lambda^{-1} v_{i,l+1} \leq v_{i,l} < \phi_{i,l}(\mathbf{b}') \leq \lambda^{-1}\phi_{i,l+1}(\mathbf{b}') \leq \lambda^{-2}\phi_{i,l+2}(\mathbf{b}') \leq \ldots$. Hence, player i has negative utility for service levels l and higher in the run on $\mathbf{b}'$, whereas when bidding truthfully, the utility for each level is non-negative.
2. $b'_{i,l} < p \leq v_{i,l}$. Then, player i obtains only $l - 1$ levels of service in the run on $\mathbf{b}'$, but she may get additional utility by accepting level l in the run on $\mathbf{b}$.

Hence, player i gets less or equal utility in the run on $\mathbf{b}'$, a contradiction. □

For binary demand cost sharing games, both Definitions 1 and 2 are always fulfilled since there is only one service level. Hence, quite remarkably, binary demand online cost sharing mechanisms are inherently weakly group-strategyproof.

Group-Strategyproofness. To ensure the stronger notion of group-strategyproofness, we must prevent that *dropping out*, i.e. underbidding in case of indifference, can help subsequent players. Towards this end, we introduce the following generalization of the well-known notion of cross-monotonicity for binary demand cost sharing games [12].

Consider a fixed instance of an online cost sharing game and let $\phi_{i,l}(\mathbf{b})$ denote the price that player i is offered for service level l when $\mathbf{b}$ is the bid vector input to the mechanism. Throughout this section, we assume λ-decreasing valuations and λ-increasing prices.

Definition 3. *An online mechanism is* cross-monotonic *if for every player $i \in U$ and service level l, the offered price does not decrease when a subset of requests are accepted in previous iterations, i.e.*

$$\phi_{i,l}(\mathbf{b}') \geq \phi_{i,l}(\mathbf{b})$$

for all bid vectors $\mathbf{b}, \mathbf{b}'$ such that $x^{t-1}(\mathbf{b}') \leq x^{t-1}(\mathbf{b})$, where (i, l) is request t.

This condition is sufficient for group-strategyproofness. The proof contains two main ideas: First, dropping out can never help others since it only increases cost shares of subsequent bidders. Second, the first member of a coalition who overbids for an additional level of service can only decrease her utility by doing this, since prices increase more than valuations in terms of service levels.

Theorem 2. *If valuations are λ-decreasing, an online cost sharing mechanism with λ-increasing prices is group-strategyproof if it is cross-monotonic.*

Proof. Fix a coalition $S \subseteq U$ and a bid vector $\mathbf{b}$ with $b_i = v_i$ for all $i \in S$. Assume that every member of the coalition increases or maintains her utility when the coalition changes their bids to $\mathbf{b}'$ (while $b_i = b'_i$ for all $i \notin S$).

We first prove that $\mathbf{x}^t(\mathbf{b}') \leq \mathbf{x}^t(\mathbf{b})$ for all $t \in T$. Assume for contradiction that there is a request which is accepted in the run on $\mathbf{b}'$ but not in the run on $\mathbf{b}$. Let (i, l) be the earliest such request, say request t. That is, $\mathbf{x}^\tau(\mathbf{b}') \leq \mathbf{x}^\tau(\mathbf{b})$ for all $\tau < t$. By cross-monotonicity, we have $\phi_{i,l}(\mathbf{b}') \geq \phi_{i,l}(\mathbf{b})$. Since players outside the coalition submit the same bids in both runs, player i must be a member of the coalition to gain service in the run on $\mathbf{b}'$. But then, $\phi_{i,l}(\mathbf{b}') \geq \phi_{i,l}(\mathbf{b}) > b_{i,l} = v_{i,l}$ and hence by λ-decreasing valuations and λ-increasing prices, player i has negative utility for service levels l and higher in the run on $\mathbf{b}'$. Since $\mathbf{x}^\tau(\mathbf{b}') \leq \mathbf{x}^\tau(\mathbf{b})$ for all $\tau < t$, by cross-monotonicity $\phi_{i,k}(\mathbf{b}') \geq \phi_{i,k}(\mathbf{b})$ for all $k < l$ as well, and therefore $u_i(\mathbf{b}') < u_i(\mathbf{b})$, contradicting the first assumption.

We can conclude that $\mathbf{x}^t(\mathbf{b}') \leq \mathbf{x}^t(\mathbf{b})$ for all $t \in T$. Hence, $\phi_{i,l}(\mathbf{b}') \geq \phi_{i,l}(\mathbf{b})$ for all i, l by cross-monotonicity. This means that

$$u_i(\mathbf{b}') = \sum_{l=1}^{x_i(\mathbf{b}')} (v_{i,l} - \phi_{i,l}(\mathbf{b}')) \leq \sum_{l=1}^{x_i(\mathbf{b})} (v_{i,l} - \phi_{i,l}(\mathbf{b})) = u_i(\mathbf{b})$$

for all i and l, hence we obtain group-strategyproofness. □

We prove next that the conditions in Theorem 2 are not only sufficient but also necessary, even in the binary demand case.

Theorem 3. *A binary demand online mechanism is not group-strategyproof if it is not cross-monotonic.*

Proof. Consider an online mechanism that is not cross-monotonic; let $L = 1$. That is, there are bid vectors $\mathbf{b}, \mathbf{b}'$ with $\mathbf{x}^{t-1}(\mathbf{b}') \leq \mathbf{x}^{t-1}(\mathbf{b})$ and $\phi_i(\mathbf{b}') < \phi_i(\mathbf{b})$ for some player i. For simplicity, assume that i is the last player in the online instance. Since the mechanism is online, $\phi_i(\mathbf{b}')$ does not depend on b'_i, so we can assume that $b'_i = \phi_i(\mathbf{b})$. We will define valuations such that there is a coalition S which has an incentive to misreport their valuations.

Define $S := \{j \in U \mid b_j \neq b'_j\} \cup \{i\}$. Assume that all $j \in U \setminus S$ bid $b_j = b'_j$. Now, define $v_j := \phi_j(\mathbf{b})$ for all $j \in S$. Observe that if all players in S bid truthfully, the outcome of the mechanism is the same as for bid vector $\mathbf{b}$. Now, if the coalition changes their bids to $\mathbf{b}'$, some players $j \in S \setminus \{i\}$ lose service but all retain their previous utility of zero. Meanwhile, player i increases her utility from zero to $\phi_i(\mathbf{b}) - \phi_i(\mathbf{b}') > 0$. Hence, the mechanism is not group-strategyproof. □

4 Incremental Online Mechanisms

We now describe a generic method to turn competitive online algorithms into online cost sharing mechanisms. Given a ρ-competitive algorithm ALG for a combinatorial optimization problem, we define an *incremental* online mechanism for the corresponding cost sharing game, which is ρ-budget balanced at all times. The mechanism is weakly group-strategyproof if the algorithm's marginal costs are increasing, which is gratuitous in the binary demand case.

Let ALG be a ρ-competitive algorithm for an online combinatorial optimization problem $\mathcal{P}$. Consider an instance $\mathcal{I}$ of $\mathcal{P}$. The incremental online mechanism induced by ALG works as follows: Requests arrive according to $\mathcal{I}$. Each time a new request arrives, we simulate ALG on the online instance induced by the requests that have previously been accepted and the new item. The price p for the additional service level is set to be the incremental cost caused by the update in the competitive algorithm. We call an online algorithm ALG *cross-monotonic* if the induced incremental online mechanism is cross-monotonic. It is straightforward to see that the budget balance factor of an incremental online mechanism is inherited from the competitive ratio of the input algorithm: In every iteration, the sum of the collected payments equals the cost inferred by the algorithm.

Lemma 1. *The incremental online mechanism is ρ-budget balanced at all times.*

4.1 Binary Demand Examples

To demonstrate the applicability of our framework, we apply it to competitive online algorithms for a number of combinatorial optimization problems. In this section, we give examples for *binary demand* cost sharing games, i.e. the maximum service level is $L = 1$ and every player has only one request.

Online Scheduling. Consider the parallel machine scheduling problem with the objective of minimizing the makespan. Any list scheduling algorithm has an approximation factor of at most 2 for this problem. Hence, the online algorithm that adds each arriving job to the machine with currently least load is 2-competitive. Unfortunately, it is not cross-monotonic as deleting jobs can cause higher or lower completion times for subsequent jobs. Thus, our framework yields a 2-budget balanced weakly group-strategyproof online mechanism.

Corollary 1. *There is a 2-budget balanced weakly group-strategyproof incremental online mechanism for the minimum makespan scheduling problem on parallel machines $P||C_{\max}$.*

Online Steiner Tree and Forest. Given an undirected graph G with edge costs, connection requests arrive online. In the Steiner forest problem, each request consists of a pair of terminals s_i, t_i; in the Steiner tree problem, all requests have one vertex in common, i.e. $s_i = s_j$ for all $i, j \in U$. The goal is to select a minimum cost set of edges such that each terminal pair is connected by a path. Let n denote the number of players (i.e., terminal pairs).

The online greedy Steiner tree algorithm picks the shortest path to the current tree each time a new terminal pair arrives. It has a competitive ratio of $\log n$, while the competitive ratio of any online algorithm is known to be at least $1/2 \log n$ [8]. Hence, our framework gives a weakly group-strategyproof $\Theta(\log n)$-budget balanced online cost sharing mechanism for the Steiner tree problem, which is asymptotically best possible. The greedy algorithm for the online Steiner forest problem achieves an approximation ratio of $O(\log^2 n)$.

Corollary 2. *There is an $O(\log^2 n)$-budget balanced weakly group-strategyproof incremental online mechanism for the Steiner forest game. This mechanism is $(\log n)$-budget balanced for the Steiner tree game.*

Unfortunately, the greedy algorithm is not cross-monotonic, as the removal of some players may cause other players to switch their paths, which in turn can have arbitrary effects on the costs incurred by subsequent players. This issue can be overcome if paths are unambiguous; e.g. if $G = (V, E)$ is a forest, the above mechanisms are group-strategyproof. Pushing this insight further, we obtain an $O(\log |V|)$-budget balanced group-strategyproof mechanism for the Steiner forest game if the underlying graph is known in advance: We use the *oblivious* online Steiner forest algorithm by Gupta et al. [6], which essentially works as follows: Given the underlying graph, the algorithm precomputes a collection of paths. When a new terminal pair arrives, it simply connects it by one of the predefined paths. The authors show that one can identify a collection of paths such that the resulting algorithm is $O(\log |V|)$-competitive. Since the used paths are defined in advance, a player can only benefit from the presence of other players, who might pay for parts of her designated path. Hence, we obtain cross-monotonicity without losing much in terms of the budget balance guarantee.

Corollary 3. *There is an $O(\log^2 |V|)$-budget balanced group-strategyproof incremental online mechanism for the Steiner forest game, where V is the vertex set of the underlying graph.*

We believe that such "universal" algorithms that determine generic solutions without knowing the upcoming instance will also yield group-strategyproof online mechanisms for several other interesting problems like e.g. the traveling salesman problem.

4.2 General Demand Examples

We now exploit the whole range of our framework by deriving incremental mechanisms for *general demand* cost sharing games. In the first example, players arrive only once with the complete list of their requests. In the second example, the arrival sequence is mixed, i.e. players take turns announcing additional requests.

Online Preemptive Scheduling. A common problem in preemptive scheduling is the parallel machine setting in which each job has a release date. The cost of a solution is given by the sum of all completion times. The single machine case is solved optimally by the *shortest remaining processing time* (SRPT) algorithm [13]. SRPT is a 2-approximation for the parallel machine case [5].

In the corresponding cost sharing game, we treat the release date of a job as its arrival time. Upon arrival, each player may request multiple executions of her job. In scheduling terms, each player owns a set of jobs which all have the same release date and processing time. E.g. imagine a student who asks a copy shop to print and bind several copies of his thesis, or a joinery is asked to produce a few of the same individual piece of furniture. In such scenarios, it is very natural to assume that the marginal valuation for each additional copy is decreasing, i.e. $v_{i,l} \geq v_{i,l+1}$ for all i, l.

SRPT schedules all of a player's jobs subsequently. Hence, each of them delays the same number of jobs, and later copies have larger completion times. Therefore, the mechanism induced by SRPT has increasing marginal prices.

Corollary 4. *There is a 2-budget balanced weakly group-strategyproof general demand incremental online mechanism for the preemptive scheduling problem with release dates* $P|r_i, pmtn|\sum C_i$. *This mechanism is budget balanced in the single machine case.*

Online Multicommodity Routing. In an online multicommodity routing problem, we are given a directed graph with monotonically increasing cost functions on each arc. Commodities arrive online and request routing of l units from some vertex to another. We assume that the routing is splittable in integer units. The greedy algorithm which routes each unit of flow separately in an optimal way is $(3 + 2\sqrt{2})$-competitive for this problem [7]. Marginal costs are increasing since the cost functions on each arc grow with increasing traffic. This is true even when players arrive in a mixed order and request to route additional units between their source-destination pair. However, this is a congestion-type game (the more players in the game, the higher the costs per request), and so we cannot expect group-strategyproofness.

Corollary 5. *There is a* $(3 + 2\sqrt{2})$*-budget balanced weakly group-strategyproof incremental online mechanism for the online multicommodity routing problem in which each player arrives multiple times.*

5 Conclusion

We characterized strategyproofness, weak group-strategyproofness and group-strategyproofness of mechanisms in a new framework for online general demand cost sharing games. Quite notably, weak group-strategyproofness comes for free for binary demand problems. Online mechanisms are group-strategyproof if and only if *dropping out* cannot help subsequent players. Consequently, we cannot expect incremental cost sharing mechanisms for problems with congestion effects like e.g. scheduling games to be group-strategyproof, while this seems easier for network design problems.

In the offline setting, finding a good order to consider players is the key to derive cost sharing mechanisms with attractive budget balance factors (see [3]). In the online case, this order is determined by an adversary and thus not under the control of the mechanism designer, which strongly constrains the possibilities of designing valuable cost sharing mechanisms. However, our results prove that there is no gap between the best possible competitive ratio of an online algorithm and the best possible budget balance factor of a weakly group-strategyproof online cost sharing mechanism.

We consider this work as a very natural and general starting point to exploit cooperative cost sharing in an online context. It would be interesting to see more applications to our framework, with or without usage of the direct derivation of incremental mechanisms from competitive algorithms. Our model restricts feasible allocations to a continuous sequence of accepts for each player, starting with their first request. This feature of the model is crucial for truthfulness as it prevents players from underbidding to reject some service request and then obtain it later for a cheaper price. One interesting line of research would be to allow for more general mechanisms which might accept further requests of players even after a request has been rejected.

References

1. Y. Bleischwitz and F. Schoppmann. Group-strategyproof cost sharing for metric fault tolerant facility location. In *Proc. of the 1st Int. Sympos. on Algorithmic Game Theory*, volume 4997 of *LNCS*, pages 350–361, 2008.
2. A. Borodin and R. El-Yaniv. *Online computation and competitive analysis*. Cambridge University Press, New York, NY, USA, 1998.
3. J. Brenner and G. Schäfer. Singleton acyclic mechanisms and their applications to scheduling problems. In *Proc. of the 1st Int. Sympos. on Algorithmic Game Theory*, volume 4997 of *LNCS*, pages 315–326, 2008.
4. N. Devanur, M. Mihail, and V. Vazirani. Strategyproof cost-sharing mechanisms for set cover and facility location games. *Decision Support Syst.*, 39(1):11–22, 2005.
5. J. Du, J. Y. T. Leung, and G. H. Young. Minimizing mean flow time with release time constraint. *Theoretical Computer Sci.*, 75(3):347–355, 1990.
6. A. Gupta, M. Hajiaghayi, and H. Räcke. Oblivious network design. In *Proc. of the 17th ACM-SIAM Sympos. on Discrete Algorithms*, pages 970–979, 2006.
7. T. Harks, S. Heinz, and M. Pfetsch. Competitive online multicommodity routing. *Theory of Computing Systems*, 2008.
8. M. Imase and B. Waxman. Dynamic Steiner tree problems. *SIAM J. Disc. Math.*, 4(3):369–384, 1991.
9. N. Immorlica, M. Mahdian, and V. S. Mirrokni. Limitations of cross-monotonic cost-sharing schemes. *ACM Trans. Algorithms*, 4(2):1–25, 2008.
10. R. Juarez. Group strategyproof cost sharing. In *Unpublished*, 2008.
11. A. Mehta, T. Roughgarden, and M. Sundararajan. Beyond Moulin mechanisms. In *Proc. of the ACM Conference on Electronic Commerce*, 2007.
12. H. Moulin. Incremental cost sharing: Characterization by coalition strategyproofness. *Soc. Choice Welfare*, 16:279–320, 1999.
13. L. Schrage. A proof of the optimality of the shortest remaining processing time discipline. *Operations Res.*, 16:687–690, 1968.