

Optimal Mechanisms for Single Machine Scheduling

Birgit Heydenreich¹, Debasis Mishra², Rudolf Müller¹, and Marc Uetz³

¹ Maastricht University, Quantitative Economics, P.O.Box 616, 6200 MD Maastricht, The Netherlands. {b.heydenreich,r.muller}@ke.unimaas.nl

² Indian Statistical Institute, Planning Unit, 7, S.J.S. Sansanwal Marg, New Delhi - 110 016, India. dmishra@isid.ac.in

³ University of Twente, Applied Mathematics, P.O. Box 217, 7500 AE Enschede, The Netherlands. m.uetz@utwente.nl

Abstract We study the design of optimal mechanisms in a setting where job-agents compete for being processed by a service provider that can handle one job at a time. Each job has a processing time and incurs a waiting cost. Jobs need to be compensated for waiting. We consider two models, one where only the waiting costs of jobs are private information (1-d), and another where both waiting costs and processing times are private (2-d). An optimal mechanism minimizes the total expected expenses to compensate all jobs, while it has to be Bayes-Nash incentive compatible. We derive closed formulae for the optimal mechanism in the 1-d case and show that it is efficient for symmetric jobs. For non-symmetric jobs, we show that efficient mechanisms perform arbitrarily bad. For the 2-d case, we prove that the optimal mechanism in general does not even satisfy IIA, the ‘independent of irrelevant alternatives’ condition. We also show that the optimal mechanism is not even efficient for symmetric agents in the 2-d case

1 Introduction

The design of optimal auctions is recognized as an intriguing issue in auction theory; first studied by Myerson (1981) for single item auctions. In that setting, the goal is to maximize the seller’s revenue. We study the design of optimal auctions (or more precisely, mechanisms) in a setting where job-agents compete for being processed by a service provider that can only handle one job at a time.

Our results. We consider two cases. In the *one-dimensional* (1-d) case, jobs’ processing times are public information and a job’s weight is only known to the job itself. Publicly known probability distributions over a finite set of possible weights represent common beliefs about the weights. In the *two-dimensional* (2-d) case, both weights and processing times are private information of the jobs. In both cases we aim at finding Bayes-Nash incentive compatible mechanisms that minimize the expected expenses of the service provider. Given jobs’ reports about their private information, a mechanism determines both an order in which jobs are served, and for each job a payment that the job receives. The

payment can be seen as a compensation for waiting. By a graph theoretic interpretation of the incentive compatibility constraints - as used e.g. by Rochet [12] and Malakhov and Vohra [7] - we are able to derive optimal mechanisms. For the one-dimensional case, we obtain closed formulae for modified job weights, and show that serving the jobs in the order of non-increasing ratios of modified weights over service times is optimal for the service provider, as long as a certain regularity condition is fulfilled. It turns out that the optimal mechanism is not necessarily efficient, i.e., in general it does not maximize total utility. But it does so if e.g. all jobs are symmetric. For non-symmetric jobs, we show by example that the objective can be arbitrarily far from optimal if we insist on efficiency. We also compare our optimal mechanism to the generalized VCG mechanism and see that expected payments differ even for the case of symmetric jobs. For the two-dimensional case, our main result is that the optimal mechanism generally does not satisfy a property called IIA, ‘independent of irrelevant alternatives’. That implies that the optimal mechanism cannot be expressed in terms of modified weights along the lines of the 1-d case. In fact, any kind of priority based list scheduling algorithm where the priorities of a job depend only on the characteristics of that job itself cannot in general be an optimal mechanism. We conclude that optimal mechanism design for the two-dimensional case is substantially more involved than two-dimensional mechanism design for auction settings, as studied in [7]. We also show that even for symmetric jobs, in the 2-d case the optimal mechanism is not efficient.

Related Work. Myerson [11] studies optimal mechanisms for single item auctions and continuous 1-dimensional type spaces. Here, optimal auctions are modified Vickrey auctions, i.e. modified efficient auctions. When regarding the seller as additional agent who bids zero in the original auction, his modified bid might become non-zero in the optimal auctions yielding a reservation price. For a comparison between Myerson’s and our results, see Section 3. In [4], the authors give an introduction to optimal mechanism design with 1-dimensional continuous types under dominant strategy incentive compatibility. Both Myerson’s and our optimal allocation rules turn out to be dominant strategy implementable as well, while they yield optimal mechanisms in the larger class of Bayes-Nash incentive compatible mechanisms. Malakhov and Vohra [7] regard optimal mechanisms for an auction setting with discrete 2-dimensional type spaces. The derived optimal mechanisms again employ the efficient allocation rule with modified bids. We show that their approach must fail in our setting. For details, we refer to Section 4. Armstrong [1] studies a multi-object auction model where valuations are additive and drawn from a binary distribution (i.e. high or low). He gives optimal auctions under specific conditions that reduce the type graph. From this paper it becomes evident that optimal mechanism design with multi-dimensional discrete types is difficult. For our model, we formalize this difficulty by showing that traditional approaches inevitably yield IIA-mechanisms and therefore must fail. Other scheduling models have been looked at from a different angle in the economic literature. See, e.g., [8] for efficient and budget-balanced mechanism

design in a 1-dimensional model and [9] for mechanisms that prevent merging and splitting of jobs.

2 Optimal Mechanisms for the 1-Dimensional Setting

Setting & Preliminaries. Consider a single machine which can handle one job at a time. Let $J = \{1, \dots, n\}$ denote the set of non-preemptive jobs. We regard jobs as selfish agents that act strategically. Each job j has a processing time p_j and a weight w_j . While p_j is publicly known, the actual w_j is private information to job j . We refer to the private information of a job as its type. Jobs share common beliefs about other jobs' types in terms of probability distributions. We assume discrete distribution of weights, that is, agent j 's weight w_j follows a probability distribution over the discrete set $W_j = \{w_j^1, \dots, w_j^{m_j}\} \subset \mathbb{R}$, where $w_j^1 \leq \dots \leq w_j^{m_j}$. Let φ_j be the probability distribution of w_j , that is, $\varphi_j(w_j^i)$ denotes the probability associated with w_j^i for $i = 1, \dots, m_j$. Let $\Phi_j(w_j^i) = \sum_{k=1}^i \varphi_j(w_j^k)$ be the cumulative probability up to w_j^i . Both φ_j and Φ_j are public information. We assume that jobs' weights are independently distributed. Let us denote by $W = \prod_{j \in J} W_j$ the set of all type profiles. For any job j , let $W_{-j} = \prod_{k \neq j} W_k$. Let φ be the joint probability distribution of $w = (w_1, \dots, w_n)$. Then $\varphi(w) = \prod_{j=1}^n \varphi_j(w_j^i)$ for $w = (w_1^i, \dots, w_n^i) \in W$. Let w_{-j} and φ_{-j} be defined analogously. For $w_j^i \in W_j$ and $w_{-j} \in W_{-j}$, we denote by (w_j^i, w_{-j}) the type profile where job j has type w_j^i and the types of all other jobs are w_{-j} .

A direct revelation mechanisms consists of an allocation rule f and a payment scheme π . Jobs have to report their weights and they might report untruthfully if it suits them. Depending on those reports, the allocation rule selects a *schedule*, i.e. an order in which jobs are processed on the machine. The payment scheme assigns a payment that is made to jobs in order to reimburse them for their waiting cost. The payments can be seen as a reimbursement for waiting.

Let $\mathfrak{S} = \{\sigma \mid \sigma \text{ is a permutation of } (1, \dots, n)\}$ denote the set of all feasible schedules. Then the allocation rule is a mapping $f: W \rightarrow \mathfrak{S}$. For any schedule $\sigma \in \mathfrak{S}$, let σ_j be the position of job j in the ordering of jobs in σ . Then, by $S_j(\sigma) = \sum_{\sigma_k < \sigma_j} p_k$, we denote the start time or waiting time of job j in σ . If job j has waiting time S_j and actual weight w_j^i , it encounters a valuation of $-w_j^i S_j$. If j additionally receives payment π_j , his total utility is $\pi_j - w_j^i S_j$, i.e., we assume quasi-linear utilities. Let us denote by $ES_j(f, w_j^i) := \sum_{w_{-j} \in W_{-j}} S_j(f(w_j^i, w_{-j})) \varphi_{-j}(w_{-j})$ the expected waiting time of job j if it reports weight w_j^i and allocation rule f is applied. Denote by $E\pi_j(w_j^i) := \sum_{w_{-j} \in W_{-j}} \pi_j(w_j^i, w_{-j}) \varphi_{-j}(w_{-j})$ the expected payment to j . We assume that jobs aim at maximizing their expected utility.

Definition 1. A mechanism (f, π) is Bayes-Nash incentive compatible if for every agent j and every two types $w_j^i, w_j^k \in W_j$

$$E\pi_j(w_j^i) - w_j^i ES_j(f, w_j^i) \geq E\pi_j(w_j^k) - w_j^k ES_j(f, w_j^k) \quad (1)$$

under the assumption that all agents apart from j report truthfully. If for allocation rule f there exists a payment scheme π such that (f, π) is Bayes-Nash incentive compatible, then f is called Bayes-Nash implementable. The payment scheme π is referred to as an incentive compatible payment scheme.

In order to account for individual rationality, we need to guarantee non-negative utilities for all agents that report their true weight. To that end, we add a dummy weight $w_j^{m_j+1}$ to the type space W_j for every agent j . We assume $ES_j(f, w_j^{m_j+1}) = 0$ and $E\pi_j(w_j^{m_j+1}) = 0$ for all $j \in J$. Furthermore, we impose the incentive constraints $E\pi_j(w_j^i) - w_j^i ES_j(f, w_j^i) \geq E\pi_j(w_j^{m_j+1}) - w_j^i ES_j(f, w_j^{m_j+1})$ implying that $E\pi_j(w_j^i) - w_j^i ES_j(f, w_j^i) \geq 0$ for any Bayes-Nash incentive compatible mechanism (f, π) . Therefore, the dummy weights together with the mentioned assumptions guarantee that individual rationality is satisfied along with the incentive constraints. The dummy weight can be interpreted as an option for any job not to take part in the mechanism.

Definition 2. An allocation rule f satisfies monotonicity w.r.t. weights or short monotonicity if for every agent $j \in J$, $w_j^i < w_j^k$ implies that $ES_j(f, w_j^i) \geq ES_j(f, w_j^k)$.

Theorem 1. An allocation rule f is Bayes-Nash incentive compatible if and only if it satisfies monotonicity w.r.t. weights.

The proof is standard and therefore omitted. We refer, e.g., to [10] for details.

The Type Graph. A useful tool for deriving optimal mechanisms is the type graph. It has been used earlier, e.g. in [6,7,10]⁴. The *type graph*⁵ T_f is defined for a fixed agent j . T_f has node set W_j and contains an arc from any node w_j^i to any other node w_j^k of length

$$\ell_{ik} = w_j^i [ES_j(f, w_j^k) - ES_j(f, w_j^i)].$$

Here, ℓ_{ik} represents the gain in expected valuation for agent j by truthfully reporting type w_j^i instead of lying type w_j^k . The incentive constraints for a Bayes-Nash incentive compatible mechanism (f, π) and job j can be read as

$$E\pi_j(w_j^k) \leq E\pi_j(w_j^i) + w_j^i [ES_j(f, w_j^k) - ES_j(f, w_j^i)] = E\pi_j(w_j^i) + \ell_{ik}.$$

That is, the expected payments $E\pi_j(\cdot)$ constitute a *node potential* in T_f . A standard result in graph theory says that these node potentials exist if and only if there is no negative cycle in the graph. That is, Bayes-Nash implementability of an allocation rule f is equivalent to the fact that the type graph T_f for any agent j has no negative cycle. We then say that the T_f 's satisfy the *non-negative cycle property*. Monotonicity is equivalent to the fact that there is no negative cycle

⁴ The exact definitions of the type graph might differ in the papers depending on the underlying model.

⁵ We suppress the dependence on agent j in the notation and simply write T_f .

of length two in T_f . We call this property the *non-negative two-cycle property*. It follows from

$$\begin{aligned}\ell_{ik} + \ell_{ki} &= w_j^i[ES_j(f, w_j^k) - ES_j(f, w_j^i)] + w_j^k[ES_j(f, w_j^i) - ES_j(f, w_j^k)] \\ &= (w_j^i - w_j^k)[ES_j(f, w_j^k) - ES_j(f, w_j^i)].\end{aligned}$$

The last term is non-negative for all jobs j and any two types w_j^i and w_j^k iff monotonicity holds.

Optimal Mechanisms. It is well known that scheduling in order of non-increasing weight over processing time ratios minimizes the sum of weighted start times $\sum_{j=1}^n w_j S_j(f(w))$ for any type profile $w \in W$, and therefore maximizes the total valuation of all agents. This allocation rule is known as Smith's rule [13]. The optimal mechanism that we derive deploys a slightly different allocation rule, namely Smith's rule with respect to certain modified weights.

Our goal is to set up a mechanism that is Bayes-Nash incentive compatible and among all such mechanisms minimizes the expected total payment that has to be made to the jobs. Given any Bayes-Nash incentive compatible mechanism (f, π) , one can obviously substitute the payment scheme by its expected payment scheme yielding $(f, E\pi(\cdot))$ without losing Bayes-Nash incentive compatibility. Moreover, the expected total payment to the agents remains unchanged under the substitution. Therefore, we restrict focus to mechanisms in which agents always receive a payment that is equal to the expected payment given the agent's report and which is independent of the specific report of the other agents and of the actual allocation.

Note that, unlike e.g. in [11], in the discrete setting considered here revenue equivalence does not hold. Therefore, there are possibly multiple payment schemes that make an allocation rule incentive compatible. Let f be an allocation rule and let $\pi^f(\cdot)$ be a payment scheme that minimizes expected expenses for the machine among all payment schemes that make f Bayes-Nash incentive compatible. More specifically, $\pi_j^f(w_j^i)$ denotes the payment to agent j declaring weight w_j^i under this optimal payment scheme. Let $P^{min}(f) = \sum_{j \in J} \sum_{w_j^i \in W_j} \varphi_j(w_j^i) \pi_j^f(w_j^i)$ be the minimum expected total expenses for allocation rule f . The following lemma specifies the optimal payment scheme for a given allocation rule.

Lemma 1. *For a Bayes-Nash implementable allocation rule f , the payment scheme defined by*

$$\pi_j^f(w_j^{m_j+1}) = 0, \quad \pi_j^f(w_j^i) = \sum_{k=i}^{m_j} w_j^k [ES_j(f, w_j^k) - ES_j(f, w_j^{k+1})] \text{ for } i = 1, \dots, m_j$$

is incentive compatible, individually rational and minimizes the expected total payment made to agents. The corresponding expected total payment is given by

$$P^{min}(f) = \sum_{j \in J} \sum_{i=1}^{m_j} \varphi_j(w_j^i) \bar{w}_j^i ES_j(f, w_j^i),$$

where the modified weights \bar{w}_j are defined as follows

$$\bar{w}_j^1 = w_j^1, \quad \bar{w}_j^i = w_j^i + (w_j^i - w_j^{i-1}) \frac{\Phi_j(w_j^{i-1})}{\varphi_j(w_j^i)} \text{ for } i = 2, \dots, m_j.$$

The proof relies on the observation that minimal expected payments can be expressed as shortest path lengths in the type graph; we refer to the full version [5] for more details.

Given the minimum payments per allocation rule, we want to specify the allocation rule f which minimizes $P^{\min}(f)$ among all Bayes-Nash implementable allocation rules.

Definition 3. If $f \in \arg \min\{P^{\min}(f) \mid f: W \rightarrow \mathfrak{S}, f \text{ Bayes-Nash implementable}\}$, then we call the mechanism (f, π^f) an optimal mechanism.

We will need the following regularity condition that ensures Bayes-Nash implementability of the allocation rule in our optimal mechanism.

Definition 4. We say that regularity is satisfied if for every agent j and $i = 2, \dots, m_j - 1$

$$w_j^i + (w_j^i - w_j^{i-1}) \frac{\Phi_j(w_j^{i-1})}{\varphi_j(w_j^i)} \leq w_j^{i+1} + (w_j^{i+1} - w_j^i) \frac{\Phi_j(w_j^i)}{\varphi_j(w_j^{i+1})}.$$

This implies that $\bar{w}_j^i < \bar{w}_j^k$ whenever $w_j^i < w_j^k$.

Note that regularity is satisfied e.g. if the differences $w_j^i - w_j^{i-1}$ are constant and the distribution has a non-increasing reverse hazard rate.

Theorem 2. Let the modified weights be defined as in Lemma 1. Let f be the allocation rule that schedules jobs in order of non-increasing ratios \bar{w}_j/p_j . If regularity holds, then (f, π^f) is an optimal mechanism.

Proof. We show that f is Bayes-Nash implementable and minimizes $P^{\min}(f)$ among all Bayes-Nash implementable allocation rules. For any allocation rule f , it is not hard to see that we can rewrite $P^{\min}(f)$ as follows, using independence of weight distributions. Let $W'_j = W_j \setminus \{w_j^{m_j+1}\}$ and $W' = \prod_{j \in J} W'_j$.

$$\begin{aligned} P^{\min}(f) &= \sum_{j \in J} \sum_{w_j^i \in W'_j} \varphi_j(w_j^i) \bar{w}_j^i ES_j(f, w_j^i) \\ &= \sum_{w \in W'} \varphi(w) \sum_{j \in J} \bar{w}_j S_j(f(w)). \end{aligned}$$

Thus, $P^{\min}(f)$ can be minimized by minimizing $\sum_{j \in J} \bar{w}_j S_j(f(w))$ for every reported type profile w . This is achieved by scheduling in order of non-increasing ratios \bar{w}_j/p_j . Under Smith's rule, the expected start time $ES_j(w_j)$ is clearly non-increasing in the modified weight \bar{w}_j . The regularity condition ensures that it is non-increasing in the original weights w_j . Therefore, Smith's rule with respect to modified weights satisfies monotonicity and is hence Bayes-Nash implementable by Theorem 1. This completes the proof. \square

It is not hard to see that the optimal allocation rule – Smith’s rule with respect to modified weights – is even dominant strategy implementable, with the same total expected payment for the mechanism.

3 Optimality versus Efficiency

For symmetric agents the optimal and the efficient allocation coincide.

Corollary 1. *If agents are symmetric, i.e. $W_1 = \dots = W_n$, $\varphi_1 = \dots = \varphi_n$ and $p_1 = \dots = p_n$ and if distributions are such that regularity holds, then the optimal mechanism is efficient.*

If weight distributions differ among agents or if agents have different processing times, then the optimal mechanism is in general not efficient. In fact, when restricting to efficient mechanisms, the total expected payment can be arbitrarily bad in comparison to the optimal one. This is illustrated by the following two examples; proofs can be found in the full version of this paper [5].

Example 1. Let there be two jobs 1 and 2 with $W_1 = \{M+1\}$ and $W_2 = \{1, M\}$ for some constant M . Let $\varphi_2(1) = 1 - 1/M$, $\varphi_2(M) = 1/M$ and $p_1 = p_2 = 1$. Let Eff be the efficient and Opt be the optimal allocation rule. Then the ratio $P^{min}(Eff)/P^{min}(Opt)$ goes to infinity as M goes to infinity.

Remark 1. In the above, the ratio of the expected payments of the efficient versus the optimal allocation rule is analyzed. It is also easy to see that the expected ratio of the payments tends to infinity as M approaches infinity.

Example 2. Let there be two jobs 1 and 2 with the same weight distribution $W_1 = W_2 = \{1, M\}$, $\varphi_j(1) = 1 - 1/M$, $\varphi_j(M) = 1/M$ for $j = 1, 2$. Let $p_1 = 1/2$ and $p_2 = M/2 + 1$. Let Eff be the efficient and Opt be the optimal allocation rule. Then the ratio $P^{min}(Eff)/P^{min}(Opt)$ goes to infinity as M goes to infinity.

Remark 2. As in the first example, it is easy to see that the expected ratio of the payments tends to infinity as M approaches infinity.

Comparison to Myerson’s result. For the single item auction and continuous type spaces, Myerson [11] has made similar observations: in his setting, the efficient auction is the Vickrey auction. The optimal auction can be seen as a modified Vickrey Auction with the seller submitting a bid himself. In our setting also, the allocation in the optimal mechanism is equivalent to the efficient allocation rule with respect to modified data. Nevertheless, in [11] the optimal and the efficient mechanism may differ. For the single item auction this can be due to the seller keeping the item (even in the symmetric case) or because a bidder that has not submitted the highest bid can get the item in the asymmetric case. In our setting, the optimal and the efficient mechanism can only differ if agents are asymmetric, see Corollary 1 and Examples 1 and 2.

On the generalized VCG Mechanism. The VCG mechanism is due to Vickrey [14], Clarke [2] and Groves [3]. The allocation rule is the efficient one. In our setting this means scheduling in order of non-increasing ratios w_j/p_j . The payment scheme can be shown to be

$$\pi_j^{VCG}(w) = p_j \sum_{\substack{k \in J \\ \sigma_k < \sigma_j}} w_k,$$

where w is the reported type profile and σ the efficient schedule. As illustrated by examples 1 and 2, the allocation of the VCG mechanism can differ from the allocation of the optimal mechanism if agents are not symmetric. Moreover, if jobs are symmetric, the VCG mechanism still can be non-optimal in terms of payments. This is illustrated by the following example.

Example 3. There are two symmetric agents with $W_1 = W_2 = \{w^1, w^2\}$, $w^1 < w^2$, and $\varphi_j(w^1) = \varphi_j(w^2) = 1/2$ for $j = 1, 2$. Processing times are equal (w.l.o.g., $p_1 = p_2 = 1$). Then the expected expenses of the VCG mechanism can be shown to be strictly higher than those of the optimal mechanism.

4 The 2-Dimensional Setting

Setting and Notation. In contrast to the 1-dimensional setting, both weight and processing time of a job are now private information of the job. Hence j 's type is the tuple (w_j, p_j) . We assume public probability distribution information, i.e. $(w_j, p_j) \in W_j \times P_j$, where $W_j = \{w_j^1, \dots, w_j^{m_j}\}$ with $w_j^1 \leq \dots \leq w_j^{m_j}$ and $P_j = \{p_j^1, \dots, p_j^{q_j}\}$ with $p_j^1 \leq \dots \leq p_j^{q_j}$. Let φ_j be the probability distribution of j 's type, that is, $\varphi_j(w_j^i, p_j^k)$ denotes the probability associated with the type (w_j^i, p_j^k) for $i = 1, \dots, m_j$ and $k = 1, \dots, q_j$. Both φ_j and Φ_j are public. Distributions are independent between agents. Denote by $T = \prod_{j \in J} (W_j \times P_j)$ the set of all type profiles. For any job j , let $T_{-j} = \prod_{r \neq j} (W_r \times P_r)$ be the set of type profiles of all jobs except j . Let φ be the joint probability distribution of $(w_1, p_1, \dots, w_n, p_n)$. Then for type profile $t = (w_1^{i_1}, p_1^{k_1}, \dots, w_n^{i_n}, p_n^{k_n}) \in T$, $\varphi(t) = \prod_{j=1}^n \varphi_j(w_j^{i_j}, p_j^{k_j})$. Let t_{-j} and φ_{-j} be defined analogously. For $(w_j^i, p_j^k) \in W_j \times P_j$ and $t_{-j} \in T_{-j}$, we denote by $((w_j^i, p_j^k), t_{-j})$ the type profile where job j has type (w_j^i, p_j^k) and the types of the other jobs are represented by t_{-j} . Denote by $ES_j(f, w_j^i, p_j^k) := \sum_{t_{-j} \in T_{-j}} S_j(f((w_j^i, p_j^k), t_{-j})) \varphi_{-j}(t_{-j})$ the expected waiting time of job j if he reports type (w_j^i, p_j^k) and allocation rule f is applied. Denote by $E\pi_j(w_j^i, p_j^k) := \sum_{t_{-j} \in T_{-j}} \pi_j((w_j^i, p_j^k), t_{-j}) \varphi_{-j}(t_{-j})$ the expected payment to j .

We assume that an agent can only report a processing time that is not lower than his true processing time and that a job is processed for his reported processing time. This is a natural assumption, since a job can add unnecessary work to achieve a longer processing time, but reporting a shorter processing time can

easily be punished by preempting the job after the declared processing time (before it is actually finished).

Note that by regarding the processing time as private information, we introduce informational externalities: job j has a different valuation for a schedule if the processing time (and hence the type) of a job scheduled before j changes. In this regard, our model differs from the auction models studied in [11] and [7].

4.1 Bayes-Nash Implementability and the Type Graph

Definition 5. A mechanism (f, π) is called Bayes-Nash incentive compatible if for every agent j and every two types $(w_j^{i_1}, p_j^{k_1})$ and $(w_j^{i_2}, p_j^{k_2})$ with $i_1, i_2 \in \{1, \dots, m_j\}$, $k_1, k_2 \in \{1, \dots, q_j\}$, $k_1 \leq k_2$,

$$E\pi_j(w_j^{i_1}, p_j^{k_1}) - w_j^{i_1} ES_j(f, w_j^{i_1}, p_j^{k_1}) \geq E\pi_j(w_j^{i_2}, p_j^{k_2}) - w_j^{i_1} ES_j(f, w_j^{i_2}, p_j^{k_2}) \quad (2)$$

under the assumption that all agents apart from j report truthfully.

Note that by defining the incentive constraints only for $k_1 \leq k_2$, we account for the fact that agents can only overstate their processing time, but cannot understate it.

In order to ensure individual rationality, again add a dummy type t_j^d to the type space for every agent j , and let $ES_j(f, t_j^d) = 0$ and $E\pi_j(t_j^d) = 0$ for all $j \in J$. As in the 1-dimensional case, the dummy types together with the mentioned extra incentive constraints guarantee that individual rationality is satisfied along with the incentive constraints. Sometimes, it will be convenient to write $(w_j^{m_j+1}, p_j^k)$ for some $k \in \{1, \dots, q_j\}$ instead of t_j^d .

In the 2-dimensional setting, the type graph T_f of agent j has node set $W_j \times P_j$ and contains an arc from any node $(w_j^{i_1}, p_j^{k_1})$ to every other node $(w_j^{i_2}, p_j^{k_2})$ with $i \in \{1, \dots, m_j\}$, $i_2 \in \{1, \dots, m_j + 1\}$, $k \in \{1, \dots, q_j\}$, $k_1 \leq k_2$ of length

$$\ell_{(i_1 k_1)(i_2 k_2)} = w_j^{i_1} [ES_j(f, w_j^{i_2}, p_j^{k_2}) - ES_j(f, w_j^{i_1}, p_j^{k_1})].$$

Note that we have arcs only in direction of increasing processing times, since agents can only overstate their processing time. Furthermore, every node has an arc to the dummy type, but there are no outgoing arcs from the dummy type.

Definition 6. An allocation rule f satisfies monotonicity w.r.t. weights if for every agent $j \in J$ and fixed $p_j^k \in P_j$, $w_j^{i_1} < w_j^{i_2}$ implies that $ES_j(f, w_j^{i_1}, p_j^k) \geq ES_j(f, w_j^{i_2}, p_j^k)$.

Theorem 3. An allocation rule f is Bayes-Nash incentive compatible in the 2-dimensional setting if and only if it satisfies monotonicity with respect to weights.

Proof. The claim reduces to showing that in the type graph of any agent j the non-negative cycle property is equivalent to the non-negative two-cycle property. Since there is an arc from a node representing type $(w_j^{i_1}, p_j^{k_1})$ to the node

representing type $(w_j^{i_2}, p_j^{k_2})$ if and only if $p_j^{k_1} \leq p_j^{k_2}$, cycles can only occur between nodes representing types with equal processing times. Hence, the proof is analogous to the 1-dimensional case. \square

Similar as in [7], one can show that some arcs in the type graph are not necessary, since the corresponding incentive constraints are implied by others. The reduced type graph of agent j contains only arcs that are necessary in that sense. A sketch of the reduced type graph is given in Figure 1. Expected payments correspond to node potentials in the reduced type graph. The reduced type graph comes handy particularly when considering our (counter) examples in the next subsection.




Figure 1. Reduced type graph 2-d case.

4.2 On Optimal Mechanisms

We start by quickly reviewing an approach to two-dimensional optimal mechanism design studied in [7]. Here, the authors regard a limited-supply multi-item auction, where each agent's type (i, j) is given by a marginal valuation i per item and a capacity j . Above that capacity, the agent has zero valuation for each additional item. The goal is revenue maximization. Bayes-Nash implementability is equivalent to the expected amount of items allocated to an agent being monotone in his reported value for i . Malakhov and Vohra [7] use the type graph approach to derive optimal mechanisms in this 2-d setting. Note, however, that the approach of [7], and also our approach for the 1-dimensional setting focus on one agent and its type graph. Hence, in terms of the scheduling model considered here, any optimal allocation rule derived this way is necessarily a modified Smith's rule with modified weights that can be computed from the characteristics (type report and distribution) of the agent itself. Such an allocation rule necessarily satisfies the following IIA property.

Definition 7. We say that an allocation rule f is independent of irrelevant alternatives (IIA) if the relative order of any two jobs j_1 and j_2 is the same in the schedules $f(t_1)$ and $f(t_2)$ for any two type profiles $t_1, t_2 \in T$ that differ only in the types of agents from $J \setminus \{j_1, j_2\}$.

In other words, the relative order of two jobs is independent of all other jobs. For the 2-d setting, this is not necessarily the case for optimal mechanisms.

Theorem 4. *The optimal allocation rule for the 2-dimensional setting does in general not satisfy IIA.*

Proof. The proof uses the following instance with three jobs. Job 1 has type $(1, 1)$, job 2 has type $(2, 2)$ and job 3 has type space $\{1.9, 2\} \times \{1, 2\}$. The probabilities for job 3's types are $\varphi_3(1.9, 1) = 0.8$, $\varphi_3(2, 2) = 0.2$ and $\varphi_3(1.9, 2) = \varphi_3(2, 1) = 0$ respectively. We show that the best allocation rule that satisfies IIA achieves a minimum expected total payment of at least 5.6, whereas there exists an allocation rule – violating IIA – with an expected total payment of 4.88. The details are contained in the full version of this paper [5]. \square

Theorem 4 shows that any list scheduling algorithm where the priority of a job can be computed from the characteristics of the job itself cannot be optimal in general. Moreover, the type graph approach must fail, since it focusses on a single agent. Hence, optimal mechanism design for our 2-dimensional setting is considerably more complicated than for the 1-dimensional setting and for traditional auction settings as described in [11] and [7].

One explanation for this complication may lie in the fact that the 2-d setting considered here in fact entails informational externalities, as opposed to the auction settings in [11] and [7]. On the other hand, the informational externalities introduced by private processing times are not the only cause for complications in the 2-dimensional setting: Consider the 1-dimensional setting, where only the processing times are private, but the weights are public information. It turns out that all allocation rules are implementable, even when we allow that jobs understate their processing times. The optimal payment to a job j that reports processing time p_j^k is equal to $w_j ES_j(f, p_j^k)$, and therefore the total payment to jobs for allocation rule f is equal to $P^{min}(f) = \sum_{j \in J} \sum_{k=1}^{q_j} \varphi_j(p_j^k) w_j ES_j(f, p_j^k)$. This is minimized by Smith's rule.

When there are only two agents present, then IIA is trivially satisfied. Recall that in the 1-dimensional case the optimal mechanism is efficient for symmetric agents and regular distributions and that the uniform distribution is regular. This is contrasted by the following theorem.

Theorem 5. *Even for two symmetric agents, 2×2 -type spaces and uniform probability distributions, the optimal mechanism is not efficient.*

Proof. We show that the efficient allocation is for some instances dominated by the w -rule, which schedules the job with the higher weight first. For details we refer to the full version of this paper [5]. \square

5 Conclusion

We have seen that the graph theoretic approach is an intuitive tool for optimal mechanism design, and yields a closed formula for the optimal mechanism in the 1-d discrete case. The same approach works for the continuous case, too.

Moreover, we have seen that in the two-dimensional case the canonical approach does not work and that optimal mechanism design seems to be considerably more complicated than in the traditional auction models. We leave it as an open problem to identify (closed formulae for) optimal mechanisms for the 2-d case. It is conceivable, however, that closed formulae don't exist.

Acknowledgements. Birgit Heydenreich was supported by NWO grant 2004/03545/MaGW 'Local Decisions in Decentralised Planning Environments'. Debasis Mishra acknowledges support by NWO (visiting grants B 46-593 and 040.11.032) and by METEOR, the Maastricht Research School of Economics of Technology and Organizations.

References

1. M. Armstrong. Optimal multi-object auctions. *Review of Economic Studies*, 67:455–481, 2000.
2. E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.
3. T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
4. J. Hartline and A. Karlin. Profit maximization in mechanism design. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007.
5. B. Heydenreich, D. Mishra, R. Müller, and M. Uetz. Optimal Mechanisms for Single Machine Scheduling. Research Memorandum RM/08/XXX, Maastricht University, 2008.
6. B. Heydenreich, R. Müller, M. Uetz, and R. Vohra. Characterization of revenue equivalence. *Econometrica*, forthcoming, 2008.
7. A. Malakhov and R. Vohra. An optimal auction for capacity constrained bidders: a network perspective. *Economic Theory*, forthcoming, 2007.
8. M. Mitra. Mechanism design in queueing problems. *Economic Theory*, 17(2):277–305, 2001.
9. H. Moulin. On scheduling fees to prevent merging, splitting, and transferring of jobs. *Mathematics of Operations Research*, 32:266–283, 2007.
10. R. Müller, A. Perea, and S. Wolf. Weak monotonicity and Bayes-Nash incentive compatibility. *Games and Economic Behavior*, 61(2):344–358, 2007. Extended version available as Research Memorandum RM/05/040, Maastricht University, 2005.
11. R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
12. J.-C. Rochet. A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of Mathematical Economics*, 16(2):191–200, 1987.
13. W. Smith. Various optimizers for single stage production. *Naval Research Logistics Quarterly*, 3:59–66, 1956.
14. W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.