# Local Computation Mechanism Design 

Avinatan Hassidim * $\quad$ Yishay Mansour ${ }^{\dagger} \quad$ Shai Vardi ${ }^{\ddagger}$


#### Abstract

We introduce the notion of local computation mechanism design - designing game theoretic mechanisms that run in polylogarithmic time and space. Local computation mechanisms reply to each query in polylogarithmic time and space, and the replies to different queries are consistent with the same global feasible solution. When the mechanism employs payments, the computation of the payments is also done in polylogarithmic time and space. Furthermore, the mechanism needs to maintain incentive compatibility with respect to the allocation and payments.

We present local computation mechanisms for a variety of classical game-theoretical problems: (1) stable matching, (2) job scheduling, (3) combinatorial auctions for unit-demand and $k$-minded bidders, and (4) the housing allocation problem.

For stable matching, some of our techniques may have implications to the global (non-LCA) setting. Specifically, we show that when the men's preference lists are bounded, we can achieve an arbitrarily good approximation to the stable matching within a fixed number of iterations of the Gale-Shapley algorithm.


## 1 Introduction

Assume that we would like to design an auction for millions of buyers and items. Alternatively, we have a cloud of hundreds of thousands of computers on which we would like to schedule several millions of jobs. In the not-so-distant past, these ideas would have been unthinkable, but today, technological advances, especially the Internet, have led us to the point where they are not only possible, but necessary. One can easily conceive a cloud computation with thousands of selfish computers, each one wanting to minimize its work load. Alternatively, an ad-auction for millions of businesses competing for advertising on millions of websites does not appear to be a far away dream. In cases like these, the data sets on which we need to work are so large, that polynomial-time tractability may not be enough. Sometimes, even computing a solution in linear time may be infeasible. Often, however, only parts of the solution to a problem are required at each point in time. In such cases, we can use local computation algorithms (LCAs).

Local computation algorithms, which were introduced by Rubinfeld et al. [RTVX11], consider the scenario in which we need to be able to respond to queries (regarding a feasible solution) quickly, but we never need the entire solution at once. For example, in most auctions, this is a reasonable assumption. When queried, we need to be able to tell each buyer which items she received and how much to pay; for a given item we need to tell the seller to whom and when to ship the item. There is no need to calculate the

[^0]entire allocation and payment at any specific time or to commit the entire solution to memory. Having an LCA to such an auction would mean that we can reply to queries in polylogarithmic time and only require polylogarithmic space. Furthermore, if all of the items and buyers are queried, combining the results will give us a complete solution that meets our requirements.

The field of algorithmic mechanism design is an area at the intersection of economic game theory and algorithm design, whose objective is to design mechanisms in decentralized strategic environments. These mechanisms need to take into account both the algorithmic efficiency considerations and the selfish behavior of the participating agents.

In this paper we propose local computation mechanism design, which shares the motivations of both local computation algorithms and algorithmic mechanism design. Our abstract model is the following: We have a large data-set and a set of allowable queries. Our goal is to implement each query locally, with polylogarithmic time and space, while maintaining the incentives of participants. It is worthwhile to give a few illustrative examples:

1. Consider the problem of assigning doctor interns to hospitals internships, the classical motivation for stable matching. We would like to be able to compute, for each doctor, her assigned hospital, without performing the entire global computation.
2. Consider a large distributed data center that has to assign jobs to machines and elicits from each machine its speed. When queried on a job, we would like to reply to which machine it is assigned, and when queried regarding a machine, we would like to reply with the set of jobs that need to run on it. Again, we would like the computation to be local, without constructing a global solution, and still be able to ensure the machines have an incentive to report their speeds truthfully.
3. Consider a large auction (for instance, an ad auction platform). When queried regarding a bidder, we would like to compute the items she receives and her payment; when queried on an item, we would like to compute which bidder won it, and it's price. Again, we would like guarantee that the bidders have an incentive to report their preferences truthfully.

The following are our main contributions. First, we formalize the notion of local computation mechanism design. A mechanism is local if, for every query, it calculates an allocation (and a payment) in polylogarithmic time and space. Furthermore, the allocation must be consistent with some global solution, and the payment must ensure truthfulness of the agents. Second, we present local computation mechanisms for several interesting problems, where our main result is an LCA for stable matching. Third, we use our techniques to show that in the general case when the men's lists have bounded length (even in cases that do not admit an LCA), we can find arbitrarily good matchings (up to both additive and multiplicative constants) by truncating the Gale-Shapley algorithm to a constant number of rounds.

We provide LCAs for the following problems:

Stable matching In the stable matching (or stable marriage) problem, introduced by Gale and Shapley [GS62], we would like to find a stable perfect matching between a group of $n$ men and a group of $n$ women. We focus on the model introduced by Immorlica and Mahdian [M05], in which the the women can have arbitrary preferences over the men, and the men have preference lists of length $k$ over the women, sampled uniformly at random.

Our main result is a local computation algorithm which matches all but an arbitrarily small fraction of the participants (this is often called an almost stable matching; see, e.g., [TL84, LZ03]). Furthermore, limited to the matched participants, the matching is stable.

[^1]Scheduling on related machines In the makespan minimization problem, we want to schedule $n$ jobs on $m$ machines so as to minimize the maximal running time (makespan) of the machines. This problem has many variations; we consider the scenario in which $m$ identical jobs need to be allocated among $n$ related machines. The machines are strategic agents, whose private information is their speed. We show:

1. A local mechanism that is truthful in expectation for scheduling on related machines, which provides an $O(\log \log n)$-approximation to the optimal makespan.
2. A local mechanism that is universally truthful for the restricted case (i.e., when each job can run on one of at most a constant number of predetermined machines), which provides an $O(\log \log n)$ approximation to the optimal makespan.
We also show some subtle and surprising results on the truthfulness of our algorithms.

Matching combinatorial auctions Combinatorial auctions (CAs) are auctions in which buyers can bid on bundles of items. We consider the following scenario: $m$ items are to be auctioned off to $n$ unitdemand buyers, where each buyer is interested in a set of at most $k$ items, sampled uniformly at random. We show universally truthful local mechanisms for the following variations, both of which provide a $1 / 2$-approximation to the optimal solution:

1. When all buyers have an identical valuation for the items in their sets, and the buyers' private information is the sets of items they are interested in.
2. When the sets are public knowledge, and the buyers' private information is their valuation.

We also show that there cannot exist an (exact) LCA for maximum matching, and therefore, there cannot exist any local mechanism which computes the optimal solution.

Combinatorial auctions with $k$-single minded bidders If each buyer is interested in a set of at most $k$ items, sampled uniformly at random, and has private valuation for this set, we show a universally truthful local mechanism which gives a $1 / k$-approximation to the optimal social welfare.

Random Serial Dictatorship (RSD) We show that in the housing allocation setting, in which each agent is interested in a constant number of houses, sampled uniformly at random, the RSD algorithm is implementable as an LCA.

### 1.1 Related Work

Local Computation Algorithms: Rubinfeld et al. [RTVX11], showed how to transform distributed algorithms to LCAs, and gave LCAs for several problems, including maximal independent set and hypergraph 2-coloring. Alon et al. [ARVX12], expanded the work of [RTVX11] and gave better space bounds for maximal independent set and hypergraph 2 -coloring, using query trees. Query trees were introduced in the local setting by Nguyen and Onak [NO08]: a random permutation of the vertices is generated, and a sequential algorithm is simulated on this order. The query tree represents the dependence of each query on the results of previous queries. [NO08] showed that if the graph has a bounded degree, the query tree has a constant expected size. [ARVX12] showed that the query tree has polylogarithmic size with high probalility, and that the space required by the algorithm can be reduced by using a random seed to generate the ordering. Mansour et al. [MRVX12], showed that the size of the query tree can be bounded, with high probability, by $O(\log n)$, and showed how it is possible to transform many on-line algorithms to LCAs. Using this technique, they showed LCAs for maximal matching and several machine scheduling problems. Mansour and Vardi [MV13], showed an LCA which finds a $(1-\epsilon)$-approximation to the maximum matching.
Mechanism Design: We divide our attention between two types of mechanisms: with and without payments. When the mechanism designer can incorporate payments (to or from participating parties), these
payments are usually used to guarantee incentive compatibility or "truthfulness" ([MCWG95], [NR99]): the agents are rationally motivated to truthfully reveal their private information. The mechanisms with payments discussed in this paper are all randomized, and there are several quantifications of truthfulness for such mechanisms in the literature, e.g., [FKKV13, HKV13]. We focus on the two most widely accepted (e.g., [NR99, AT01, LOS02, DD09, DNS12]): truthfulness in expectation, in which the expectation of each agent's utility is maximized by being truthful (where the expectation is taken over the coin flips of the mechanism), and universal truthfulness, where each agent's utility is maximized by being truthful, regardless of the realization of the randomness of the mechanism. When the mechanism does not support payments, it is sometimes impossible to guarantee truthfulness without crippling restrictions to the mechanism Arr50]. In such cases, it is common to look at (Bayesian) ex-ante truthfulnes:2] in which the expectation of each agent's utility is maximized by being truthful (where the expectation is taken over the prior distribution of the other agents' private information).

Mechanisms are generally required to run in time (and space) polynomial in the size of the input. In cases when the optimal solution can be calculated in polynomial time and space, one can apply the well-known VCG payments [Vic61, Cla71, Gro73]. Unfortunately, these payments can only be applied when the optimal solution can be computed [NR99, LOS02], and in many cases, it has been shown that computing an optimal solution is NP-hard [NR99, Dob11]. In these cases, we can only hope to design algorithms that approximate the optimal solution in polynomial time.
Additional Related Work: Because of the large variety of game-theoretic settings considered in this paper, instead of listing the entire glossary of related game-theoretic works here, we provide a short subsection dedicated to related work pertaining to each topic at the start of the relevant sections.

### 1.2 Paper Organization

In Section 2/we give some general definitions and notation which we will use in the rest of the paper, and present our model for local computation game theoretic mechanisms. In Section 3 we present our LCA for stable matching. In Section 4 we show some properties of the (global) Gale-Shapley algorithm that can be derived using our proof techniques. In Section 5 we show that our LCA for stable matching is ex-ante truthful. In Section 6we present LCAs for machine scheduling; in Section 7, we give LCAs for combinatorial auctions with unit demand bidders, and prove that there cannot be an LCA for maximum matching; in Section 8 we extend the results of Section 7 to combinatorial auctions with single-minded bidders. In Section 9 we give an LCA for random serial dictatorship, and finally, in Section 10 we provide several interesting open questions for future research.

## 2 Our Model and Preliminaries

### 2.1 Local Computation Algorithms

The model we use is a generalization of the model of local computation algorithms (LCAs) introduced in [RTVX11] $3^{3} \mathrm{~A}(t(n), s(n), \delta(n))$-local computation algorithm $\mathcal{L A}$ for a computational problem is a (randomized) algorithm that receives an input of size $n$, and a query $x$. Algorithm $\mathcal{L} \mathcal{A}$ replies to query $x$ in time $t(n)$ and uses at most $s(n)$ memory, with probability at least $1-\delta(n)$. Furthermore, the replies to all of the possible queries are consistent and combine to a feasible solution to the problem. That is,

[^2]the algorithm always replies correctly, but there is a $\delta(n)$ probability that the time and/or space bounds will be violated.

### 2.2 Mechanism Design

We use the standard notation of game theoretic mechanisms. There is a set $N$ of $n$ rational agents and a set $M$ of $m$ objects. In some settings, e.g., the stable marriage setting, there are no objects, only rational agents. Each agent $i \in N$ has a valuation function $v_{i}$ that maps subsets $S \subseteq M$ of the items to non-negative numbers. The utilities of the agents are quasi-linear, namely, when agent $i$ receives subset $S$ of items and pays $p$, her utility is $u_{i}(S, p)=v_{i}(S)-p$. Agents are rational in the sense that they select actions to maximize their utility. We would like to allocate items to agents (or possibly agents to other agents), in order to meet global goal, e.g., maximize the sum of the valuations of allocated objects (see, e.g., [NRTV05]).

A mechanism with payments $\mathcal{M}=(\mathcal{A}, \mathcal{P})$ is composed of an allocation function $\mathcal{A}$, which allocates items to agents, and a payment scheme $\mathcal{P}$, which assigns each agent a payment. A mechanism without payments consists only of an allocation function. Agents report their bids to the mechanism. Given the bids $b=\left(b_{1}, \ldots, b_{n}\right)$, the mechanism allocates the item subset $\mathcal{A}_{i}(b) \subseteq M$ to agent $i$, and, if the mechanism is with payments, charges her $\mathcal{P}_{i}(b)$; the utility of agent $i$ is $u_{i}(b)=v_{i}\left(\mathcal{A}_{i}(b)\right)-\mathcal{P}_{i}(b)$.

A randomized mechanism is universally truthful if for every agent $i$, for every random choice of the mechanism, reporting her true private valuation maximizes her utility. A randomized mechanism is truthful in expectation, if for every agent $i$, reporting her true private valuation maximizes her expected utility. That is, for all agents $i$, any bids $b_{-i}$ and $b_{i}, \mathbb{E}\left[u_{i}\left(v_{i}, b_{-i}\right)\right] \geq \mathbb{E}\left[u_{i}\left(b_{i}, b_{-i}\right)\right]$.

We say that an allocation function $\mathcal{A}$ admits a truthful payment scheme if there exists a payment scheme $\mathcal{P}$ such that the mechanism $\mathcal{M}=(\mathcal{A}, \mathcal{P})$ is truthful.

A mechanism $\mathcal{M}=(\mathcal{A}, \mathcal{P})$ fulfills voluntary participation if, when the agent bids truthfully, the utility of every agent is always non-negative, regardless of the other agents' bids, i.e., for all agents $i$ and bids $b_{-i}, u_{i}\left(v_{i}, b_{-i}\right) \geq 0$.

### 2.3 Local Computation Mechanisms

Definition 2.1 (Mechanisms without payments). We say that a mechanism $\mathcal{M}$ is $(t(n), s(n), \delta(n))$-local if its allocation function is computed by a $(t(n), s(n), \delta(n))$-local computation algorithm.

Definition 2.2 (Mechanisms with payments). We say that a mechanism $\mathcal{M}=(\mathcal{A}, \mathcal{P})$ is $(t(n), s(n), \delta(n))$ local if both the allocation function $\mathcal{A}$ and the payment scheme $\mathcal{P}$ are computed by $(t(n), s(n), \delta(n))$ local computation algorithms.

In other words, given a query $x, \mathcal{A}$ computes an allocation and $\mathcal{P}$ computes a payment, and both run in time $t(n)$ and space $s(n)$ with probability at least $1-\delta(n)$. Furthermore, the replies of $\mathcal{A}$ to all of the queries combine to a feasible allocation.

A truthful local mechanism $\mathcal{M}=(\mathcal{A}, \mathcal{P})$ is a local mechanism which is also truthful. Namely, each agent's dominant bid is her true valuation, regardless of the fact that the mechanism is local.

## 3 Stable Matching

The stable matching problem is represented by a tuple $A=(M, W, P)$, where $M$ is the set of men, $W$ is set of women, and $P$ is the set of preference relations of the men and the women: each man $m \in M$ has a preference relation over the women: if $m$ prefers $w$ to $w^{\prime}$ we denote this by $w \stackrel{m}{\succ} w^{\prime}$. Similarly there is a preference relationship $\stackrel{w}{\succ}$ for each woman $w$.

Related work Stable matching has been at the center of game-theoretic research since the seminal paper of Gale ans Shapley [GS62] (see, e.g., [Rot03] for an introduction and a summary of many important results). Roth and Rothblum [RR99] examined the scenario in which the preference lists are of bounded length; in most real-life scenarios, this is indeed the case. For example, a medical student will not submit a preference list for internship over all of the hospitals in the United States. Furthermore, the mechanism designer in most of these cases de facto decides on the list length because the mechanism will usually require the men to submit a list of some predetermined length. It is known that a linear number of iterations of the Gale-Shapley algorithm is necessary to attain stability [GI89], and several works address the situation when we are interested in a sublinear number of queries, for example Feder et al. [FMP00] propose a parallel sublinear time algorithm for stable matching. Unfortunately, it is not possible to convert their algorithm to a local computation algorithm. Several experimental works on parallel algorithms for the stable matching problem provide evidence that after a constant number of rounds, the matching is almost stable (e.g., [TL84, Qui85, LZO3]). Floreen et al. [FKPS10] show that in the special case when both the mens' and the womens' preference lists are bound by a constant, there exists a distributed version of the Gale-Shapley algorithm, which can be run for a constant number of rounds and finds an almost stable matching.

We examine the variant in which each man $m \in M$ is interested in at most $k$ women, (and prefers to be unmatched than to be matched to anyone not on their list; cf. [RR99]). We limit our attention to the setting in which the men's preference are assumed to be uniformly distributed; cf. [IM05, KP09].

Definition 3.1. A matching $H: M \rightarrow W \cup\{\perp\}$ is a function which is injective over $W$ (but it is possible that for several men $m_{i}, H\left(m_{i}\right)=\perp$ - these men are said to be unmatched).
A matching $H$ is said to be stable iffor every man $m$ such that $H(m)=w$ (possibly $w=\perp$ ), and every $w^{\prime}$ such that $w^{\prime} \stackrel{m}{\succ} w$, then $H^{-1}\left(w^{\prime}\right) \stackrel{w^{\prime}}{\succ} m$. A couple $\left(m, w^{\prime}\right)$ is said to be unstable if $\exists w \in W \cup\{\perp\}$ : $H(m)=w, w^{\prime} \succ_{\succ} w^{2}$, and for $m^{\prime}=H^{-1}\left(w^{\prime}\right), m \stackrel{w^{\prime}}{\succ} m^{\prime}$ (possibly $m^{\prime}=\perp$ ).
The stable matching problem $(M, W, P)$ where each man has a preference list of length $k$ and each woman is chosen uniformly at random is called $k$-uniform.

The Gale-Shapley algorithm finds a stable matching in the $k$-uniform setting (e.g., GS85]). To ensure the locality of our algorithm, we allow our mechanism to "disqualify" men, in which case they remain unmatched, but are unable to contest the matching. We try, however, to keep the number of disqualified men to a minimum. Our main result is the following.

Theorem 3.2. Let $A=(M, W, P)$ be a stable matching problem, $|M|=|W|=n$, in the $k$-uniform setting. Then there is an $(O(\log n), O(\log n), 1 / n)$-local computation mechanism for $A$ in which at most an $O(1 / k)$ fraction of the men remain unmatched.

We begin by describing a non-local algorithm, ABRIDGEDGS, and then show how to simulate it locally by a local algorithm, LOCALAGS.

### 3.1 AbridgedGS

Let AbridgedGS be the Gale-Shapley men's courtship algorithm, where, in addition to the preference lists being of constant length, the algorithm is stopped after $\ell$ rounds. That is, in each round, each unassigned man goes to the highest ranked woman who has not yet rejected him. Each woman then keeps the man she prefers out of the men who approached her, and rejects the rest. This continues until the $\ell^{t h}$ round, and the men who were rejected on the $\ell^{t h}$ round are disqualified. (Note that the men who were rejected $k$ times are left unmatched as well, but they are not said to be disqualified.) We simulate AbridgedGS using an LCA.

### 3.2 LOCALAGS - an LCA implementation of AbridgedgS

Consider a graph which represents the problem, where the men and the women are represented by vertices, and an edge exists between two vertices if and only if they are on each other's list. Define the distance between two agents to be the length of the shortest path between them in the graph. Define the $d$-neighborhood of a person $v$ to be everyone at a distance at most $d$ from $v$, denoted $N_{d}(v)$

Assume we are queried on a specific man, $m_{1}$. We simulate AbridgedGS locally as follows: Let the number of rounds be $\ell=2 k^{2}$ (see Lemma 3.7). We look at the $2 \ell$-neighborhood of $m_{1}$. For each man $m_{i} \in N_{2 \ell}\left(m_{1}\right)$, we simulate round 1 . Then, for each man $m_{i} \in N_{2 \ell-2}\left(m_{1}\right)$, we simulate round 2 . And so on, until for $m_{i} \in N_{2}\left(m_{1}\right)$, that is, $m_{1}$ and his closest male neighbors, we simulate round $\ell$. We return the woman to whom $m_{1}$ is paired, "unassigned" if he was rejected by $k$ women, and "disqualified" if he was rejected by a woman in round $\ell$. We denote this algorithm by LocalaGS.
Claim 3.3. For any two men, $m_{i}$ and $m_{j}$, whose distance from each other is greater than $2 \ell, m_{i}$ 's actions cannot affect $m_{j}$ if Algorithm ABRIDGEDGS terminates after $\ell$ rounds.

Proof. The proof is by induction. For $\ell=1$, let $w_{1}$ be $m_{j}$ 's first choice. Only men for whom $w_{1}$ is their first choice can affect $m_{j}$, and these are a subset of the men at distance 2 from $m_{j}$. For the inductive step, assume that the claim holds for $\ell-1$. Assume that there is a man $m_{i}$ whose actions can affect $m_{j}$ within $\ell$ rounds, who is at a distance of at least $2 \ell+2$ from $m_{j}$. From the inductive claim, none of $m_{i}$ 's actions can affect any of $m_{j}$ 's neighbors within $\ell-1$ rounds. As their actions in round $\ell-1$ (or any previous round) will not be affected by $m_{i}$, and they are the only ones who can affect $m_{j}$ in round $\ell$, it follows that $m_{i}$ cannot affect $m_{j}$ within $\ell$ rounds.

Corollary 3.4. The query of a man m's status in round $\ell$ only needs to consider men at distance at most $2 \ell$ from $m$.

Lemma 3.5. The running time and space of algorithm Localags is $O(\log n)$ with probability at least $1-1 / n^{2}$.

We prove the following claim, which immediately implies Lemma 3.5
Claim 3.6. For sufficiently large $n$, for any integer $i>0$, there exists a constant $c_{i}$ such $\operatorname{Pr}\left[\left|N_{i}(v)\right| \leq\right.$ $\left.c_{i} \log n\right] \geq 1-1 / n^{2}$.
Proof. Let $N_{v}^{i}$ be the random variable representing the number of vertices in the $i$-neighborhood of vertex $v$. Note that as the degree of a women $v$ is distributed binomially $N_{v}^{1} \sim B(n, k / n)$ and $\mathbb{E}\left[N_{v}^{1}\right]=$ $k$. We prove by induction that $\operatorname{Pr}\left[N_{v}^{i} \leq c_{i}(\log n)\right] \geq 1-\frac{i}{n^{3}}$.

For the base, $i=1$, if $v$ is a man, $N_{v}^{1}=k$. If $v$ is a woman, we employ the Chernoff bound with $\lambda>2 e-1 \cdot 4 \operatorname{Pr}\left[N_{v}^{1}>(1+\lambda) k\right]<2^{-k \lambda}$. Therefore, for $c_{1}=4$ and for $n \geq 2^{k}$,

$$
\operatorname{Pr}\left[N_{v}^{1}>c_{1} \log n\right] \leq 2^{-c_{1} \log n+k}<\frac{2^{k}}{n^{c_{1}}} \leq \frac{1}{n^{3}}
$$

Assuming that the claim holds for all integers smaller than $i$, we show that it holds for $i$. If the outermost vertices of the neighborhood are men, then $N_{v}^{i} \leq k N_{v}^{i-1}$ and we can take $c_{i}=k c_{i-1}$. Otherwise, we use the law of total probability.

$$
\begin{aligned}
\operatorname{Pr}\left[N_{v}^{i}>c_{i} \log n\right]= & \operatorname{Pr}\left[N_{v}^{j}>c_{i} \log n \mid N_{v}^{i-1} \leq c_{i-1} \log n\right] \operatorname{Pr}\left[N_{v}^{i-1} \leq c_{i-1} \log n\right] \\
& +\operatorname{Pr}\left[N_{v}^{j}>c_{i} \log n \mid N_{v}^{j-1}>c_{i-1} \log n\right] \operatorname{Pr}\left[N_{v}^{j-1}>c_{i-1} \log n\right] \\
\leq & \operatorname{Pr}\left[N_{v}^{j}>c_{i} \log n \mid N_{v}^{j-1} \leq c_{i-1} \log n\right]+\operatorname{Pr}\left[N_{v}^{j-1}>c_{i-1} \log n\right] \\
\leq & \operatorname{Pr}\left[N_{v}^{j}>c_{i} \log n \mid N_{v}^{j-1} \leq c_{i-1} \log n\right]+\frac{i-1}{n^{3}} .
\end{aligned}
$$

[^3]where the last inequality uses the inductive hypothesis.
The probability that the degree of any node $u$ is exactly $z$ is at most
$$
\operatorname{Pr}[\operatorname{deg}(u)=z] \leq\binom{ n}{z}\left(\frac{k}{n}\right)^{z} \leq\left(\frac{e k}{z}\right)^{z}
$$
using the inequality $\binom{n}{k} \leq\left(\frac{n e}{k}\right)^{k}$. Hence, for $z \geq e^{2} k$ we have that $\operatorname{Pr}[\operatorname{deg}(u)=z] \leq e^{-z}$ and $\operatorname{Pr}[\operatorname{deg}(u)=z] \leq e^{-\bar{z}}$, where $\widehat{z}=\max \left\{0, z-e^{2} k\right\}$.

We like to bound the probability that $N_{v}^{i}$ is larger than $c_{i} \log n$ although $N_{v}^{i-1}$ is less than $c_{i-1} \log n$. We define a new random variable $\widehat{N}_{v}^{i}$ as follows. Let $y \leq c_{i-1} \log n$ be the number of nodes at distance $i-1$ from $v$ and let $z=\left(z_{1}, z_{2}, \ldots, z_{y}\right)$ be their degrees. We define the truncated degrees as $\widehat{z}=$ $\left\{\widehat{z}_{1}, \widehat{z}_{2}, \ldots, \widehat{z}_{y}\right\}$ such that $\widehat{z}_{j}=\max \left\{0, z_{j}-e^{2} k\right\}$. The value of $\widehat{N}_{v}^{i}$ is the sum of the truncated degrees at distance $i-1$ from $v$, i.e., $\widehat{N}_{v}^{i}=\sum_{i=1}^{y} \widehat{z}_{i}$. Clearly $N_{v}^{i} \leq \widehat{N}_{v}^{i}+e^{2} k y \leq \widehat{N}_{v}^{i}+c_{i-1} e^{2} k \log n$. Therefore it is sufficient to bound $\widehat{N}_{v}^{i}$.

Let $\widehat{x}=\sum_{i=1}^{y} \widehat{z_{i}}$. The probability that the truncated degrees of the vertices at distance $i-1$ are exactly $\widehat{z}=\left(\widehat{z}_{1}, \widehat{z}_{2}, \ldots, \widehat{z}_{y}\right)$ is at most $\prod_{i=1}^{y} e^{-\widehat{z}_{i}}=e^{-\widehat{x}}$. There are $\binom{\widehat{x}+y}{y}$ vectors $\widehat{z}$ that can realize $\widehat{x}$. We bound $\operatorname{Pr}\left[\widehat{N}_{v}^{i}=\widehat{x} \mid N_{v}^{i-1} \leq y\right]$, for $\widehat{x} \geq 7 y$ as follows:

$$
\operatorname{Pr}\left[\widehat{N}_{v}^{i}=\widehat{x} \mid N_{v}^{i-1} \leq y\right] \leq\binom{\widehat{x}+y}{y} e^{-\widehat{x}} \leq\left(\frac{e \cdot(\widehat{x}+\widehat{x} / 7)}{\widehat{x} / 7}\right)^{\widehat{x} / 7} e^{-\widehat{x}}=e^{-(1-(1+\ln (8)) / 7) \widehat{x}} \leq e^{-\widehat{x} / 2}
$$

It follows that

$$
\operatorname{Pr}\left[\widehat{N}_{v}^{i} \geq 7 y \mid N_{v}^{i-1} \leq y\right] \leq \sum_{\widehat{x}=7 y}^{\infty} e^{-\widehat{x} / 2}=\frac{e^{-7 y / 2}}{1-e^{-1 / 2}} \leq e^{-y} \leq 1 / n^{3}
$$

which follows since $c_{i-1} \geq 3$. Therefore for $c_{i}=\left(e^{2} k+7\right) c_{i-1} \leq(16 k)^{i}$ we have,

$$
\operatorname{Pr}\left[N_{v}^{i}>c_{i} \log n\right] \leq \frac{1}{n^{3}}+\frac{i-1}{n^{3}}=\frac{i}{n^{3}}
$$

Claim 3.6 implies Algorithm LocalAGS makes $O(\log n)$ queries with probability at least $1 / n^{2}$, and so Lemma 3.5 follows.

Lemma 3.7. In Algorithm LocalAGS, setting $\ell=2 k^{2}$ ensures at most $4 n / k$ men remain unmatched with probability at least $1-\frac{1}{n^{2}}$.

Note: This implies that the mechanism can limit the number of unmatched pairs with a high degree of certainty by specifying the length of the list.

Lemmas 3.5 and 3.7 imply Theorem 3.2 We prove Lemma 3.7 in the following subsection; the proof follows from Claims 3.9 and 3.11

### 3.3 Bounding the number of men removed

To prove Lemma 3.7 we bound the number of men remaining unassigned due to the fact that the lists are short, and the number of men disqualified due to the number of rounds being bounded.

### 3.3.1 Removal due to short lists

We bound the number of unpaired women as a result of the fact that the lists are short, using the principle of deferred decisions $\boxed{\text { Knu76 }}{ }^{5}$. Note that the number of unpaired women equals the number of unpaired men.

Lemma 3.8. In the $k$-uniform setting, the Gale-Shapley algorithm results in at most $\frac{2 n}{k}$ men being unpaired, with probability at least $1-1 / 2 n^{2}$.

Proof. Consider the following stochastic process: in the first round, each of the $n$ men chooses a woman independently and uniformly at random. For each consecutive round, for each woman that has been chosen by at least one man, one of the men remains married to her (arbitrarily chosen), and the others remain single and choose again. This process repeats for $k$ rounds. This is modeled by the functions $f^{t}: M \rightarrow W$, where in round $t, f^{t}$ maps each single man to a woman uniformly at random and each married man to the same woman. Let $\mathcal{F}$ be the set of all possible allocation functions $f: M \rightarrow W$.

For ease of analysis, we assume that each man can choose the same woman again, as the number of free women in this case is an upper bound to the number of free women in the system where he can not. Note that the number of unmatched men after $t$ rounds is identical to the number of unmatched men after $t$ rounds of the Gale-Shapley process. This stochastic process, however, terminates after $k$ rounds, whereas the Gale-Shapley process can continue. As the number of matched men can only increase when more rounds are added, the number of unmatched men created by the process is an upper bound to the number of unmatched men created by the Gale-Shapley process.

Let $X_{j}^{t}$ be the indicator variable which is 1 if woman $j$ is unassigned at the end of round $t$. Let $X^{t}=\sum_{j=1}^{n} X_{j}^{t}$. The following claim implies Lemma 3.8
Claim 3.9. For any constant $t, \operatorname{Pr}\left[X^{t}>\frac{2 n}{t}\right] \leq \frac{t}{n^{3}}$.
Proof. The proof is by induction. The base of the induction, $t=1$, is immediate. For the inductive step, assume that after round $t, X^{t}=n / z$ (for some $z>0$ ). In round $t+1, \mathbb{E}\left[X^{t+1} \left\lvert\, X^{t}=\frac{n}{z}\right.\right]=$ $\frac{n}{z}(1-1 / n)^{n / z}$. By the inductive hypothesis

$$
\begin{equation*}
\operatorname{Pr}\left[X^{t}>\frac{2 n}{t}\right] \leq \frac{t}{n^{3}} \tag{1}
\end{equation*}
$$

For the rest of the proof, assume $X^{t} \leq \frac{2 n}{t}$, and fix $X^{t}$ to be some such value. We get

$$
\begin{equation*}
\mathbb{E}\left[X^{t+1} \left\lvert\, X^{t} \leq \frac{2 n}{t}\right.\right] \leq \frac{2 n}{t}(1-1 / n)^{2 n / t}<\frac{n}{t / 2 \cdot e^{2 / t}}<\frac{2 n}{t+2} \tag{2}
\end{equation*}
$$

using $e^{x}>1+x$.
Order the women arbitrarily, $(W=\{1,2, \ldots n\})$, and let $W_{i}=\{1,2, \ldots i\}$. For $h \in F$, define the martingale

$$
Y_{i}^{t+1}(h)=\mathbb{E}\left[X^{t+1}\left(f^{t+1}\right) \mid f^{t+1}(j)=h(j) \text { for all } j \in W_{i}\right]
$$

where $X^{t+1}\left(f^{t+1}\right)$ is the realization of $X^{t+1}$ given that the allocation vector is $f^{t+1}$. Note that $Y_{0}^{t+1}(h)$ is the expected value of $X^{t+1}$ over all possible functions $f^{t+1}$; that is, the expected number of unmatched women after $t+1$ rounds. $Y_{n}^{t+1}(h)$ is the number of unmatched women after $t+1$ rounds when the allocation function is $h . X^{t+1}$ satisfies the Lipschitz condition, because if $h$ and $h^{\prime}$ only differ on the allocation of one man, $\left|X^{t+1}(h)-X^{t+1}\left(h^{\prime}\right)\right| \leq 1$. Therefore

$$
\left|Y_{i+1}^{t+1}(h)-Y_{i}^{t+1}(h)\right| \leq 1
$$

[^4](see, e.g. AS08]), and so we can apply Azuma's inequality, from which we get
$$
\operatorname{Pr}\left[\left|X^{t+1}-\mathbb{E}\left[X^{t+1}\right]\right|>\lambda \sqrt{n}\right]<2 e^{-\lambda^{2} / 2}
$$

Setting $\lambda=\frac{2 \sqrt{n}}{(t+1)(t+2)}$, we have that

$$
\operatorname{Pr}\left[\left|X^{t+1}-\mathbb{E}\left[X^{t+1}\right]\right|>\frac{2 n}{(t+1)(t+2)}\right]<2 e^{-n /\left(81 t^{4}\right)}
$$

for constant $t \geq 2$.
Therefore, since we assume that $\mathbb{E}\left[X^{t+1}\right]<\frac{2 n}{t+2}$,

$$
\begin{equation*}
\operatorname{Pr}\left[X^{t+1}>\frac{2 n}{t+1}\right]<2 e^{-n /\left(81 t^{4}\right)}<\frac{1}{n^{3}} \tag{3}
\end{equation*}
$$

Note that

$$
\begin{aligned}
\operatorname{Pr}\left[X^{t+1}>\frac{2 n}{t+1}\right] & =\operatorname{Pr}\left[\left.X^{t+1}>\frac{2 n}{t+1} \right\rvert\, X^{t} \leq \frac{2 n}{t}\right] \operatorname{Pr}\left[X^{t} \leq \frac{2 n}{t}\right] \\
& +\operatorname{Pr}\left[\left.X^{t+1}>\frac{2 n}{t+1} \right\rvert\, X^{t}>\frac{2 n}{t}\right] \operatorname{Pr}\left[X^{t}>\frac{2 n}{t}\right]
\end{aligned}
$$

From the induction hypothesis and Equations (1), (2) and (3), using the union bound,

$$
\operatorname{Pr}\left[X^{t+1}>\frac{2 n}{t+1}\right] \leq \frac{t}{n^{3}}+\frac{1}{n^{3}} \leq \frac{t+1}{n^{3}}
$$

The stochastic process for which Claim 3.9 holds ends at least as early as the Gale-Shapley algorithm with short lists; therefore Claim 3.9 implies Lemma 3.8

### 3.3.2 Removal due to the number of rounds being limited

Because we stop the LocalAGS algorithm after a constant $(\ell)$ number of rounds, it is possible that some men who "should have been" matched are disqualified because they were rejected by their $i^{\text {th }}$ choice in round $\ell(i<k)$. We show that this number cannot be very large.

Let $R_{i}$ denote the number of men rejected in round $i \geq 1$.
Observation 3.10. $R_{i}$ is monotone decreasing in $i$.
Claim 3.11. The number of men rejected in round $i$ is at most $n \frac{k}{i}$.
Proof. As each man can be rejected at most $k$ times, the total number of rejections possible is $k n$. The number of men who can be rejected in round $i$ is at most

$$
\begin{equation*}
R_{i} \leq k n-\sum_{j=1}^{i-1} R_{j} \Rightarrow R_{i} \leq k n-(i-1) R_{i} \quad \Rightarrow \quad R_{i} \leq n \frac{k}{i} \tag{4}
\end{equation*}
$$

Where Inequality 4 is due to Observation 3.10

## 4 Some general properties of the Gale-Shapley algorithm

We use the results and ideas of Section 3to prove some interesting features of the (general) Gale-Shapley stable matching algorithm, when the mens' lists are of length at most $k$. (These results immediately extend to our local version of the algorithm, LocalAGS.) Note that the proof of Claim 3.11 makes no assumption on how the men's selection is made, and therefore, Claim 3.11 implies that as long as each man's list is bounded by $k$, if we run the Gale-Shapley for $\ell$ rounds, at most $\frac{n k}{\ell}$ men will be rejected in that round. This immediately gives us an additive approximation bound for the algorithm if we stop after $\ell$ rounds.

Corollary 4.1 (to Claim3.11). Assume that the output of the Gale-Shapley algorithm on a stable matching problem, where the preference lists of the men are of length at most $k$, is a matching of size $M^{*}$. Then, stopping the Gale Shapley algorithm after $\ell$ rounds will result in a matching of size at least $M^{*}-\frac{n k}{\ell}$.

We would like to also provide a multiplicative bound. Henceforth, we assume that the mens' list length is bounded by $k$, but make no other assumptions. For each round $i$, let $M_{i}$ be the size of the current matching; let $D_{i}$ be the number of men who have already approached all $k$ women on their list and have been rejected by all of them; let $C_{i}$ be the number of men who were rejected by women in round $i$, but have approached fewer than $k$ women so far; recall that $R_{i}$ is the number of men rejected in round $i$. Denote the size of the matching returned by the un-truncated Gale-Shapley algorithm by $M^{*}$.

Claim 4.2. $C_{k+1} \leq k M^{*}$.
Proof. Note that $R_{i}=C_{i}+D_{i}-D_{i-1}$. For $i<k, D_{i}=0$. As $M_{i}$ is monotonically increasing, $\forall i \leq k, R_{i} \geq n-M^{*}$.

$$
\sum_{i=1}^{k} R_{i} \geq k n-k M^{*}
$$

Hence,

$$
C_{k+1} \leq k n-\sum_{i=1}^{k} R_{i} \leq k M^{*}
$$

Corollary 4.3. For every $\epsilon>0$, there exists a constant $\ell>0$ such that $C_{\ell} \leq \epsilon M^{*}$.
Proof. Denote the maximum number of rejections possible from round $i$ onwards by $L_{i}$. Clearly,

$$
L_{i} \leq k\left(M_{i}+C_{i}\right) \leq k\left(M^{*}+C_{i}\right)
$$

For all $i$ such that $C_{i} \geq \epsilon M^{*}$, we have

$$
L_{i} \leq\left(1+\frac{1}{\epsilon}\right) k C_{i}
$$

Therefore, from Claim 4.2

$$
L_{k+1} \leq\left(1+\frac{1}{\epsilon}\right) k^{2} M^{*}
$$

Putting everything together, we have,

$$
\begin{aligned}
L_{i+1} & \leq L_{i}-C_{i} \\
\Rightarrow L_{i+1} & \leq L_{i}\left(1-\frac{1}{k\left(1+\frac{1}{\epsilon}\right)}\right) \\
\Rightarrow L_{k+i+1} & \leq L_{k}\left(1-\frac{1}{k\left(1+\frac{1}{\epsilon}\right)}\right)^{i} \\
& \leq\left(1+\frac{1}{\epsilon}\right) k^{2} M^{*}\left(1-\frac{1}{k\left(1+\frac{1}{\epsilon}\right)}\right)^{i} \\
& \leq 2 k^{2} M^{*} e^{-\frac{i}{k\left(1+\frac{1}{\epsilon}\right)}}
\end{aligned}
$$

Taking $i=k\left(1+\frac{1}{\epsilon}\right) \log \frac{2 k^{2}}{\epsilon}$ gives $C_{k+i+1} \leq L_{k+i+1} \leq \epsilon M^{*}$.
This gives us,
Theorem 4.4. Consider a stable matching problem. Let each man's list be bounded by $k$. Denote the size of the stable matching returned by the Gale-Shapley algorithm by $M^{*}$. Then, if the process is stopped after $O\left(\frac{k}{\epsilon} \log \frac{k}{\epsilon}\right)$ rounds, the matching returned is at most $a(1+\epsilon)$-approximation to $M^{*}$, and has at most $\epsilon M^{*}$ unstable couples.

Corollary 4.5. If both men and women have lists of length at most $k$, then for any $\epsilon$ there is an $(O(1), O(1), 0)$-LCA for stable matching which returns a matching that is at most a $(1+\epsilon)$-approximation to the matching returned by the Gale-Shapley algorithm, and with at most an $\epsilon$-fraction of the edges being unstable.

## 5 Ex-ante truthfulness of LocalAGS

It is known that the Gale-Shapley men's courtship algorithm is strategy-proof for the men but not for the women (see, e.g., [MSZ13]). Unfortunately, Algorithm LocalAGS is neither strategy-proof for the men nor for the women: for the women, this follows immediately from the fact that the men's courtship algorithm is not strategy-proof; for the men - a man who was rejected in round $\ell$ might prefer to not declare the woman from which he was rejected in that round. It is possible, though, to show that LocalagS is ex-ante truthful. The men have a preference relation over the women (and women over the men). However, it is not clear how to calculate the expected utility in this case, so for the purposes of the proof, assume the men have a utility function $u: W \rightarrow \mathbb{R}$. The utility of the men for women not among their first $k$ preferences is 0 . Note that LocalAGS has no access to the utilities themselves, but only to the preferences, therefore the only way for a man to manipulate the algorithm is by misrepresenting his preference vector.

Assume that man $m$ 's real utility function is $t_{m}$, but he declares $b_{m}$. Denote the set of all the possible preference vectors of the other players by $\mathcal{B}_{-m}$. Let $b_{-m} \in \mathcal{B}_{-m}$ be a single realization of the preference vectors of the other players. Let $b=\left(b_{m}, b_{-m}\right)$. The outcome of instance $b$ is deterministically determined; denote the utility of $m$ for instance $b$ by $u_{m}(b) .\left(u_{m}(b)=t_{m}(w)\right.$ if $m$ is paired with woman $w$ and 0 otherwise.) Denote the expected utility of $m$ when he bids $b_{m}$ by $\overline{u_{m}}\left(b_{m}\right)=\sum_{b_{-m} \in \mathcal{B}_{-m}} u_{m}(b) \cdot \operatorname{Pr}\left[b_{-m}\right]$.

Definition 5.1. A transformation of a utility function by swapping the utilities of two women is called a swap.

Observation 5.2. It is possible to convert any preference vector $v$ to any other preference vector $v^{\prime}$ by a finite series of swaps.

This implies that given that a utilities function has a fixed list of outcomes, it is possible to reach any permutation of the utilities from any other by a finite series of swaps.

Claim 5.3. Algorithm LocalagS is ex-ante truthful for the men.
Proof. Assume that man $m$ gains by declaring his utility function to be $b_{m}$, over his true function, $t_{m}$, where $b_{m}$ is a permutation of $t_{m}$. That is, $\overline{u_{m}}\left(b_{m}\right)>\overline{u_{m}}\left(t_{m}\right)$. (Recall that $m$ can make no better manipulations.)

Take any legal chain of swaps from $t_{m}$ to $\left.b_{m}:<t_{m}=u^{1}, u^{2}, \ldots u^{i}, \ldots, u^{n}=b_{m}\right\rangle$. There must be two consecutive functions, $u^{i}$ and $u^{i+1}$ on the chain, for which

$$
\begin{equation*}
\overline{u_{m}}\left(u^{i+1}\right)>\overline{u_{m}}\left(u^{i}\right), \tag{5}
\end{equation*}
$$

from the transitivity of the relation " $>$ ". Functions $u^{i}$ and $u^{i+1}$ differ in the utilities of two women, say $w_{x}$ and $w_{y}: u^{i}\left(w_{x}\right)=u^{i+1}\left(w_{y}\right), u^{i}\left(w_{y}\right)=u^{i+1}\left(w_{x}\right)$.

Denote by $\mathcal{B}_{-m}^{\prime}$ the set of all preferences which is reached by taking $\mathcal{B}_{-m}$ and interchanging women $w_{x}$ and $w_{y}$ wherever they appear. Likewise the preference vectors of women $w_{x}$ and $w_{y}$ are interchanged. Note that $\mathcal{B}_{-m}=\mathcal{B}_{-m}^{\prime}$, because they are both uniform distributions over all possible preference vectors. By symmetry, it must hold that

$$
\overline{u_{m}}\left(u^{i}\right)>\overline{u_{m}}\left(u^{i+1}\right),
$$

in contradiction to Equation (5).
Similarly,
Claim 5.4. Algorithm LocalagS is ex-ante truthful for the women.
The proof is similar to the proof of Claim5.3 and is omitted.

## 6 Local machine scheduling

In this section we consider the following job scheduling setting. There is a set $\mathcal{M}$ of $n$ machines (or "bins") and a set $\mathcal{J}$ of $m$ uniform jobs (or "balls"). Each machine $i \in \mathcal{M}$ has an associated capacity $c_{i}$ (also referred to as its "speed"). We assume that the capacities are positive integers. Given that $h_{i}$ jobs are allocated to machine $i$, its load is $\ell_{i}=h_{i} / c_{i}$. ( $h_{i}$ is referred to as the height of machine $i$.) The utility of machine $i$ is quasi-linear, namely, when it has load $\ell_{i}$ and receives payment $p_{i}$ then its utility is $u_{i}\left(\ell_{i}, p_{i}\right)=p_{i}-\ell_{i}$.

The makespan of an allocation is $\max _{i}\left\{\ell_{i}\right\}=\max _{i}\left\{h_{i} / c_{i}\right\}$. In our setting, the players are the machines and their private information is their true capacities. Each machine $i$ submits a bid $b_{i}$ (which represents its capacity). Our goal is to elicit from the machines the information about their capacities, in order to minimize the makespan of the resulting allocation.

For any allocation algorithm $\mathcal{A}$, define $\mathcal{A}(b)$ to be the allocation vector, which, given bid vector $b$, assigns each job $j \in \mathcal{J}$ to a unique machine $i \in \mathcal{M}$. Let $\mathcal{A}^{j}(b)$ be the machine to which job $j$ is allocated in $\mathcal{A}$. When the bids $b_{-i}$ are fixed, we sometimes omit them from the notation for clarity.

Related work Azar et al., ABKU99] proposed the GREEDY [d] algorithm in the online setting, where $m$ balls need to be allocated to $n$ bins, with the objective of minimizing the makespan: each ball chooses, uniformly at random, $d$ bins, and allocates itself to the least loaded bin among its $d$ choices at the time of its arrival. They also showed that the maximal load is $\Theta(m / n)+(1+o(1)) \ln \ln n / \ln d$. A large volume of work has been devoted to variations on this problem, such as adding weights to the balls [TW07]; and variations on the algorithm, such as the non-uniform ball placement strategies of Vöcking [Vöc03]. Of particular relevance to this work is the case of non-uniform bins: Berenbrink et al. [BBFN14], showed that in this case the maximum load can also be bounded by $\Theta(m / n)+O(\ln \ln n)$. Hochbaum and Shmoys HS88 showed a PTAS for scheduling on related machines. Lenstra et al. [LST87], presented a 2-approximation algorithm for scheduling on unrelated machines and showed that the optimal allocation is not approximable to within $\frac{3}{2}-\epsilon$. The problem of finding a truthful mechanism for scheduling (on unrelated machines) was introduced by Nisan and Ronen, [NR99], who showed an $m$-approximation to the problem, and a lower bound of 2. Archer and Tardos [AT01] were the first to tackle the related machine case; they showed a randomized 3-approximation polynomial algorithm and a polynomial pricing scheme to derive a mechanism that is truthful in expectation. Since then, much work has gone into finding mechanisms with improved approximation ratios, until Christodoulou and Kovács [CK10] recently settled the problem by showing a deterministic PTAS, and a corresponding mechanism that is deterministically truthful. Babaioff et al. [BKS10], showed how to transform any monotone allocation rule for single parameter agents to a truthful-in-expectation mechanism.

Definition 6.1. (Monotonicity) $A$ randomized allocation function $\mathcal{A}$ is monotone in expectation if for any machine $i$, and any bids $b_{-i}$ of the other machines, the expected load of machine $i, E\left[\ell_{i}\left(b_{i}, b_{-i}\right)\right]$, is a non-decreasing function of $b_{i}$.
A randomized allocation function $\mathcal{A}$ is universally monotone if for any machine $i$, and any bids $b_{-i}$ of the other machines, the load of machine $i, \ell_{i}\left(b_{i}, b_{-i}\right)$, is a non-decreasing function of $b_{i}$ for any realization of the randomization of the allocation function.

Given an allocation function $\mathcal{A}$, we would like to provide a payment scheme $\mathcal{P}$ to ensure that our mechanism $\mathcal{M}=(\mathcal{A}, \mathcal{P})$ is truthful. It is known that a necessary and sufficient condition is that the allocation function $\mathcal{A}$ is monotone.

Theorem 6.2. Mye81, AT01]) The allocation algorithm $\mathcal{A}$ admits a payment scheme $\mathcal{P}$ such that the mechanism $\mathcal{M}=(\mathcal{A}, \mathcal{P})$ is truthful-in-expectation (universally truthful) if and only if $\mathcal{A}$ is monotone in expectation (universally monotone).

We differentiate between two settings. The standard setting (cf. [BBFN14, Wie07]) is a slight variation on the basic power-of- $d$ choices setting proposed in ABKU99], for some constant $d \geq 2$. For each job $j$, the mechanism chooses a subset $M_{j} \subseteq \mathcal{M},\left|M_{j}\right|=d$ of machines that the job can be allocated to. The probability that machine $i \in M_{j}$ is proportional to $b_{i}$.

In the restricted setting (cf. ANR95]), each job can be allocated to a subset of at most $d$ machines, where the subsets $M_{j}$ are given as an input to the allocation algorithm. The restricted setting models the case when the jobs have different requirements, and there is only a small subset of machines that can run each job.

Mansour et at. [MRVX12], showed that it is possible to transform any on-line algorithm on a graph of bounded degree (or whose degree is distributed binomially) to an LCA. The idea behind the reduction is simple: generate a random permutation on the vertices and simulate the on-line algorithm on this permutation. They show that, with high probability, this results in at most $O(\log n)$ queries. We require the following theorem from [MRVX12].

Theorem 6.3. MRVX12] Consider a generic on-line algorithm $\mathcal{L B}$ which requires constant time per query, for $n$ balls and $m$ bins, where $n=c m$ for some constant $c>0$. There exists an $\left(O\left(\log ^{4} n\right)\right.$,
$\left.O\left(\log ^{3} n\right), 1 / n\right)$-local computation algorithm which, on query of a (ball) vertex $v \in V$, allocates $v$ a (bin) vertex $u \in U$, such that the resulting allocation is identical to that of $\mathcal{L B} \cdot 6$

### 6.1 A truthful in expectation mechanism for the standard setting

For the presentation of the allocation algorithm, we regard the allocation field as consisting of slots of size 1. Each machine $i$ 's bid $b_{i}$ represents the number of slots it "owns". Given the bids $b$ of the machines, let $B=\sum_{i=1}^{n} b_{i}$. Thus, we can treat this as $B$ slots of size 1 each, where each machine $i$ owns $b_{i}$ slots. In the allocation, we will allocate jobs to slots. When a job $j$ is allocated to a specific slot, this means that the machine that owns the slot receives $j$. We provide the following simple on-line allocation algorithm $\mathcal{A}_{S L M S}$ (cf. [BBFN14]):

1. Choose for job $j$ a subset $M_{j}$ of $d$ slots out of $B$, where each slot has equal probability. (Note that $M_{j}$ may include different slots owned by the same machine.)
2. Given $M_{j}$, job $j$ is allocated to the least loaded slot in $M_{j}$ (breaking ties uniformly at random). Slots are treated as being independent of their machines. That is, it is possible that if a job chooses two slots $a$ and $b$, which belong to machines $A$ and $B, a$ has fewer jobs than $b$, but $B$ is more loaded than $A$, in terms of the average of the loads of its slots.
Note: Although it may not be possible to compute $B$ and $M_{j}$ locally exactly, it has been shown in that an approximate calculation suffices (e.g., [BCM03, Wie07]). The reader is invited to peruse the references herein for a more in-depth discussion.

Lemma 6.4. The randomized allocation function $\mathcal{A}_{\text {SLMS }}$ is monotone in expectation.
Proof. Let $B=\sum_{i^{\prime}} b_{i^{\prime}}$ and $B_{-i}=\sum_{i^{\prime} \neq i} b_{i^{\prime}}$. Since all the slots are identical, by symmetry the expected number of jobs allocated to each slot is exactly $m / B$. Since machine $i$ owns $b_{i}$ slots, the expected height of machine $i$ is

$$
\mathbb{E}\left[h_{i}\left(b_{i}\right)\right]=\frac{b_{i}}{B_{-i}+b_{i}} m,
$$

which is monotone increasing in $b_{i}$ (for $b_{i}, B_{-i} \geq 0$ ).
From Theorem 6.2 we immediately get:
Lemma 6.5. The randomized allocation function $\mathcal{A}_{\text {SLMS }}$ admits a payment scheme $\mathcal{P}_{\text {SLMS }}$ such that the mechanism $\mathcal{M}_{\text {SLMS }}=\left(\mathcal{A}_{S L M S}, \mathcal{P}_{\text {SLMS }}\right)$ is truthful in expectation.

It is interesting to note that the above algorithm does not admit a universally truthful mechanism. To show this, we prove a slightly stronger claim: that the GREEDY algorithm of [ABKU99], (in which each job chooses $d$ machines at random, and is allocated to the least loaded among them, post placemen $\sqrt{7}$, breaking ties arbitrarily), does not admit a universally truthful mechanism.

Claim 6.6. Algorithm Greedy is not universally monotone.
Proof. Assume we have 4 machines: $A, B, C$, and $D$, with bids $4,4,8$ and 1 respectively. The first 2 jobs choose machines $A$ and $D$ (which we abbreviate to $A D$ ), the next 2 jobs choose $B D$, and the next 6 jobs choose $C D$. After these 10 jobs, the heights of the machines are $(2,2,6,0)$ (recall that the Greedy algorithm allocates according to the post-placement load). The 11th job chooses $A B$, and the 12th job chooses $A C$. As ties are broken at random, assume machine $A$ receives job 11. Machine $C$ then receives job 12 , making the capacities $(3,2,7,0)$.

[^5]Now assume machine $C$ bids 9 , and the choices of the first 10 jobs and the 12 th job remain the same, but because $C$ bid higher, now the 11th job chooses $C$ instead of $A$ (so job 11 chooses $B C$ ). Now machine $B$ receives the 11 th job and machine $A$ receives the 12 th job, making the capacities $(3,3,6,0)$. Machine $C$ received less jobs although it was bidding higher!

Corollary 6.7. Algorithm $\mathcal{A}_{S L M S}$ is not universally monotone.
By Theorem 6.3, the allocation function $\mathcal{A}_{S L M S}$ can be transformed to a $\left(O\left(\log ^{4} n\right), O\left(\log ^{3} n\right)\right.$, $1 / n)$ LCA. We would now like to show a payment scheme $\mathcal{P}_{S L M S}$ such that the mechanism $\mathcal{M}_{S L M S}=$ $\left(\mathcal{A}_{S L M S}, \mathcal{P}_{S L M S}\right)$ is a local mechanism. (We overload the notation, letting $\mathcal{A}_{S L M S}$ represent both the on-line allocation algorithm and its respective LCA.) We therefore need to show a payment scheme which can be implemented in an LCA and guarantees truthfulness. Our payment schemes are similar in idea to the payments schemes of [APTT03] and [BBFN14].
Lemma 6.8. There exists a randomized local payment scheme $\mathcal{P}_{S L M S}$ such that the mechanism $\mathcal{M}_{\text {SLMS }}=$ $\left(\mathcal{A}_{S L M S}, \mathcal{P}_{S L M S}\right)$ is truthful in expectation. Furthermore, if all the bids are bounded by a polylogarithmic function, there exists a deterministic local payment scheme $\mathcal{P}_{\text {SLMS }}$ such that the mechanism $\mathcal{M}_{\text {SLMS }}=\left(\mathcal{A}_{\text {SLMS }}, \mathcal{P}_{\text {SLMS }}\right)$ is truthful in expectation.
Proof. Archer and Tardos AT01] showed that the following payment scheme makes for a truthful mechanism fulfilling voluntary participation. For bid $b_{i}$ :

$$
p_{i}\left(b_{i}, b_{-i}\right)=b_{i} h_{i}\left(b_{i}, b_{-i}\right)+\int_{0}^{b_{i}} h_{i}\left(x, b_{-i}\right) d x .
$$

This has to be the expected payment, and we can (deterministically) take $h_{i}(b)$ to be the expected height of machine $i$ when it bids $b_{i}$. Since $E\left[h_{i}(b)\right]=m \frac{b_{i}}{B_{-i}+b_{i}}$, where $B_{-i}=\sum_{i^{\prime} \neq i} b_{i^{\prime}}$, the payment is of the form:

$$
\begin{aligned}
p_{i}\left(b_{i}, b_{-i}\right) & =m \frac{b_{i}^{2}}{B_{-i}+b_{i}}+\sum_{x=0}^{b_{i}} h_{i}\left(x, b_{-i}\right) \\
& =m \frac{b_{i}^{2}}{B_{-i}+b_{i}}+m \sum_{x=0}^{b_{i}} \frac{x}{B_{-i}+x} .
\end{aligned}
$$

Our remaining challenge is to compute the payment in polylogarithmic time and space.
If $b_{i}$ is bounded by some polylogarithmic function of $n$, we can calculate the sum in polylogarithmic time and space, and take this to be the payment. If, however, $b_{i}$ is larger, it is not apparent how to calculate this sum in a straightforward fashion 8 We provide the following (randomized) payment scheme:

Choose, uniformly at random, $k \in\left[1, b_{i}\right]$, and take the payment to be

$$
m \frac{b_{i}^{2}}{B}+m b_{i} \cdot \frac{k}{B_{-i}+k}
$$

This gives the correct expected payment, and takes $O(1)$ time.
[BBFN14], showed that $\mathcal{A}_{S L M S}$ provides an $O(\log \log n)$ approximation to the optimal makespan. Therefore, by Theorem 6.3, the LCA of $\mathcal{A}_{S L M S}$ provides the same approximation ratio. Combining Lemma 6.5, and Lemma 6.8, we state our main result for the standard setting:
Theorem 6.9. There exists an $\left(O\left(\log ^{4} n\right), O\left(\log ^{3} n\right), 1 / n\right)$ - local mechanism to scheduling on related machines in the standard setting that is truthful in expectation, and provides an $O(\log \log n)$-approximation to the makespan.

[^6]
### 6.2 A universally truthful mechanism for the restricted setting

In the restricted setting, each job can only be allocated to one of a set $M_{j} \subseteq \mathcal{M}$ of $d$ machines. As opposed to the standard setting, $M_{j}$ is not selected by the mechanism, but is part of the input. We assume that the allocation is random (this is necessary to bound the query tree size), and the probability of machine $i$ to be in $M_{j}$ is proportional to its capacity $c_{i}$ (necessary to guarantee the approximation ratio of the allocation algorithm $\mathcal{A}_{R L M S}$ 旬. The latter requirement can be relaxed slightly, but for clarity of the proofs, we will assume that it holds exactly. Furthermore, we assume that the capacity of each machine is not too large, that is, each machine has capacity at most $O(\log n)$. We restrict our attention to the case where we would like to allocate $m=\Omega(C)$ jobs where $C$ is the total capacity of the machines. This is usually considered to be the worst case scenario (see e.g., ABKU99, BBFN14, Wie07]). We define the following (on-line) algorithm $\mathcal{A}_{R L M S}$ for assigning jobs to machines as follows. Initially, we select a permutation $\pi$ of the machines, for tie-breaking. Place job $t$ in the machine $i \in M_{j}$ for which the post-placement load, $l p_{i}^{t+1}\left(b_{i}\right)=\left\lfloor\frac{h_{i}^{t}\left(b_{i}\right)+1}{b_{i}}\right\rfloor$ is smallest, breaking ties according to $\pi$. The following claim shows why it is necessary to take the floor of the load, as the simple Greedy algorithm does not admit a universally truthful mechanism in this case.

Claim 6.10. The Greedy algorithm (unmodified) is not universally monotone in the restricted case.
Proof. Assume we have 3 machines $A, B, C$, with bids $(4,8,36)$ respectively, and a tie-breaking permutation: $A<B<C$ (Jobs always prefer machine $A$ to machines $B$ and $C$, and machine $B$ to machine $C$ ). The allocation until time $t$ has caused the heights to be $(1,3,18)$. The next job's choices are machines $A$ and $B$ (which we abbreviate to $A B$ ), and the following two jobs' choices are $B C$ and $A B$ respectively. The first job is allocated to $A$ (since the post-placement loads on $A$ and $B$ are $2 / 4$ and $4 / 8$ respectively, and $2 / 4=4 / 8$, we use the tie-breaking rule). The second job is allocated to $B$ $(4 / 8<19 / 36)$ and the third job to $B(5 / 8>3 / 4)$. The heights of the machines are now $(2,5,18)$.

Now let $B$ declare its capacity to be 9 , and assume that until time $t$, there is no difference in the allocation. The loads at time $t$ in this case are: $1 / 4,3 / 9,18 / 36$. The jobs' choices are part of the input to the mechanism, so are unaffected by the bids, and remain $A B, B C, A B$. The first job is allocated to $B(2 / 4>4 / 9)$, the second job to $C(19 / 36<20 / 36=5 / 9)$, and the third job to $A(2 / 4<5 / 9)$. The heights of the machines are now $(2,4,19)$. Thus, $B$ gets fewer jobs after bidding higher.

Theorem 6.11. For any permutation $\pi$ of the machines and any job arrival order, the allocation function $\mathcal{A}_{\text {RLMS }}$ is universally monotone increasing in the machines' bids.

From Definition 6.11 it suffices to prove the following lemma:
Lemma 6.12. For any machine $i$, fixing $b_{-i}$, for any $b_{i}^{\prime}>b_{i}$, we have that $h_{i}\left(\mathcal{A}_{R L M S}\left(b_{i}^{\prime}, b_{-i}\right)\right) \geq$ $h_{i}\left(\mathcal{A}_{R L M S}\left(b_{i}, b_{-i}\right)\right)$.

To prove Lemma 6.12, define $D^{t}\left(k, b_{i}^{\prime}, b_{i}\right)$ to be the difference in the number of jobs allocated to machine $k$ between $\mathcal{A}_{\text {RLMS }}\left(b_{i}^{\prime}\right)$ and $\mathcal{A}_{\text {RLMS }}\left(b_{i}\right)$ up to and including time $t$. We abbreviate this to $D^{t}(k)$ when $b_{i}^{\prime}$ and $b_{i}$ are clear from the context. (If machine $k$ received less jobs, then $D^{t}(k)$ is negative.) We say that machine $k$ steals a job from machine $l$ at time $t$ if $\mathcal{A}_{R L M S}^{t}\left(b_{i}\right)=l$ and $\mathcal{A}_{R L M S}^{t}\left(b_{i}^{\prime}\right)=k$. We will show that the only machine for which $D^{t}(k)$ can be positive at some time $t$ is machine $i$, therefore, as $\sum_{j=1}^{n} D^{t}(j)=0$, we have that $D^{t}(i)$ can never be negative.

Proposition 6.13. For any machine $i$, fixing $b_{-i}$, if $b_{i}^{\prime}>b_{i}$ then at all times $t$, for any machine $k \neq i$, $D^{t}(k) \leq 0$.

[^7]Informally, Proposition 6.13 says that if bin $i$ claims its capacity is larger than it actually is, no bin except for $i$ can receive more balls. The following corollary follows immediately from Proposition 6.13, and implies Lemma 6.12

Corollary 6.14. For any machine $i$, fixing $b_{-i}$, if $b_{i}^{\prime}>b_{i}$ then at all times $t, D^{t}(i) \geq 0$.
Before proving Proposition 6.13 we first will make the following simple observation
Observation 6.15. For any machine $k$, if $D^{t}(k) \leq 0$ then $l p_{k}^{t}\left(b_{i}^{\prime}\right) \leq l p_{k}^{t}\left(b_{i}\right)$.
Proof. For $k \neq i$, as $k$ 's bid is the same in both allocations, if it received less jobs by time $t$ in $\mathcal{A}_{R L M S}\left(b_{i}\right)$ then the claim follows. If $k=i$, the claim follows since $b_{i}^{\prime}>b_{i}$.

We now prove Proposition 6.13 .
Proof. The proof is by induction on $t$. At $t=1, D^{1}(k)=0$ for every $k$.
Assume the proposition is true for times $t=1, \ldots, \tau-1$. We show it holds for $t=\tau$, by contradiction. Assume that we have a machine $k \neq i$ such that $D^{\tau}(k)>0$. At time $\tau-1$, for all $k \neq i$, by the induction hypothesis, it holds that $D^{\tau-1}(k) \leq 0$. The only way that $D^{\tau}(k)>0$ is if machine $k$ has $D^{\tau-1}(k)=0$ and at time $\tau$ steals a job. Assume first that machine $k$ steals a job from machine $l \neq i$. This means that in $\mathcal{A}_{R L M S}\left(b_{i}\right)$, machine $l$ received job $\tau$, therefore

$$
\begin{equation*}
l p_{l}^{\tau}\left(b_{i}\right) \leq l p_{k}^{\tau}\left(b_{i}\right) \tag{6}
\end{equation*}
$$

By Observation 6.15 $l p_{l}^{\tau}\left(b_{i}^{\prime}\right) \leq l p_{l}^{\tau}\left(b_{i}\right)$, and so

$$
l p_{l}^{\tau}\left(b_{i}^{\prime}\right) \leq l p_{l}^{\tau}\left(b_{i}\right) \leq l p_{k}^{\tau}\left(b_{i}\right)=l p_{k}^{\tau}\left(b_{i}^{\prime}\right)
$$

If machine $k$ steals job $\tau$ from machine $l$, then $l p_{k}^{\tau}\left(b_{i}^{\prime}\right) \leq l p_{l}^{\tau}\left(b_{i}^{\prime}\right)$. This is in contradiction to Equation (6) because there cannot be an equality both here and in Equation (6), as the tie-breaking permutation $\pi$ is fixed. More precisely, if $l p_{l}^{\tau}\left(b_{i}^{\prime}\right)=l p_{l}^{\tau}\left(b_{i}\right)=l p_{k}^{\tau}\left(b_{i}\right)=l p_{k}^{\tau}\left(b_{i}^{\prime}\right)$, then job $\tau$ will be allocated to the same machine in $b_{i}$ and $b_{i}^{\prime}$, according to the permutation $\pi$.

Therefore, machine $k$ must steal job $\tau$ from machine $i$, which gives us

$$
\begin{equation*}
l p_{i}^{\tau}\left(b_{i}\right) \leq l p_{k}^{\tau}\left(b_{i}\right)=l p_{k}^{\tau}\left(b_{i}^{\prime}\right) \leq l p_{i}^{\tau}\left(b_{i}^{\prime}\right) \tag{7}
\end{equation*}
$$

The first inequality is due to the fact that machine $i$ receives job $\tau$ in $\mathcal{A}_{R L M S}\left(b_{i}\right)$. The equality is due to the fact that $D^{\tau-1}(k)=0$, and the second inequality is because machine $k$ receives job $\tau$ in $\mathcal{A}_{R L M S}\left(b_{i}^{\prime}\right)$. And so,

$$
\begin{equation*}
l p_{i}^{\tau}\left(b_{i}\right)<l p_{i}^{\tau}\left(b_{i}^{\prime}\right) \tag{8}
\end{equation*}
$$

because one of the inequalities in Equation (7) must be strict, as the tie-breaking permutation $\pi$ is fixed.
Assume that the last time before $\tau$ that machine $i$ stole a job is time $\rho$, and label by $z$ the machine that $i$ stole from at that time. We have

$$
l p_{i}^{\rho}\left(b_{i}^{\prime}\right) \leq l p_{z}^{\rho}\left(b_{i}^{\prime}\right) \leq l p_{z}^{\rho}\left(b_{i}\right) \leq l p_{i}^{\rho}\left(b_{i}\right) .
$$

The first inequality is because machine $i$ received job $\rho$ in $\mathcal{A}_{R L M S}\left(b_{i}^{\prime}\right)$. The middle inequality is because $D^{\rho}(z) \leq 0$. The last inequality is because machine $z$ received job $\rho$ in $\mathcal{A}_{R L M S}\left(b_{i}\right)$. Again, at least one inequality must be strict, giving

$$
l p_{i}^{\rho}\left(b_{i}^{\prime}\right)<l p_{i}^{\rho}\left(b_{i}\right)
$$

which implies, for all $\alpha \geq 0$,

$$
\begin{equation*}
\left\lfloor\frac{h_{i}^{\rho}\left(b_{i}^{\prime}\right)+\alpha+1}{b_{i}^{\prime}}\right\rfloor \leq\left\lfloor\frac{h_{i}^{\rho}\left(b_{i}\right)+\alpha}{b_{i}}\right\rfloor, \tag{9}
\end{equation*}
$$

since $b_{i}^{\prime}>b_{i} \geq 1$.
Because job $\rho$ was the last job that machine $i$ stole, it received at least as many jobs between $\rho$ and $\tau$ in $\mathcal{A}_{R L M S}\left(b_{i}\right)$ as in $\mathcal{A}_{\text {RLMS }}\left(b_{i}^{\prime}\right)$. Label the number of jobs $i$ received between $\rho$ and $\tau$ (including $\rho$ but excluding $\tau$ ) in $\mathcal{A}_{R L M S}\left(b_{i}\right)$ by $\beta$ and in $\mathcal{A}_{R L M S}\left(b_{i}^{\prime}\right)$ by $\beta^{*}$.
Observation 6.16. $\beta^{*} \leq \beta+1$.
Proof. Machine $i$ received at least as many jobs in $\mathcal{A}_{\text {RLMS }}\left(b_{i}\right)$ as in $\mathcal{A}_{R L M S}\left(b_{i}^{\prime}\right)$ after $\rho$. This must be true because $\rho$ was the last time machine $i$ stole a job. However, machine $i$ received the job at time $\rho$ in $\mathcal{A}_{R L M S}\left(b_{i}^{\prime}\right)$ but not in $\mathcal{A}_{R L M S}\left(b_{i}\right)$, and so we cannot claim that $\beta^{*} \leq \beta$, but only that $\beta^{*} \leq \beta+1$.

Proof of Proposition 6.13 continued. From the definition of $l p$ and equation (99), we get:

$$
\begin{align*}
l p_{i}^{\tau}\left(b_{i}^{\prime}\right) & =\left\lfloor\frac{h_{i}^{\tau}\left(b_{i}^{\prime}\right)+1}{b_{i}^{\prime}}\right\rfloor \\
& =\left\lfloor\frac{h_{i}^{\rho}\left(b_{i}^{\prime}\right)+\beta^{*}+1}{b_{i}^{\prime}}\right\rfloor  \tag{10}\\
& \leq\left\lfloor\frac{h_{i}^{\rho}\left(b_{i}^{\prime}\right)+\beta+2}{b_{i}^{\prime}}\right\rfloor  \tag{11}\\
& \leq\left\lfloor\frac{h_{i}^{\rho}\left(b_{i}\right)+\beta+1}{b_{i}}\right\rfloor  \tag{12}\\
& =l p_{i}^{\tau}\left(b_{i}\right) . \tag{13}
\end{align*}
$$

Equality (10) stems from the definition of $\beta^{*}$, Inequality (11) is due to Observation 6.16 Inequality (12) is due to Equation (9), and Equality (12) is from the definition of $\beta$.

This is in contradiction to Equation (8), and therefore $D^{\tau}(k) \leq 0$. This concludes the proof of the proposition.

Lemma 6.17. The allocation algorithm $\mathcal{A}_{\text {RLMS }}$ provides an $O(\log \log n)$-approximation to the optimal allocation.

The proof is similar to the proof for the unmodified Greedy algorithm in the case of non-uniform bins of [BBFN14]. We provide it in Appendix $\Delta$ for completeness.

Lemma 6.18. There exists a local payment scheme $\mathcal{P}_{\text {RLMS }}$ such that the mechanism $\mathcal{M}_{2}=\left(\mathcal{A}_{2}, \mathcal{P}_{2}\right)$ is universally truthful.

Proof. Having shown that $\mathcal{A}_{R L M S}$ is universally monotone, we can use the payment scheme of [AT01]:

$$
\begin{equation*}
p_{i}\left(b_{i}, b_{-i}\right)=b_{i} h_{i}\left(b_{i}, b_{-i}\right)+\sum_{x=0}^{b_{i}} h_{i}\left(x, b_{-i}\right) \tag{14}
\end{equation*}
$$

Unfortunately, the height of machine $i$ is not an easily computable function of $i$ 's bid. We therefore need to explicitly compute the value of $h_{i}\left(x, b_{-i}\right)$ for every $x \in\left[0, b_{i}\right]$. That is, we need to run the allocation algorithm $\mathcal{A}_{R L M S}$ again with each possible bid $x \in\left[0, b_{i}\right]$. As the running time of $\mathcal{A}_{R L M S}$ is $O(\log n)$, and $b_{i}=O(\log n)$, this will take at most $O\left(\log ^{2} n\right)$. Note that as the time to compute the permutation is $O\left(\log ^{3} n\right)$, computing the payment this way will not affect the asymptotic running time of the mechanism.

We conclude:
Theorem 6.19. There exists an $\left(O\left(\log ^{4} n\right), O\left(\log ^{3} n\right), 1 / n\right)$-local mechanism to scheduling on related machines in the restricted setting that is universally truthful and gives an $O(\log \log n)$-approximation to the makespan.

## 7 Local auctions for unit demand buyers

We propose local truthful mechanisms for auctions with unit-demand buyers, in which each buyer is interested in at most $k$ items, and each item is desired by at most a polylogarithmic number of buyers. First, we tackle the case where all buyers have the same valuation for the items in their sets, and the buyers' private information is their sets. Then we examine the case in which the buyers' sets are public knowledge and the buyers' private information is their valuations for their items, with the restriction that buyers have the same valuation for all items in their set.

Related work Combinatorial auctions are an extremely well-studied problem in algorithmic game theory. The general premise is the following: we wish to allocate $m$ goods to $n$ players, who have valuations for subsets of goods, with the goal of maximizing the social welfare. The general problem, where each player may have an an arbitrary valuation for each subset of the goods is known to be $N P$ hard; indeed, even approximating the optimal solution for single-minded bidders to within $\sqrt{m}(1-\epsilon)$ is $N P$-hard [LOS02]. Therefore, in order to obtain useful approximation algorithms, we must relax some of our demands. One such relaxation is limiting ourselves to identical items. An example is the case of $k$-minded bidders: There are $m$ identical indivisible goods, and $n$ buyers with $k$ valuations each - each buyer $i$ has a valuation $v_{i}(j)$ for obtaining $j$ items (where $j$ is between 1 and $k$ ). [DN10] gave a PTAS for this problem, and showed that (under certain restrictions), obtaining an FPTAS is $N P$-hard. [KV12] gave universally truthful mechanisms for combinatorial auctions in an on-line model.

We use the following theorem from [MRVX12] for finding maximal matchings in undirected graphs.
Theorem 7.1. $M R V X 12]$ Let $G=(V, E)$ be an undirected graph with $n$ vertices and maximum degree d. Then there is an $\left(O\left(\log ^{4} n\right), O\left(\log ^{3} n\right), 1 / n\right)$ - local computation algorithm for maximal matching.

As in the case of load balancing, the idea behind Theorem 7.1 is that one can simulate the wellknown Greedy on-line algorithm for maximal matching, using a random permutation of the edges.

### 7.1 Unit-demand buyers with uniform value

We first consider the following scenario. We have a set $\mathcal{I}$ of $n$ unit-demand buyers, and a set $\mathcal{J}$ of $m$ indivisible items. There is a fixed, identical value for all items, which we normalize to 1 . Each buyer $i$ is interested in a set $J_{i}$ of at most $k$ items (where $k$ is a constant). We can treat this auction as a graph $G=(V, E)$, in which $V=\mathcal{J} \cup \mathcal{I}$, and $E=\left\{(i, j): i \in \mathcal{I}, j \in J_{i}\right\}$. The value of a subset $S$ to buyer $i$ is $v_{i}(S)=1$ if $S \cap J_{i} \neq \emptyset$ and 0 otherwise. Namely, the buyers are indifferent between the items in their set (they all have the same valuation for the items in their set, and a zero valuation for all other items). The utility of buyer $i$ is quasi-linear, that is, when she receives items $S$ and pays $p$ her utility is $u_{i}(S, p)=v_{i}(S)-p$. We assume that the subsets $J_{i}$ are selected uniformly at random and that $k n / m=O(1) 10$

Our goal is to design a local mechanism that maximizes the social welfare. In order to do this, we would like to satisfy as many buyers as possible, allocating each buyer a single item from her set. We call this type of auction an $k$-UDUV (unit demand, uniform value) auction.

Ideally, we would like to find a maximum matching between the buyers and items, as this will maximize the social welfare. However, it is not possible to solve the maximum matching problem locally, as the following theorem shows. This is one important challenge that the local setting adds to the algorithmic mechanism design.

Theorem 7.2. There does not exist an LCA for the maximum matching problem.

[^8]Proof. To see why it is not always possible to solve the maximum matching locally, consider the following family of homomorphic graphs: $\mathcal{G}=\left\{G_{i}\right\}$. All $G_{i} \in \mathcal{G}$ have $2 n$ vertices: $\left\{v_{1}, v_{2}, \ldots, v_{2 n}\right\}$. In each $G_{i}$, vertices $v_{-i}$ comprise an (odd) cycle, and vertex $v_{i}$ is connected to vertex $v_{i-1}$ (modulo $2 n$ ). Each $G_{i}$ has a unique maximum matching. We are given as input a graph $G \in \mathcal{G}$, (i.e., we know it is $G_{i}$ for some $i$, but we don't know the value of $i$ ). We would like to know whether the edge $e=\left(v_{1}, v_{2}\right)$ is in the maximum matching. Note that the edge $e$ will be in exactly half of the maximum matchings.

Assume that the graph is either $G_{n}$ or $G_{n+1}$. In the distributed model, this implies that the distance between the edge $e=\left(v_{1}, v_{2}\right)$ and the distinguishing place of $G_{n}$ and $G_{n+1}$ is $n$ edges, which will be a lower bound on the time to detect the correct graph.

In the local computation model, we can write the edges in a random order. This implies that one needs to query, on average, $n$ edges to distinguish between $G_{n}$ and $G_{n+1}$.

Therefore there cannot exist an LCA for maximum matching.

## Corollary 7.3. There does not exist an LCA for the maximum matching problem in bipartite graphs.

Proof. The proof is similar to the general case. In the bipartite case, though, in each $G_{i}$, the vertices $v_{1}, v_{2}, \ldots, v_{i-1}, v_{i+2}, \ldots, v_{2 n}$ comprise an even cycle. Vertex $v_{i}$ is connected to vertex $v_{i-1}$ and vertex $v_{i+1}$ is connected to vertex $v_{i+2}$ (modulo $2 n$ ). Note that the edges $\left(v_{i-1}, v_{i}\right)$ and $\left(v_{i+1}, v_{i+2}\right)$ must be in the maximum matching, because the maximum matching in this case is of size $n$, meaning all vertices must be matched, including $v_{i}$ and $v_{i+1}$.

Since there is no local algorithm for computing the maximum matching ${ }^{111}$, we will content ourselves with finding an approximation to the maximum matching.

### 7.1.1 A $\frac{1}{2}$-approximation to the maximum matching

To obtain a $\frac{1}{2}$-approximation, we use the local greedy matching algorithm of [MRVX12], which we denote by $\mathcal{A}_{U D U V}$. Algorithm $\mathcal{A}_{U D U V}$ simulates the well-known greedy on-line algorithm - as an edge arrives, it is added to the matching, if possible.

Our mechanism $\mathcal{M}_{U D U V}=\left(\mathcal{A}_{U D U V}, \mathcal{P}_{U D U V}\right)$ works as follows. The mechanism receives from each buyer $i$ a subset $J_{i}^{\prime} \subset J$. For the allocation algorithm $\mathcal{A}_{U D U V}$, the mechanism decides on a random order in which it considers the items. Specifically, the mechanism assigns each item $j$ a real number $r_{j} \sim_{u}[0,1]$, sampled independently and uniformly. The order of the items is determined by $r_{j}$ (higher $r_{j}$ items are considered first). Notice that because the $r_{j}$ are allocated independently, buyers cannot influence the order in which the items as considered. Given this order induced by $r_{j}$, the mechanism considers items one at a time. When item $j$ is considered, if there is some buyer $i$ such that $j \in J_{i}^{\prime}$ and buyer $i$ was not allocated any item yet, then $j$ is allocated to buyer $i$. If there is more than one such buyer, ties are broken lexicographically.

As the values of all items are identical, any payment scheme $\mathcal{P}_{U D U V}$ that fulfills voluntary participation is adequate, i.e., the payment can be any value in the range $[0,1]$. For example, charge $p=1 / 2$ from any buyer that receives an item and $p=0$ from any buyer that does not receive an item.

In order to show that our mechanism is truthful, we need only to show that buyers cannot profit by bidding $J_{i}^{\prime} \neq J_{i}$.

Theorem 7.4. In the $k$-UDUV auction, the mechanism $\mathcal{M}_{U D U V}=\left(\mathcal{A}_{U D U V}, \mathcal{P}_{U D U V}\right)$ is universally truthful and provides a $\frac{1}{2}$-approximation to the optimal allocation.

Proof. The proof will be done in two steps. First, we show that for any $J_{i}^{\prime}$, bidding $J_{i} \cap J_{i}^{\prime}$ weakly dominates bidding $J_{i}^{\prime}$. Second, we show that bidding $J_{i}$ weakly dominates bidding any $J_{i}^{*} \subseteq J_{i}$.

[^9]To show that $J_{i} \cap J_{i}^{\prime}$ weakly dominates bidding $J_{i}^{\prime}$, label the items in $J_{i} \cap J_{i}^{\prime}$ as good, and those in $J_{i}^{\prime} \backslash J_{i}$ by bad. If a good item is allocated to buyer $i$ when she bids $J_{i}^{\prime}$, it will also be allocated to her when bidding $J_{i} \cap J_{i}^{\prime}$. Therefore the value of buyer $i$ cannot decrease by bidding $J_{i} \cap J_{i}^{\prime}$, and hence $J_{i} \cap J_{i}^{\prime}$ weakly dominates bidding $J_{i}^{\prime}$.

To show that bidding $J_{i}$ weakly dominates bidding any $J_{i}^{*} \subseteq J_{i}$, consider the following. If buyer $i$ does not receive any items when bidding $J_{i}^{*}$, the claim trivially holds. Assume buyer $i$ receives item $j$ when bidding $J_{i}^{*}$. Then, when bidding $J_{i}$, if she has not received any item from $J_{i} \backslash J_{i}^{*}$ before considering item $j$, then she will receive item $j$. Therefore, if she receives an item when bidding $J_{i}^{*}$, she will also receive an item when bidding $J_{i}$, and have the same valuation and utility. This proves that bidding $J_{i}$ weakly dominates bidding $J_{i}^{*} \subseteq J_{i}$.

The reasoning that the allocation is a $\frac{1}{2}$-approximation is similar to the proof of maximal versus maximum matching. Consider a buyer which is not allocated an item in $\mathcal{A}_{U D U V}$ and is allocated an item in the optimal allocation. Her item is allocated to a unique different buyer in $\mathcal{A}_{U D U V}$. This bounds the number of buyers allocated items in the optimal allocation and not in $\mathcal{A}_{U D U V}$ by the number of buyers that are allocated items in $\mathcal{A}_{U D U V}$, giving the factor of $\frac{1}{2}$ approximation, and completing the proof of the theorem.

The fact that $\mathcal{A}_{U D U V}$ is a $\left(O\left(\log ^{4} n\right), O\left(\log ^{3} n\right), 1 / n\right)$ - LCA for maximal matching on graph of bounded degree $k$ was shown in [MRVX12]. The result also implicitly holds for bipartite graphs in which the degrees are bounded on one side and distributed binomially on the other. This is exactly the case when buyers are interested in $k$ items each and the buyers' choices can be seen as sampled uniformly from the items. Therefore, we derive the following theorem:

Theorem 7.5. The $k$-UDUV auction has an $\left(O\left(\log ^{4} n\right), O\left(\log ^{3} n\right), 1 / n\right)$ - local computation mechanism which is universally truthful and provides a $\frac{1}{2}$-approximation to the optimal social welfare.

### 7.2 Unit demand buyers, uniform-buyer-value

We have a set $\mathcal{I}$ of $n$ buyers, and a set $\mathcal{J}$ of $m$ items. Each buyer $i$ is interested in a set of at most $k$ items, $J_{i} \subseteq \mathcal{J}$, which is public knowledge, and has a private valuation, $t_{i}$ (which represents the value of any item from $J_{i}$ to buyer $i$ ). Buyer $i$ 's valuation for subset $S$ is $v_{i}(S)=t_{i}$ if $S \cap J_{i} \neq \emptyset$, and 0 otherwise. The utility of buyer $i$ is quasi-linear, namely her utility of receiving subset $S$ and paying $p$ is $u_{i}(S, p)=v_{i}(S)-p$.

We can treat this auction as a weighted graph $G=(V, E)$, in which $V=\mathcal{J} \cup \mathcal{I}$, and $E=\cup_{i} E_{i}$ where $E_{i}=\left\{(i, j): j \in J_{i}\right\}$. Every edge $e \in E_{i}$ has weight $w(e)=t_{i}$.

We assume that $J_{i}$ has a uniform or binomial distribution and that $k n / m=O(1){ }^{12}$ In addition, we make the simplifying assumption that the buyers are Bayesian - the valuations $t_{i}$ are randomly drawn from some prior (not necessarily known) distribution, that is identical to all buyers. We call this type of auction an $k$-UDUBV (unit demand, uniform buyer value) auction.

We require that if buyer $i$ does not receive an item, she pays nothing. If buyer $i$ receives an item, the mechanism charges her $p_{i}(b)$, where $b$ is the bid vector. (Any buyer will receive at most one item in the allocation of the mechanism). We would like to ensure that bidding truthfully is a dominant strategy for all buyers. Hence, we need to show that, for all $b_{i}$ and $b_{-i}$, we have $u_{i}\left(t_{i}, b_{-i}\right) \geq u_{i}\left(b_{i}, b_{-i}\right)$.

The allocation algorithm, $\mathcal{A}_{U D U B V}$, is as follows. First, $\mathcal{A}_{U D U B V}$ orders the buyers by their bids. Starting with the buyer with the highest bid, each buyer $i$ is allocated an item $j_{i} \in J_{i}$ such that $j_{i}$ has not yet been allocated. If more than one such item exists, we allocate the (lexicographically) first $j_{i} \in J_{i}$. (We assume the items have lexicographic order.) If there is no such item, then buyer $i$ is not allocated any item. We continue until we cannot allocate any more items.

[^10]First, we claim that the resulting allocation is a $\frac{1}{2}$-approximation.
Claim 7.6. The allocation algorithm $\mathcal{A}_{U D U B V}$ provides a $\frac{1}{2}$-approximation to the optimal allocation, with respect to the bids $b$.

Proof. The proof is similar to the proof that any maximal matching is a 2-approximation to a maximum matching. Regard the auction as a bipartite graph $G=(U, W, E)$, with $U$ representing the buyers and $W$ representing the items. There is a weighted edge between each buyer $i$ and every item $j_{i} \in J_{i}$. The weight of each edge $e=\left(i, j_{i}\right)$ is the bid of buyer $i, b_{i}$. The optimal allocation is a maximum weighted matching, while $\mathcal{A}_{U D U B V}$ considers the buyers in the order of their $b_{i}$ 's and finds a maximal matching.

If an edge $e=\left(i, j_{i}\right)$ is added in $\mathcal{A}_{U D U B V}$ but not in the optimal matching, then it is allocated instead of at most 2 edges in the optimal matching (an edge $e^{\prime}$ containing $i$ and an edge $e^{\prime \prime}$ containing $\left.j_{i}\right)$. Because $\mathcal{A}_{U D U B V}$ considers edges according to their weights, we know that $w(e) \geq w\left(e^{\prime}\right)$ and $w(e) \geq w\left(e^{\prime \prime}\right)$. Therefore $2 w(e) \geq w\left(e^{\prime}\right)+w\left(e^{\prime \prime}\right)$ and so the ratio between $\mathcal{A}_{U D U B V}$ and the optimal allocation is at least $\frac{1}{2}$.

We now need to specify the payment mechanism. To calculate buyer $i$ 's payment when she receives an item, we run $\mathcal{A}_{U D U B V}$ without buyer $i$. Buyer $i$ pays the smallest value for which any of her items is sold when the auction is run without her. (This is exactly the minimal value of $b_{i}$ which would still gain her an item). We label buyer $i$ 's payment by $p_{i}$, hence, the payments are $\mathcal{P}_{U D U B V}=\left\{p_{1}, \ldots p_{n}\right\}$.

Claim 7.7. In mechanism $\mathcal{M}_{U D U B V}=\left(\mathcal{A}_{U D U B V}, \mathcal{P}_{U D U B V}\right)$, for all buyers $i$ and all $b_{i}$, bidding $t_{i}$ weakly dominates bidding $b_{i}$.

Proof. We will show that, fixing the bids of all other buyers at $b_{-i}$,

1. Buyer $i$ has no incentive to over-bid, i.e., bid $b_{i}>t_{i}$.
2. Buyer $i$ has no incentive to under-bid i.e., bid $b_{i}<t_{i}$.

To prove (1), we notice that if buyer $i$ receives an item, then she has no incentive to bid higher, as she has no preference between items. Furthermore, bidding higher cannot change her payment, as her payment is independent of her bid. If she does not receive an item, then $p_{i} \geq b_{i}\left(=t_{i}\right) \sqrt{13}$, and so if she bids more, she might receive an item, but will have to pay at least $t_{i}$ if she does, which will result in a non-positive utility.

To prove (2), we notice that if buyer $i$ does not receive an item, she cannot obtain an item by bidding lower, because the algorithm allocates first to higher bids. If she is allocated an item, then bidding lower will not make a difference, unless she bids under $p_{i}$, in which case she will not receive any item, and hence have zero utility.

Claims 7.6 and 7.7 imply the following.
Theorem 7.8. The mechanism $\mathcal{M}_{U D U B V}$ is universally truthful and provides a $\frac{1}{2}$-approximation to the optimal social welfare.

Algorithm $\mathcal{A}_{U D U B V}$ is a $\left(O\left(\log ^{4} n\right), O\left(\log ^{3} n\right), 1 / n\right)$ - LCA for maximal matching on graph of bounded degree $k$, by [MRVX12]. Notice, however, that we need to run $\mathcal{A}_{U D U B V}$ once for calculating the allocation, and $k$ more times for calculating the payment. Hence, we have the following.

Theorem 7.5, (2) There is an $\left(O\left(\log ^{4} n\right), O\left(\log ^{3} n\right), 1 / n\right)$ - local mechanism for $k-U D U B V$ auction which is universally truthful and provides a $\frac{1}{2}$-approximation to the optimal social welfare.

[^11]
## 8 Single minded bidders

We extend the results of Section 7.2 to the case of combinatorial auctions with single-minded bidders: There is a set $\mathcal{I}$ of $n$ buyers, and a set $\mathcal{J}$ of $m$ items. Each buyer $i$ is interested in a set of at most $k$ items, $J_{i} \subseteq \mathcal{J}$, which is public knowledge, and has a private valuation, $t_{i}$, which represents the value of the entire subset $J_{i}$ to buyer $i$. Buyer $i$ 's valuation for subset $S$ is $v_{i}(S)=t_{i}$ if $J_{i} \subseteq S$, and 0 otherwise. The utility of buyer $i$ is quasi-linear, namely her utility of receiving subset $S$ and paying $p$ is $u_{i}(S, p)=v_{i}(S)-p$.

As in Subsection 7.2, we assume that $J_{i}$ has a uniform or binomial distribution, and that the valuations $t_{i}$ are randomly drawn from some prior (not necessarily known) distribution, and $k n / m=O(1)$.

The allocation algorithm, $\mathcal{A}_{k S M B}$, is as follows. First, $\mathcal{A}_{k S M B}$ orders the buyers by their bids. Starting with the buyer with the highest bid, each buyer $i$ is allocated subset $J_{i}$ such that no item $j_{i} \in J_{i}$ has been allocated yet. We continue until we cannot allocate any more subsets.

Claim 8.1. The allocation algorithm $\mathcal{A}_{k S M B}$ provides a $\frac{1}{k}$-approximation to the optimal allocation, with respect to the values $b$.

Proof. Compare the allocation of Algorithm $\mathcal{A}_{k S M B}, J^{*}$, to the optimal allocation, OPT. Each set $J \in J^{*}$ is chosen by $\mathcal{A}_{k S M B}$ instead of at most $k$ sets in $O P T$, but its weight is greater than each of their weights, because $\mathcal{A}_{k S M B}$ is a greedy algorithm.

The payment scheme is as follows. To calculate buyer $i$ 's payment when she receives an item, we run $\mathcal{A}_{k S M B}$ without buyer $i$. Buyer $i$ pays the highest value of the allocated sets $J_{x}$ for which $J_{i} \cap J_{x} \neq \emptyset$. (This is exactly the minimal value of $b_{i}$ which would still gain her an item). We label buyer $i$ 's payment by $p_{i}$, and let $\mathcal{P}_{k S M B}=\left\{p_{1}, \ldots p_{n}\right\}$.

Claim 8.2. In mechanism $\mathcal{M}_{k S M B}=\left(\mathcal{A}_{k S M B}, \mathcal{P}_{k S M B}\right)$, for all buyers $i$ and all $b_{i}$, bidding $t_{i}$ weakly dominates bidding $b_{i}$.

The proof is similar to the proof of Claim 7.7 and is omitted.
Combining Claims 8.1 and 8.2 we get
Theorem 8.3. There is an $\left(O\left(\log ^{4} n\right), O\left(\log ^{3} n\right), 1 / n\right)$ - local mechanism for combinatorial auctions with known $k$-single minded bidders (where the sets are sampled uniformly at random), which is universally truthful and provides a $\frac{1}{k}$-approximation to the optimal social welfare.

## 9 Random serial dictatorship

We would like to allocate $n$ houses to $n$ agents. We assume that each agent is interested in a constant number of houses, $d$, and that the preferences are drawn from the uniform distribution. Each agent $i$ has a complete preference relation $R_{i}$ over the $d$ houses. In the random serial dictatorship algorithm, a permutation over the agents is generated, and then each agent chooses her most preferred house out of the unallocated houses. We would like to simulate this algorithm locally: Each agent $i$ is allocated, uniformly and at random, an integer $r_{i} \in\left[n^{4}\right]$. $r$ determines the permutation: if $r_{i}<r_{j}$ then agent $i$ makes a choice before agent $j$. If $r_{i}=r_{j}$, the tie is broken lexicographically. Agent $i$ checks, for each of its $d$ housing choices in the order induced by $R_{i}$, whether it has been allocated already. If it has not, $i$ chooses the house. For $i$ to check whether or not house $j$ has already been allocated, it needs to check whether any of the agents interested in $j$ had already chosen it. $i$ therefore needs to recursively check the allocation of all agents that arrived before her, on which her allocation depends. As in ARVX12, MRVX12], we model this using a query tree. We recall the following lemma from [MRVX12]:

Lemma 9.1 ([|MRVX12]). Let $G=(\{V, U\}, E)$ be a bipartite graph, $|V|=n$ and $|U|=m$ and $n=c m$ for some constant $c \geq 1$, such that for each vertex $v \in V$ there are d edges chosen independently and at random between $v$ and $U$. For any constant $\alpha$ there is a constant $C$ which depends only on $d$ and $\alpha$ such that

$$
\operatorname{Pr}[|T|<C \log n]>1-1 / n^{\alpha},
$$

where the probability is taken over all of the possible permutations $\pi \in \Pi$ of the vertices of $G$, and $T$ is a random query tree in $G$ under $\pi$.

This implies that in order to reply to any query, we will require more that $O(\log n)$ queries to the graph with probability at most $1 / n^{2}$, which implies the following theorem:

Theorem 9.2. Let $k$ be some constant integer $k>0$. Consider a house allocation problem with $n$ agents and $n$ houses, and let each agent preference list length be bounded by $k$, where each list is drawn uniformly at random from the set of all possible lists of length $k$. Then there is an $\left(O\left(\log ^{4} n\right), O\left(\log ^{3} n\right), 1 / n\right)$ - LCA which simulates, and whose output is identical to, the Random Serial Dictatorship allocation algorithm.

## 10 Open questions

We provide several interesting open questions.

Open question 1 How restrictive is the local computation setting? It remains an interesting open question in general, and specifically for the load balancing problem, while the intractability of the nonlocal problem carries over to the local computation setting, there exists a truthful PTAS for the non-local setting. How well can we approximate the optimal solution using LCAs? Furthermore, it would be interesting to quantify the added difficulty imposed by the requirement of truthfulness.

Open question 2 An intriguing open problem is whether there exists a truthful local mechanism in the unit-demand combinatorial auction setting, where both the sets and the valuations are private information. We conjecture that this is impossible, especially considering a similar result of [FLSS11], albeit in a different setting (auctions with budgets).

Open question 3 In this paper, we show an LCA which provides a $1 / 2$-approximation to the optimal solution for unit-demand buyers, using a maximal matching LCA. [MV13], showed that there exists a $(1-\epsilon)$-approximation algorithm to maximum matching. Unfortunately, that algorithm does not yield a monotonic allocation. Is there an LCA which provides a better than $1 / 2$-approximation that can be transformed to a local computation mechanism?

Open question 4 If we view the $k$-single minded bidder combinatorial auction as a hypergraph $H=$ $(V, E)$, where each item is represented by a vertex $v \in V$, and each player by a weighted hyperedge $e \in E$, the problem is reduced to maximal weighted independent set in $k$-regular hypergraphs, (equivalently weighted $k$-set packing). [CH01] show how to approximate weighted $k$-set packing to within $2(k+1) / 3$, via local improvements. It would be interesting to see if it would be possible to apply of their techniques to obtain a local mechanism with a better approximation ratio.

Open question 5 In this paper, we often need to make the assumption that the demands of the bidders are drawn from some distribution in order to ensure the locality of our mechanisms. This is quite a strong assumption, and it would be interesting to see how far it can be relaxed.

Acknowledgements We would like to thank Amos Fiat, Alon Naor and Amit Weinstein for their useful input.

## References

[ABKU99] Yossi Azar, Andrei Z. Broder, Anna R. Karlin, and Eli Upfal. Balanced allocations. SIAM Journal on Computing, 29(1):180-200, 1999.
[ANR95] Yossi Azar, Joseph Naor, and Raphael Rom. The competitiveness of on-line assignments. J. Algorithms, 18(2):221-237, 1995.
[APTT03] Aaron Archer, Christos H. Papadimitriou, Kunal Talwar, and Éva Tardos. An approximate truthful mechanism for combinatorial auctions with single parameter agents. Internet Mathematics, 1(2), 2003.
[Arr50] Kenneth J. Arrow. A difficulty in the concept of social welfare. Journal of Political Economy, 58(4):328-346, 1950.
[ARVX12] Noga Alon, Ronitt Rubinfeld, Shai Vardi, and Ning Xie. Space-efficient local computation algorithms. In SODA, pages 1132-1139. SIAM, 2012.
[AS08] Noga Alon and Joel Spencer. The Probabilistic Method. John Wiley, 3rd edition, 2008.
[AT01] Aaron Archer and Éva Tardos. Truthful mechanisms for one-parameter agents. In Proc. 42nd Annual IEEE Symposium on Foundations of Computer Science (FOCS), pages 482491, 2001.
[BBFN14] Petra Berenbrink, André Brinkmann, Tom Friedetzky, and Lars Nagel. Balls into nonuniform bins. J. Parallel Distrib. Comput., 74(2):2065-2076, 2014.
[BCM03] John W. Byers, Jeffrey Considine, and Michael Mitzenmacher. Simple load balancing for distributed hash tables. In IPTPS, pages 80-87, 2003.
[BKS10] Moshe Babaioff, Robert D. Kleinberg, and Aleksandrs Slivkins. Truthful mechanisms with implicit payment computation. In ACM Conference on Electronic Commerce, pages 43-52, 2010.
[CH01] Barun Chandra and Magnús M. Halldórsson. Greedy local improvement and weighted set packing approximation. J. Algorithms, 39(2):223-240, 2001.
[CK10] George Christodoulou and Annamária Kovács. A deterministic truthful ptas for scheduling related machines. In Proc. 21st ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 1005-1016, 2010.
[Cla71] Edward H. Clarke. Multipart pricing of public goods. Public Choice, 11(1):17-33, 1971.
[DD09] Shahar Dobzinski and Shaddin Dughmi. On the power of randomization in algorithmic mechanism design. In Proc. 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pages 505-514, 2009.
[DN10] Shahar Dobzinski and Noam Nisan. Mechanisms for multi-unit auctions. J. Artif. Intell. Res. (JAIR), 37:85-98, 2010.
[DNS12] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Truthful randomized mechanisms for combinatorial auctions. J. Comput. Syst. Sci., 78(1):15-25, 2012.
[Dob11] Shahar Dobzinski. An impossibility result for truthful combinatorial auctions with submodular valuations. In Proc. 43rd Annual ACM Symposium on the Theory of Computing (STOC), pages 139-148, 2011.
[FKKV13] Amos Fiat, Anna R. Karlin, Elias Koutsoupias, and Angelina Vidali. Approaching utopia: Strong truthfulness and externality-resistant mechanisms. In ITCS, pages 221-230, 2013.
[FKPS10] Patrik Floréen, Petteri Kaski, Valentin Polishchuk, and Jukka Suomela. Almost stable matchings by truncating the gale-shapley algorithm. Algorithmica, 58(1):102-118, 2010.
[FLSS11] Amos Fiat, Stefano Leonardi, Jared Saia, and Piotr Sankowski. Single valued combinatorial auctions with budgets. In ACM Conference on Electronic Commerce, pages 223-232, 2011.
[FMP00] Tomás Feder, Nimrod Megiddo, and Serge A. Plotkin. A sublinear parallel algorithm for stable matching. Theor. Comput. Sci., 233(1-2):297-308, 2000.
[GI89] Dan Gusfield and Robert W. Irving. The Stable Marriage Problem - Structure and Algorithms. Foundations of computing series. MIT Press, 1989.
[Gro73] Theodore Groves. Incentives in teams. Econometrica, 41 (4):617-631, 1973.
[GS62] David Gale and Lloyd S. Shapley. College admissions and the stability of marriage. American Mathematical Monthly, 69:9-14, 1962.
[GS85] David Gale and Marilda Sotomayor. Some remarks on the stable matching problem. Discrete Applied Mathematics, 11(3):223-232, 1985.
[HKV13] Martin Hoefer, Thomas Kesselheim, and Berthold Vöcking. Truthfulness and stochastic dominance with monetary transfers. In ACM Conference on Electronic Commerce, pages 567-582, 2013.
[HS88] Dorit S. Hochbaum and David B. Shmoys. A polynomial approximation scheme for machine scheduling on uniform processors: Using the dual approximation approach. SIAM Journal on Computing, 17(3):539-551, 1988.
[IM05] Nicole Immorlica and Mohammad Mahdian. Marriage, honesty, and stability. In SODA, pages 53-62, 2005.
[Knu76] Donald E. Knuth. Mariages stables. Les Presses de l'Université de Montréal, 1976.
[KP09] Fuhito Kojima and Parag A. Pathak. Incentives and stability in large two-sided matching markets. American Economic Review, 99(3):608-627, 2009.
[KV12] Piotr Krysta and Berthold Vöcking. Online mechanism design (randomized rounding on the fly). In ICALP (2), pages 636-647, 2012.
[LOS02] Daniel J. Lehmann, Liadan O'Callaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. J. ACM, 49(5):577-602, 2002.
[LST87] Jan Karel Lenstra, David B. Shmoys, and Éva Tardos. Approximation algorithms for scheduling unrelated parallel machines. In Proc. 28 th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pages 217-224, 1987.
[LZ03] Enyue Lu and S. Q. Zheng. A parallel iterative improvement stable matching algorithm. In $H i P C$, pages 55-65, 2003.
[MCWG95] Andreu Mas-Collel, Michael Whinston, and Jerry Green. Microeconomic Theory. Oxford University Press, 1995.
[MRVX12] Yishay Mansour, Aviad Rubinstein, Shai Vardi, and Ning Xie. Converting online algorithms to local computation algorithms. In Proc. 39th International Colloquium on Au tomata, Languages and Programming (ICALP), pages 653-664, 2012.
[MSZ13] Michael Maschler, Eilon Solan, and Shmuel Zamir. Game Theory. Cambridge Press, 2013.
[MV13] Yishay Mansour and Shai Vardi. A local computation approximation scheme to maximum matching. In APPROX-RANDOM, pages 260-273, 2013.
[Mye81] Roger B. Myerson. Optimal auction desing. Mathematics of Operations Research, 6:5874, 1981.
[NO08] Huy N. Nguyen and Krzystof Onak. Constant-time approximation algorithms via local improvements. In Proc. 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pages 327-336, 2008.
[NR99] Noam Nisan and Amir Ronen. Algorithmic mechanism design (extended abstract). In Proc. 31st Annual ACM Symposium on the Theory of Computing (STOC), pages 129-140, 1999.
[NRTV05] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani. Algorithmic Game Theory. Cambridge University Press, 2005.
[Qui85] Michael J. Quinn. A note on two parallel algorithms to solve the stable marriage problem. BIT Numerical Mathematics, 25(3):473-476, 1985.
[Rot03] Alvin E. Roth. The origins, history, and design of the resident match. Journal of the American Medical Association, 289(7):909-912, 2003.
[RR99] Alvin E. Roth and Uriel G. Rothblum. Truncation strategies in matching markets - in search of advice for participants. Econometrica, 67(1):21-43, 1999.
[RTVX11] Ronitt Rubinfeld, Gil Tamir, Shai Vardi, and Ning Xie. Fast local computation algorithms. In Proc. 2nd Symposium on Innovations in Computer Science (ICS), pages 223-238, 2011.
[TL84] S. S. Tseng and Richard C. T. Lee. A parallel algorithm to solve the stable marriage problem. BIT, 24(3):308-316, 1984.
[TW07] Kunal Talwar and Udi Wieder. Balanced allocations: the weighted case. In Proc. 39th Annual ACM Symposium on the Theory of Computing (STOC), pages 256-265, 2007.
[Vic61] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. The Journal of Finance, 16 (1):8-37, 1961.
[Vöc03] Berthold Vöcking. How asymmetry helps load balancing. J. ACM, 50(4):568-589, 2003.
[Wie07] Udi Wieder. Balanced allocations with heterogenous bins. In Proc. 19th ACM Symposium on Parallel Algorithms and Architectures (SPAA), pages 188-193, 2007.

## A Proof of Lemma 6.17

Lemma6.17, The allocation algorithm $\mathcal{A}_{\text {RLMS }}$ provides an $O(\log \log n)$-approximation to the optimal allocation.

We prove the theorem for the case $d=2$ (each job can be assigned to one of 2 machines). The proof is easily expandable for the case of $d>2$. For the proof (not the algorithm), we regard each machine $i$ of capacity $c_{i}$ as having $c_{i}$ slots of capacity 1 . Before presenting the proof we need several definitions:

The load vector of an allocation of jobs into $n$ machines is $L=\left(\ell_{1}, \ldots, \ell_{n}\right)$, where $\ell_{i}=h_{i}$ is the load of machine $i$. The normalized load vector $\bar{L}$ consists of the members of $L$ in non-increasing order (where the order among machines with the same load is arbitrary). For the case of non-uniform machines of capacities $c_{1}, \ldots c_{n}$, and total capacity $C=\sum_{i=1}^{n} c_{i}$, we define the slot-load vector $S=$ $\left(h_{1,1}, \ldots h_{1, c_{1}}, h_{2,1}, \ldots h_{2, c_{2}}, \ldots h_{n, 1}, \ldots h_{n, c_{n}}\right.$ ), where if machine $i$ contains $r$ jobs, the first $r \bmod c$ slots will have $\lceil r / c\rceil$ jobs, and the remaining slots will have $\lfloor r / c\rfloor$ jobs. If a machine has an uneven allocation of jobs, we call the slots with more jobs heavy, and the slots with less jobs light. If the load on a machine is an integer (all of the slots of the machine have an identical number of jobs assigned to them), we call all the slots light. When we add a job to a machine, we add the job to one of the light slots, arbitrarily. The normalized slot load vector $\bar{S}$ is $S$ sorted in non-increasing order (slots of the same machine may be separated in $\bar{S}$ ). We add a subscript $t$ to these vectors, i.e., $L_{t}, \bar{L}_{t}, S_{t}$ and $\bar{S}_{t}$ to indicate the vector after the allocation of the $t$-th job.

Definition A. 1 (Majorization, $\succeq)$. We say that a vector $P=\left(p_{1}, \ldots, p_{n}\right)$ majorizes vector $Q=$ $\left(q_{1}, \ldots, q_{m}\right)$ (denoted $P \succeq Q$ ) if and only if for all $1 \leq k \leq \min (m, n)$,

$$
\sum_{i=1}^{k} \bar{p}_{i} \geq \sum_{i=1}^{k} \bar{q}_{i},
$$

where $\bar{p}_{i}$ and $\bar{q}_{i}$ are the $i$-th entries of the normalized vectors $\bar{P}$ and $\bar{Q}$.

$$
\text { For } n \in \mathbb{N} \text {, let }[n] \text { denote }\{1, \ldots, n\} \text {. }
$$

Definition A. 2 (System Majorization). Let $A$ and $B$ be two processes allocating $m$ jobs into machines with the same total capacity $n$. Let $\tau=\left(\tau_{1} \ldots \tau_{2 m}\right), \tau_{i} \in[n]$ be a vector representing the (slot) choices of the $m$ jobs ( $\tau_{2 i-1}$ and $\tau_{2 i}$ are the choices of the $i$-th job). Let $S^{A}(\tau)$ and $S^{B}(\tau)$ be the slot load vectors using $A$ and $B$ respectively with the random choices specified by $\tau$. Then we say

1. A majorizes $B$ (denoted by the overloaded notation $A \succeq B$ ) if there is a bijection $f:[n]^{2 m} \rightarrow$ $[n]^{2 m}$ such that for all possible random choices $\tau \in[n]^{2 m}$, we have

$$
L^{A}(\tau) \succeq L^{B}(f(\tau))
$$

2. The maximum load of $A$ majorizes the maximum load of $B$ (denoted by $A \succeq_{m} B$ ) if there is a bijection $f:[n]^{2 m} \rightarrow[n]^{2 m}$ such that for all possible random choices $\tau \in[n]^{2 m}$, we have

$$
\ell_{1}^{A}(\tau) \geq \ell_{1}^{B}(f(\tau)),
$$

where $\ell_{1}^{A}(\tau)$ and $\ell_{1}^{B}(f(\tau))$ are the loads of the most loaded bins in $A$ and $B$ respectively with the random choices specified by $\tau$ and $f(\tau)$ respectively.

Observation A.3. $A \succeq B \Rightarrow A \succeq{ }_{m} B$.

We now turn to the proof of Lemma6.17
First, notice that if we have an system of $m$ identical machines, each of capacity 1 , both the unmodified Greedy algorithm and the allocation algorithm $\mathcal{A}_{R L M S}$ will behave in exactly the same way - the load and the $\lfloor$ load $\rfloor$ are the same if the capacity is 1 . From [ABKU99], we know that the maximal load on any machine when allocating $m=n$ jobs (to $n$ machines with capacity 1 ) with the Greedy algorithm, is $\Theta(\log \log n)$. Therefore, the maximal load when allocating $m=m$ jobs with $\mathcal{A}_{R L M S}$ is also $\Theta(\log \log n)$ in this setting. We would like to show that the maximal load of a system with non-uniform machines of total capacity $C$ is majorized by the maximal load of a system with $C$ machines of capacity 1 , when the allocating algorithm is $\mathcal{A}_{R L M S}$. We will show that the first system majorizes the second, and deduce the required result from Observation A. 3

We restate Claim 2.4 of [Wie07]:
Claim A. 4 (Wie07]). Let $P$ and $Q$ be two normalized integer vectors such that $P \succeq Q$. If $i \leq j$ then $P+e_{i} \succeq Q+e_{j}$ where $e_{i}$ is the $i$-th unit vector and $P+e_{i}$ and $Q+e_{j}$ are normalized.

Lemma A.5. For allocation algorithm $\mathcal{A}_{\text {RLMS }}$, let $A$ be a system with non-uniform machines of total capacity $C$, and $B$ be a system with $C$ uniform machines of capacity 1 each. Then $B \succeq A$.

Proof. We use the slot load vectors of systems $A$ and $B$ (in $B$ the load vector and slot load vector are identical), and show that $S^{B}(f(\tau)) \succeq S^{A}(\tau)$. The bijection is such that the jobs in both processes choose the same $k_{1}<k_{2} \in\{1, \ldots, C\}$ in the normalized slot load vectors, and the choice corresponds to machines $k_{1}, k_{2}$ in $B$ and the machines associated with those specific slots in system $A$. We use induction: for $t=0$, the claim is trivially true.

From the inductive hypothesis, before the allocation of the $t$-th job, $S_{t-1}^{B}(f(\tau)) \succeq S_{t-1}^{A}(\tau)$. In system $B$, the $t$-th job goes to machine $k_{2}$. In system $A$, if the $\lfloor$ load $\rfloor$ of the machine of $k_{1}$ is greater than that of the machine of $k_{2}$, the job goes to $k_{2}$ if $k_{2}$ is a light slot, or to a slot to the right of $k_{2}$ (a lighter slot of the same machine), if $k_{2}$ is a heavy slot. If the ไloads」of the machines of $k_{1}$ and $k_{2}$ are the same, again, the job goes to $k_{2}$ if $k_{2}$ is a light slot, or to a slot to the right of $k_{2}$ (again, a lighter slot of the same machine), if $k_{2}$ is a heavy slot. In all cases, by Claim A.4 it follows that $S^{B}(f(\tau)) \succeq S^{A}(\tau)$.


[^0]:    *Department of Computer Science, Bar Ilan University. E-mail: avinatan@cs.biu.ac.il. This research was supported by supported by ISF grant number 1241/12, and by BSF grant number 2012344.
    ${ }^{\dagger}$ Blavatnik School of Computer Science, Tel Aviv University. E-mail: mansour@tau.ac.il. This research was supported in part by the Google Inter-university Center for Electronic Markets and Auctions, the Israeli Centers of Research Excellence program, the Israel Science Foundation, the United States-Israel Binational Science Foundation, and the Israeli Ministry of Science .
    ${ }^{\ddagger}$ Blavatnik School of Computer Science, Tel Aviv University. E-mail: shaivar1@post.tau.ac.il. This research was supported in part by the Google Europe Fellowship in Game Theory.

[^1]:    ${ }^{1}$ A blocking pair is a man $m$ and a woman $w$ such that $m$ prefers $w$ to the woman he is matched to, and $w$ prefers $m$ to the man she is matched to. A matching is stable if there is no blocking pair.

[^2]:    ${ }^{2}$ Usually referred to in Economic literature as ex-ante efficiency.
    ${ }^{3}$ Our model differs from the model of [RTVX11] in that their model requires that the LCA always obeys the time and space bounds, and returns an error with some probability. It is easy to see that any algorithm which conforms to our model can be modified to conform to the model of [RTVX11] by forcing it to return an error if the time or space bound is violated (the other direction does not necessarily hold). Note however, that using this translation, a truthful mechanism in our model would not necessarily translate to a truthful mechanism in the model of [RTVX11].

[^3]:    ${ }^{4}$ This bound is reached by substituting $\lambda \geq 2 e-1$ into the standard Chernoff bound $\operatorname{Pr}[X>(1+\lambda) \mu] \leq\left(\frac{e^{\lambda}}{(1+\lambda)^{1+\lambda}}\right)^{\mu}$

[^4]:    ${ }^{5}$ Instead of "deciding" on the preference lists in advance, each man chooses the $(i+1)^{t h}$ woman on his list only if he is rejected from the $i^{t h}$ - this mechanism is known to be equivalent to the mechanism which we use (see, e.g., (Knu76).

[^5]:    ${ }^{6}$ We need an assumption that each ball restricted to a constant number of machines, either explicitly or implicitly through the algorithm $\mathcal{L B}$.
    ${ }^{7}$ That is, if bin $A$ has capacity 4 and height 2 and bin $B$ has capacity 8 and height 5 , the job will go to machine $B$.

[^6]:    ${ }^{8}$ We can approximate the payment up to an additive factor of $1 / B_{-i}$ using $p_{i}\left(b_{i}, b_{-i}\right) \approx m \frac{b_{i}^{2}}{B_{-i}+b_{i}}+m\left(b_{i}-B_{-i} \ln (1+\right.$ $\left.b_{i} / B_{-i}\right)$ ), but this might change the incentives of the machines and is not guaranteed to be truthful.

[^7]:    ${ }^{9}$ Although this assumption is somewhat strong for a theoretical discussion, from a practical viewpoint it usually holds that a machines' capacity is somewhat proportional to the number of jobs that is usually scheduled on it.

[^8]:    ${ }^{10}$ We require this for ease of analysis. However, it suffices that the sets are distributed in such a way as to resemble a uniform or binomial distribution MRVX12].

[^9]:    ${ }^{11}$ Actually, we only require a maximum matching in the special case of bipartite graphs. Corollary 7.3 shows that this is not possible either.

[^10]:    ${ }^{12}$ As in the previous section, we require this for ease of analysis. However, it suffices that the sets are distributed in such a way as to resemble a uniform or binomial distribution [MRVX12].

[^11]:    ${ }^{13}$ As specified, if buyer $i$ does not receive an item, she pays 0 . However, if the mechanism were to compute the payment, i.e., run the mechanism without her, the payment would be $p_{i} \geq b_{i}\left(=t_{i}\right)$.

