Settling the Communication Complexity of VCG-Based Mechanisms for All Approximation Guarantees

Frederick Qiu
Princeton University
Princeton, USA
fqiu@princeton.edu

S. Matthew Weinberg
Princeton University
Princeton, USA
smweinberg@princeton.edu

ABSTRACT

We consider truthful combinatorial auctions with items $M := [m]$ for sale to n bidders, where each bidder i has a private monotone valuation function $v_i : 2^M \to \mathbb{R}_+$. Among truthful mechanisms, *maximal-in-range* (MIR) mechanisms (sometimes called *VCG-based*) achieve the best-known approximation guarantees among all poly-communication deterministic truthful mechanisms in all previously-studied settings. Our work settles the communication complexity necessary to achieve any approximation guarantee via an MIR mechanism. Specifically:

Let $\mathsf{MIR}_{\mathsf{SubMod}}(m,k)$ denote the best approximation guarantee achievable by an MIR mechanism using 2^k communication between bidders with submodular valuations over m items. Then for all $k = \Omega(\log(m))$, $\mathsf{MIR}_{\mathsf{SubMod}}(m,k) = \Omega(\sqrt{m/(k\log(m/k))})$. When we set $k = \Theta(\log(m))$, this improves the previous best lower bound for polynomial communication maximal-in-range mechanisms from $\Omega(m^{1/3}/\log^{2/3}(m))$ to $\Omega(\sqrt{m}/\log(m))$. Additionally, $\mathsf{MIR}_{\mathsf{SubMod}}(m,k) = O(\sqrt{m/k})$. Moreover, our mechanism can be implemented with 2^k simultaneous value queries and computation, and is optimal with respect to the value query and computational/succinct representation models. The mechanism also works for bidders with subadditive valuations. When $k = \Theta(\log(m))$, this improves the previous best approximation guarantee for polynomial communication maximal-in-range mechanisms from $O(\sqrt{m})$ to $O(\sqrt{m/\log(m)})$.

Let also $\mathsf{MIR}_{\mathsf{Gen}}(m,k)$ denote the best approximation guarantee achievable by an MIR mechanism using 2^k communication between bidders with general valuations over m items. Then for all $k = \Omega(\log(m))$, $\mathsf{MIR}_{\mathsf{Gen}}(m,k) = \Omega(m/k)$. When $k = \Theta(\log(m))$, this improves the previous best lower bound for polynomial communication maximal-in-range mechanisms from $\Omega(m/\log^2(m))$ to $\Omega(m/\log(m))$. Additionally, $\mathsf{MIR}_{\mathsf{Gen}}(m,k) = O(m/k)$. Moreover, our mechanism can be implemented with 2^k simultaneous value queries and computation, and is optimal with respect to the value query and computational/succinct representation models. When $k = \Theta(\log(m))$, this improves the previous best approximation guarantee for polynomial communication maximal-in-range mechanisms from $O(m/\sqrt{\log(m)})$ to $O(m/\log(m))$.

STOC '24, June 24–28, 2024, Vancouver, BC, Canada

ACM ISBN 979-8-4007-0383-6/24/06
https://doi.org/10.1145/3618260.3649706

CCS CONCEPTS

• **Theory of computation → Computational pricing and auctions**; *Approximation algorithms analysis.*

KEYWORDS

Combinatorial Auction, Communication Complexity, Maximal-in-Range, VCG-Based

ACM Reference Format:
Frederick Qiu and S. Matthew Weinberg. 2024. Settling the Communication Complexity of VCG-Based Mechanisms for All Approximation Guarantees. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24), June 24–28, 2024, Vancouver, BC, Canada.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3618260.3649706

1 INTRODUCTION

In a *combinatorial auction*, a central designer has a set of items $M := [m]$, and bidders $N := [n]$. Each bidder i has a monotone valuation $v_i : 2^M \to \mathbb{R}_+$, unknown to the designer. The designer interacts with the bidders to produce an allocation $A = (A_1, \dots, A_n)$ of the items (where $A_i \cap A_j = \emptyset$ for all $i \neq j$), and their goal is to select one maximizing the welfare, defined $\sum_{i \in N} v_i(A_i)$.

As an algorithmic resource allocation problem, combinatorial auctions are extremely well-studied – see further discussion in Section 1.3. Combinatorial auctions are similarly well-studied in economic settings, where the bidders' incentives are now relevant. That is, while an efficient communication protocol suffices in a purely algorithmic setting, that protocol must also be incentive compatible and incentivize all bidders to follow it. Here, the designer may also charge each bidder i a price p_i, and the bidder aims to optimize their utility: $v_i(A_i) - p_i$.

When considering either desiderata separately, the state of affairs is well-understood. For example, with polynomial communication, a tight $\Theta(\sqrt{m})$-approximation for monotone valuations [5, 32, 36], a tight 2-approximation for subadditive valuations [24, 25], and a tight $e/(e-1)$-approximation for XOS valuations [17, 25] are known. Additionally, the optimal achievable guarantee for submodular valuations is known to lie in $[2e/(2e-1), e/(e-1) - 10^{-5}]$ [21, 26]. However, these protocols are not incentive compatible. Similarly, the classical *Vickrey-Clark-Groves* (VCG) mechanism is incentive compatible and finds the welfare-maximizing allocation [8, 27, 37], but requires exponential communication for any of the above-referenced valuation classes.

As such, a central open problem within Economics and Computation is understanding the extent to which communication-efficient truthful mechanisms can match the approximation guarantees of communication-efficient (not necessarily incentive compatible) protocols. A key framework to tackle this agenda is *maximal-in-range* (MIR) mechanisms. For example, state-of-the-art deterministic truthful mechanisms for monotone, subadditive, XOS, and submodular valuations are maximal-in-range [17, 28]. Our main results settle the approximation guarantees achievable by MIR mechanisms. We overview this agenda and our results below.

1.1 Maximal-in-Range Mechanisms

MIR mechanisms leverage the VCG mechanism to trade off approximation guarantees for efficiency. In particular, the VCG mechanism implies a truthful mechanism that maximizes welfare over any set of (possibly unstructured) outcomes. Specifically, one can define any *allocation bank* $\mathcal{A}$ of allocations, and the VCG mechanism will truthfully optimize welfare *over all allocations in* $\mathcal{A}$. Formally:

Theorem 1.1 (VCG Mechanism [8, 27, 37]). *Let $\mathcal{A}$ be a collection of allocations, and let $\mathcal{P}$ be any communication protocol among the n bidders to find the welfare-maximizing allocation in $\mathcal{A}$. Then there is a deterministic truthful mechanism that selects a welfare-maximizing allocation in $\mathcal{A}$ using $n+1$ black-box calls to $\mathcal{P}$.*

The resulting mechanism is termed a Maximal-in-Range Mechanism *for $\mathcal{A}$ [35].*

MIR mechanisms therefore provide a structured algorithmic framework to design deterministic truthful mechanisms: one selects an allocation bank $\mathcal{A}$ and designs a protocol $\mathcal{P}$ to optimize over it. This framework induces a tradeoff between efficiency and optimality: richer allocation banks may contain allocations whose welfare better approximates the true optimal welfare, but smaller allocation banks may require less communication to optimize over.

For the following reasons, there is significant interest in understanding the approximation guarantees achievable by MIR Mechanisms:

- In all settings, the best-known approximation guarantees achieved by polynomial communication deterministic truthful mechanisms are achieved by MIR mechanisms [17, 28]. Moreover, this claim has held for the entire duration of the study of combinatorial auctions (that is, no polynomial communication deterministic truthful mechanisms that outperform the best-known MIR mechanism at the time have ever been discovered).[1]
- All deterministic truthful mechanisms satisfying four natural properties are affine maximizers, a generalization of MIR mechanisms [30].[2] Moreover, if an affine maximizer guarantees an α-approximation using β communication on all submodular/XOS/subadditive/monotone valuations, then the MIR mechanism with the same allocation bank does so for the same valuation class as well.
- Some conjecture that indeed MIR mechanisms are the optimal deterministic mechanisms (and therefore settling the approximation guarantees of poly-communication MIR mechanisms will eventually settle the approximation guarantees of all poly-communication deterministic truthful mechanisms), although this conjecture remains far from settled. However, we show that MIR mechanisms indeed achieve optimal approximation guarantees for *all* deterministic truthful mechanisms in the value query model and the computational/succinct representation model (see Table 1).

Indeed, MIR mechanisms have been studied since the start of Economics and Computation as a field [35], and several works over the past two decades make significant progress understanding their strengths and limitations. Our main results close these gaps. We now state our main results, and afterwards discuss their context. In the following theorem statements (and the rest of this paper), define $\mathsf{MIR}_{\mathsf{SubMod}}(m,k)$ (and $\mathsf{MIR}_{\mathsf{Gen}}(m,k)$, respectively) to be the optimal approximation guarantee that can be achieved by an MIR mechanism using at most 2^k communication between bidders with submodular (and general, respectively) valuations over m items.

Theorem 1.2. *For all $k=\Omega(\log(m))$ and for all $n=\Omega(\sqrt{m/k})$, $\mathsf{MIR}_{\mathsf{SubMod}}(m,k)=\Omega(\sqrt{m/(k\log(m/k))})$. In particular, the best possible approximation guarantee for submodular valuations by MIR mechanisms with $\text{poly}(m)$ communication is $\Omega(\sqrt{m}/\log(m))$.*

This improves prior work, beginning with the impossibility of $m^{1/6}$ with polynomial communication for MIR mechanisms for submodular valuations by [14], which was later improved to $m^{1/3}/\log^{2/3}(m)$ by [9].

The particular constant of $1/2$ in the exponent is significant because we now know that the MIR mechanism of [17] achieving an $O(\sqrt{m})$-approximation using polynomial communication is essentially tight. Still, our second main result improves their guarantee slightly.

Theorem 1.3. *For all $k=\Omega(\log(m))$ and for all n, we have $\mathsf{MIR}_{\mathsf{SubMod}}(m,k)=O(\sqrt{m/k})$. In particular, our mechanism guarantees a $O(\sqrt{m/\log m})$ approximation in $\text{poly}(m)$ communication.*

Moreover, the mechanism we construct can be implemented using $2^{O(k)}$ simultaneous value queries, or in time $2^{O(k)}$ in the succinct representation model. Our mechanism guarantees an $O(\sqrt{m/k})$-approximation for subadditive valuations as well.

Theorem 1.3 is a mild improvement over [17] (it saves a $\sqrt{\log(m)}$ factor). Still, we note that the MIR mechanism of [17] is exceptionally simple, and no better guarantee was previously known. Together, Theorems 1.2 and 1.3 nail down the achievable approximation guarantees for submodular valuations by MIR mechanisms with $\text{poly}(m)$ communication up to a factor of $\Theta(\sqrt{\log(m)})$, exponentially improved over the prior gap of $\tilde{\Theta}(m^{1/6})$. Additionally, it is worth noting that in the value query and computational/succinct representation models, the MIR mechanism we construct for subadditive valuations is optimal for *all* deterministic truthful mechanisms; see [22] and the appendix of the full paper.[3]

[1] Note that [14] discover *exponential-communication* non-MIR deterministic truthful mechanisms that outperform the best *poly-communication* MIR mechanisms in multi-unit domains [14], but *poly-communication* deterministic truthful mechanisms have never outperformed MIR mechanisms.

[2] An affine maximizer also is defined by an allocation bank $\mathcal{A}$, scalars $\vec{c}\in\mathbb{R}^n_{\geq 0}$, and an adjustment $v_0:\mathcal{A}\to\mathcal{A}$. The affine maximizer selects an allocation $(A_1,\ldots,A_n)$ optimizing $v_0(A_1,\ldots,A_n)+\sum_{i\in N}c_i\cdot v_i(A_i)$. Affine maximizers are also truthful using the VCG payment scheme.

[3] https://arxiv.org/pdf/2404.00831

We also consider general valuations. Here, nearly-tight bounds were previously known, and we improve both to be tight up to constant factors.

Theorem 1.4. *For all $k = \Omega(\log(m))$, $\mathsf{MIR}_{\mathsf{Gen}}(m,k) = \Theta(m/k)$. In particular, the best possible approximation guarantee for monotone valuations by MIR mechanisms with* $\text{poly}(m)$ *communication is* $\Theta(m/\log m)$. *Moreover, the mechanism we construct can be implemented using $2^{O(k)}$ simultaneous value queries, or in time $2^{O(k)}$ in the succinct representation model.*

When considering $\text{poly}(m)$ communication, this gives logarithmic improvements for the previous best impossibility result of $\Omega(m/\log^2(m))$ by [9], and the previous best approximation guarantee of $O(m/\sqrt{\log(m)})$ by [28].[4] Still, the minor improvements are significant as both bounds are now tight.

Table 1 places our work alongside prior work for all three considered models. Our main results consider the communication model, but also imply minor improvements in the value query or computational model for free (or modest additional work). We contextualize the key takeaways from this table.

- In the communication model with arbitrary monotone valuations, we improve both the state-of-the-art mechanism and lower bound, reducing the gap between them from $\Theta(\log^{1.5}(m))$ to $\Theta(1)$. In other words, $\Theta(m/\log(m))$ is the best achievable guarantee in $\text{poly}(m)$ communication.
- In the communication model with subadditive/XOS/submodular valuations, we improve the state-of-the-art lower bounds and slightly improve the state-of-the-art mechanism, reducing the gap between them from $\tilde{\Theta}(m^{1/6})$ to $\Theta(\sqrt{\log(m)})$. In other words, $\tilde{\Theta}(\sqrt{m})$ is the best achievable guarantee in $\text{poly}(m)$ communication.
- In the value queries model and computational model with either arbitrary monotone or subadditive/XOS/submodular valuations, we slightly improve state-of-the-art mechanisms. In the value queries model, these improvements match pre-existing lower bounds on *any* $\text{poly}(m)$*-query deterministic truthful mechanism*.[5] In the computational model, we further slightly improve prior-best lower bounds on *any* $\text{poly}(m)$-*time deterministic truthful mechanism* to match our MIR mechanisms. That is, we now know that MIR mechanisms achieve the optimal approximation guarantees among all $\text{poly}(m)$-query deterministic truthful mechanisms, and all $\text{poly}(m)$-time deterministic truthful mechanisms in the succinct representation model. These MIR mechanisms follow by observing that our new communication-efficient mechanisms can be implemented with $\text{poly}(m)$ value queries (in fact, simultaneous value queries). Our slight improvement on lower bounds follows a similar outline as prior work, but is more careful with lower order terms.[6]

[4] Although the $\Theta(\sqrt{\log(m)})$ improvement in the upper bound seems small, the mechanism of [28] is actually far from optimal for large k; for example, if we allow $2^{O(m)}$ communication for a sufficiently small constant, then the mechanism of [28] is only $O(\sqrt{m})$-approximate, whereas ours is $O(1)$-approximate.

[5] For arbitrary monotone, subadditive, and XOS valuations, our mechanisms even match pre-existing lower bounds on *any* $\text{poly}(m)$*-query deterministic algorithm*.

[6] Our computational lower bounds also use the stronger assumption of the randomized Exponential Time Hypothesis instead of the assumption RP ≠ NP used in prior work. However, this is essentially necessary if we care about lower order terms.

1.2 Technical Highlights

Below, we overview one technical highlight from our algorithms, and one technical highlight from our lower bounds.

Technical Background: Prior Algorithms. The state-of-the-art MIR mechanism for general valuations [28] and for submodular valuations [17] are quite different. Parameterizing their approaches by k, the MIR mechanism of [28] partitions the items into k chunks of size m/k, and considers $\mathcal{A}$ to be the set of allocations that keeps together all items in the same chunk. Observe that there are only 2^k sets that any bidder might possibly receive in $\mathcal{A}$, so optimizing over $\mathcal{A}$ can be done with 2^k communication.

On the other hand, [17] considers $\mathcal{A}$ to be the set of allocations that either give all items to the same bidder, or that gives each bidder a set of at most $O(k/\log(m/k))$ items. Again, there are only 2^k sets that any bidder might receive, so optimizing over $\mathcal{A}$ can be done with 2^k communication.

Technical Highlight: Our Algorithms. Our algorithm for monotone valuations adds just one new idea to that of [28]: consider multiple partitions. Specifically, repeat the process z times of partitioning M into k chunks, calling the chunks $B_s^{(\ell)}$ for $\ell \in [z]$ and $s \in [k]$. Let $\mathcal{A}$ be the set of allocations such that each bidder receives a set of the form $\bigcup_{s \in S} B_s^{(\ell)}$ for some $\ell \in [z]$ and $S \subseteq [k]$ (that is, each bidder receives a set that picks a single partition, and then is a union of chunks for that partition). Taking $z = 2^{O(k)}$ still requires just $2^{O(k)}$ communication. The main difference to [28] is that while any single partition can only give an $m/\sqrt{k}$-approximation, the best of $2^{O(k)}$ partitions improve the guarantee to m/k. Informally, this is because by taking many partitions, it is likely that for every set of k items, there exists a partition where the items go to different chunks. This allows any k items to be optimally allocated, so allocating the "most important" k items will yield an m/k-approximation.

Our algorithm for subadditive/XOS/submodular valuations is in some sense more like [28] than [17]. Sticking exclusively to an approach like [17] is almost optimal, but doomed to lose a $\sqrt{\log(m/k)}$ factor due to the fact that we can only exhaust over sets of size $O(k/\log(m/k))$ in 2^k communication. On the other hand, sticking exclusively to an approach like [28] cannot guarantee an $o(m/k)$ approximation since at most k bidders are given items, so significant changes are needed to leverage this approach. Our algorithm is as follows:

- Take $z = 2^{O(k)}$ partitions of the items into $\sqrt{m/k}$ buckets of $\sqrt{mk}$ items (so there are $2^{O(k)}\sqrt{m/k}$ total buckets).
- Within each bucket, take $z' = 2^{O(k)}$ partitions of the items into k chunks of $\sqrt{m/k}$ items. This induces sets of the form $C_{s',\ell',s,\ell}$, where $\ell \in [z]$ determines which bucketing we use, $s \in [\sqrt{m/k}]$ labels the bucket, $\ell' \in [z']$ determines which partition of that bucket into chunks we use, and $s' \in [k]$ labels the chunk.
- Finally, $\mathcal{A}$ denotes the set of allocations where each bidder either receives all the items M, or receives a set of the form $\cup_{s' \in S} C_{s',\ell',s,\ell}$ for some $S \subseteq [k]$ and ℓ', s, ℓ. In other words, each bidder chooses a bucket (ℓ chooses the bucketing and s chooses a bucket in that bucketing), chooses a partition of

Table 1: Summary of our results (bolded) compared to prior work (unbolded) when $k = \Theta(\log(m))$. All models referenced above are for deterministic mechanisms.

Communication	General	Subadditive/XOS/Submodular
Prior Best Mechanisms	$O(m/\sqrt{\log(m)})$ [28]	$O(\sqrt{m})$ [17]
Our MIR Mechanisms	$\boldsymbol{O(m/\log(m))}$ (Thm. 1.4)	$\boldsymbol{O(\sqrt{m/\log(m)})}$ (Thm. 1.3)
Our MIR Lower Bounds	$\boldsymbol{\Omega(m/\log(m))}$ (Thm. 1.4)	$\boldsymbol{\Omega(\sqrt{m}/\log(m))}$ (Thm. 1.2)
Prior MIR Lower Bounds	$\Omega(m/\log^2(m))$ [9]	$\Omega(m^{1/3}/\log^{2/3}(m))$ [9]
Value Queries	**General**	**Subadditive/XOS/Submodular**
Prior Best Mechanisms	$O(m/\sqrt{\log(m)})$ [28]	$O(\sqrt{m})$ [17]
Our MIR Mechanisms	$\boldsymbol{O(m/\log(m))}$ (Thm. 1.4)	$\boldsymbol{O(\sqrt{m/\log(m)})}$ (Thm. 1.3)
Truthful Lower Bounds	$\downarrow$	$\Omega(\sqrt{m/\log(m)})$ [22]
Algorithmic Lower Bounds	$\Omega(m/\log(m))$ [5]	$\Omega(\sqrt{m/\log(m)})$ for Subadditive/XOS [17]
Computation (Succ. Rep.)	**General**	**Subadditive/XOS/Submodular**
Prior Best Mechanisms	$O(m/\sqrt{\log(m)})$ [28]	$O(\sqrt{m})$ [17]
Our MIR Mechanisms	$\boldsymbol{O(m/\log(m))}$ (Thm. 1.4)	$\boldsymbol{O(\sqrt{m/\log(m)})}$ (Thm. 1.3)
Our Truthful Lower Bounds	$\boldsymbol{\Omega(m/\log(m))}$ (Full Paper)	$\boldsymbol{\Omega(\sqrt{m/\log(m)})}$ (Full Paper)
Prior Truthful Lower Bounds	$m^{1-\Theta(1)}$ [9]	$\sqrt{m^{1-\Theta(1)}}$ [22]

that bucket into chunks, and then receives a subset of the chunks for the chosen partition of the chosen bucket.

Observe that there are again only $2^{O(k)} \cdot \sqrt{mk} \cdot 2^{O(k)} \cdot 2^k = 2^{O(k)}$ ways to choose such a set that a bidder might get in $\mathcal{A}$, so $\mathcal{A}$ can be optimized over in $2^{O(k)}$ communication.

Broadly, the idea is to use subadditivity and binomial tail bounds to argue that a bucketing exists where we only need to allocate k items optimally within each bucket to get a $\sqrt{m/k}$-approximation. Then we leverage our mechanism for general valuations within each bucket to allocate the k items.

The key high-level technical takeaway we wish to emphasize is that exhausting over collections of large chunks of items, versus exhausting over collections of a few items, seems to be "the right" way to achieve optimal approximation guarantees for MIR mechanisms. This is because each partition into large chunks of items can achieve the desired approximation ratios for many configurations of bidders simultaneously, which allows us to overcome the barrier that there are too many configurations of bidders to try satisfying them by asking for small sets.

Technical Background: Prior Lower Bounds. Prior lower bounds on MIR mechanisms follow from an argument of the following form: (1), derive structure on $\mathcal{A}$, using the fact that $\mathcal{A}$ guarantees a good approximation, then (2), show that this structure embeds a hard communication problem. At a very high-level, the initial approach of [14] could be described as using first-principles for (1), and then a non-trivial reduction from SetDisjointness for (2). The state-of-the-art approaches of [7, 9] instead use advanced machinery based on generalizations of the VC-dimension for (1), so that a trivial argument for (2) suffices.

In slightly more detail, [7, 9] find a set N' of bidders and M' of items such that $\mathcal{A}$ must contain every possible allocation of items in M' to bidders in N' (in this case, we say that $\mathcal{A}$ *shatters* (M', N')). Then if the valuation class requires exponential communication to exactly optimize welfare, a communication lower bound of $2^{\Omega(|M'|)}$ follows immediately for (2), because when restricting attention to M' and N', $\mathcal{A}$ considers all allocations and is exactly optimal.

Technical Highlight: Our Lower Bounds. Our lower bounds leverage *some* of the advanced machinery developed in [7, 9] to understand the structure of any $\mathcal{A}$ that achieves a good approximation, but stops short of going all the way to shattering. Indeed, the bounds in [9] are tight (up to perhaps logarithmic factors) for approaches that insist on fully shattering some (M', N'). Instead, we leverage just enough structure to move towards a communication lower bound. For general valuations, the structure established in prior work actually suffices for a direct reduction from SetDisjointness that saves a $\log(m)$ factor. For submodular valuations, we derive a novel structure on $\mathcal{A}$, but ultimately avoid a full shattering argument to save a $\tilde{\Theta}(m^{1/6})$ factor.

The high-level takeaway is that our lower bounds improve over prior results by uncovering "the right" structure on $\mathcal{A}$ to enable a simple-but-not-trivial communication lower bound, rather than pushing all the way towards fully shattering.

1.3 Related Work

We've previously discussed the most-related work to ours: [35] introduces the concept of MIR mechanisms, based off principles of the VCG mechanism [8, 27, 37]. [28] provides the first (and until our work, state-of-the-art) approximation for general valuations via an MIR mechanism, which is also the best previous truthful deterministic mechanism. [17] provides the first (and until our work, state-of-the-art) approximation for submodular valuations via an MIR mechanism, which is also the best previous truthful deterministic mechanism. [14] provide the first lower bounds on poly(m)-communication MIR mechanisms for submodular valuations. [9], building on tools developed in [7], improve their bounds for submodular valuations and provide the first bounds for general valuations.

Beyond these directly-related works, there is a rich body of works on combinatorial auctions broadly. These works provide context for the study of MIR mechanisms specifically. For example, [30] establish that all mechanisms satisfying four natural properties are affine maximizers (a generalization of MIR mechanisms that achieve identical approximation guarantees). Therefore, our results on MIR mechanisms also immediately bound approximation guarantees achievable by this class of mechanisms. There are, however, deterministic mechanisms that are not affine maximizers (posted-price mechanisms are one such example). As such, there is also an active body of research aiming to understand the approximation guarantees of deterministic truthful mechanisms with bounded communication. Our mechanisms are now the state-of-the-art deterministic truthful mechanisms with poly(m) communication for submodular, XOS, subadditive, and general valuations. On the other hand, there are significantly fewer lower bounds that hold for all deterministic truthful mechanisms. Specifically, the only such result for any of these four classes is a $4/3 + \varepsilon$ lower bound for two XOS bidders [3, 6, 13].

The discussion of the previous paragraph considers a protocol to be truthful if it is an ex-post Nash equilibrium for bidders to follow it. That is, as long as every other bidder is following the protocol *for some plausible valuations* $\vec{v}_{-i}$, it is in bidder i's best interest to follow the protocol as well (for all i). One could instead seek mechanisms that are dominant strategy truthful: even if the other players use bizarre strategies that are not prescribed for any $\vec{v}_{-i}$, it is still in bidder i's best interest to follow the protocol. On this front, [18] recently establish that no dominant strategy truthful mechanism can achieve an $m^{1-\varepsilon}$ approximation for general valuations in poly(m) communication. This means that, up to lower-order terms, the MIR mechanisms we develop are also optimal among dominant strategy truthful mechanisms (our mechanisms are dominant strategy truthful because they can be implemented using poly(m) simultaneous communication).

Finally, there is significant related work on the communication complexity of combinatorial auctions broadly, considering protocols (without incentives) [1, 15, 17, 21, 24–26, 36], deterministic truthful mechanisms [32], and randomized truthful mechanisms [2, 4, 10, 12, 16, 29]. There is also significant related work on the computational complexity of combinatorial auctions broadly, again considering protocols without incentives [31, 33, 38], and strong inapproximability results for truthful and computationally efficient mechanisms [11, 19, 20, 22, 23].

2 PRELIMINARIES AND NOTATION

2.1 Shattering

When convenient, we may think of an allocation $A : M \to N \cup \{*\}$ as a function from items to bidders, where $*$ denotes an item not allocated to any bidder. As such, we may use notation such as $A|_{M'} := (A_1 \cap M', \dots, A_n \cap M')$ and $\mathcal{A}|_{M'} := \{A|_{M'} : A \in \mathcal{A}\}$ to denote the restriction of an allocation/allocation bank to the items M'. For allocation banks on disjoint sets of items $\mathcal{A}$ and $\mathcal{B}$, we may also use the notation $\mathcal{A} \times \mathcal{B} := \{(A_1 \cup B_1, \dots, A_n \cup B_n) : A \in \mathcal{A}, B \in \mathcal{A}\}$ to denote an allocation bank where each combination of $A \in \mathcal{A}$ and $B \in \mathcal{B}$ is possible.

For a set of bidders $v_1, \dots, v_n : 2^M \to \mathbb{R}_+$, define $A^*(\vec{v})$ to be an optimal allocation, and let $\mathrm{OPT}(\vec{v})$ be the welfare under $A^*(\vec{v})$. For any $N' \subseteq N$, define $\mathrm{OPT}(\vec{v}, N')$ to be the welfare of bidders N' under $A^*(\vec{v})$. Let $\mathrm{MIR}_{\mathcal{A}}(\vec{v})$ be the welfare from an optimal allocation in $\mathcal{A}$.

One concept that will repeatedly appear in our arguments is that of an allocation bank $\mathcal{A}$ *shattering* a collection of items/bidders.

Definition 2.1 (Shattering). An allocation bank $\mathcal{A}$ *d-shatters* a pair (M', N') if for all items $j \in M'$, there exists a set $T_j \subseteq N'$ with $|T_j| = d$ such that $\bigtimes_{j \in M'} T_j \subseteq \mathcal{A}|_{M'}$. That is, for each of the $d^{|M'|}$ ways to allocate each item $j \in M'$ to a bidder in T_j, there exists an allocation in $\mathcal{A}$ that allocates the items in M' in this manner.

If $\mathcal{A}$ $|N'|$-shatters (M', N'), we will simply say that $\mathcal{A}$ shatters (M', N').

Prior lower bounds of [7, 9] use this concept extensively, and eventually find a large set of items that are shattered. Our lower bound for submodular functions leverage this machinery for $d < |N'|$ instead of $d = |N'|$ to achieve a $\tilde{\Theta}(m^{1/6})$ improvement. This concept is also helpful for understanding intuitively how our mechanisms provide good approximation guarantees.

2.2 Formal Statement of Models

Our main results consider the communication model, where each player i holds the valuation function $v_i(\cdot)$ and we consider only the communication cost of the protocol (for concreteness, in the blackboard model). Our new mechanisms (like the previous-best mechanisms) can be implemented simultaneously using only value queries. As such, these also imply results in the value query model, and the succinct representation model. In the succinct representation model, each player has a $v_i(\cdot)$ that can be represented by an explicit circuit of size at most poly(m). Because our main results are parameterized by k, we will further refer to the 2^k*-succinct representation model* as the case where each $v_i(\cdot)$ can be represented by an explicit circuit of size at most $2^{O(k)}$.

Additionally, we make the simplifying assumption that when considering the class of mechanisms that can be run in $2^{O(k)}$ communication/value queries/computation, all numbers are integers less than $2^{2^{O(k)}}$ (and therefore can be represented in $2^{O(k)}$ bits in a standard fashion). We make this assumption to avoid any strangeness with things like arithmetic, value queries representing arbitrary precision numbers, etc.

3 AN OPTIMAL MAXIMAL-IN-RANGE MECHANISM FOR GENERAL VALUATIONS

In this section, we will prove the upper bound in Theorem 1.4 by constructing an m/k-approximate MIR mechanism which uses $2^{O(k)}$ communication.

Definition 3.1 (Chunking Mechanism). Let $B^{(1)}, \dots, B^{(z)} \in [t]^M$ partition M into t *chunks* each. The allocation bank $\mathcal{A}$ contains every allocation where each bidder gets a set of the form $C(S, \ell) := \bigcup_{j \in S} B_j^{(\ell)}$ for some $\ell \in [z], S \subseteq [t]$, and the *chunking mechanism for* $B^{(1)}, \dots, B^{(z)}$ is MIR over $\mathcal{A}$.

The prior state-of-the-art for general valuations is simply a chunking mechanism for a single partition into k chunks of equal size [28] (i.e. a chunking mechanism with $z = 1$). They prove their mechanism guarantees an $m/\sqrt{k}$ approximation ratio for all monotone valuations (and this is tight – no chunking mechanism with $z = 1$ can guarantee better than $m/\sqrt{k}$).

Our only new idea is to instead consider a chunking mechanism for *multiple* carefully chosen partitions which satisfy the following property.

Definition 3.2 (r-Itemizing). A partition $B^{(\ell)}$ *itemizes* a set S if each chunk of $B^{(\ell)}$ contains at most one item of S. A list of partitions $B^{(1)}, \dots, B^{(z)}$ is r*-itemizing* if for any set S of size at most r, there exists $B^{(\ell)}$ which itemizes S.

Lemma 3.3. *For all $r = \Omega(\log\log(m))$ and some $z = 2^{\Theta(r)}$, there exists a list of partitions $B^{(1)}, \dots, B^{(z)} \in [r]^M$ into r chunks which is r-itemizing.*

Proof. Suppose we randomly sample the partitions such that $B^{(1)}, \dots, B^{(z)} \in [r]^M$ are independent and uniformly random. Then for a fixed set S of size r, the partition $B^{(\ell)}$ itemizes r w.p. $r!/r^r = 2^{-\Theta(r)}$. Therefore, by independence, no partition itemizes S w.p. at most $(1 - 2^{-\Theta(r)})^z = 2^{-2^{\Theta(r)}}$. By a union bound over the $\binom{m}{r} \le 2^{r\log(m)} \le 2^{2^{\Theta(r)}}$ sets of size r, $B^{(1)}, \dots, B^{(z)}$ is r-itemizing w.p. > 0. Thus, there exists a fixed list of partitions $B^{(1)}, \dots, B^{(z)}$ that is r-itemizing. ◻

Theorem 3.4. *Let $k = \Omega(\log(m))$ and let $z = 2^{\Theta(k)}$. Additionally, let $B^{(1)}, \dots, B^{(z)} \in [r]^M$ be a $(4k)$-itemizing list of partitions, which exists by Lemma 3.3. Then the chunking mechanism for $B^{(1)}, \dots, B^{(z)}$ is m/k-approximate and can be implemented using $2^{O(k)}$ communication.*

Moreover, the mechanism can be implemented simultaneously with $2^{O(k)}$ value queries, and in time $2^{O(k)}$ in the 2^k-succinct representation model.

Proof. The first step in our analysis for general valuations is similar to the analysis of [17] for subadditive valuations, which separately analyzes the bidders who receive many items versus few items in the optimal allocation.

Let $t = m/(2k)$. We will partition bidders into sets $N_0, N_1, \dots, N_t$ such that $i \in N_0$ if and only if $|A_i^*(\vec{v})| > 2k$, and for all $s > 0$, $\sum_{i \in N_s} |A_i^*(\vec{v})| \le 4k$. Observe that the condition on all N_s is possible because the bidders not in N_0 all get at most $2k$ items each.[7]

Let $\mathcal{A}$ be the allocation bank which defines the chunking mechanism and observe that $\mathcal{A}$ can allocate all items to a single bidder. Then since there are at most $m/(2k)$ bidders who get more than $2k$ items, $\mathrm{MIR}_{\mathcal{A}}(\vec{v}) \ge (2k/m)\,\mathrm{OPT}(\vec{v}, N_0)$.

Now, observe that any set S which is itemized by some $B^{(\ell)}$ is shattered by $\mathcal{A}$, as the chunking mechanism can assign any combination of the chunks (and hence any combination of the items) to the bidders. Thus, $\mathcal{A}$ shatters every set of size $4k$. Since each set of bidders N_s gets at most $4k$ items in $A^*(\vec{v})$, $\mathrm{MIR}_{\mathcal{A}}(\vec{v}) \ge \max_{s \in [t]} \mathrm{OPT}(\vec{v}, N_s) \ge (2k/m)\,\mathrm{OPT}(\vec{v}, N \setminus N_0)$.

Therefore, we get that for all valuations $v_1, \dots, v_n$,

$$\begin{aligned} \mathrm{MIR}_{\mathcal{A}}(\vec{v}) &\ge \max\left\{\frac{2k}{m}\,\mathrm{OPT}(\vec{v}, N_0),\ \frac{2k}{m}\,\mathrm{OPT}(\vec{v}, N \setminus N_0)\right\} \\ &\ge \frac{k}{m}\,\mathrm{OPT}(\vec{v})\ . \end{aligned}$$

Communication and Computation. Each bidder can only receive at most $2^{4k}z = 2^{\Theta(k)}$ possible sets in $\mathcal{A}$, so optimizing over $\mathcal{A}$ can be done with just $2^{\Theta(k)}$ simultaneous value queries per bidder. On the computation side, we will make use of the following lemma.

Lemma 3.5. *In the $2^{m'}$-succinct representation model with items $M' := [m']$ and bidders $N' := [n']$, a welfare-maximizing allocation can be found in time $2^{O(m')} \cdot n'$.*

Proof. For $T \subseteq N'$ and $S \subseteq M'$, define $v_T(S)$ to be the optimal welfare for bidders T and items S, and define $A_T(S)$ to be an optimal allocation of items S to bidders T. Suppose the functions $v_{[i]}$ and $A_{[i]}$ are known for some $i \in [n' - 1]$. Then $v_{[i+1]}$ and $A_{[i+1]}$ can be computed in $2^{O(m')}$ computation by brute forcing over all $2^{m'}$ sets $S \subseteq M$ and all of at most $2^{m'}$ allocations of S between $v_{[i]}$ and v_{i+1}. Hence, we can iteratively compute $v_{N'}$ and $A_{N'}$ in $2^{O(m')}n'$ time, and the optimal allocation is $A_{N'}(M')$. ◻

Observe that to run our mechanism, we only need to solve z welfare maximization problems over $4k$ "items," where we interpret a chunk as a single item. Therefore, by Lemma 3.5, the total computation needed is $2^{O(k)}nz = 2^{O(k)}$. ◻

Remark 3.6. Note that while we can run the above mechanism in polynomial time *given* an r-itemizing list of partitions, we do not know how to *explicitly* find such a list in polynomial time. Therefore, if we want an explicit mechanism, then we can only achieve an $m/\log(m)$ approximation in polynomial time w.h.p. by sampling a random list of partitions. Note that this is still stronger than a mechanism which achieves the desired approximation with constant probability/in expectation since not all truthful mechanisms can have their success probability amplified by repetition.

A similar statement holds true for the subadditive mechanism in the next section.

4 A MAXIMAL-IN-RANGE MECHANISM FOR SUBADDITIVE VALUATIONS

In this section, we will prove Theorem 1.3 by giving a $\sqrt{m/k}$-approximate MIR mechanism for subadditive valuations which uses $2^{O(k)}$ communication.

The prior state-of-the-art for subadditive valuations asks each bidder for their value for the entire set of items, and each set of

[7]That is, this partition can be created by first placing all bidders with $|A_i^*(\vec{v})| > 2k$ in N_0, and then greedily filling N_i with remaining bidders without exceeding the cap of $4k$. Because each bidder not in N_0 gets at most $2k$ items, each non-empty bidder set will have at least $2k$ items.

$O(k/\log(m/k))$ items [17]. Then we can either give one influential bidder every item (not too many influential bidders $\implies$ good approximation), or we can give every bidder their $k/\log(m/k)$ favorite items (subadditivity $\implies$ good approximation).

Our mechanism for subadditive valuations deviates from this approach of asking for smaller sets, and is instead more closely related to our approach for general mechanisms.

Definition 4.1 (Bucketing Mechanism). Let $B^{(1)}, \dots, B^{(z)} \in [t]^M$ partition M into t *buckets* each, and for each $\ell \in [z]$, $s \in [t]$, and $T \subseteq N$, let $\mathcal{A}^{(\ell)}_{s,T} \subseteq T^{B^{(\ell)}_s}$ be an allocation bank for items $B^{(\ell)}_s$ among bidders T. Also, let $A^{(i)}$ denote the allocation that awards all items to bidder i. Then the *bucketing mechanism for* $\{\mathcal{A}^{(\ell)}_{s,T} : \ell \in [z], s \in [t], T \subseteq N\}$ is MIR over the allocation bank

$$\mathcal{A} \quad := \quad \left(\bigcup_{i\in N}\{A^{(i)}\}\right) \cup \left(\bigcup_{\ell\in[z]}\bigcup_{P\in[t]^N}\bigtimes_{s\in[t]}\mathcal{A}^{(\ell)}_{s,P_s}\right).$$

In other words, $\mathcal{A}$ includes all allocations that award all items to the same bidder. All other allocations in $\mathcal{A}$ first choose a bucketing $\ell \in [z]$, then partitions bidders among buckets (with bidders P_s going to bucket s), and then finally chooses an allocation in $\mathcal{A}^{(\ell)}_{s,P_s}$ of items in bucket $B^{(\ell)}_s$ to bidders P_s. The bucketing mechanism for $\mathcal{A}$ is MIR over $\mathcal{A}$.

Definition 4.2 (Regular). A partition $B^{(\ell)} \in [t]^M$ is *regular* for a (possibly incomplete) partition of the items B into t buckets if for all $s \in [t]$ where $|B_s| = O(m/t)$, $|B^{(\ell)}_{s'} \cap B_s| = O(m/t^2)$ for all $s' \in [t]$. A list of partitions $B^{(1)}, \dots, B^{(z)} \in [t]^M$ is *regular* if for all B, some $B^{(\ell)}$ is regular for B.

Lemma 4.3. *For $t = O(\sqrt{m/\log(m)})$, there exists a regular list of partitions $B^{(1)}, \dots, B^{(m)} \in [t]^M$.*

Proof. Suppose we randomly sample the partitions such that $B^{(1)}, \dots, B^{(m)} \in [t]^M$ are independent and uniformly random. Then for a fixed (possibly incomplete) partition B into t buckets and any $s \in [t]$ where $|B_s| \le Cm/t$ for some constant C, $|B^{(\ell)}_{s'} \cap B_s|$ is stochastically dominated by $X \sim \mathrm{Binom}(Cm/t, 1/t)$. Thus, noting that $m/t^2 = \Omega(\log(m))$, we have

$$\begin{aligned}
&\Pr\left[B^{(1)}, \dots, B^{(m)} \text{ is not regular for } B\right] \\
&\qquad\le \left(t^2 \Pr[X = \omega(m/t^2)]\right)^m \\
&\qquad\le \left(mt^2 \Pr[X = 3Cm/t^2]\right)^m \\
&\qquad\le \left(mt^2 \binom{Cm/t}{3Cm/t^2}\left(\frac{1}{t}\right)^{3Cm/t^2}\right)^m \\
&\qquad\le \left(mt^2 \left(\frac{Cem/t}{3Cm/t^2}\right)^{3Cm/t^2}\left(\frac{1}{t}\right)^{3Cm/t^2}\right)^m \\
&\qquad= \left(mt^2\left(\frac{e}{3}\right)^{3Cm/t^2}\right)^m \quad\le\quad \frac{1}{m^m}.
\end{aligned}$$

Hence, by a union bound over $\le 2^m t^m < m^m$ (possibly incomplete) partitions of the items into t buckets, $B^{(1)}, \dots, B^{(m)}$ is regular w.p. > 0, and therefore there exists a fixed list of partitions $B^{(1)}, \dots, B^{(m)}$ which is regular. □

We now present a $\sqrt{m/k}$-approximate mechanism for subadditive valuations using $2^{O(k)}$ simultaneous value queries.

Definition 4.4 (Bucket-Shattering Mechanism). Let $k = \Omega(\log(m))$ and let $t = \sqrt{m/k}/2$. Additionally, let $B^{(1)}, \dots, B^{(m)} \in [t]^M$ be a regular list of partitions, which exists by Lemma 4.3. For each $\ell \in [z], s \in [t]$, fix a $\Theta(k)$-itemizing list of $2^{\Theta(k)}$ partitions for $B^{(\ell)}_s$ (which exists by Lemma 3.3), and for each $T \subseteq N$, let $\mathcal{A}^{(\ell)}_{s,T}$ be its chunking mechanism (i.e. the $\Theta(m/k)$-approximate MIR mechanism from Theorem 3.4 for this specific list of partitions). The *bucket-shattering mechanism for* k is the bucketing mechanism for this choice of $\mathcal{A}^{(\ell)}_{s,P_s}$.

Example 4.5. Suppose we have $m = 8$ and $n = 4$. Then one bucketing and chunking is

Buckets 1: $(\{1,2,3,4\}, \{5,6,7,8\})$

$$\begin{cases}\text{Chunks 1a: } (\{1,2\},\{3,4\}), & \text{Chunks 2a: } (\{1,3\},\{2,4\})\\ \text{Chunks 1b: } (\{5,6\},\{7,8\}), & \text{Chunks 2b: } (\{5,7\},\{6,8\})\end{cases}$$

Buckets 2: $(\{1,3,5,7\}, \{2,4,6,8\})$

$$\begin{cases}\text{Chunks 1a: } (\{1,3\},\{5,7\}), & \text{Chunks 2a: } (\{1,5\},\{3,7\})\\ \text{Chunks 1b: } (\{2,4\},\{6,8\}), & \text{Chunks 2b: } (\{2,6\},\{4,8\})\end{cases}$$

Then we could

- Choose bucketing 1
- Choose chunking 1a for bucket a and chunking 2b for the bucket b.
- Assign bidder 1 to the bucket a and bidders 2, 3, 4 to the bucket b.
- Give bidder 1 chunks $\{1,2\}$ and $\{3,4\}$ from bucket a.
- Give bidder 2 chunk $\{5,7\}$ and bidder 4 chunk $\{6,8\}$ from bucket b.

This results in the allocation $(\{1,2,3,4\}, \{5,7\}, \emptyset, \{6,8\})$. Any allocation resulting from a similar procedure would be in the allocation bank of the bucket-shattering mechanism.

On the other hand, the allocation $(\{1,2,5,6\}, \{3,4\}, \{7,8\}, \emptyset)$ is impossible, because while we could choose bucketing 1, and chunkings 1a for bucket a and 1b for bucket b, bidder 1 can only receive chunks from a single bucket. This is a crucial restriction that saves a large factor of communication.

Remark 4.6. One can think of the bucket-shattering mechanism as adding an additional layer of shattering to the mechanism for general valuations: we first shatter the bidders among the buckets, in the sense that any allocation of the bidders to buckets is possible. Then we run the general mechanism within each bucket, which shatters the items among the bidders for that bucket.

Note that once the general mechanism within each bucket is solved, no additional communication is needed to find the optimal bucketing. However, $2^{\Omega(n)}$ *computation* is needed (at least naively), which is too much when $n = \omega(k)$. This can be avoided by restricting consideration to only $\mathrm{poly}(m)$ random allocations of bidders to buckets instead of all t^n allocations. We first analyze the more elegant, computationally inefficient version.

Theorem 4.7. *For $k = \Omega(\log(m))$, the bucket-shattering mechanism for k (Definition 4.4) is $\sqrt{m/k}$-approximate for subadditive valuations and can be implemented using $2^{O(k)}$ communication.*

Moreover, the mechanism can be implemented simultaneously with $2^{O(k)}$ value queries.

Proof. We again first partition bidders into sets $N_0, N_1, \dots, N_t$ such that $i \in N_0$ if and only if $|A_i^*(\vec{v})| \geq 2\sqrt{mk}$, and for all $s \in [t]$, $\sum_{i \in N_s} |A_i^*(\vec{v})| \leq 4\sqrt{mk}$.

Observe that $\mathcal{A}$ can allocate all items to a single bidder, so since there are at most $\sqrt{m/k}/2$ bidders who get more than $2\sqrt{mk}$ items, $\mathrm{MIR}_{\mathcal{A}}(\vec{v}) \geq 2\sqrt{k/m}\,\mathrm{OPT}(\vec{v}, N_0)$.

Define $B_s^* := \bigcup_{i \in N_s} A_i^*(\vec{v})$, meaning that B_s^* is the items that bidders in bucket s get in an optimal allocation. By construction of $N_1, \dots, N_t$, $|B_s^*| \leq 4\sqrt{mk} = 2m/t$ for all $s \in [t]$. Then if we interpret $B^* := (B_1^*, \dots, B_t^*)$ as an incomplete partition into t buckets, there exists $\ell \in [m]$ such that $B^{(\ell)}$ is regular for B^*.

Now, define $B_s^{(\ell)}$ for $s \in [t+1, 2t]$ by $B_{s-t}^{(\ell)}$, and observe that by subadditivity of $v_1, \dots, v_n$,

$$\sum_{\Delta \in [t]} \sum_{s \in [t]} \sum_{s \in N_s} v_i(B_{s+\Delta}^{(\ell)} \cap A_i^*(\vec{v})) \geq \sum_{s \in [t]} \sum_{i \in N_s} v_i(A_i^*(\vec{v})) = \mathrm{OPT}(\vec{v}, N \setminus N_0) \ ,$$

because $\bigcup_{\Delta \in [t]} (B_{s+\Delta}^{(\ell)} \cap A_i^*(\vec{v})) = A_i^*(\vec{v})$ for any s, i. Hence,

$$\max_{\Delta \in [t]} \left\{ \sum_{s \in [t]} \sum_{i \in N_s} v_i(B_{s+\Delta}^{(\ell)} \cap A_i^*(\vec{v})) \right\} \geq 2\sqrt{\frac{k}{m}}\,\mathrm{OPT}(\vec{v}, N \setminus N_0) \ . \tag{1}$$

Now, observe that:

- $\bigtimes_{s \in [t]} \mathcal{A}_{s+\Delta, N_s}^{(\ell)} \subseteq \mathcal{A}$ for all $\Delta \in [t]$.
- $\mathcal{A}_{s+\Delta, N_s}^{(\ell)}$ shatters every subset of $B_s^{(\ell)}$ of size $O(k)$ (see proof of Theorem 3.4).
- $\sum_{i \in N_s} |B_{s+\Delta}^{(\ell)} \cap A_i^*(\vec{v})| = |B_{s+\Delta}^{(\ell)} \cap B_s^*| = O(k)$, since $B^{(\ell)}$ is regular for B^*.

By the above points, there exists $A \in \mathcal{A}$ such that $B_{s+\Delta}^{(\ell)} \cap A_i^*(\vec{v}) \subseteq A_i$ for all $i \in N$ and $\Delta \in [t]$. Therefore, by (1),

$$\mathrm{MIR}_{\mathcal{A}}(\vec{v}) \geq 2\sqrt{\frac{k}{m}}\,\mathrm{OPT}(\vec{v}, N \setminus N_0) \ ,$$

and thus,

$$\mathrm{MIR}_{\mathcal{A}}(\vec{v}) \geq \max\left\{ 2\sqrt{\frac{k}{m}}\,\mathrm{OPT}(\vec{v}, N_0),\ 2\sqrt{\frac{k}{m}}\,\mathrm{OPT}(\vec{v}, N \setminus N_0) \right\} \geq \sqrt{\frac{k}{m}}\,\mathrm{OPT}(\vec{v}) \ .$$

Communication. Observe that for each $\ell \in [m]$, $s \in [t]$, and $T \subseteq N$, $\mathcal{A}_{s,T}^{(\ell)} \subseteq \mathcal{A}_{s,N}^{(\ell)}$, and each bidder can only receive at most $2^{\Theta(k)}$ possible sets in $\mathcal{A}_{s,N}^{(\ell)}$, as it is a Chunking mechanism with $2^{\Theta(k)}$ partitions into $\Theta(k)$ chunks. Hence, each bidder can only receive at most $2^{\Theta(k)} mt = 2^{\Theta(k)}$ sets in $\mathcal{A}$, so optimizing over $\mathcal{A}$ can be done with just $2^{\Theta(k)}$ simultaneous value queries per bidder. □

4.1 Computational Efficiency

Naively, the bucket-shattering mechanism requires $2^{\Omega(n)}$ computation to implement, as there are exponentially many allocations of bidders to buckets. We can resolve this by randomly sampling polynomially-many allocations of bidders to buckets, and optimizing over this restricted subset instead. We define the mechanism below, but leave the proofs to the appendix of the full paper.

Definition 4.8 ($\mathcal{P}$-Bucketing Mechanism). Let $\mathcal{P} \subseteq [t]^N$ be a set of partitions of the bidders into t buckets, and for each $\ell \in [z]$, $s \in [t]$, and $T \subseteq N$, define $\mathcal{A}_{s,T}^{(\ell)}$ as in Definition 4.1. Similarly, let $A^{(i)}$ denote the allocation that awards all items to bidder i. Then the *$\mathcal{P}$-bucketing mechanism for* $\{\mathcal{A}_{s,T}^{(\ell)} : \ell \in [z], s \in [t], T \subseteq N\}$ is MIR over the allocation bank

$$\mathcal{A} := \left(\bigcup_{i \in N} \{A^{(i)}\} \right) \cup \left(\bigcup_{\ell \in [z]} \bigcup_{P \in \mathcal{P}} \bigtimes_{s \in [t]} \mathcal{A}_{s,P_s}^{(t)} \right) \ .$$

In other words, the $\mathcal{P}$-bucketing mechanism is the bucketing mechanism with a restricted range for the assignment of bidders to buckets.

Definition 4.9 (Balanced). Let $\vec{v} := (v_1, \dots, v_n)$, and let $N_1(\vec{v}) := \{i \in N : |A_i^*(\vec{v})| \leq m/t\}$. For any bucketing of the bidders $P \in [t]^N$ and $s \in [t]$, let $B_s^*(P) := \bigcup_{i \in P_s \cap N_1(\vec{v})} A_i^*(\vec{v})$, let $S_P(\vec{v}) := \{s \in [t] : |B_s^*(\vec{v}, P)| = O(m/t)\}$, and let $N_P(\vec{v}) := \bigcup_{s \in S_P(\vec{v})} (P_s \cap N_1)$. In other words, $N_P(\vec{v})$ is the set of bidders which belong to buckets that do not receive many items, when we restrict attention only to items awarded in OPT to bidders who do not individually receive many items (in OPT). Then a bucketing $P^{(\ell)} \in [t]^N$ is *balanced* for $\vec{v}$ if $\mathrm{OPT}(\vec{v}, N_P(\vec{v})) = \Theta(\mathrm{OPT}(\vec{v}, N_1(\vec{v})))$. A list of bucketings $P^{(1)}, \dots, P^{(z)} \in [t]^N$ is *balanced* if for all $\vec{v}$, some $P^{(\ell)}$ is balanced for $\vec{v}$.

Lemma 4.10. *For $y = \mathrm{poly}(m)$, there exists a balanced list of bucketings $P^{(1)}, \dots, P^{(y)} \in [t]^N$.*

Definition 4.11 (Efficient Bucket-Shattering Mechanism). For $y = \mathrm{poly}(m)$, let $P^{(1)}, \dots, P^{(y)}$ be a balanced list of bucketings, which exists by Lemma 4.10. Let $\mathcal{P}$ be the set of all $P^{(x)}$ and their shifts (e.g., $(P_{1+\Delta}^{(x)}, \dots, P_{t+\Delta}^{(x)})$ for all $\Delta \in [t]$). The *efficient bucket-shattering mechanism for k* is defined as the $\mathcal{P}$-bucketing mechanism for $\{\mathcal{A}_{s,T}^{(\ell)} : \ell \in [z], s \in [t], T \subseteq N\}$ as defined by the bucket-shattering mechanism for k.

Theorem 4.12. *For $k = \Omega(\log(m))$, the efficient bucket-shattering mechanism for k (Definition 4.11) is $\sqrt{m/k}$-approximate for subadditive valuations and can be implemented simultaneously with $2^{O(k)}$ value queries, and in time $2^{O(k)}$ in the 2^k-succinct representation model.*

5 COMMUNICATION LOWER BOUNDS

In this section, we will prove the lower bound in Theorem 1.4 and Theorem 1.2.

Our approach can be broken down into the same two parts as prior approaches: we first show that guarantees on the approximation ratio implies a rich allocation bank, then show that optimizing over the allocation bank requires lots of communication. While the first part primarily adapts existing results, the second part uses novel techniques to close the $m^{1/6}$ gap between the upper and lower bounds for submodular valuations.

5.1 Part 1: Approximation Implies Rich Allocation Bank

We will show that a rich allocation bank 2-shatters a large set of items. To do so, we make use of the following results.

Proposition 5.1 ([9], Theorem 1.5). *Suppose an allocation bank $\mathcal{A} \subseteq N^M$ does not 2-shatter any set of size d. Then*

$$|\mathcal{A}| \quad \leq \quad \sum_{i=0}^{d} \binom{m}{i}\binom{n}{2}^i \quad \leq \quad \left(\frac{emn^2}{d}\right)^d .$$

For an allocation bank $\mathcal{A}$, there exists $S \subseteq M$ such that $\mathcal{A}$ 2-shatters (S, N) and

$$|\mathcal{A}| \quad \leq \quad (mn^2)^{|S|} .$$

Proposition 5.2 ([7], Lemma 3.2). *If the MIR mechanism for $\mathcal{A}$ is $n/3$-approximate for additive valuations, then there exists $S \subseteq M$ such that $|\mathcal{A}|_S| \geq 2^{m/n}$.*

Lemma 5.3. *Let $k = \Omega(\log(m))$ and $n = 3\sqrt{m/k}$. Then if the MIR mechanism for $\mathcal{A}$ is $\sqrt{m/k}$-approximate for submodular valuations, there exists $S \subseteq M$ of size $\Omega(\sqrt{mk}/\log(m/k))$ such that $\mathcal{A}$ 2-shatters (S, N).*

Proof. Suppose that $\mathcal{A}$ does not 2-shatter any sufficiently large set. Then by Proposition 5.1, choosing the correct constants yields

$$\begin{aligned} |\mathcal{A}| \quad &\leq \quad \left(\frac{emn^2}{\Theta(\sqrt{mk}/\log(m/k))}\right)^{\Theta(\sqrt{mk}/\log(m/k))} \\ &= \quad O\left(\frac{m^2}{k^2}\right)^{\Theta(\sqrt{mk}/\log(m/k))} \quad < \quad 2^{\sqrt{mk}/3} . \end{aligned}$$

However, by Proposition 5.2, $|\mathcal{A}| \geq 2^{\sqrt{mk}/3}$, a contradiction. Thus, $\mathcal{A}$ must 2-shatter some set of size $\Omega(\sqrt{mk}/\log(m/k))$. ◻

Lemma 5.4. *Let $k = \Omega(\log(m))$ and $n \geq 2m/k$. Then if the MIR mechanism for $\mathcal{A}$ is m/k-approximate for all monotone valuations, there exists $S \subseteq M$ of size $\Omega(k)$ and $T \subseteq N$ of size 2 such that $\mathcal{A}$ shatters (S, T).*

Proof. Suppose for simplicity that $n = 2m/k$ is an integer and k is even. Let $\mathcal{B} \subseteq N^M$ be the collection of partitions such that each part has $k/2$ items. The valuations *induced* by $B \in \mathcal{B}$ are $v_i(R) := \mathbb{1}_{R \supseteq B_i}$. Consider all partitions $B \in N^M$ where $|B_i| = k/2$ for all $i \in N$, and consider the valuations $v_i(S) := \mathbb{1}_{S \supseteq B_i}$ induced by each such B.

Since the MIR mechanism for $\mathcal{A}$ is m/k-approximate, for every set of valuations induced by $B \in \mathcal{B}$, there exist two bidders who get value 1. Since there are only $\binom{n}{2} < m^2$ pairs of bidders, there exists a fixed pair of bidders $T = \{i, j\}$ such that the MIR mechanism for $\mathcal{A}$ gives bidder i items B_i and bidder j items B_j for a $1/m^2$ fraction of the $B \in \mathcal{B}$.

Now, partition $\mathcal{B}$ into $\{\mathcal{B}_R : R \subseteq M,\ |R| = k\}$ so that $B \in \mathcal{B}_S$ if $B_i \cup B_j = R$. It follows that there exists some $\mathcal{B}_R$ such that the MIR mechanism for $\mathcal{A}$ gives bidder i items B_i and bidder j items B_j for a $1/m^2$ fraction of the $B \in \mathcal{B}_R$.

Therefore, there exists $R \subseteq M$ of size k such that $|T^R \cap \mathcal{A}|_R| \geq |\mathcal{B}_R|_R|/m^2 = \binom{k}{k/2}/m^2 = 2^{\Omega(k)}$. Applying Sauer's Lemma to $T^R \cap \mathcal{A}|_R$, there exists S of size $\Omega(k)$ such that $\mathcal{A}$ shatters (S, T). ◻

Remark 5.5. We have demonstrated that good approximations in either the submodular or arbitrary monotone valuation setting imply a rich allocation bank in the sense that there is a large set of 2-shattered items. This unified view then allows us (in the next section) to derive, from 2-shattering, a useful structure for a lower bound, which then implies lower bounds for both submodular and arbitrary monotone valuations. However, we note that enough structure is already recovered by Lemma 5.4 to give a direct lower bound for monotone valuations, e.g., by noting that it requires $2^{\Omega(k)}$ communication to maximize the sum of two monotone functions over the k items shattered between 2 bidders [36].

5.2 Part 2a: Rich Allocation Bank Contains Structure

For simplicity, let $S := [s]$ and $T := [t]$. Our aim now is to find a structure within the 2-shattered (S, T) which will be suitable for embedding SetDisjointness. Let T_j for $j \in S$ be the pair of bidders that item j can go to. In other words, $\bigtimes_{j \in S} T_j$ is the allocation bank witnessing the 2-shattering. Our end goal will be to prove the following:

Proposition 5.6. *For $t = O(s/\log(s))$ and $z = 2^{\Theta(s/t)}$, there exist $B^{(1)}, \dots, B^{(z)} \in \bigtimes_{j \in S} T_j$ such that*

- *There exists $V \subseteq T$ such that any $f : V \to [z]$ is constant if and only if $B_1^{(f(1))}, \dots, B_t^{(f(t))}$ are pairwise disjoint.*
- *$|B_i^{(1)}| = \cdots = |B_i^{(z)}|$ for all $i \in T$.*

The next subsection makes clear why this structure allows us to embed SetDisjointness. At a high level, the first bullet ensures that we can encode valuations for the bidders using sets $X_1, \dots, X_t$ such that a certain welfare can be attained if and only if $\bigcap_{i \in V} X_i \neq \emptyset$. The first bullet alone suffices to give a lower bound for XOS valuations, and captures the main idea behind the lower bound. The second bullet introduces highly non-trivial technical challenges, but is required to extend the lower bound to submodular valuations. We only provide the proof of the claim in its entirety, but will comment on when simplifications can be made by not satisfying the second bullet.

Proof of Proposition 5.6. The first step is (with a slight abuse of notation) to consider $\mathcal{A} \subseteq \bigtimes_{j \in S} T_j$ such that there exist $a_1, \dots, a_t$ such that for all $A \in \mathcal{A}$, $|A_i| = a_i$ for all i. In other words, $\mathcal{A}$ satisfies the second bullet. Since there are $(s+1)^t \leq 2^{\varepsilon s}$ possible values for $a_1, \dots, a_t$ for any constant $\varepsilon > 0$, there exists such $\mathcal{A}$ such that $|\mathcal{A}| \geq 2^{-\varepsilon s}|\bigtimes_{j \in S} T_j| = 2^{(1-\varepsilon)s}$. Thus, we can achieve the second bullet without losing too many allocations from our bank.

Step I: A Helpful Interpretation. Intuitively, we may think of the first bullet as saying that there exist z allocations such that there is no way to combine two or more allocations into another valid allocation. In other words, the function f chooses which allocation each bidder in V receives a set from, and f will always cause some item to be allocated to two bidders unless f assigns each bidder to the same allocation.

Because each item can only go to one of two bidders, only items that go between bidders who are assigned different allocations under f are able to cause the allocation induced by f to be invalid. As such, it is helpful to interpret the 2-shattering structure as a graph, and reason about the viability of mixing allocations across cuts in the graph.

Definition 5.7. For $V \subseteq T$ and $E \subseteq S$ such that $\bigcup_{j \in E} T_j \subseteq V$, define $G(V, E)$ to be the graph where each bidder denotes a vertex and each item $j \in E$ denotes an edge between bidders T_j.

Our eventual goal is to sample z allocations independently and uniformly from $\mathcal{A}$, and show that the probability that the first bullet is satisfied is nonzero. Had we instead considered $\bigtimes_{j \in S} T_j$ rather than $\mathcal{A}$ (i.e., disregarded the second bullet), then observe that sampling uniformly from $\bigtimes_{j \in S} T_j$ assigns each item j to a bidder in T_j independently and uniformly, and hence the probability that independently sampled allocations can be combined is exponentially small in the number of edges crossing the cut separating the bidders who are assigned to different allocations. Then, it roughly suffices to take a subgraph with min-cut $\Omega(k)$, and apply some clever union bounds.

However, when we instead consider $\mathcal{A}$, restricting to a $2^{-\varepsilon s}$ fraction of the original 2-shattered structure introduces correlations between items when we sample uniformly from $\mathcal{A}$, and hence a more complex argument is needed. Importantly, the analysis focuses on the following quantity.

Definition 5.8. For a graph $G := G(V, E)$, edges $C \subseteq E$ (typically, we will take C to be the edges across a cut), and allocation bank $\mathcal{A} \subseteq \bigtimes_{j \in E} T_j$, define $p(G, C, \mathcal{A})$ to be the maximum fraction of allocations that remain after fixing the allocation of the items in C. In other words,

$$p(G, C, \mathcal{A}) \quad := \quad \max_{A \in \mathcal{A}|_C} \frac{|\{B \in \mathcal{A} : A = B|_C\}|}{|\mathcal{A}|} \ .$$

Observe that $p(G, C, \bigtimes_{j \in E} T_j) = 2^{-|C|}$. To eventually apply a similar probabilistic argument as we would apply to a 2-shattered structure, we will try to find a subgraph $H := H(V, E)$ of $G(T, S)$ such that the min-cut of H is $\Omega(k)$ and $p(H, C, \mathcal{A}|_E) \approx 2^{-|C|}$ for all cuts C.

Step II: Finding a Good Subgraph. From here on, we will denote cuts by the set of their edges. Additionally, for a graph G, let $\gamma_r(G)$ denote its r-way min-cut, i.e., the smallest cut which partitions the bidders into r non-empty parts.

Lemma 5.9. *Let $\mathcal{A} \subseteq \bigtimes_{j \in S} T_j$ and $|\mathcal{A}| \geq 2^{(1-\varepsilon)s}$ for some constant $\varepsilon > 0$. Then there exists $V \subseteq T$ where $|V| \geq 2$ and $E \subseteq S$ such that $H := H(V, E)$ satisfies $\gamma_r(H) \geq \varepsilon(r-1)s/t$ for all r, and for any edges C across a cut in H, $p(H, C) \leq 2^{-(1-2\varepsilon)s}$.*

Proof. Initialize $S' := S$, and while there exists an r-way cut C across some connected component of $G := G(T, S')$ such that $p(G, C, \mathcal{A}|_{S'}) > 2^{-\varepsilon(r-1)s/t} 2^{-(1-2\varepsilon)|C|}$, set $S' := S' \setminus C$. By definition of $p(G, C, \mathcal{A}|_{S'})$, we have $|\mathcal{A}|_{S' \setminus C}| \geq p(G, C, \mathcal{A}|_{S'}) |\mathcal{A}|_{S'}|$.

Suppose for contradiction that the process results in a disconnected graph after removing the $r_1, \ldots, r_\ell$-way cuts $C_1, \ldots, C_\ell$. Then since removing an r-way cut increases the number of connected components by $r-1$, we have $\sum_{i \in [\ell]} (r_i - 1) < t$. Further, there are no loops in G, so $S' = \emptyset$ and $\sum_{i \in [\ell]} |C_i| = s$. Therefore,

$$\begin{aligned} 1 &= |\mathcal{A}|_\emptyset| \\ &> 2^{-\varepsilon s/t \sum_{i \in [\ell]} (r_i - 1)} 2^{-(1-2\varepsilon) \sum_{i \in [\ell]} |C_i|} |\mathcal{A}| \\ &> 2^{-\varepsilon s} 2^{-(1-2\varepsilon)s} 2^{(1-\varepsilon)s} \\ &= 1 \ , \end{aligned}$$

a contradiction. Thus, the process terminates with a connected component $H := H(V, E)$ (which is a subgraph of the original $G(T, S)$) such that $|V| \geq 2$ and every r-way cut C across H satisfies $p(H, C, \mathcal{A}|_E) \leq 2^{-\varepsilon(r-1)s/t} 2^{-(1-2\varepsilon)|C|} \leq 2^{-(1-2\varepsilon)|C|}$. Additionally, $p(H, C, \mathcal{A}|_E) \geq 2^{-|C|}$, so the first exponential ensures that $|C| \geq \varepsilon(r-1)s/t$ for all C. Thus, H satisfies the desired properties. □

Step III: A Probabilistic Construction. Before we proceed to the final probabilistic construction, we prove a property of the subgraph H promised by Lemma 5.9 that makes clear why we need a structure like H to be contained in $G(T, S)$.

Lemma 5.10. *Let $H := H(V, E)$ satisfy the condition in Lemma 5.9 for $\varepsilon = 1/50$. Let $f : V \to [r]$ map bidders to parts of an r-way cut C. Then for independent $B^{(1)}, \ldots, B^{(r)} \sim \mathcal{A}|_E$, the probability that $\{B^{(f(i))}(i) : i \in V\}$ are pairwise disjoint is at most $2^{-|C|/3}$.*

Proof. Let $V_1, \ldots, V_r$ be the partition of V where $f(i) = \ell$ for any bidder $i \in V_\ell$. Let C_ℓ for $\ell \in [r]$ be the 2-way cut between V_ℓ and $V \setminus V_\ell$.

Observe that there are $3^{|C|}$ r-tuples of $B^{(1)}|_{C_1}, \ldots, B^{(r)}|_{C_r}$ such that $B^{(f(i))}(i)$ for $i \in V$ are pairwise disjoint, because for each $j \in C$ which can go to either V_ℓ or $V_{\ell'}$, we need either $B^{(\ell)}$ to allocate j to $V_{\ell'}$, or $B^{(\ell')}$ to allocate j to V_ℓ, or both.

On the other hand, by definition of $p(H, C_\ell, \mathcal{A}|_E)$ and the upper bound on it given by Lemma 5.9, the probability of any such r-tuple being sampled is at most

$$\begin{aligned} \prod_{\ell \in [r]} p(H, C_\ell, \mathcal{A}|_E) &\leq \prod_{\ell \in [r]} 2^{-(1-2\varepsilon)|C_\ell|} \\ &= 2^{-2(1-2\varepsilon)|C|} \ , \end{aligned}$$

so the probability that the $B^{(f(i))}(i)$ are pairwise disjoint is at most $3^{|C|} 2^{-2(1-2\varepsilon)|C|} \leq 2^{-|C|/3}$. □

In other words, H has the property that for every r-way cut in H, the probability that r independently sampled allocations from $\mathcal{A}|_E$ can be combined into a feasible allocation is exponentially small in the size of the cut, which is exactly the property we wished to replicate from the 2-shattered structure! We can now proceed to the probabilistic construction. First, we give a graph theoretic result that will be needed for the union bound (proof in the appendix of the full paper)

Lemma 5.11. *For any graph $G = (V, E)$ (possibly with parallel edges, but no loops) and $c \in \mathbb{Z}_{\geq 1}$, the number of r-way cuts with at most $c\gamma_2(G)$ edges is at most $|V|^{4c}$.*

Lemma 5.12. *Let $t = O(s/\log(s))$ and let $z = 2^{\Theta(s/t)}$. Additionally, let $B^{(1)}, \dots, B^{(z)} \sim \mathcal{A}|_E$. Then w.p. > 0, a function $f : V \to [z]$ is constant if and only if $\{B_i^{(f(i))} : i \in V\}$ are pairwise disjoint.*

Proof. Since $B^{(1)}, \dots, B^{(z)}$ are valid allocations, the forward direction holds trivially.

For the reverse direction, fix any function $f : V \to [z]$ which takes on a fixed set of $r \geq 2$ values, and let the C be the r-way cut induced by f. Then by Lemma 5.10, the probability that $\{B^{(f(i))}(i) : i \in V\}$ are pairwise disjoint is at most $2^{-|C|/3}$.

By Lemma 5.11, the number of distinct cuts C with $|C| \leq c\varepsilon s/t \leq c\gamma_2(H)$ is at most t^{4c}. Thus, by a union bound over all possible values of r, all possible fixed subsets of r values in $[z]$, all possible values of c, and all possible cuts satisfying these parameters, the probability that the desired condition is not satisfied is at most

$$\begin{aligned}
\sum_{r=2}^{t} \left(\binom{z}{r} \sum_{c=r-1}^{s} t^{4(c+1)} 2^{-(c\varepsilon s/t)/3} \right) & \\
\leq \quad & \sum_{r=2}^{t} \left(z^r \sum_{c=r-1}^{s} 2^{8c(\log t - \varepsilon s/(6t))} \right) \\
\leq \quad & \sum_{r=2}^{t} \left(2^{\Theta(rs/t)} \sum_{c=r-1}^{s} 2^{-\Theta(cs/t)} \right) \\
\leq \quad & \sum_{r=2}^{t} 2^{-\Theta(rs/t)} \quad < \quad 1 \ ,
\end{aligned}$$

where the third line follows because $\log t \leq \log s = O(s/t)$. □

Extending the allocations $B^{(1)}, \dots, B^{(z)} \in \mathcal{A}|_E$ to any allocations in $\mathcal{A}$ which agree on the allocation of the items E completes the proof of Proposition 5.6.

5.3 Part 2b: Structure Yields Set Disjointness Embedding

We now show that the MIR mechanism for $\mathcal{A}$ can solve SetDisjointness. The bidder valuations we will use are the following.

Definition 5.13. A *mild-desires bidder for $\mathcal{F} \subseteq 2^M$*, where all $F \in \mathcal{F}$ are the same size a, has the valuation function

$$v(G) \quad = \quad \begin{cases} 2|G| & |G| < a \\ 2|G| - \mathbb{1}_{G \notin \mathcal{F}} & |G| = a \\ 2a & |G| > a \end{cases} .$$

We say that such a bidder is *satisfied* if their allocation gives them value $2a$, which occurs when they receive items $F \in \mathcal{F}$, or any $a+1$ items. Mild-desires bidders have submodular valuations [36].

For each bidder $i \in V$, associate an input set $X_i \subseteq [z]$. For the bidders $i \notin V$, associate the set $X_i = [z]$. Fix $B^{(1)}, \dots, B^{(z)} \in \mathcal{A}$ promised by Proposition 5.6. By bullet two, we can let each bidder $i \in T$ be mild-desires for $\{B_i^{(\ell)} : \ell \in X_i\}$.

Note that if we were only interested in a communication lower bound for MIR mechanisms for XOS valuations, then we would not have needed bullet two to hold, as we could instead let the valuation for bidder i be the rank function of the downward-closed set family defined by $\{B_i^{(\ell)} : \ell \in X_i\}$ (which is an XOS function).

Then to solve SetDisjointness with the MIR mechanism for $\mathcal{A}$, we only need to show that $\bigcap_{i \in V} X_i \neq \emptyset$ if and only if the optimal welfare over $\mathcal{A}$ is $2s$. The forward direction holds because the allocation $B^{(x)}$ for some $x \in \bigcap_{i \in V} X_i$ satisfies every bidder.

For the reverse direction, by Proposition 5.6, we know that any collection of sets desired by the bidders V are pairwise disjoint (i.e., results in a valid allocation) if and only if those collection of sets belong to the same allocation. To achieve $2s$ welfare, we need to satisfy every bidder (since $\sum_{i \in T} a_i = s$), and this can only be done if every bidder receives a desired set (if we satisfy bidder i by giving them $a_i + 1$ items, some bidder j can only receive at most $a_j - 1$ items and cannot be satisfied). Hence, if the optimal welfare is $2s$, it must be the case that $\bigcap_{i \in V} X_i \neq \emptyset$.

Thus, the MIR mechanism for $\mathcal{A}$ is capable of solving SetDisjointness over a universe of size z. Since the communication complexity of SetDisjointness is $\Omega(z)$ [34], we conclude that maximizing the welfare of submodular valuations over $\mathcal{A}$ requires $2^{\Omega(s/t)}$ communication.[8]

Wrapping Up. By Lemmas 5.3 and 5.4,

- An m/k-approximate MIR mechanism for general valuations must use $2^{\Omega(k)}$ communication.
- A $\sqrt{m/k}$-approximate MIR mechanism for submodular valuations must use $2^{\Omega(k/\log(m/k))}$ communication. Therefore, a $\sqrt{m/(k\log(m/k))}$-approximate MIR mechanism for submodular valuations must use $2^{\Omega(k)}$ communication.

Thus, we have $\mathsf{MIR}_{\mathsf{Gen}}(m,k) = \Omega(m/k)$ and $\mathsf{MIR}_{\mathsf{SubMod}}(m,k) = \Omega(\sqrt{m/(k\log(m/k))})$.

6 CONCLUSION

For all amounts of communication, we improve both upper and lower bounds for approximation guarantees of MIR mechanisms over submodular, XOS, subadditive, and general valuations. This resolves the approximation guarantees of MIR mechanisms for general valuations up to a constant factor, and for submodular, etc. valuations up to a $\Theta(\sqrt{\log m})$ factor. In addition, the mechanisms which witness the upper bounds use only value queries, demonstrating that using arbitrary communication instead of the far more restrictive regime of value queries does not give a mechanism much power. Even so, there are a few open questions for future work.

Closing the Logarithmic Gap for Submodular Valuations. Although we were able to significantly improve existing lower bounds for submodular, XOS, and subadditive valuations (reducing the gap from $\tilde{\Theta}((m/k)^{1/6})$ to $\Theta(\sqrt{\log(m/k)})$), we still started from the same 2-shattering argument of [9] in order to embed a hard communication game.

If one conjectures that our lower bound can be slightly improved (as we do), this unfortunately cannot follow after a 2-shattering argument – there exist mechanisms which are $\sqrt{m/k}$-approximate (see appendix of full paper) and which do not 2-shatter any pair (S,T) such that $|S|/|T| = \Omega(k)$.

If instead one conjectures that our MIR mechanisms can be slightly improved, then the must be neither implementable $2^{O(k)}$

[8]Note that because the randomized communication complexity of SetDisjointness is also $\Omega(b)$, even randomized protocols which maximize the welfare of monotone submodular valuations over $\mathcal{A}$ must use $2^{\Omega(s/t)}$ communication in expectation.

value queries (by [22]), nor implementable with $2^{O(k)}$ simultaneous communication (see appendix of full paper). Such a mechanism would be fundamentally different than all prior MIR mechanisms, which can be implemented with simultaneous value queries.

Beyond MIR mechanisms. The major open problem is to understand communication lower bounds that hold for all deterministic truthful mechanisms, and not just MIR mechanisms. There is significantly less progress in this direction – only [18] for dominant strategy truthful mechanisms, and [3] for two-player mechanisms.

REFERENCES

[1] Sepehr Assadi. 2020. Combinatorial Auctions Do Need Modest Interaction. *ACM Trans. Economics and Comput.* 8, 1, 3:1–3:23. https://doi.org/10.1145/3381521

[2] Sepehr Assadi, Thomas Kesselheim, and Sahil Singla. 2021. Improved Truthful Mechanisms for Subadditive Combinatorial Auctions: Breaking the Logarithmic Barrier. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, Dániel Marx (Ed.). SIAM, 653–661. https://doi.org/10.1137/1.9781611976465.40

[3] Sepehr Assadi, Hrishikesh Khandeparkar, Raghuvansh R. Saxena, and S. Matthew Weinberg. 2020. Separating the communication complexity of truthful and non-truthful combinatorial auctions. In *Proccedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy (Eds.). ACM, 1073–1085. https://doi.org/10.1145/3357713.3384267

[4] Sepehr Assadi and Sahil Singla. 2019. Exponentially Improved Truthful Combinatorial Auctions with Submodular Bidders. In *Proceedings of the Sixtieth Annual IEEE Foundations of Computer Science (FOCS).*

[5] Liad Blumrosen and Noam Nisan. 2005. On the computational power of iterative auctions. In *Proceedings 6th ACM Conference on Electronic Commerce (EC-2005), Vancouver, BC, Canada, June 5-8, 2005*, John Riedl, Michael J. Kearns, and Michael K. Reiter (Eds.). ACM, 29–43. https://doi.org/10.1145/1064009.1064013

[6] Mark Braverman, Jieming Mao, and S. Matthew Weinberg. 2018. On Simultaneous Two-player Combinatorial Auctions. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018.* 2256–2273. https://doi.org/10.1137/1.9781611975031.146

[7] David Buchfuhrer, Shaddin Dughmi, Hu Fu, Robert Kleinberg, Elchanan Mossel, Christos H. Papadimitriou, Michael Schapira, Yaron Singer, and Christopher Umans. 2010. Inapproximability for VCG-Based Combinatorial Auctions. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA).*

[8] Edward H. Clarke. 1971. Multipart Pricing of Public Goods. *Public Choice* 11, 1 (1971), 17–33.

[9] Amit Daniely, Michael Schapira, and Gal Shahaf. 2015. Inapproximability of Truthful Mechanisms via Generalizations of the VC Dimension. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015.* 401–408. https://doi.org/10.1145/2746539.2746597

[10] Shahar Dobzinski. 2007. Two Randomized Mechanisms for Combinatorial Auctions. In *Proceedings of the 10th International Workshop on Approximation and the 11th International Workshop on Randomization, and Combinatorial Optimization. Algorithms and Techniques.* 89–103.

[11] Shahar Dobzinski. 2011. An Impossibility Result for Truthful Combinatorial Auctions with Submodular Valuations. In *Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC).*

[12] Shahar Dobzinski. 2016. Breaking the Logarithmic Barrier for Truthful Combinatorial Auctions with Submodular Bidders. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing* (Cambridge, MA, USA) *(STOC 2016).* ACM, New York, NY, USA, 940–948. https://doi.org/10.1145/2897518.2897569

[13] Shahar Dobzinski. 2016. Computational Efficiency Requires Simple Taxation. In *FOCS.*

[14] Shahar Dobzinski and Noam Nisan. 2011. Limitations of VCG-Based Mechanisms. *Comb.* 31, 4 (2011), 379–396. https://doi.org/10.1007/s00493-011-2528-4

[15] Shahar Dobzinski, Noam Nisan, and Sigal Oren. 2014. Economic Efficiency Requires Interaction. In *the 46th annual ACM symposium on Theory of computing (STOC).*

[16] Shahar Dobzinski, Noam Nisan, and Michael Schapira. 2006. Truthful randomized mechanisms for combinatorial auctions. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing, Seattle, WA, USA, May 21-23, 2006*, Jon M. Kleinberg (Ed.). ACM, 644–652. https://doi.org/10.1145/1132516.1132607

[17] Shahar Dobzinski, Noam Nisan, and Michael Schapira. 2010. Approximation Algorithms for Combinatorial Auctions with Complement-Free Bidders. *Math. Oper. Res.* 35, 1 (2010), 1–13. https://doi.org/10.1287/moor.1090.0436 https://www.cs.yale.edu/homes/schapira/cfjournal.pdf.

[18] Shahar Dobzinski, Shiri Ron, and Jan Vondrák. 2022. On the hardness of dominant strategy mechanism design. In *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, Stefano Leonardi and Anupam Gupta (Eds.). ACM, 690–703. https://doi.org/10.1145/3519935.3520013

[19] Shahar Dobzinski and Jan Vondrák. 2012. From query complexity to computational complexity. In *Proceedings of the 44th Symposium on Theory of Computing (STOC).*

[20] Shahar Dobzinski and Jan Vondrak. 2012. The Computational Complexity of Truthfulness in Combinatorial Auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC).*

[21] Shahar Dobzinski and Jan Vondrák. 2013. Communication Complexity of Combinatorial Auctions with Submodular Valuations. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, Sanjeev Khanna (Ed.). SIAM, 1205–1215. https://doi.org/10.1137/1.9781611973105.87 https://theory.stanford.edu/~jvondrak/data/submod-comm.pdf.

[22] Shahar Dobzinski and Jan Vondrák. 2016. Impossibility Results for Truthful Combinatorial Auctions with Submodular Valuations. *J. ACM* 63, 1 (2016), 5:1–5:19. https://doi.org/10.1145/2786754

[23] Shaddin Dughmi and Jan Vondrak. 2011. Limitations of Randomized Mechanisms for Combinatorial Auctions. In *52nd Annual Symposium on Foundations of Computer Science (FOCS).*

[24] Tomer Ezra, Michal Feldman, Eric Neyman, Inbal Talgam-Cohen, and S. Matthew Weinberg. 2019. Settling the Communication Complexity of Combinatorial Auctions with Two Subadditive Buyers. In *the 60th Annual IEEE Symposium on Foundations of Computer Science (FOCS).* https://arxiv.org/pdf/1811.09871.pdf.

[25] Uriel Feige. 2009. On Maximizing Welfare When Utility Functions Are Subadditive. *SIAM J. Comput.* 39, 1 (2009), 122–142. https://doi.org/10.1137/070680977 http://www.wisdom.weizmann.ac.il/~feige/mypapers/subadditive.pdf.

[26] Uriel Feige and Jan Vondrák. 2010. The Submodular Welfare Problem with Demand Queries. *Theory of Computing* 6, 1 (2010), 247–290. https://doi.org/10.4086/toc.2010.v006a011 https://theory.stanford.edu/~jvondrak/data/submod-improve-ToC.pdf.

[27] Theodore Groves. 1973. Incentives in Teams. *Econometrica* 41, 4 (1973), 617–631.

[28] Ron Holzman, Noa E. Kfir-Dahav, Dov Monderer, and Moshe Tennenholtz. 2004. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior* 47, 1 (2004), 104–123. https://doi.org/10.1016/S0899-8256(03)00184-2

[29] Piotr Krysta and Berthold Vöcking. 2012. Online mechanism design (randomized rounding on the fly). In *Automata, Languages, and Programming.* Springer, 636–647.

[30] Ron Lavi, Ahuva Mu'alem, and Noam Nisan. 2003. Towards a Characterization of Truthful Combinatorial Auctions. In *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003, Cambridge, MA, USA, Proceedings.* IEEE Computer Society, 574–583. https://doi.org/10.1109/SFCS.2003.1238230

[31] Benny Lehmann, Daniel Lehmann, and Noam Nisan. 2001. Combinatorial Auctions with Decreasing Marginal Utilities. In *the 3rd Annual ACM Conference on Electronic Commerce (EC).*

[32] Daniel Lehmann, Liadan O'Callaghan, and Yoav Shoham. 2002. Truth revelation in approximately efficient combinatorial auctions. *J. ACM* 49, 5 (2002), 577–602. https://doi.org/10.1145/585265.585266

[33] Vahab S. Mirrokni, Michael Schapira, and Jan Vondrák. 2008. Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. In *Proceedings 9th ACM Conference on Electronic Commerce (EC-2008), Chicago, IL, USA, June 8-12, 2008.* 70–77. https://doi.org/10.1145/1386790.1386805

[34] Noam Nisan. 2002. The Communication Complexity of Approximate Set Packing and Covering. In *Automata, Languages and Programming, 29th International Colloquium, ICALP 2002, Malaga, Spain, July 8-13, 2002, Proceedings (Lecture Notes in Computer Science, Vol. 2380)*, Peter Widmayer, Francisco Triguero Ruiz, Rafael Morales Bueno, Matthew Hennessy, Stephan J. Eidenbenz, and Ricardo Conejo (Eds.). Springer, 868–875. https://doi.org/10.1007/3-540-45465-9_74

[35] Noam Nisan and Amir Ronen. 2007. Computationally Feasible VCG Mechanisms. *J. Artif. Intell. Res.* 29 (2007), 19–47. https://doi.org/10.1613/jair.2046

[36] Noam Nisan and Ilya Segal. 2006. The communication requirements of efficient allocations and supporting prices. *J. Economic Theory* 129, 1 (2006), 192–224. https://doi.org/10.1016/j.jet.2004.10.007

[37] William Vickrey. 1961. Counterspeculations, Auctions, and Competitive Sealed Tenders. *Journal of Finance* 16, 1 (1961), 8–37.

[38] Jan Vondrák. 2008. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008.* 67–74. https://doi.org/10.1145/1374376.1374389

Received 13-NOV-2023; accepted 2024-02-11