Mechanism Design For Covering Problems

by

Hadi Minooei

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Combinatorics and Optimization

Waterloo, Ontario, Canada, 2014

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Algorithmic mechanism design deals with efficiently-computable algorithmic constructions in the presence of strategic players who hold the inputs to the problem and may misreport their input if doing so benefits them. Algorithmic mechanism design finds applications in a variety of internet settings such as resource allocation, facility location and e-commerce, such as sponsored search auctions.

There is an extensive amount of work in algorithmic mechanism design on packing problems such as single-item auctions, multi-unit auctions and combinatorial auctions. But, surprisingly, covering problems, also called *procurement auctions*, have almost been completely unexplored, especially in the multidimensional setting.

In this thesis, we systematically investigate *multidimensional covering mechanism-design* problems, wherein there are m items that need to be covered and n players who provide covering objects, with each player i having a private cost for the covering objects he provides. A feasible solution to the covering problem is a collection of covering objects (obtained from the various players) that together cover all items.

Two widely considered objectives in mechanism design are: (i) *cost-minimization* (CM) which aims to minimize the total cost incurred by the players and the mechanism designer; and (ii) *payment minimization* (PayM), which aims to minimize the payment to players.

Covering mechanism design problems turn out to behave quite differently from packing mechanism design problems. In particular, various techniques utilized successfully for packing problems do not perform well for covering mechanism design problems, and this necessitates new approaches and solution concepts. In this thesis we devise various techniques for handling covering mechanism design problems, which yield a variety of results for both the CM and PayM objectives.

In our investigation of the CM objective, we focus on two representative covering problems: uncapacitated facility location (UFL) and vertex cover. For multi-dimensional UFL, we give a black-box method to transform any *Lagrangian-multiplier-preserving* ρ-approximation algorithm for UFL into a truthful-in-expectation, ρ-approximation mechanism. This yields the first result for multi-dimensional UFL, namely a truthful-in-expectation

2-approximation mechanism. For multi-dimensional VCP (Multi-VCP), we develop a *decomposition method* that reduces the mechanism-design problem into the simpler task of constructing *threshold mechanisms*, which are a restricted class of truthful mechanisms, for simpler (in terms of graph structure or problem dimension) instances of Multi-VCP. By suitably designing the decomposition and the threshold mechanisms it uses as building blocks, we obtain truthful mechanisms with approximation ratios (n is the number of nodes): (1) $O(r^2 \log n)$ for r-dimensional VCP; and (2) $O(r \log n)$ for r-dimensional VCP on any proper minor-closed family of graphs (which improves to $O(\log n)$ if no two neighbors of a node belong to the same player). These are the first truthful mechanisms for Multi-VCP with non-trivial approximation guarantees.

For the PayM objective, we work in the oft-used Bayesian setting, where players' types are drawn from an underlying distribution and may be correlated, and the goal is to minimize the expected total payment made by the mechanism. We consider the problem of designing incentive compatible, ex-post individually rational (IR) mechanisms for covering problems in the above model. The standard notion of incentive compatibility (IC) in such settings is Bayesian incentive compatibility (BIC), but this notion is over-reliant on having precise knowledge of the underlying distribution, which makes it a rather non-robust notion. We formulate a notion of IC that we call *robust Bayesian IC* (robust BIC) that is substantially more robust than BIC, and develop black-box reductions from robust BIC-mechanism design to algorithm design. This black-box reduction applies to single-dimensional settings even when we only have an LP-relative *approximation algorithm* for the algorithmic problem. We obtain near-optimal mechanisms for various covering settings including single- and multi-item procurement auctions, various single-dimensional covering problems, and multidimensional facility location problems.

Finally, we study the notion of *frugality*, which considers the PayM objective but in a worst-case setting, where one does not have prior information about the players' types. We show that some of our mechanisms developed for the CM objective are *also* good with respect to certain oft-used frugality benchmarks proposed in the literature. We also introduce an alternate benchmark for frugality, which more directly reflects the goal that the mechanism's payment be close to the best possible payment, and obtain some preliminary results with respect to this benchmark.

Acknowledgements

At first, I would like to thank my supervisor, Chaitanya Swamy, without whom I couldn't have achieved this stage of my life. His brilliant ideas and invaluable comments has made a lot of contributions to every part of this thesis. He has spent a lot of time to help me in many areas, such as teaching courses, writing papers, proposal, and thesis, and also preparing presentations and talks. He was available for advice or academic help whenever I needed it. Thank you!

I thank Joseph Cheriyan, Jochen Konemann, Shreyas Sundaram and Jose Correa for accepting to serve on my defence committee and also their insightful comments and suggestions on this thesis.

I had the pleasure of having friends and colleagues like Marcel de Carli Silva, Cristiane Maria Sato, Konstantinos Georgiou, Linda Farczadi, Nishad Kothari and Zachary Friggstad whose feedback and inspiring discussions have helped me in different aspects of my academic life.

I have some incredible friends who have always been there for me during my stay at Waterloo. I would like to thank Tohid Ahadifard, Sohrab Yarbakht, Mojtaba Rajabi, Hamid Molawian, Azad Qazi Zadeh, Samad Bazargan, Hamid Faramarzi, Amin Yazdani, Payman Charkh Zarrin and Bahador Biglari.

Lastly, I am grateful to my parents Majid and Khadijeh, and also my siblings and their spouses Maryam, Mojtaba, Amin and Najmeh for their support and encouragement during all these years.

Dedication

This thesis is dedicated to US consuls in Canada. Despite their lack of background in the subject, they have always shown interest in my research and more importantly have asked me various good questions about it's applications. They have given me the confidence that no matter what I think about my research, it looks like rocket-science!

Table of Contents

List of Tables

List of Figures

Chapter 1

Overview and Preliminaries

1.1 Introduction

Traditional algorithm design deals with problems in the standard Computer Science input-output model, where the input is readily available to the algorithm, and the goal is to efficiently compute an output satisfying one or more objectives. Such algorithm design problems constitute the bread-and-butter of Computer Science, and have been extensively studied; see, for example, Algorithm Design by Kleinberg and Tardos [30] and Approximation algorithms by Vazirani [49]. In recent years, with the advent of the Internet and especially E-commerce platforms such as Google Ad-auctions, eBay, etc., we see various settings, where the inputs to the underlying algorithmic problem are held by *strategic* entities or players, who are motivated by their individual self-interests. This has naturally led to the interaction of algorithm design with *mechanism design*, which is the field of economics that traditionally studies such strategic settings using game theory to model strategic behaviour, resulting in the emergence of the highly-active research area of *algorithmic mechanism design* lying at the intersection of Computer Science and Economics.

Algorithmic mechanism design (AMD) deals with efficiently-computable algorithmic constructions in the presence of strategic players who hold the input to the problem, and may misreport their input if doing so benefits them. Thus, one seeks to suitably incentivize the players to truthfully report their private inputs. In order to achieve this task, a *mechanism* is a protocol that specifies both an algorithm and a pricing or payment scheme that can be used to incentivize players to reveal their true inputs by suitably affecting their utilities. A mechanism is said to be *truthful*, if each player maximizes his utility by revealing his true input regardless of the other players' declarations. The challenge is to design the algorithm and the relevant pricing or payment schemes such that they hold desired properties from both the mechanism-design perspective, e.g. truthfulness, and algorithm-design perspective, e.g. computational efficiency.

Motivated by resource-allocation problems, there is an extensive amount of literature in AMD that deals with packing problems such as single-item auctions, multi-unit auctions and combinatorial auctions. These problems have been well studied under two widely considered objectives in mechanism design: *social-welfare maximization* (SWM), where the goal is to maximize the total value received by the players, and *revenue maximization*, where the goal is to maximize the total revenue of the mechanism. Furthermore, these problems have been investigated both from the standpoint of designing efficient truthful mechanisms with provable performance guarantees (e.g. see [17, 18, 20, 26, 32, 16, 21]) and from the perspective of proving lower bounds on the performance guarantees achievable by efficient truthful mechanisms (e.g. see [19, 39, 45, 31, 22]). In contrast, surprisingly, covering problems, also called *procurement auctions* or *reverse auctions*, have been almost completely unexplored, especially in the multidimensional settings (where dimensionality is, roughly speaking, a measure of the complexity of the players' private information; see Section 1.2 for a precise definition).

In this thesis, we systematically investigate *covering mechanism-design* problems. In a covering problem there are m items that need to be covered and n players who provide covering objects, with each player i having a private cost for the covering objects he provides. A feasible solution to the covering problem is a collection of covering objects (obtained from the various players) that together cover all items.

In the context of covering problems, the aforementioned objectives of social-welfare- and revenue- maximization translate respectively to the following: (i) *cost minimization* (CM), (CM) which aims to minimize the total cost incurred by the players and the mechanism designer; and (ii) *payment minimization* (PayM), which aims to minimize the payment to players. Our goal in AMD (as mentioned earlier) is to devise mechanisms for covering mechanism design problems that satisfy various desirable mechanism- and algorithm-design criteria, such as truthfulness, computational efficiency, and near-optimality for the underlying objective(s).

Covering mechanism design problems turn out to behave quite differently from packing mechanism design problems. In particular, various techniques utilized successfully for packing problems do not perform well for covering mechanism design problems, and this necessitates new approaches and solution concepts. In this thesis we devise various techniques for handling covering mechanism design problems, which yield a variety of results for both the CM and PayM objectives.

1.2 Preliminaries

In order to describe our contributions in a more meaningful way, and also lay some groundwork and notation for the rest of this thesis, we describe some basic concepts from algorithm and mechanism design. In subsequent chapters, we will supplement this with more problem-specific information as and when needed.

Covering mechanism-design problems. In a *multidimensional covering mechanism-design problem*, we have m items that need to be covered, and n players who provide covering objects. Each player i provides a set $\mathcal{T}_i$ of covering objects where each covering object $v \in \mathcal{T}_i$ covers a subset of items. All this information is public knowledge. We use $[k]$ to denote the set $\{1, \ldots, k\}$. Each player i has a *private cost* (or type) vector $c_i = \{c_{i,v}\}_{v \in \mathcal{T}_i}$, where $c_{i,v}$ is the cost he incurs for providing object $v \in \mathcal{T}_i$; for $T \subseteq \mathcal{T}_i$, we use $c_i(T)$ to denote $\sum_{v \in T} c_{i,v}$. A solution or allocation selects a subset $T_i \subseteq \mathcal{T}_i$ for each player i, denoting that i provides the objects in T_i. Given this solution, each player i incurs the private cost $c_i(T_i)$. Also, the mechanism designer incurs a publicly-known cost $pub(T_1, \ldots, T_n)$. Note that we can encode any feasibility constraints in the covering problem by simply setting $pub(a) = \infty$ if a is not a feasible allocation.

Cost minimization problem. A problem that we will often encounter is the *cost-minimization* (CM) problem, where the goal is to minimize the total cost $\sum_i c_i(T_i) + pub(T_1, \ldots, T_n)$ incurred. Observe that if we view the mechanism designer also as a player, then the CM problem is equivalent to maximizing the *social welfare*, which is the total value received by the players, and is given by $\sum_i -c_i(T_i) - pub(T_1, \ldots, T_n)$.

Let C_i denote the set of all possible types of player i, and $C = \prod_{i=1}^n C_i$. Let $\Omega := \{(T_1, \ldots, T_n) : pub(T_1, \ldots, T_n) < \infty\}$ be the (finite) set of all possible feasible allocations. For a tuple $x = (x_1, \ldots, x_n)$, we use x_{-i} to denote $(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n)$. Similarly, let $C_{-i} = \prod_{j \neq i} C_j$. For an allocation $\omega = (T_1, \ldots, T_n)$, we sometimes use ω_i to denote T_i, $c_i(\omega)$ to denote $c_i(\omega_i) = c_i(T_i)$, and $pub(\omega)$ to denote $pub(T_1, \ldots, T_n)$. We make the mild assumption that $pub(\omega') \leq pub(\omega)$ if $\omega_i \subseteq \omega'_i$ for all i; so in particular, if ω is feasible, then adding covering objects to the ω_is preserves feasibility.

A (direct revelation) *mechanism* $M = (\mathcal{A}, p_1, \ldots, p_n)$ for a covering problem consists of an allocation algorithm $\mathcal{A} : C \mapsto \Omega$ and a payment function $p_i : C \mapsto \mathbb{R}$ for each player i. Each player i reports a cost function c_i (that might be different from his true cost function). The mechanism computes the allocation $\mathcal{A}(c) = (T_1, \ldots, T_n) = \omega \in \Omega$, and pays $p_i(c)$ to each player i. The *utility* $u_i(c_i, c_{-i}; \overline{c}_i)$ that player i derives when he reports c_i and the others report c_{-i} is $p_i(c) - \overline{c}_i(\omega_i)$ where $\overline{c}_i$ is his true cost function, and each player i aims to maximize his own utility.

Definition 1.2.1 *We refer to* $\max_i |\mathcal{T}_i|$ *as the dimension of a covering problem. Thus, for a* single-dimensional *problem, each player i's cost can be specified as* $c_i(\omega) = c_i \alpha_{i,\omega}$*, where* $c_i \in \mathbb{R}_+$ *is his private type and* $\alpha_{i,\omega} = 1$ *if* $\omega_i \neq \emptyset$ *and 0 otherwise.*

Note that, $|\mathcal{T}_i|$ is an upper bound on the *type-domain dimension* of player i which is defined as the lowest dimension of an affine subspace of $\mathbb{R}^{|\Omega|}$ containing the set $\{(c_i(\omega))_{\omega \in \Omega} | c_i \in C_i\}$, and the maximum type-domain dimension among players is the dimension of the covering problem. But the above simpler definition holds true for most of general settings which are our focus in this thesis. Observe that in single-dimensional covering problems, by Definition 1.2.1, we can assume player i owns exactly one covering object i and incurs cost c_i for providing it.

A desirable property for a mechanism to satisfy is *truthfulness*, wherein every player i maximizes his utility by reporting his true cost function. Truthfulness is also (more precisely) called *dominant strategy incentive compatible (DSIC)* since it is a dominant strategy for each player to report his cost function truthfully. This explains one of the main benefits of designing truthful mechanisms: since truth-telling is the best strategy for each player regardless of the other players' strategies, it is straightforward for him to play according to his best interests; in particular, he does not need to perform any expensive computations to figure out his best strategy (if one exists). All our mechanisms will also satisfy the natural property of *individual rationality* (IR), which means that every player has nonnegative utility if he reports his true cost.

Definition 1.2.2 *A mechanism $M = \big(\mathcal{A}, \{p_i\}\big)$ is truthful (DSIC) if for every player i, every $c_{-i} \in C_{-i}$, and every $\bar{c}_i, c_i \in C_i$, we have $u_i(\bar{c}_i, c_{-i}; \bar{c}_i) \geq u_i(c_i, c_{-i}; \bar{c}_i)$. M is IR if for every i, every $\bar{c}_i \in C_i$ and every $c_{-i} \in C_{-i}$, we have $u_i(\bar{c}_i, c_{-i}; \bar{c}_i) \geq 0$.*

To ensure that truthfulness and IR are compatible, we consider *monopoly-free* settings: for every player i, there is a feasible allocation ω (i.e., $pub(\omega) < \infty$) with $\omega_i = \emptyset$. (Otherwise, if there is no such allocation, then i needs to be paid at least $\min_{v \in \mathcal{T}_i} c_{i,v}$ for IR, so he can lie and increase his utility arbitrarily.)

For a *randomized mechanism* M, where $\mathcal{A}$ or the p_i's are randomized, we say that M is *truthful in expectation* if each player i maximizes his expected utility by reporting his true cost. We now say that M is IR if for *every* coin toss of the mechanism, the utility of each player is nonnegative upon bidding truthfully. There are other weaker notions of incentive compatibility (IC) that are used in certain settings, e.g. Bayesian settings, where there is an underlying distribution on the players' types; we will define these in Chapter 3 when we consider such settings.

Since the CM problem is often *NP*-hard, our goal is to design a mechanism $M = \big(\mathcal{A}, \{p_i\}\big)$ that is truthful (or truthful in expectation), and where $\mathcal{A}$ is a ρ-approximation algorithm; that is, for every input c, the solution $\omega = \mathcal{A}(c)$ satisfies $\sum_i c_i(\omega) + pub(\omega) \leq \rho \cdot \min_{\omega' \in \Omega}\big(\sum_i c_i(\omega') + pub(\omega')\big)$. We call such a mechanism a *truthful, ρ-approximation mechanism.*

A well-known truthful mechanism is the VCG mechanism, which is optimal for the CM problem (i.e., achieves a 1-approximation). Moreover, in monopoly-free settings, the mechanism is IR and has the additional desirable property that the mechanism only makes positive payments to players from whom covering objects are procured.

Definition 1.2.3 *The VCG mechanism $M = \big(\mathcal{A}, \{p_i\}\big)$ for a covering problem is defined as follows: for any $c \in C$ we have $\mathcal{A}(c) = \omega^* := \text{argmin}_{\omega\in\Omega}\big(\sum_i c_i(\omega) + pub(\omega)\big)$ and the payment to each player i is $p_i(c) = \min_{\omega'\in\Omega:\omega'_i=\emptyset}(\sum_{j\neq i} c_j(\omega') + pub(\omega')) - (\sum_{j\neq i} c_j(\omega^*) + pub(\omega^*))$.*

Theorem 1.2.4 *VCG is a truthful and IR mechanism that only makes positive payments to players from whom covering objects are procured.*

Proof : Fix a player i, his true cost $\bar{c}_i \in C_i$ and reported cost of other players $c_{-i} \in C_{-i}$.

Truthfulness requires that for any given $c_i \in C_i$ we have $u_i(\bar{c}_i, c_{-i}; \bar{c}_i) \geq u_i(c_i, c_{-i}; \bar{c}_i)$ or if $\omega^* := \text{argmin}_{\omega\in\Omega}\big(\sum_{j\neq i} c_j(\omega) + \bar{c}_i(\omega) + pub(\omega)\big)$ and $\hat{\omega} := \text{argmin}_{\omega\in\Omega}\big(\sum_{j\neq i} c_j(\omega) + c_i(\omega) + pub(\omega)\big)$ then it requires $\min_{\omega'\in\Omega:\omega'_i=\emptyset}(\sum_{j\neq i} c_j(\omega') + pub(\omega')) - (\sum_{j\neq i} c_j(\omega^*) + pub(\omega^*)) - \bar{c}_i(\omega^*) \geq \min_{\omega'\in\Omega:\omega'_i=\emptyset}(\sum_{j\neq i} c_j(\omega') + pub(\omega')) - (\sum_{j\neq i} c_j(\hat{\omega}) + pub(\hat{\omega})) - \bar{c}_i(\hat{\omega})$ which is equivalent to $(\sum_{j\neq i} c_j(\hat{\omega}) + pub(\hat{\omega})) + \bar{c}_i(\hat{\omega}) \geq (\sum_{j\neq i} c_j(\omega^*) + pub(\omega^*)) + \bar{c}_i(\omega^*)$ and this is true according to the definition of ω^*.

IR property is straightforward from the definition of $p_i(c)$. Note that if $\omega^*_i = \emptyset$ then also $\omega^* = \text{argmin}_{\omega'\in\Omega:\omega'_i=\emptyset}(\sum_{j\neq i} c_j(\omega') + pub(\omega'))$ which leads to $p_i(c) = 0$ and proves the last desired property of VCG. ■

The following theorem gives a necessary and sometimes sufficient condition for when an algorithm $\mathcal{A}$ is *implementable*, that is, admits suitable payment functions $\{p_i\}$ such that $\big(\mathcal{A}, \{p_i\}\big)$ is a truthful mechanism. Before stating the theorem, we need to define *weak monotonicity*.

Definition 1.2.5 *An algorithm $\mathcal{A}$ is said to satisfy* **weak monotonicity (WMON)** *if for all i, and all c_{-i}, we have that if $\mathcal{A}(c_i, c_{-i}) = \omega$, $\mathcal{A}(c'_i, c_{-i}) = \omega'$, then $c_i(\omega) - c_i(\omega') \leq c'_i(\omega) - c'_i(\omega')$.*

It is easy to see that for a single-dimensional covering problem WMON is equivalent to the following simpler condition: say that $\mathcal{A}$ is *monotone* (MON) if for all i, all $c_i, c'_i \in C_i,\ c_i \leq c'_i$, and all $c_{-i} \in C_{-i}$, if $\mathcal{A}(c_i, c_{-i}) = \omega,\ \mathcal{A}(c'_i, c_{-i}) = \omega'$ then $\omega'_i \subseteq \omega_i$.

Theorem 1.2.6 (Theorems 9.29 and 9.36 in [43]) *If a mechanism $\big(\mathcal{A}, \{p_i\}\big)$ is truthful, then $\mathcal{A}$ satisfies WMON. Conversely, if the problem is single-dimensional, or if C_i is convex for all i, then every WMON algorithm $\mathcal{A}$ is implementable.*

In single-dimensional settings, given a MON algorithm $\mathcal{A}$, one can specify precisely the unique payments that when combined with it yield a truthful mechanism.

Definition 1.2.7 *Let $\mathcal{A}$ be a MON algorithm for an instance of single-dimensional covering problem. Given players' cost vector c, for each player i, its critical value is defined as $b_i(\mathcal{A}, c_{-i}) := \sup\{c'_i \in C_i | \mathcal{A}(c'_i, c_{-i})_i \neq \emptyset\}$ whenever the RHS is finite.*

Theorem 1.2.8 (Theorem 9.36 in [43]) *Let $\mathcal{A}$ be a MON algorithm such that $b_i(\mathcal{A}, c_{-i})$ is well defined for all i, all $c \in C$. Consider the mechanism $\mathcal{M} = (\mathcal{A}, p_1, ..., p_n)$ where $p_i(c) = b_i(\mathcal{A}, c_{-i})$ if player i wins, and 0 otherwise. This mechanism $\mathcal{M}$ is truthful, IR, and makes zero payment to a player whose covering object is not chosen. Moreover, if $b_i(\mathcal{A}, c_{-i}) = \inf\{c'_i \in C_i | \mathcal{A}(c'_i, c_{-i}) = \emptyset\}$, then these are the unique payments that yield a truthful, IR mechanism that makes zero payments to a player whose covering object is not chosen.*

Note that if a mechanism for a single-dimensional covering problem has a bounded approximation ratio, then by the monopoly-free nature of the single-dimensional setting, the critical values (Definition 1.2.7) are always defined. This fact is implicitly used in the subsequent chapters.

1.3 Summary of Our Contributions

In Chapter 2, we initiate a study of multidimensional *covering mechanism-design* problems from a cost-minimization (CM) perspective. We consider two representative covering problems, namely vertex cover (VCP) and uncapacitated facility location (UFL), and devise polytime, truthful, approximation mechanisms for these problems. For multidimensional UFL we present a *black-box reduction from truthful mechanism design to algorithm design*. Such reductions from mechanism- to algorithm- design are highly sought after, but quite rare. Whereas some such reductions are known for multidimensional packing problems, e.g., [32, 20, 4], this is the *first* such black-box reduction for a multidimensional covering problem, and it leads to the first result for multidimensional UFL, namely, a truthful-in-expectation, 2-approximation mechanism.

We devise two main techniques for multidimensional vertex cover problem (Multi-VCP). We introduce a simple class of truthful mechanisms called *threshold mechanisms* (Section 2.4.1), and show that despite their restrictions, threshold mechanisms can achieve non-trivial approximation guarantees. We next develop a *decomposition method* for Multi-VCP (Section 2.4.2) that provides a general way of reducing the mechanism-design problem for Multi-VCP into simpler—either in terms of graph structure, or problem dimension—mechanism-design problems by using threshold mechanisms as building blocks. By leveraging the decomposition method along with threshold mechanisms, we obtain various truthful, approximation mechanisms for Multi-VCP, which yield the *first* truthful mechanisms for multidimensional vertex cover with non-trivial approximation guarantees. We believe that these techniques will also find use in other mechanism-design problems.

In Chapter 3, we consider the payment-minimization (PayM) objective. We investigate covering problems in the *Bayesian setting*, where players' types are drawn from an underlying distribution and may be *correlated* and we seek to design incentive-compatible (IC) and individually rational (IR) mechanisms with minimum expected total payment. We develop black-box reductions from mechanism design to algorithm design whose application yields a variety of optimal and near-optimal mechanisms. As we elaborate in Chapter 3, covering problems turn out to behave quite differently in certain respects from packing problems, which necessitates new approaches (and solution concepts).

The standard solution concept in Bayesian settings is *Bayesian incentive compatibility* (BIC). However, this turns out to be a rather *non-robust* concept in that it is overly-reliant on having precise knowledge of the underlying distribution. We formulate a notion of incentive compatibility that we call *robust Bayesian IC* (robust BIC) that on the one hand is substantially more robust than BIC, and on the other is flexible enough that it allows one to obtain various polytime near-optimal mechanisms satisfying this notion. A *robust-(BIC, IR)* mechanism (see Section 3.2) ensures that truthful participation in the mechanism is in the best interest of every player even *after* when the other players' (randomly-chosen) types are revealed to him; thus, such a mechanism retains its desirable (IC and IR) properties for a wide variety of distributions, including those having the same support as the actual distribution making robust BIC a significantly more robust notion.

We show that for a variety of settings, one can reduce the robust-(BIC, IR) payment-minimization (PayM) mechanism-design problem to the algorithmic cost-minimization (CM) problem. We emphasize that our definition of additive types (see Section 3.2) should not be confused with, and is more general than, additive valuations in combinatorial auctions (CAs).

Our reduction yields near-optimal mechanisms for a variety of covering settings such as (a) various single-dimensional covering problems including single-item procurement auctions (b) multi-item procurement auctions; and (c) multidimensional facility location.

In Chapter 4, we investigate payment-minimization from a worst-case perspective, when we do not have any prior distribution on the players' types, by focusing on the vertex cover problem. This has been called the *frugality objective* in the literature to distinguish it from the PayM objective (which usually implies a Bayesian setup). We obtain various results regarding frugal truthful mechanism design for vertex cover. We show that some of the mechanisms that we devise for Multi-VCP in Chapter 2 for the CM problem also enjoy good frugality properties with respect to certain benchmarks used in the literature. Thus, we obtain polytime mechanisms that are simultaneously good with respect to both the approximation and frugality measures. Finally, we consider an alternate frugality benchmark to address some of the limitations of existing benchmarks, and obtain some preliminary results with respect to this benchmark.

1.4 Related Work

As mentioned earlier, although packing mechanism design problems have been studied from various perspectives in the AMD literature, there is a general lack of work on covering mechanism design problems, especially multidimensional covering CM problems (a recent result of [20] is an exception). In fact, to our knowledge, the only multidimensional problem with a covering flavour that has been studied in the AMD literature is the makespan-minimization problem on unrelated machines [42, 33, 3], which is not an CM problem.

For the PayM objective discussed extensively in Chapter 3, whereas the analogous revenue-maximization problem for packing domains, such as combinatorial auctions (CAs), has been extensively studied in the algorithmic mechanism design (AMD) literature, both in the case of independent and correlated (even interdependent) player-types (see, e.g., [6, 7, 5, 1, 9, 26, 17, 44, 5, 47] and the references therein), surprisingly, there are almost no results on the payment-minimization problem in the AMD literature (see however [7]). The economics literature does contain various general results that apply to both covering and packing problems. However much of this work focuses on characterizing special cases; see, e.g., [51]. An exception is the work of Crémer and McLean [12, 13], which shows that under certain conditions, one can devise a Bayesian-incentive-compatible (BIC) mechanism whose expected total payment is exactly equal to the expected cost incurred by the players, albeit one where players may incur negative utility under certain type-profile realizations. Our work on the PayM objective was inspired by the work of Dobzinski et al. [17], who address similar questions for the revenue-maximization objective. However, as noted earlier, fundamental differences between packing and covering problems necessitate different approaches. We elaborate on the benchmarks proposed in this body of work in Chapter 4.

The frugality objective considered in Chapter 4 is perhaps the only objective pertaining to covering problems that has been actively investigated. In particular, single-dimensional covering problems have been well studied from the perspective of frugality. Starting with the work of Archer and Tardos [2], various benchmarks for frugality have been proposed and investigated for various problems including VCP, k-edge-disjoint paths, spanning tree, s-t cut; see [28, 23, 29, 10] and the references therein.

1.5 Outline of the Thesis

The rest of this thesis is organized as follows.

In Chapter 2 we develop various techniques to design efficient, truthful and IR mechanisms for multi-dimensional uncapacitated facility location and vertex cover problems, from a cost-minimization standpoint. A preliminary version of the results of this chapter appeared as [37].

In Chapter 3 we investigate payment-minimization objective for covering problems in Bayesian setting. A preliminary version of the results of this chapter appeared as [38].

In Chapter 4 we study mechanism design for covering problems from a frugality viewpoint.

In Chapter 5 we briefly list some open questions regarding the study of multidimensional covering mechanism design problems.

Chapters 2, 3 and 4 are mostly independent of each other. Chapter 3 uses a lemma proved in Section 2.3, and Chapter 4 explores two classes of mechanisms for Multi-VCP introduced in Section 2.4.1.

Chapter 2

Cost-Minimization for Multidimensional Covering Problems

2.1 Introduction

In this chapter, we initiate a study of multidimensional *covering mechanism-design* problems from a cost-minimization (CM) perspective. Recall the definition from Chapter 1: there are m items that need to be covered and n players who provide covering objects, with each player i having a private cost for the covering objects he provides. The goal is to select (or buy) a suitable set of covering objects from each player so that their union covers all the items, and the total covering cost incurred is minimized. This *cost-minimization* (CM) problem is equivalent to the *social-welfare maximization* (SWM) (where the social welfare is $-$ (total cost incurred by the players and the mechanism designer)), so ignoring computational efficiency, the classical VCG mechanism [50, 11, 24] yields a truthful mechanism that always returns an optimal solution. However, the CM problem is often *NP*-hard, so we seek to design a *polytime* truthful mechanism where the underlying algorithm returns a near-optimal solution to the CM problem.

Although multidimensional packing mechanism-design problems have received much attention in the AMD literature, multidimensional covering CM problems are conspicuous by their absence in the literature. For example, the packing SWM problem of combinatorial auctions has been studied (in various flavors) in numerous works both from the viewpoint of designing polytime truthful, approximation mechanisms [18, 32, 16, 21], and from the perspective of proving lower bounds on the capabilities of computationally- (or query-) efficient truthful mechanisms [31, 22, 19]. In contrast, the lack of study of multidimensional covering CM problems is aptly summarized by the blank table entry for results on truthful approximations for procurement auctions in Fig. 11.2 in [43] (see "Related work" below). As stated in Chapter 1, to our knowledge, the only multidimensional problem with a covering flavor that has been studied in the AMD literature is the makespan-minimization problem on unrelated machines [42, 33, 3], which is not an SWM problem.

2.1.1 Summary of Results and Techniques

We study two representative multidimensional covering problems, namely (metric) *uncapacitated facility location (UFL)*, and *vertex cover* (VCP), and develop various techniques to devise polytime, truthful, approximation mechanisms for these problems.

In UFL, there is a set of clients who need to be served, and a set of players. Each player provides a (known) set of facilities for serving clients and the opening costs for these facilities is private information. Each client has to be assigned to one facility for service and doing so incurs a public client-assignment cost. The goal is to open a subset of facilities so as to minimize the total facility-opening and client-assignment costs.

In VCP, the edges of a graph need to be covered and each player provides a subset of nodes of the graph. The goal is to choose a minimum-cost vertex cover of the graph. (See detailed definitions in Section 2.2.)

For multidimensional UFL (Section 2.3), we present a *black-box reduction from truthful mechanism design to algorithm design*. We show that any ρ-approximation algorithm for UFL satisfying an additional *Lagrangian-multiplier-preserving* (LMP) property (that indeed holds for various algorithms) can be converted in a black-box fashion to a truthful-in-expectation ρ-approximation mechanism (Theorem 2.3.1). This is the *first* such black-box reduction for a multidimensional covering problem, and it leads to the first result for multidimensional UFL, namely, a truthful-in-expectation, 2-approximation mechanism.

Our result builds upon the convex-decomposition technique in [32]. Lavi and Swamy [32] primarily focus on packing problems, but remark that their convex-decomposition idea also yields results for *single-dimensional* covering problems, and leave open the problem of obtaining results for multidimensional covering problems. Our result for UFL identifies an interesting property under which a ρ-approximation algorithm for a covering problem can be transformed into a truthful, ρ-approximation mechanism in the multidimensional setting.

In Section 2.4, we consider multidimensional VCP. Although, algorithmically, VCP is one of the simplest covering problems, it becomes a surprisingly challenging mechanism-design problem in the *multidimensional* mechanism-design setting, and, in fact, seems significantly more difficult than multidimensional UFL. This is in stark contrast with the single-dimensional setting, where each player owns a single node. Before detailing our results and techniques, we mention some of the difficulties encountered. We use Multi-VCP to distinguish the multidimensional mechanism-design problem from the algorithmic problem.

For *single-dimensional* problems, a simple monotonicity condition characterizes the *implementability* of an algorithm, that is, whether it can be combined with suitable payments to obtain a truthful mechanism (see Theorem 1.2.8). This condition allows for ample flexibility and various algorithm-design techniques can be leveraged to design monotone algorithms for both covering and packing problems (see, e.g., [4, 32]). For single-dimensional VCP, many of the known 2-approximation algorithms for the algorithmic problem (based on LP-rounding, primal-dual methods, or combinatorial methods) are either already monotone, or can be modified in simple ways so that they become monotone, and thereby yield truthful 2-approximation mechanisms [14]. However, the underlying algorithm-design techniques fail to yield algorithms satisfying *weak monotonicity* (WMON)—a necessary condition for implementability (see Theorem 1.2.6)—even for the simplest multidimensional setting, namely, 2-dimensional VCP, where *every player owns at most two nodes*. We show this for various LP-rounding methods in Section 2.5, and for primal-dual algorithms in Section 2.6.

Furthermore, various techniques that have been devised for designing polytime truthful mechanisms for multidimensional packing problems (such as combinatorial auctions) do not seem to be helpful for Multi-VCP. For instance, the well-known technique of constructing a *maximal-in-range*, or more generally, a *maximal-in-distributional-range* (MIDR) mechanism—fix some subset of outcomes and return the best outcome in this set—does not work for Multi-VCP [20] (and more generally, for multidimensional covering problems). More precisely, any algorithm for Multi-VCP whose range is a proper subset of the collection of minimal vertex covers, cannot have bounded approximation ratio (see footnote 1 in Section 2.4). This also rules out the convex-decomposition technique of [32], which we exploit for multidimensional UFL, because, as noted in [32], this yields an MIDR mechanism.

Thus, we need to develop new techniques to attack Multi-VCP (and multidimensional covering problems in general). We devise two main techniques for Multi-VCP. We introduce a simple class of truthful mechanisms called *threshold mechanisms* (Section 2.4.1), and show that despite their restrictions, threshold mechanisms can achieve non-trivial approximation guarantees. We next develop a *decomposition method* for Multi-VCP (Section 2.4.2) that provides a general way of reducing the mechanism-design problem for Multi-VCP into simpler—either in terms of graph structure, or problem dimension—mechanism-design problems by using threshold mechanisms as building blocks. We believe that these techniques will also find use in other mechanism-design problems.

By leveraging the decomposition method along with threshold mechanisms, we obtain various truthful, approximation mechanisms for Multi-VCP, which yield the *first* truthful mechanisms for multidimensional vertex cover with non-trivial approximation guarantees. Let n be the number of nodes. Our decomposition method shows that any instance of r-dimensional VCP can be broken up into $O(r^2 \log n)$ instances of *single-dimensional VCP*; this in turn leads to a truthful, $O(r^2 \log n)$-approximation mechanism for r-dimensional VCP (Theorem 2.4.9). In particular, for any fixed r, we obtain an $O(\log n)$-approximation for any graph. We give another decomposition method that yields an improved truthful, $O(r \log n)$-approximation mechanism (Theorem 2.4.11) for any proper minor-closed family of graphs (such as planar graphs). This guarantee improves to $O(\log n)$ for any proper minor-closed family, when no two neighbors of a node belong to the same player.

2.1.2 Related Work

As mentioned earlier, there is little prior work on the CM problem for multidimensional covering problems. Dughmi and Roughgarden [20] give a general technique to convert an FPTAS for an SWM problem to a truthful-in-expectation FPTAS by constructing an MIDR mechanism. However, for covering problems, they obtain an additive approximation, which does not translate to a (worst-case) multiplicative approximation. In fact, as observed in [20] and noted earlier, a multiplicative approximation ratio is impossible (in polytime) using their technique, or any other technique that constructs an MIDR mechanism whose range is a proper subset of all outcomes.

For single-dimensional covering problems, various other results, including black-box results, are known. Briest et al. [4] consider a closely-related generalization, which one may call the "single-value setting"; although this is a multidimensional setting, it admits a simple monotonicity condition sufficient for implementability, which makes this setting easier to deal with than our multidimensional settings. They show that a pseudopolynomial time algorithm (for covering and packing problems) can be converted into a truthful FPTAS. Lavi and Swamy [32] mainly consider packing problems, but mention that their technique also yields results for single-dimensional covering problems.

Our decomposition method, where we combine mechanisms for simpler problems into a mechanism for the given problem, is somewhat in the same spirit as the construction in [40]. They give a toolkit for combining truthful mechanisms, identifying sufficient conditions under which this combination preserves truthfulness. But they work only with the single-dimensional setting, which is much more tractable to deal with.

2.2 Problem Definitions and Preliminaries

Various covering problems can be cast in the framework defined in Section 1.2. For example, in the mechanism-design version of *vertex cover* (Section 2.4), the items are edges of a graph. Each player i provides a subset $\mathcal{T}_i$ of the nodes of the graph and incurs a private cost $c_{i,v}$ if node $v \in T_i$ is used to cover an edge. We can set $pub(T_1, \ldots, T_n) = 0$ if $\bigcup_i T_i$ is a vertex cover, and ∞ otherwise, to encode that the solution must be a vertex cover. It is also easy to see that the mechanism-design version of *uncapacitated facility location* (UFL; Section 2.3), where each player provides some facilities and has private facility-opening costs, and the client-assignment costs are public, can be modeled by letting $pub(T_1, \ldots, T_n)$ be the total client-assignment cost given the set $\bigcup_i T_i$ of open facilities.

Below we formally define both problems and related notations.

Multi-VCP mechanism design problem In this problem, we have a graph $G = (V, E)$ with n nodes. Each player i provides a subset $\mathcal{T}_i$ of nodes. For simplicity, we first assume that the $\mathcal{T}_i$s are disjoint, and given a cost-vector $\{c_{i,u}\}_{i\in[n],u\in\mathcal{T}_i}$, we use c_u to denote $c_{i,u}$ for $u \in \mathcal{T}_i$. Notice that our assumption of monopoly-free then means that each $\mathcal{T}_i$ is an independent set. In Remark 2.4.7 we argue that many of the results obtained in this disjoint-$\mathcal{T}_i$s setting (in particular, Theorems 2.4.11 and 2.4.9) also hold when the $\mathcal{T}_i$s are not disjoint (but each $\mathcal{T}_i$ is still an independent set). The goal is to choose a minimum-cost *vertex cover*, i.e., a min-cost set $S \subseteq V$ such that every edge is incident to a node in S.

UFL mechanism design problem In this problem we have a set $\mathcal{D}$ of clients that need to be serviced by facilities, and a set $\mathcal{F}$ of locations where facilities may be opened. Each player i may provide facilities at the locations in $\mathcal{T}_i \subseteq \mathcal{F}$. By making multiple copies of a location if necessary, we may assume that the $\mathcal{T}_i$s are disjoint. Hence, we will simply say "facility ℓ" to refer to the facility at location $\ell \in \mathcal{F}$. For each facility $\ell \in \mathcal{T}_i$ that is opened, i incurs a private opening cost of $\overline{f}_{i,\ell}$, and assigning client j to an open facility ℓ incurs a publicly known assignment/connection cost $c_{\ell j}$. To simplify notation, given a tuple $\{f_{i,\ell}\}_{i\in[n],\ell\in\mathcal{T}_i}$ of facility costs, we use f_ℓ to denote $f_{i,\ell}$ for $\ell \in \mathcal{T}_i$. The goal is to open a subset $F \subseteq \mathcal{F}$ of facilities, so as to minimize $\sum_{\ell\in F} \overline{f}_\ell + \sum_{j\in\mathcal{D}} \min_{\ell\in F} c_{\ell j}$. We will assume throughout that the $c_{\ell j}$s form a metric. It will be notationally convenient to allow our algorithms to have the flexibility of choosing the open facility $\sigma(j)$ to which a client j is assigned (instead of $\text{argmin}_{\ell\in F} c_{\ell j}$); since assignment costs are public, this does not affect truthfulness, and any approximation guarantee achieved also clearly holds when we drop this flexibility.

2.3 A Black-Box Reduction for Multidimensional Metric UFL

In this section, we consider the multidimensional metric *uncapacitated facility location* (UFL) problem and present a *black-box reduction* from truthful mechanism design to algorithm design. We show that any ρ-approximation algorithm for UFL satisfying an additional property can be converted in a black-box fashion to a truthful-in-expectation ρ-approximation mechanism (Theorem 2.3.1). This is the first such result for a multidimensional covering problem. As a corollary, we obtain a truthful-in-expectation, 2-approximation mechanism (Corollary 2.3.3).

We can formulate (metric) UFL as an integer program as follows. Throughout, we use ℓ to index facilities in $\mathcal{F}$ and j to index clients in $\mathcal{D}$.

$$\begin{aligned}
\min \quad & \sum_{\ell} f_\ell y_\ell + \sum_{j,\ell} c_{\ell j} x_{\ell j} & \\
\text{s.t.} \quad & \sum_{\ell} x_{\ell j} \geq 1 & \text{for all } j, \\
& x_{\ell,j} \leq y_\ell & \text{for all } \ell, j, \\
& x_{\ell,j}, y_\ell \in \{0,1\} & \text{for all } \ell, j.
\end{aligned}$$

Here, $\{f_\ell\}_\ell = \{f_{i,\ell}\}_{i\in[n],\ell\in\mathcal{T}_i}$ is the vector of reported facility costs. Variable y_ℓ denotes if facility ℓ is opened, and $x_{\ell j}$ denotes if client j is assigned to facility ℓ; the constraints encode that each client is assigned to a facility, and that this facility must be open.

By relaxing the integrality constraints and removing redundant constraints we obtain the following LP

$$\min \quad \sum_{\ell} f_\ell y_\ell + \sum_{j,\ell} c_{\ell j} x_{\ell j} \qquad \text{s.t.} \qquad \sum_{\ell} x_{\ell j} \geq 1 \quad \forall j, \qquad 0 \leq x_{\ell j} \leq y_\ell \leq 1 \quad \forall \ell, j. \tag{FL-P}$$

Say that an algorithm $\mathcal{A}$ is a *Lagrangian multiplier preserving* (LMP) ρ-approximation algorithm for UFL if for every instance, it returns a solution $\big(F, \{\sigma(j)\}_{j\in\mathcal{D}}\big)$ such that $\rho\sum_{\ell\in F} f_\ell + \sum_j c_{\sigma(j)j} \leq \rho \cdot OPT_{\text{(FL-P)}}$.

The main result of this section is the following black-box reduction.

Theorem 2.3.1 *Given a polytime, LMP ρ-approximation algorithm $\mathcal{A}$ for UFL, one can construct a polytime, truthful-in-expectation, individually rational, ρ-approximation mechanism M for multidimensional UFL.*

Proof : We build upon the convex-decomposition idea used in [32]. The randomized mechanism M works as follows. Let $f = \{f_\ell\}$ be the vector of reported facility-opening costs, and c be the public connection-cost metric.

1. Compute the optimal solution (y^*, x^*) to (FL-P) (for the input (f, c)). Let $\{p_i^* = p_i^*(f)\}$ be the payments made by the "*fractional* VCG" mechanism that outputs the optimal LP solution for every input. That is, $p_i^* = \big(\sum_\ell f_\ell y'_\ell + \sum_{\ell,j} c_{\ell j} x'_{\ell j}\big) - \big(\sum_{\ell\notin\mathcal{T}_i} f_\ell y^*_\ell + \sum_{\ell,j} c_{\ell j} x^*_{\ell j}\big)$, where (y', x') is the optimal solution to (FL-P) with the additional constraints $y_\ell = 0$ for all $\ell \in \mathcal{T}_i$.

2. Let $\mathbb{Z}(P) = \{(y^{(q)}, x^{(q)})\}_{q\in\mathcal{I}}$ be the set of all integral solutions to (FL-P) (i.e. all feasible solutions to UFL). In Lemma 2.3.2, we prove the key technical result that using the LMP algorithm $\mathcal{A}$, one can compute, in polynomial time, nonnegative multipliers $\{\lambda^{(q)}\}_{q\in\mathcal{I}}$ such that $\sum_q \lambda^{(q)} = 1$, $\sum_q \lambda^{(q)} y_\ell^{(q)} = y^*_\ell$ for all ℓ, and $\sum_{q,\ell,j} \lambda^{(q)} c_{\ell j} x_{\ell j}^{(q)} \leq \rho\sum_{\ell,j} c_{\ell j} x^*_{\ell j}$.

3. With probability $\lambda^{(q)}$: (a) output the solution $\big(y^{(q)}, x^{(q)}\big)$; (b) pay $p_i^{(q)}$ to player i, where $p_i^{(q)} = 0$ if $\sum_{\ell\in\mathcal{T}_i} f_\ell y^*_\ell = 0$, and $\sum_{\ell\in\mathcal{T}_i} f_\ell y_\ell^{(q)} \cdot \frac{p_i^*}{\sum_{\ell\in\mathcal{T}_i} f_\ell y^*_\ell}$ otherwise.

Clearly, M runs in polynomial time. Fix a player i. Let $\overline{f}_i$ and f_i be the true and reported cost vector of i. Let f_{-i} be the reported cost vectors of the other players. Let (y^*, x^*) be an optimal solution to (FL-P) for (f, c). Note that $\mathrm{E}[p_i(f)] = p_i^*(f)$. If $\sum_{\ell\in\mathcal{T}_i} f_\ell y^*_\ell = 0$ then this follows since $p_i^*(f) = 0$ (because then (y^*, x^*) is also an optimal solution to (FL-P) when player i does not participate). Otherwise, this follows since $\sum_q \lambda^{(q)} y^{(q)} = y^*_\ell$ for all ℓ. So $\mathrm{E}[u_i(f_i, f_i; \overline{f}_i)] = \mathrm{E}[p_i] - \sum_q \lambda^{(q)} \sum_{\ell\in\mathcal{T}_i} \overline{f}_\ell y_\ell^{(q)} = p_i^*(f) - \sum_{\ell\in\mathcal{T}_i} \overline{f}_\ell y^*_\ell$ where the last equality is again because $\sum_q \lambda^{(q)} y^{(q)} = y^*_\ell$ for all ℓ. Since p^*_i and y^* are respectively the payment to i and the assignment computed for input (f_i, f_{-i}) by the fractional VCG mechanism, which is truthful, it follows that player i maximizes his utility in the VCG mechanism, and hence, his expected utility under mechanism M, by reporting his true opening costs. Thus, M is truthful in expectation.

Note that the above argument (on truthfulness in expectation of M) is heavily dependent on the fact that $\sum_q \lambda^{(q)} y^{(q)} = y^*_\ell$ for all ℓ, and it is the LMP ρ-approximation algorithm (see Lemma 2.3.2) that enables us to efficiently decompose (y^*, x^*) in this way (or precisely as stated in Step 2).

This also implies the ρ-approximation guarantee because the convex decomposition obtained in Step 2 shows that the expected cost of the solution computed by M for input (f, c) (where we may assume that f is the true cost vector) is at most $\rho \cdot OPT_{(\text{FL-P})}(f, c)$. Finally, since the fractional VCG mechanism is IR, for any player i, the VCG payment $p^*_i(f)$ satisfies $p^*_i(f) \geq \sum_{\ell \in \mathcal{T}_i} f_\ell y^*_\ell$, and therefore $p^{(q)}_i \geq \sum_{\ell \in \mathcal{T}_i} f_\ell y^{(q)}_\ell$. So M is IR. ■

Lemma 2.3.2 *The convex decomposition in step 2 can be computed in polytime.*

Proof : It suffices to show that the LP (P) can be solved in polynomial time and its optimal value is 1. Recall that $\{(y^{(q)}, x^{(q)})\}_{q \in \mathcal{I}}$ is the set of all integral solutions to (FL-P).

$$\begin{aligned}
\max \quad & \sum_q \lambda^{(q)} && \text{(P)} \\
\text{s.t.} \quad & \sum_q \lambda^{(q)} y^{(q)}_\ell = y^*_\ell \quad \forall \ell && (2.1) \\
& \sum_{j,\ell,q} \lambda^{(q)} c_{\ell j} x^{(q)}_{\ell j} \leq \rho \sum_{j,\ell} c_{\ell j} x^*_{\ell j} && (2.2) \\
& \sum_q \lambda^{(q)} \leq 1 && (2.3) \\
& \lambda \geq 0.
\end{aligned}$$

$$\begin{aligned}
\min \quad & \sum_\ell y^*_\ell \alpha_\ell + \big(\rho \sum_{j,\ell} c_{\ell j} x^*_{\ell j}\big)\beta + z && \text{(D)} \\
\text{s.t.} \quad & \sum_\ell y^{(q)}_\ell \alpha_\ell + \big(\sum_{j,\ell} c_{\ell j} x^{(q)}_{\ell j}\big)\beta + z \geq 1 \quad \forall q && (2.4) \\
& z, \beta \geq 0.
\end{aligned}$$

Since (P) has an exponential number of variables, we consider the dual (D). Here the α_ℓs, β and z are the dual variables corresponding to constraints (2.1), (2.2), and (2.3) respectively. Clearly, $OPT_{(\mathrm{D})} \leq 1$ since $z = 1$, $\alpha_\ell = 0 = \beta$ for all ℓ is a feasible dual solution. If there is a feasible dual solution (α', β', z') of value smaller than 1, then the rough idea is that by running $\mathcal{A}$ on the UFL instance with facility costs $\{\frac{\alpha'_\ell}{\rho}\}$ and connection costs $\{\beta' c_{\ell j}\}$, we can obtain an integral solution whose constraint (2.4) is violated. (This idea needs be modified a bit since α'_ℓ could be negative; see below.) Hence, we can solve (D) efficiently via the ellipsoid method using $\mathcal{A}$ to provide the separation oracle. This also yields an equivalent dual LP consisting of only the polynomially many violated inequalities found during the ellipsoid method. The dual of this compact LP gives an LP equivalent to (P) with polynomially many $\lambda^{(q)}$ variables whose solution yields the desired convex decomposition.

We now fill in the details. Suppose (α', β', z') is feasible to (D) where $\sum_\ell y^*_\ell \alpha'_\ell + (\rho \sum_{j,\ell} c_{\ell j} x^*_{\ell j})\beta' + z' < 1$. Define $a^+ := \max(0, a)$; for a vector $v = (v_1, \ldots, v_n)$, define $v^+ := (v_1^+, \ldots, v_n^+)$. Consider the UFL instance with facility costs $\{f'_\ell = \alpha'^+_\ell / \rho\}$ and connection costs $\{c'_{\ell j} = \beta' c_{\ell j}\}$. (Clearly c' is also a metric.) Running $\mathcal{A}$ on this input, we can obtain an integral solution $(y^{(q)}, x^{(q)})$ such that

$$\rho \sum_\ell \tfrac{\alpha'^+_\ell}{\rho} y^{(q)}_\ell + \sum_{j,\ell} \beta' c_{\ell j} x^{(q)}_{\ell j} \leq \rho \cdot OPT_{(\text{FL-P})}(f', c') \leq \rho\Big(\sum_\ell \tfrac{\alpha'^+_\ell}{\rho} y^*_\ell + \sum_{j,\ell} \beta' c_{\ell j} x^*_{\ell j}\Big).$$

Clearly the facilities ℓ with $\alpha'_\ell \leq 0$ contribute 0 to the LHS and RHS of the above inequality. Now consider the integer solution $\hat{y}^{(q)}$ where $\hat{y}^{(q)}_\ell$ is 1 if $\alpha'_\ell \leq 0$ and is $y^{(q)}_\ell$ otherwise. Adding $\sum_{\ell:\alpha'_\ell \leq 0} \alpha'_\ell \hat{y}^{(q)}_\ell$ to the LHS and $\sum_{\ell:\alpha'_\ell \leq 0} \alpha'_\ell y^*_\ell$ to the RHS of the above inequality, since $y^*_\ell \leq 1$ for all ℓ and $\alpha'^+_\ell = \alpha'_\ell$ when $\alpha'_\ell > 0$, we infer that

$$\sum_\ell \alpha'_\ell \hat{y}^{(q)}_\ell + \sum_{j,\ell} \beta' c_{\ell j} x^{(q)}_{\ell j} \leq \sum_\ell \alpha'_\ell y^*_\ell + \big(\rho \sum_{j,\ell} c_{\ell j} x^*_{\ell j}\big)\beta' < 1 - z'$$

which contradicts that (α', β', z') is feasible to (D). Hence, $OPT_{(\mathrm{D})} = OPT_{(\mathrm{P})} = 1$.

Thus, we can add the constraint $\sum_\ell y^*_\ell \alpha_\ell + (\rho \sum_{j,\ell} c_{\ell j} x^*_{\ell j})\beta + z \leq 1$ to (D) without altering anything. If we solve the resulting LP using the ellipsoid method, and take the inequalities corresponding to the violated inequalities (2.4) found by $\mathcal{A}$ during the ellipsoid method, then we obtain a compact LP with only a polynomial number of constraints that is equivalent to (D). The dual of this compact LP yields an LP equivalent to (P) with a polynomial number of $\lambda^{(q)}$ variables which we can solve to obtain the desired convex decomposition. ∎

By using the polytime LMP 2-approximation algorithm for UFL devised by Jain et al. [27], we obtain the following corollary of Theorem 2.3.1.

Theorem 2.3.3 *There is a polytime, IR, truthful-in-expectation, 2-approximation mechanism for multidimensional* UFL.

2.4 Truthful Mechanisms for Multidimensional VCP

We now consider the multidimensional vertex-cover problem (VCP), and devise various polytime, truthful, approximation mechanisms for it. We often use Multi-VCP to distinguish multidimensional VCP from its algorithmic counterpart.

As mentioned earlier, VCP becomes a rather challenging mechanism-design problem in the *multidimensional* mechanism-design setting. Whereas for *single-dimensional VCP*, many of the known 2-approximation algorithms for VCP are implementable, none of these underlying techniques yield implementable algorithms even for the simplest multidimensional setting, 2-dimensional VCP, where *every player owns at most two nodes*; see Section 2.5 and 2.6 for examples. Moreover, no maximal-in-distributional-range (MIDR) mechanism whose range is a proper subset of all outcomes can achieve a bounded multiplicative approximation guarantee [20].[1] This also rules out the convex-decomposition technique of [32], which yields MIDR mechanisms.

We develop two main techniques for Multi-VCP in this section. In Section 2.4.1, we introduce a simple class of truthful mechanisms called *threshold mechanisms*, and show that although seemingly restricted, threshold mechanisms can achieve non-trivial approximation guarantees. In Section 2.4.2, we develop a *decomposition method* for Multi-VCP

[1]If $\mathcal{A}$ is a randomized MIDR algorithm and S is an inclusion-wise minimal vertex cover such that the range of $\mathcal{A}$ does not include a distribution that returns S with probability 1, then $\mathcal{A}$ incurs non-zero cost on the instance where the cost of a node u is 0 if $u \in S$ and is 1 (say) otherwise, and so its approximation ratio is unbounded.

that uses threshold mechanisms as building blocks and gives a general way of reducing the mechanism-design problem for Multi-VCP into simpler mechanism-design problems.

By leveraging the decomposition method along with threshold mechanisms, we obtain various truthful, approximation mechanisms for Multi-VCP, which yield the *first* truthful mechanisms for multidimensional vertex cover with non-trivial approximation guarantees. (1) We show that any instance of r-dimensional VCP can be decomposed into $O(r^2 \log n)$ *single-dimensional VCP* instances; this leads to a truthful, $O(r^2 \log n)$-approximation mechanism for r-dimensional VCP (Theorem 2.4.9). In particular, for any fixed r, we obtain an $O(\log n)$-approximation. (2) For any proper minor-closed family of graphs (such as planar graphs), we obtain an improved truthful, $O(r \log n)$-approximation mechanism (Theorem 2.4.11); this improves to an $O(\log n)$-approximation if no two neighbors of a node belong to the same player (Corollary 2.4.12).

Theorem 4.3.3 shows that our mechanisms also enjoy good frugality properties. We obtain the first mechanisms for Multi-VCP that are polytime, truthful, and achieve bounded approximation ratio *and* bounded frugality ratio. This nicely complements a result of [10], who devise such mechanisms for single-dimensional VCP.

2.4.1 Threshold Mechanisms

Definition 2.4.1 *A* threshold mechanism M *for Multi-VCP works as follows. On input* c, *for every* i *and every node* $u \in \mathcal{T}_i$, M *computes a threshold* $t_u = t_u(c_{-i})$ *(i.e.,* t_u *does not depend on* i*'s reported costs).* M *then returns the solution* $S = \{v \in V : c_v \leq t_v\}$ *as the output, and pays* $p_i = \sum_{u \in S \cap \mathcal{T}_i} t_u$ *to player* i.

If t_u only depends on the costs in the neighbor-set $N(u)$ of u, for all $u \in V$ (note that $N(u) \cap \mathcal{T}_i = \emptyset$ if $u \in \mathcal{T}_i$), we call M a *neighbor-threshold mechanism*. A special case of a neighbor-threshold mechanism is an *edge-threshold mechanism*: for every edge $uv \in E$ we have edge thresholds $t_u^{(uv)} = t_u^{(uv)}(c_v)$, $t_v^{(uv)} = t_v^{(uv)}(c_u)$, and the threshold of a node u is given by $t_u = \max_{v \in N(u)}(t_u^{(uv)})$.

In general, threshold mechanisms may not output a vertex cover, however it is easy to argue that threshold mechanisms are always truthful and IR.

Lemma 2.4.2 *Every threshold mechanism for* Multi-VCP *is IR and truthful.*

Proof : IR is immediate from the definition of payments. To see truthfulness, fix a player i. For every $\bar{c}_i, c_i \in C_i, c_{-i} \in C_{-i}$ we have $u_i(c_i, c_{-i}; \bar{c}_i) = \sum_{v \in \mathcal{T}_i : c_v \leq t_v} (t_v - \bar{c}_v)$. It follows that i's utility is maximized by reporting $c_i = \bar{c}_i$. ■

Inspired by [29, 10], we define an x*-scaled* edge-threshold mechanism as follows: fix a vector $(x_u)_{u \in V}$, where $x_u > 0$ for all u, and set $t_u^{(uv)} := x_u c_v / x_v$ for every edge (u, v). We abuse notation and use $\mathcal{A}_x$ to denote both the resulting edge-threshold mechanism and its allocation algorithm. Also, define $\mathcal{B}_x$ to be the neighbor-threshold mechanism where we set $t_u := \sum_{v \in N(u)} x_u c_v / x_v$. Define $\alpha(G; x) := \max_{u \in V} \left(\max_{S \subseteq N(u) : S \text{ independent}} \frac{x(S)}{x_u} \right)$.

Lemma 2.4.3 $\mathcal{A}_x$ *and* $\mathcal{B}_x$ *output feasible solutions and have a tight approximation ratio* $\alpha(G; x) + 1$.

Proof : Clearly, every node selected by $\mathcal{A}_x$ is also selected by $\mathcal{B}_x$. So it suffices to show that $\mathcal{A}_x$ is feasible, and to show the approximation ratio for $\mathcal{B}_x$. For any edge (u, v), either $c_u \leq x_u c_v / x_v$ and u is output, or $c_v \leq x_v c_u / x_u$ and v is output. So $\mathcal{A}_x$ returns a vertex cover.

Let S be the output of $\mathcal{B}_x$ on input c, and let S^* be a min-cost vertex cover. We have $c(S) = c(S \cap S^*) + c(S \setminus S^*) \leq c(S^*) + \sum_{u \in S \setminus S^*} t_u = c(S^*) + \sum_{u \in S \setminus S^*} \sum_{v \in N(u)} x_u c_v / x_v$. Note that $S \setminus S^*$ is an independent set since S^* is a vertex cover, so $\sum_{u \in S \setminus S^*} \sum_{v \in N(u)} x_u c_v / x_v \leq \sum_{v \in S^*} \frac{c_v}{x_v} \sum_{u \in N(v) \setminus S^*} x_u \leq \sum_{v \in S^*} c_v \cdot \alpha(G; x)$. Hence $c(S) \leq (\alpha(G; x) + 1) c(S^*)$. The tightness of the approximation guarantee follows from Example 2.4.5 below. ■

Corollary 2.4.4 *(i) Setting* $x = \vec{1}$ *gives* $\alpha(G; x) \leq \Delta(G)$*, which is the maximum degree of a node in* G*, so* $\mathcal{A}_{\vec{1}}$ *has approximation ratio at most* $\Delta(G) + 1$.

(ii) Taking x *to be the eigenvector corresponding to the largest eigenvalue* $\lambda_{\max}$ *of the adjacency matrix of* G *(*$x > 0$ *by the Perron-Frobenius theorem) gives* $\alpha(G; x) \leq \lambda_{max}$ *(see [10]), so* $\mathcal{A}_x$ *has approximation ratio* $\lambda_{max} + 1$.

Example 2.4.5 (Tightness of approximation ratio of $\mathcal{A}_x$ and $\mathcal{B}_x$) *Let u and $S \subseteq N(u)$ achieve the maximum in the definition of $\alpha(G; x)$. Now consider the instance (G, c) where $c_u = x_u$, $c_v = x_v$ for all $v \in S$ and $c_w = 0$ for all $w \in V \setminus (\{u\} \cup S)$. The mechanism $\mathcal{A}_x$ will include $\{u\} \cup S$ in the output, whereas $V \setminus S$ is a vertex cover of cost $c_u = x_u$. So, $\mathcal{A}_x$ has approximation ratio at least $\frac{x_u + x(S)}{x_u} = 1 + \alpha(G; x)$.*

Although neighbor-threshold mechanisms are more general than edge-threshold mechanisms, Lemma 2.4.6 shows that this yields limited dividends in the approximation ratio. Define $\alpha'(G) = \min_{\text{orientations of } G}\big(\max_{u \in V, S \subseteq N^{\text{in}}(u): S \text{ independent}} |S|\big)$, where $N^{\text{in}}(u) = \{v \in N(u) : (u, v) \text{ is directed into } u\}$. Note that $\alpha'(G) \leq \alpha(G; \vec{1}) \leq \Delta(G)$. If $G = (V, E)$ is *everywhere γ-sparse*, i.e., $|\{(u, v) \in E : u, v \in S\}| \leq \gamma |S|$ for all $S \subseteq V$, then $\alpha'(G) \leq \gamma$; this follows from Hakimi's theorem [25]. A well-known result in graph theory states that for every proper family $\mathcal{G}$ of graphs that is closed under taking minors (e.g., planar graphs), there is a constant γ, such that every $G \in \mathcal{G}$ has at most $\gamma |V(G)|$ edges [35] (see also [15], Chapter 7, Exer. 20); since $\mathcal{G}$ is minor-closed, this also implies that G is *everywhere* γ-sparse, and hence $\alpha'(G) \leq \gamma$ for all $G \in \mathcal{G}$.

Lemma 2.4.6 *A (feasible) neighbor-threshold mechanism M for graph G with approximation ratio ρ, yields an $O\big(\rho \log(\alpha'(G))\big)$-approximation edge-threshold mechanism for G. This implies an approximation ratio of (i) $O(\rho \log \gamma)$ if G is an everywhere γ-sparse graph; (ii) $O(\rho)$ if G belongs to a proper minor-closed family of graphs (where the constant in the $O(.)$ depends on the graph family).*

Proof : Statements (i) and (ii) follow from the statement for general graphs and the graph-theoretic facts mentioned before Lemma 2.4.6, so we focus on proving the statement for an arbitrary graph G. Let $\alpha' = \alpha'(G)$.

Consider an arbitrary vertex $v \in V$. For any $u \in N(v)$ define $x_v^{(uv)} := \inf\{\sigma \geq 0 : t_u(c_v = \beta, c_{-v} = \vec{0}) \geq 1 \ \ \forall \beta \geq \sigma\}$.

Claim 1: $x_v^{(uv)} < \infty$. If not, then for any $p > 0$, there exists $q \geq p$ such that $t_u(c_v = q, c_{-v} = \vec{0}) < 1$. So, let $p = \rho + \epsilon$ for some small $\epsilon > 0$ and $q \geq p$ be such that $t_u(c_v = q, c_{-v} = \vec{0}) < 1$. Consider the cost vector c where $c_u = 1$, $c_v = q$, and $c_z = 0$ for

$z \neq u, v$, we see that the approximation ratio ρ is contradicted for the instance (G, c) (i.e., graph G with the cost vector c): $V \setminus v$ is an optimal vertex cover of cost 1 but the threshold mechanism does not choose u so it chooses v as it is feasible and incurs cost $q > \rho$.

Claim 2: $x_v^{(uv)} > 0$. If $x_v^{(uv)} = 0$, then similar to the above, by considering c where $c_u = 1$, $c_v = \epsilon$, $c_z = 0$ for $z \neq u, v$, where ϵ is very small, we see that M outputs u, which means M does not have the approximation ratio ρ.

Now orient the edges of G according to the orientation that determines $\alpha'(G)$ to obtain the directed graph D. For any arc (u, v) in D, consider linear edge-threshold functions $t_v^{(uv)}(c_u) = x_v^{(uv)} c_u$, and $t_u^{(uv)}(c_v) = (1/x_v^{(uv)}) c_v$. Using these edge-thresholds we obtain an edge-threshold mechanism M'. M' is feasible since for any arc (u, v) if u is not chosen by M', we should have $c_u > t_u^{(uv)}(c_v) = (1/x_v^{(uv)}) c_v$ which implies $t_v^{(uv)}(c_u) = x_v^{(uv)} c_u > c_v$ hence v is chosen by M'.

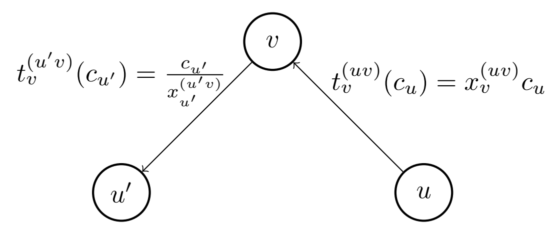

We assert that M' has approximation ratio $O(\rho \log(\alpha'))$. Note that if T is the outcome of M' and T^* is the optimal outcome, then we have

$$\begin{aligned}
c(T) &= c(T \cap T^*) + c(T \setminus T^*) \leq c(T^*) + \sum_{w \in T \setminus T^*} \max_{u \in N(w)} t_w^{(uw)}(c_u) \\
&\leq c(T^*) + \sum_{w \in T \setminus T^*} \sum_{u \in N(w)} t_w^{(uw)}(c_u) = c(T^*) + \sum_{\substack{w \in T \setminus T^* \\ u \in N(w)}} c_u t_w^{(uw)}(1) \\
&= c(T^*) + \sum_{u \in T^*} c_u \sum_{w \in N(u) \cap (T \setminus T^*)} t_w^{(uw)}(1) \qquad (\text{since } N(w) \subseteq T^* \text{ for } w \notin T^*)
\end{aligned}$$

Note that $T \setminus T^*$ is an independent set, so it suffices to show for any $u \in V(G)$, if $S \subseteq N(u)$ forms an independent set then $\sum_{w \in S} t_w^{(uw)}(1) \leq \rho(\log(\alpha') + 2)$.

Let $\delta^{out}(u) = \{v : (u,v) \in D\}$, $S_1 := S \cap \delta^{out}(u)$, and $S_2 := S \setminus S_1$. So, we have

$$\sum_{w \in S} t_w^{(uw)}(1) = \sum_{w \in S_1} t_w^{(uw)}(1) + \sum_{w \in S_2} t_w^{(uw)}(1) = \sum_{w \in S_1} x_w^{(uw)} + \sum_{w \in S_2} \tfrac{1}{x_u^{(uw)}} \tag{2.5}$$

Choose an arbitrary $w \in S_1$. By definition of $x_w^{(uw)}$, for every $\epsilon_w \geq 0$, there is some $0 \leq \delta_w \leq \epsilon_w$ such that $t_u(c_w = x_w^{(uw)} - \epsilon_w + \delta_w, \vec{0}) < 1$. Hence, $u \notin M(G, \hat{c})$ where $\hat{c}_w = x_w^{(uw)} - \epsilon_w + \delta_w$, $\hat{c}_u = 1$, and $\hat{c}_z = 0$ otherwise. So, since $M(G,\hat{c})$ is a vertex cover, we should have $w \in M(G,\hat{c})$ which means $t_w(c_u = 1, \vec{0}) \geq x_w^{(uw)} - \epsilon_w + \delta_w$. Thus, as S_1 is an independent set, for the cost vector c' where $c'_u = 1$, $c'_w = x_w^{(uw)} - \epsilon_w + \delta_w$ if $w \in S_1$, and $c'_z = 0$ otherwise, we have $S_1 \subseteq M(G, c')$ (since $t_w(c'_{N(w)}) = t_w(c_u = 1, \vec{0})$). Letting ϵ_w tend to 0, we get that $\rho \geq \sum_{w \in S_1} x_w^{(uw)}$, as $V \setminus N(u)$ is a vertex cover of cost 1.

Let $S_2 = \{v_1, \ldots, v_k\}$ where $x_u^{(uv_1)} \leq x_u^{(uv_2)} \leq \ldots \leq x_u^{(uv_k)}$. Consider c'' where $c''_u = x_u^{(uv_q)}$, $c''_z = 1$ if $z \in S_2$, and $c''_z = 0$ otherwise. Then, $\{v_1, \ldots, v_q\} \subseteq M(G, c'')$ hence $\rho \geq q/x_u^{(uv_q)}$ for each $q = 1, \ldots, k$. So, $\sum_{q=1}^{k} \frac{1}{x_u^{(uv_q)}} \leq \sum_{q=1}^{k} \rho/q \leq \rho(\log(|S_2|)+1) \leq \rho\log(\alpha') + \rho$. Therefore, (2.5) gives

$$\sum_{w \in S} t_w^{(uw)}(1) \leq \rho + \rho\log(\alpha') + \rho = \rho(\log(\alpha') + 2). \qquad \blacksquare$$

$\blacksquare$

Remark 2.4.7 *Any neighbor-threshold mechanism M with approximation ratio ρ that works under the disjoint-$\mathcal{T}_i$s assumption can be modified to yield a truthful, ρ-approximation mechanism when we drop this assumption. Let $A_u = \{i : u \in \mathcal{T}_i\}$. Set $\hat{c}_u = \min_{i \in A_u} c_{i,u}$ for each $u \in V$ and let $\hat{t}_u$ be the neighbor-threshold of u for the input $\hat{c}$. Note that $\hat{t}_u$ depends only on c_{-i} for every $i \in A_u$. Set $t_u^i := \min\{\hat{t}_u, \min_{j \neq i : u \in \mathcal{T}_j} c_{j,u}\}$ for all $i, u \in \mathcal{T}_i$. Consider the threshold mechanism M' with $\{t_u^i\}$ thresholds, where we use a fixed tie-breaking rule to ensure that we pick u for at most one player $i \in A_u$ with $c_{i,u} = t_u^i$. Then the outputs of M on c, and of M' on input $\hat{c}$ coincide. Thus, M' is a truthful, ρ-approximation mechanism.*

2.4.2 A Decomposition Method

We now propose a general reduction method for Multi-VCP that uses threshold mechanisms as building blocks to reduce the task of designing truthful mechanisms for Multi-VCP to the task of designing threshold mechanisms for simpler (in terms of graph structure or the dimensionality of the problem) Multi-VCP problems. This reduction is useful because designing good threshold mechanisms appears to be a much more tractable task for Multi-VCP. By utilizing the threshold mechanisms designed in Section 2.4.1 in our decomposition method, we obtain an $O(r \log n)$-approximation mechanism for any proper minor-closed family of graphs, and an $O(r^2 \log n)$-approximation mechanism for r-dimensional VCP.

A *decomposition mechanism* M for $G = (V, E)$ is constructed as follows.

– Let $G_1, \ldots, G_k$ be subgraphs of G such that $\bigcup_{q=1}^{k} E(G_q) = E$,

– Let $M_1, \ldots, M_k$ be threshold mechanisms for $G_1, \ldots, G_k$ respectively. For any $v \in V$, let t_v^q be v's threshold in M_q if $v \in V(G_i)$, and 0 otherwise.

– Define M to be the threshold mechanism obtained by setting the threshold for each node v to $t_v := \max_{q=1,\ldots,k}(t_v^q)$ for any $v \in V$. The payments of M are then as specified in Definition 2.4.1. Notice that if all the M_is are neighbor threshold mechanisms, then so is M.

Lemma 2.4.8 *The decomposition mechanism M described above is IR and truthful. If $\rho_1, \ldots, \rho_k$ are the approximation ratios of $M_1, \ldots, M_k$ respectively, then M has approximation ratio $\left(\sum_q \rho_q\right)$.*

Proof : Since M is a threshold mechanism, it is IR and truthful by Lemma 2.4.2. The optimal vertex cover for G induces a vertex cover for each subgraph G_q. So M_q outputs a vertex cover S_q of cost at most $\rho_q \cdot OPT$, where OPT is the optimal vertex-cover cost for G. It is clear that M outputs $\bigcup_q S_q$, which has cost at most $\left(\sum_q \rho_q\right) \cdot OPT$. ■

Theorem 2.4.9 *For any r-dimensional instance of Multi-VCP on $G = (V, E)$, one can obtain a polytime, $O(r^2 \log |V|)$-approximation, decomposition mechanism, even when the $\mathcal{T}_i$s are not disjoint.*

Proof : We decompose G into single-dimensional subgraphs, by which we mean subgraphs that contain at most one node from each $\mathcal{T}_i$. Initialize $j = 1$, $V_j = \emptyset$. While, $\bigcup_{q=1}^{j-1} E(G_q) \neq$

E, we do the following: for every player i, we pick one of the nodes of $\mathcal{T}_i$ uniformly at random and add it to V_j. We also add all the nodes in $V \setminus \left(\bigcup_{i=1}^{n} \mathcal{T}_i\right)$ to V_j. Let G_j be the induced subgraph on V_j; set $j \leftarrow j+1$.

For any edge $e \in E$, the probability that both of its ends appear in some subgraph G_j is at least $1/r^2$. So, the expected value of $|E \setminus \bigcup_{q=1}^{j-1} E(G_q)|$ decreases by a factor of at least $(1-1/r^2)$ with j. Hence, the expected number of subgraphs produced above is $O\left(\frac{\log|E|}{\log(r^2/(r^2-1))}\right) = O(r^2 \log |V|)$ (this also holds with high probability). Each G_j yields a single-dimensional VCP instance (where a node may be owned by multiple players). Any truthful mechanism for a 1D-problem is a threshold mechanism (recall Theorem 1.2.8). So we can use any truthful, 2-approximation mechanism for single-dimensional VCP for the G_js and obtain an $O(r^2 \log n)$-approximation for r-dimensional Multi-VCP. ■

The following lemma shows that the decomposition into single-dimensional subgraphs, obtained above, is essentially the best that can hope for, when $r = 2$.

Lemma 2.4.10 *There are instances of 2-dimensional VCP that require* $\Omega(\log |V(G)|)$ *single-dimensional subgraphs in any decomposition of* G.

Proof : Define G^n to be the bipartite graph with vertices $\{u_1, \ldots, u_n, v_1, \ldots, v_n\}$ and edges $\{(u_i, v_j) : i \neq j\}$. Each player $i = 1, \ldots, n$ owns vertices u_i and v_i.

For $n = 2$ the claim is obvious. Let q_n be the minimum number of single-dimensional subgraphs needed to decompose G^n. Suppose the claim is true for all $j < n$ and we have decomposed G^n into single-dimensional subgraphs $D = \{G_1, \ldots, G_{q_n}\}$. We may assume that $V(G_1) = \{u_1, \ldots, u_k, v_{k+1}, \ldots, v_n\}$ (if G_1 has less than n nodes, pad it with extra nodes). Let H_1 and H_2 be the subgraphs of G induced by $\{u_1, \ldots, u_k, v_1, \ldots, v_k\}$ and $\{u_{k+1}, \ldots, u_n, v_{k+1}, \ldots, v_n\}$, respectively. The graphs in $D \setminus \{G_1\}$ must contain a decomposition of H_1 and a decomposition of H_2. So $q_n \geq 1 + \max(q_k, q_{n-k})$, and hence, by induction, we obtain that $q_n \geq 1 + (1 + \log_2(n/2)) = 1 + \log_2 n$. ■

Complementing Theorem 2.4.9, we next present another decomposition mechanism that exploits the graph structure to obtain an improved approximation guarantee. Given a graph $G = (V, E)$ and a set $S \subseteq V$, we use $E[S]$ to denote the set of edges having both end points in S, and $N(S) = \{u \in V \setminus S : \exists v \in S \text{ s.t. } (u,v) \in E\}$ to denote the neighbors of S. Also, let $\delta(S,T)$ denote the set of edges of G having one end point each in S and T. When we subscript a quantity (e.g., $\delta(S)$ or $N(S)$) with a specific graph, we are referring to the quantity in that specific graph.

Theorem 2.4.11 *If $G = (V,E)$ is everywhere γ-sparse, then one can devise a polytime, $O(\gamma r \log |V|)$-approximation decomposition mechanism for r-dimensional VCP on G. Hence, there is a polytime, truthful, $O(r \log n)$-approximation mechanism for r-dimensional VCP on any proper minor-closed family of graphs. These guarantees also hold when the $\mathcal{T}_i$s are not disjoint.*

Proof : Set $G = G_0 = (V_0, E_0)$, and let $n_0 = |V_0|$. Since $|E_0| \leq \gamma n_0$, there are at most $n_0/2$ nodes in V_0 with degree larger than 4γ. Let $T_1 = \{u \in V_0 : \delta(u) \leq 4\gamma\}$. Let $H_1 = \big(T_1, E[T_1]\big)$ be the subgraph of G_0 induced by T_1. Also, consider the bipartite subgraph $B_1 = \big(T_1 \cup N_{G_1}(T_1), \delta_{G_1}(T_1, N_{G_1}(T_1))\big)$ where $G_1 := G_0 \setminus T_1$ (i.e., we delete the nodes in T_1 and the edges incident to them to obtain G_1) is also γ-sparse. So, we can similarly find a subgraph H_2 that contains at least half of the nodes of G_1, and the bipartite subgraph B_2 of G_1. Continuing this process, we obtain subgraphs $H_1, B_1, H_2, B_2, \ldots, H_k, B_k$ that partition G, where for every q, each node of H_q and each node on one of the sides of B_q has degree (in that subgraph) at most 4γ, and $|V(H_q)| \geq |V(G \setminus (T_1 \cup \ldots T_{q-1})|/2$. Hence, $k \leq \log n$. Using the (edge-threshold) mechanism $\mathcal{A}_{\vec{1}}$ defined in Corollary 2.4.4, for each H_q subgraph gives a $(4\gamma+1)$-approximation for each H_q. Let $B_q = \big(T_q \cup R_q, F_q\big)$, where $R_q = N_{G_q}(T_q)$, and $F_q = \delta_{G_q}(T_q, R_q)$.

Let $T = \bigcup_q T_q$, $R = \bigcup_q R_q$. Note that a node u could lie in $T \cap R$. We replace each such node $u \in T \cap R$ with two distinct "copies" u_1 and u_2, and place u_1 in T and u_2 in R. If $u \in \mathcal{T}_i$ for some player i, then we include both u_1, u_2 in $\mathcal{T}_i$, and set $c_{i,u_1} = c_{i,u_2} = c_{i,u}$. The understanding is that if any of u_1 or u_2 is picked, then we pick u; in other words, the threshold of u is the maximum of the thresholds of u_1 and u_2. Let $T \uplus R$ denote the resulting set of nodes (with bipartition T, R). We create a bipartite graph $B = (T \uplus R, F)$ representing the union of all the B_qs, where F is defined as follows. For notational simplicity, if a node u is in exactly one of T and R (so it has only one copy in $T \uplus R$), we set $u_1 = u_2 = u$. For every $q = 1, \ldots, k$, and every edge $(u,v) \in F_q$, where $u \in R_q$, $v \in T_q$, we include the edge (u_2, v_1) in F. Note that: (a) B is bipartite; (b) the maximum degree of T (in B) is at most 4γ; and, (c) every edge in $E \setminus \bigcup_q E(H_q)$ maps to exactly one edge of F. We show that one can obtain an $O(r\gamma \log n)$-approximation decomposition mechanism for B. Thus, we obtain an $O(r\gamma \log n)$-approximation decomposition mechanism for G.

We obtain $O(r \log n)$ bipartite graphs whose edges cover F, with the property that in each resulting bipartite subgraph Z, for each node $u \in R \cap V(Z)$, and each player i, *at most one* of u's neighbors in Z is in $\mathcal{T}_i$. We use a procedure similar to that in the proof of Theorem 2.4.9. For each i, we pick one node from $T \cap \mathcal{T}_i$ uniformly at random; let X be the set of nodes picked from T. We create the bipartite graph Z^j consisting of all edges between X and $N_B(X)$. We increment j and continue this process until all edges of F have been covered. Since the probability that an edge $(u, v) \in F$ is covered in an iteration is at least $\frac{1}{r}$, $O(r \log n)$ subgraphs suffice, in expectation and with high probability, to cover F.

Now, for each bipartite graph Z^j with bipartition $X^j \cup Y^j$, where $X^j \subseteq T$, $Y^j \subseteq R$, we use the following threshold mechanism. Assume for now that the $\mathcal{T}_i$s are disjoint, and set $c_u = c_{i,u}$ if $u \in \mathcal{T}_i$. For each $u \in Y^j$, we pick u if $c_u \leq \sum_{v \in N_{Z^j}(u)} c_v$, and we pick $N_{Z^j}(u)$ if $\sum_{v \in N_{Z^j}(u)} c_v \leq c_u$. Note that since $|X^j \cap \mathcal{T}_i| \leq 1$ for every i, this is a valid threshold mechanism. The cost of the solution S output by this mechanism for Z^j is at most $2 \sum_{u \in Y^j} c(S_u^*)$, where S_u^* is the optimal vertex cover for the star consisting of u and $N_{Z^j}(u)$. Since every node in X^j has degree at most 4γ, it is not hard to see that $\sum_{u \in Y^j} c(S_u^*) \leq 4\gamma \cdot OPT(Z^j)$, where $OPT(Z^j)$ is the value of an optimal vertex cover for Z^j. This follows since, for example, concatenating the optimal dual solutions corresponding to the S_u^*s and scaling by 4γ yields a feasible solution to the dual of the vertex-cover LP for Z^j. Therefore, the threshold mechanism for Z^j is an 8γ-approximation, and hence we obtain an $O(r\gamma \log n)$-approximation for B.

If the $\mathcal{T}_i$s are not disjoint, then by Remark 2.4.7, the $O(\gamma)$-approximation for the H_qs still holds. When constructing Z^j, we set the "owners" of a node $v \in T$ included in Z^j to be all the players i who picked v as the random node from their $\mathcal{T}_i$-set (and hence caused v to be included in Z^j); the owners of a node $u \in Y^j$ are unchanged, that is, $\{i : u \in \mathcal{T}_i\}$. Now, as in Remark 2.4.7, we can move from this to an instance where each node is owned by at most one player. Although the mechanism described above for

Z^j is *not* a neighbor-threshold mechanism, it is not hard to see that since the threshold for a node $v \in T \cap V(Z^j)$ depends only on nodes that are at hop-distance at most 2 from v, none of which are owned by any player owning v in Z^j, the same reasoning as in Remark 2.4.7 shows that the $O(\gamma)$-approximation threshold mechanism obtained above for Z^j holds even when a node is owned by multiple players. Thus, we still obtain an $O(\gamma r \log |V|)$-approximation mechanism.

As noted in Section 2.4.1, every proper minor-closed family of graphs is everywhere γ-sparse for some $\gamma > 0$. Thus, the above result implies a truthful, $O(r \log^2 n)$-approximation for any proper minor-closed family (where the constant in the $O(.)$ depends on the graph family; e.g., for planar graphs $\gamma \leq 4$). ∎

Given a graph $G = (V, E)$, define a *3-hop-far* instance of Multi-VCP on G to be one that satisfies $|N(u) \cap \mathcal{T}_i| \leq 1$ for every $u \in V$ and every player i; that is no two neighbors of a node are owned by the same player. On such instances, one can improve the guarantee of Theorem 2.4.11 by removing the dependence on $\max_i |\mathcal{T}_i|$.

Corollary 2.4.12 *Let $G = (V, E)$ be an everywhere γ-sparse graph. One can devise a polytime $O(\gamma \log |V|)$-approximation decomposition mechanism for 3-hop-far instances of Multi-VCP on G. Hence, one obtains a polytime, truthful $O(\log n)$-approximation mechanism for 3-hop far Multi-VCP on any proper minor-closed family of graphs. These guarantees also hold when the $\mathcal{T}_i$s are not disjoint.*

Proof : The proof follows from that of Theorem 2.4.11. The only change is that we no longer need to decompose the bipartite graph B into the Z^j subgraphs: since the input is a 3-hop-far Multi-VCP instance, the Multi-VCP instance on B already satisfies the property required of the Z^j graphs. Thus, we obtain an $O(\gamma)$-approximation for B, and an $O(\gamma)$-approximation for each H_q, and hence an $O(\gamma \log |V|)$-approximation for G. The consequences when the $\mathcal{T}_i$s are not necessarily disjoint, and for a proper minor-closed family of graphs follow as in the proof of Theorem 2.4.11. ∎

2.5 LP-rounding Does Not Work for Multi-VCP

A common method for designing approximation algorithms for VCP (and in general) is to solve the following LP-relaxation and then round the optimal solution.

$$\min \quad \sum_v c_v x_v \qquad \text{s.t.} \qquad x_u + x_v \geq 1 \quad \forall (u,v) \in E. \tag{VC-P}$$

We show that *any* LP-rounding algorithm that always includes nodes with $x_u \geq \frac{1}{2}$ and does not include any node u with $x_u = 0$ is not WMON.

Example 2.5.1 *Consider the graph G shown below where u and v belong to player 1. For the cost-vector $(c_u, c_a, c_b, c_v, c_d) = (5/4, 1, 1, 1, 1)$, the unique optimal solution to the LP is $(x_u, x_a, x_b, x_v, x_d) = (1/2, 1/2, 1/2, 1/2, 1/2)$. Therefore, the algorithm includes both u and v in the output.*

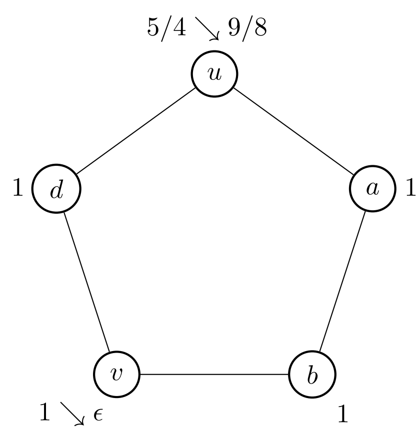

G

Figure 2.1: Example showing that a natural LP-rounding algorithm is not WMON.

Consider the cost vector $c' = (c'_1, c_{-1})$ where player 1 reduces the costs for u and v to $c'_u = 9/8$ and $c'_v = \epsilon < 1/16$ (all other costs are unchanged). Then WMON dictates that both u and v must still be chosen. However, the unique optimal solution to the LP with the new costs is $x_a = x_d = x_v = 1,\ x_u = x_b = 0$ with cost $2 + \epsilon$. (This follows because if

$x_u = 1$ then the cost of an LP solution is at least $1 + 9/8$; if $x_u = 1/2$, then the cost of an LP solution is at least $9/16 + 1 + 1/2$; both are greater than $2 + \epsilon$ as $\epsilon < 1/16$.) So M will not output u, which contradicts WMON.

The example mentioned above also shows that the following well-known combinatorial 2-approximation algorithm for VCP does not satisfy WMON: Given a graph $G = (V, E)$, construct a bipartite graph G' having two copies of V, say V_1, V_2, and having edges $(u_1, v_2), (u_2, v_1)$ for every edge $(u, v) \in E$; solve VCP on G' and if any of the copies of a node are chosen in this solution, then pick that node in the solution for G.

In the above example, for the cost-vector c, every optimal vertex cover for G' includes exactly one copy of u and one copy of v, so both u and v will be chosen in the solution for G. For the cost-vector c', no optimal vertex cover for G' includes any copies of u, so u will not be chosen in the solution for G. This contradicts WMON.

2.6 Primal-dual Methods Do Not Work for Multi-VCP

The dual of (VC-P) is as follows.

$$\max \quad \sum_e y_e \qquad \text{s.t.} \qquad \sum_{e \in \delta(v)} y_e \le c_v \quad \forall v \in V. \tag{VC-D}$$

Various primal-dual algorithm based on dual ascent are known to yield 2-approximation algorithms. All of these start with $y = \vec{0}$, raise dual variables while maintaining dual feasibility, and return the nodes whose costs are completely "paid" by the dual variables.

The two most common variants are where one fixes an ordering of the edges in which to raise dual variables, and where one raises all (unfrozen) dual variables simultaneously. We show that neither of these lead to WMON algorithms.

Example 2.6.1 *Consider the graph shown in Fig. 2.2, where the dual variables are increased in the order ux, xy, yv, and u and v belong to one player.*

Let $c_u = 1$, $c_x = 1.5$, $c_y = 1.05$, $c_v = 0.5$. The primal-dual algorithm will output $\{u, x, v\}$. Now, if we reduce c_u to 0.5 and c_v to 0.3, and keep c_x and c_y unchanged, the algorithm outputs $\{u, x, y\}$ which contradicts WMON.

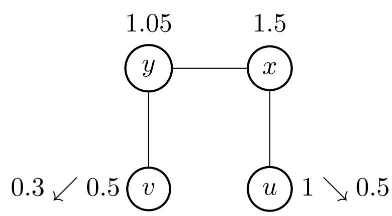

Figure 2.2: Example showing Primal-Dual algorithm is not WMON; dual variables are increased in a fixed order.

Example 2.6.2 *Now consider the simultaneous-dual-ascent primal-dual algorithm. Consider again the same graph as in Example 2.6.1 but with a different assignment of costs, as shown in Fig. 2.3. Let* $c_u = 1,\ c_x = 3,\ c_y = 4.6,\ c_v = 2.5$. *The primal-dual algorithm outputs* $\{u, x, v\}$. *Now, if we reduce* c_u *to* 0.5 *and* c_v *to* 2.4 *and keep* c_x *and* c_y *unchanged, the algorithm outputs* $\{u, y\}$, *which contradicts WMON.*

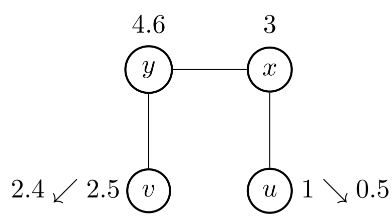

Figure 2.3: Example showing Primal-Dual algorithm is not WMON; dual variables are increased simultaneously.

Chapter 3

Near-Optimal and Robust Mechanisms for Payment Minimization

3.1 Introduction

In this chapter we consider the payment-minimization (PayM) problem of designing incentive-compatible, ex-post individually rational (IR) mechanisms for *covering problems* in the *Bayesian setting*, where players' types are drawn from an underlying distribution and may be *correlated*, and the goal is to *minimize the expected total payment made by the mechanism*. This kind of objective is indeed considered in various real-life settings, for example Ariba.com is providing online procurement auctions as supply management tool for various industries. Consider the simplest such setting of a *single-item procurement auction*, where a buyer wants to buy an item from any one of n sellers. Myerson's seminal result [41] solves this problem (and other single-dimensional problems) when players' private types are *independent*. However, no such result (or characterization) is known when players' types are correlated. This is the question that motivates the work in this chapter.

Whereas the analogous revenue-maximization problem for packing domains, such as combinatorial auctions (CAs), has been extensively studied in the algorithmic mechanism design (AMD) literature, both in the case of independent and correlated (even interdependent) player-types (see, e.g., [6, 7, 5, 1, 9, 26, 17, 44, 5, 47] and the references therein), surprisingly, there are *almost no results* on the payment-minimization problem in the AMD literature (see however [7]). The economics literature does contain various general results that apply to both covering and packing problems. However much of this work focuses on characterizing special cases; see, e.g., [51]. An exception is the work of Crémer and McLean [12, 13], which shows that under certain conditions, one can devise a Bayesian-incentive-compatible (BIC) mechanism whose expected total payment is exactly equal to the expected cost incurred by the players, albeit one where players may incur negative utility under certain type-profile realizations.

3.1.1 Summary of Results

We initiate a study of payment-minimization (PayM) problems from the AMD perspective of designing computationally efficient, near-optimal mechanisms. We develop black-box reductions from mechanism design to algorithm design whose application yields a variety of optimal and near-optimal mechanisms. As we elaborate below, covering problems turn out to behave quite differently in certain respects from packing problems, which necessitates new approaches (and solution concepts).

Formally, we consider the setting of *correlated players* in the explicit model, that is, where we have an explicitly-specified *arbitrary discrete* joint distribution of players' types. The most common solution concept in Bayesian settings is, as hinted briefly in Chapter 1, Bayesian incentive compatibility (BIC) and interim individual rationality (interim IR), wherein at the *interim* stage when a player knows his type but is oblivious of the random choice of other players' types, truthful participation in the mechanism by all players forms a Bayes-Nash equilibrium. Two serious drawbacks of this solution concept (which are exploited strikingly and elegantly in [12, 13]) are that: (i) a player may regret his decision of participating and/or truthtelling *ex post*, that is, after observing the realization of other players' types; and (ii) it is overly-reliant on having precise knowledge of the true underlying distribution making this a rather *non-robust* concept: if the true distribution differs, possibly even slightly, from the mechanism designer and/or players' beliefs or information about it, then the mechanism could lose its IC and IR properties.

We formulate a notion of incentive compatibility (IC) that we call *robust Bayesian IC* (robust BIC) that on the one hand is substantially more robust than BIC, and on the other is flexible enough that it allows one to obtain various polytime near-optimal mechanisms satisfying this notion. A *robust-(BIC, IR)* mechanism (see Section 3.2) ensures that truthful participation in the mechanism is in the best interest of every player (i.e. a "no-regret" choice) *even at the ex-post stage* when the other players' (randomly-chosen) types are revealed to him. Thus, a robust-(BIC, IR) mechanism is significantly more robust than a (BIC, interim-IR) mechanism since it retains its IC and IR properties for a wide variety of distributions, including those having the same support as the actual distribution. In other words, in keeping with Wilson's doctrine of detail-free mechanisms, the mechanism functions robustly even under fairly limited information about the type-distribution.

We show that for a variety of settings, one can reduce the robust-(BIC, IR) payment-minimization (PayM) mechanism-design problem to the algorithmic cost-minimization (CM) problem of finding an outcome that minimizes the total cost incurred. Moreover, this black-box reduction applies to: (a) single-dimensional settings even when we only have an LP-relative *approximation algorithm* for the CM problem (that is required to work only with nonnegative costs) (Theorem 3.6.2); and (b) *multidimensional problems* with additive types (Corollary 3.5.3).

Our reduction yields near-optimal robust-(BIC-in-expectation, IR) mechanisms for a variety of covering settings such as (a) various single-dimensional covering problems including single-item procurement auctions (Table 3.1); (b) multi-item procurement auctions (Theorem 3.7.1); and (c) multidimensional facility location (Theorem 3.7.3). (Robust BIC-in-expectation means that the robust-BIC guarantee holds for the *expected utility* of a player, where the expectation is over the random coin tosses of the mechanism.) Our techniques can be adapted to yield *truthful-in-expectation* mechanisms with the same guarantees for single-dimensional problems with a constant number of players. These are the *first* results for the PayM mechanism-design problem with correlated players under a notion stronger than (BIC, interim IR). To our knowledge, our results are new even for the simplest covering setting of single-item procurement auctions.

On a side note, we note that we can leverage our ideas to also expand upon the results in [17] for revenue-maximization with correlated players and make significant progress on a research direction proposed in [17]. We show that any "integrality-gap verifying" ρ-approximation algorithm for the SWM problem (as defined in [32]) can be used to obtain a truthful-in-expectation mechanism whose revenue is at least a ρ-fraction of the optimum revenue (see Appendix A).

In comparison with [17], which is the work most closely-related to ours, our reduction from robust-BIC mechanism design to the algorithmic CM problem is stronger than the reduction in [17] in two ways. First, for single-dimensional settings, it applies even with LP-relative *approximation algorithms*, and the approximation algorithm is required to work only for "proper inputs" with nonnegative costs. (Note that whereas for packing problems, allowing negative-value inputs can be benign, this can change the character of a covering problem considerably; in particular, the standard notion of approximation becomes meaningless since the optimum could be negative.) In contrast, Dobzinski et al. [17] require an exact algorithm for the analogous social-welfare-maximization (SWM) problem. Second, our reduction also applies to multidimensional settings with additive types (see Section 3.2), albeit we now require an exact algorithm for the CM problem.

3.1.2 Related Work

The AMD literature has concentrated mostly on the independent-players setting [6, 7, 5, 1, 9, 26]. There has been some, mostly recent, work that also considers correlated players [46, 17, 44, 5, 47]; as noted earlier, all of this work pertains to the revenue-maximization setting. Ronen [46] considers the single-item auction setting in the *oracle model*, where one samples from the distribution conditioned on some players' values. He proposes the (1-) lookahead auction and shows that it achieves a $\frac{1}{2}$-approximation. [44] shows that the optimal (DSIC, IR) mechanism for the single-item auction can be computed efficiently with at most 2 players, and is *NP*-hard otherwise. Cai et al. [5] give a characterization of the optimal auction under certain settings. [47] considers interdependent types, which generalizes the correlated type-distribution setting, and develop an analog of Myerson's theory for certain such settings.

Various reductions from revenue-maximization to SWM are given in [6, 7, 5]. These reductions also apply to covering problems and the PayM objective, but they are incomparable to our results. These works focus on the (BIC, interim-IR) solution concept, which is a rather weak/liberal notion for correlated distributions. Most (but not all) of these consider independent players and additive valuations, and often require that the SWM-algorithm also work with negative values, which is a benign requirement for downwards-closed environments such as CAs but is quite problematic for covering problems when only has an approximation algorithm. [5] considers correlated players and obtains mechanisms having running time polynomial in the maximum support-size of the marginal distribution of a player, which could be substantially smaller than the support-size of the entire distribution. This savings can be traced to the use of the (BIC, interim-IR) notion which allows [5] to work with a compact description of the mechanism. It is unclear if these ideas are applicable when one considers robust-(BIC, IR) mechanisms. A very interesting open question is whether one can design robust-(BIC-in-expectation, IR) mechanisms having running time polynomial in the support-sizes of the marginal player distributions (as in [5, 17]).

3.2 Problem Definition and Preliminaries

The setup in Section 1.2 yields a *multidimensional covering mechanism-design problem* with *additive types*, where additivity is the property that if $c_i, c'_i \in C_i$, then the type $c_i + c'_i$ defined by $(c_i+c'_i)(\omega) = c_i(\omega)+c'_i(\omega)$ for all $\omega \in \Omega$, is also in C_i. It is possible to define more general multidimensional settings, but additive type spaces is a reasonable starting point to explore the multidimensional covering mechanism-design setting. (As noted earlier, there has been almost no work on designing polytime, near-optimal mechanisms for covering problems.)

The Bayesian setting. We consider *Bayesian* settings where there is an underlying publicly-known *discrete* and possibly *correlated* joint type-distribution on C from which the players' types are drawn. We consider the so-called explicit model, where the players' type distribution is explicitly specified. We use $\mathcal{D} \subseteq C$ to denote the support of the type distribution, and $\Pr_{\mathcal{D}}(c)$ to denote the probability of realization of $c \in C$. Also, we define $\mathcal{D}_i := \{c_i \in C_i : \exists c_{-i} \text{ s.t. } (c_i, c_{-i}) \in \mathcal{D}\}$, and $\mathcal{D}_{-i}$ to be $\{c_{-i} : \exists c_i \text{ s.t. } (c_i, c_{-i}) \in \mathcal{D}\}$.

Solution concepts. A mechanism sets up a game between the players, and the solution concept dictates certain desirable properties that this game should satisfy, so that one can reason about the outcome that results when rational are presented with a mechanism satisfying the solution concept. As mentioned earlier, the two chief properties that one seeks to capture relate to incentive compatibility (IC) and individual rationality (IR). Differences and subtleties arise in Bayesian settings depending on the stage at which we impose these properties and how robust we would like these properties to be with respect to the underlying type distribution.

Definition 3.2.1 *A mechanism $M = \big(\mathcal{A}, \{p_i\}\big)$ is* Bayesian incentive compatible *(BIC) and* interim IR *if for every player i and every $\bar{c}_i, c_i \in C_i$, we have $\mathrm{E}_{c_{-i}}[u_i(\bar{c}_i, c_{-i}; \bar{c}_i)|\bar{c}_i] \geq \mathrm{E}_{c_{-i}}[u_i(c_i, c_{-i}; \bar{c}_i)|\bar{c}_i]$ (BIC) and $\mathrm{E}_{c_{-i}}[u_i(\bar{c}_i, c_{-i}; \bar{c}_i)|\bar{c}_i] \geq 0$ (interim IR), where $\mathrm{E}_{c_{-i}}[.|\bar{c}_i]$ denotes the expectation over the other players' types conditioned on i's type being $\bar{c}_i$.*

As mentioned in the Introduction of this chapter, the (BIC, interim-IR) solution concept may yet lead to ex-post "regret", and is quite non-robust in the sense that the mechanism's IC and IR properties rely on having detailed knowledge of the type-distribution; consequently, in order to be confident that a BIC mechanism achieves its intended functionality, one must be confident about the "correctness" of the underlying distribution, and learning this information might entail significant cost. To remedy these weaknesses, we propose and investigate the following stronger IC and IR notions.

Definition 3.2.2 *A mechanism $M = \big(\mathcal{A}, \{p_i\}\big)$ is* robust BIC *and* robust IR*, if for every player i, every $\bar{c}_i, c_i \in C_i$, and every $c_{-i} \in \mathcal{D}_{-i}$, we have $u_i(\bar{c}_i, c_{-i}; \bar{c}_i) \geq u_i(c_i, c_{-i}; \bar{c}_i)$ (robust BIC) and $u_i(\bar{c}_i, c_{-i}; \bar{c}_i) \geq 0$ (robust IR).*

Robust (BIC, IR) ensures that participating truthfully in the mechanism is in the best interest of every player even at the ex-post stage when he knows the realized types of all players. Recall that we focus on *monopoly-free* settings where for every player i, there is some $\omega \in \Omega$ with $\omega_i = \emptyset$, to ensure that robust BIC and robust IR are compatible.

Notice that robust (BIC, IR) is subtly weaker than the notion of (DSIC, IR), wherein the IC and IR conditions of Definition 3.2.2 must hold for all $c_{-i} \in C_{-i}$, ensuring that truth-telling and participation are no-regret choices for a player even if the other players' reports are outside the support of the underlying type-distribution. We focus on robust BIC because it forms a suitable middle-ground between BIC and DSIC: it inherits the desirable robustness properties of DSIC, making it much more robust than BIC (and closer to a worst-case notion), and yet is flexible enough that one can devise polytime mechanisms satisfying this solution concept.

The above definitions are stated for a deterministic mechanism, but they have analogous extensions to a randomized mechanism M; the only change is that each $u_i(.)$ and $p_i(.)$ term is now replaced by the expected utility $\mathrm{E}_M[u_i(.)]$ and expected price $\mathrm{E}_M[p_i(.)]$ over the random coin tosses of M. We denote the analogous solution concept for a randomized mechanism by appending "in expectation" to the solution concept, e.g., a (BIC, interim IR)-in-expectation mechanism denotes a randomized mechanism whose expected utility satisfies the BIC and interim-IR requirements stated in Definition 3.2.1.

A robust-(BIC, IR)-in-expectation mechanism $M = \big(\mathcal{A}, \{p_i\}\big)$ can be easily modified so that the IR condition holds *with probability 1* (with respect to M's coin tosses) while the expected payment to a player (again over M's coin tosses) is unchanged: on input c, if $\mathcal{A}(c) = \omega \in \Omega$ with probability q, the new mechanism returns, with probability q, the allocation ω, and payment $c_i(\omega) \cdot \frac{\mathrm{E}_M[p_i(c)]}{\mathrm{E}_M[c_i(\omega)]}$ to each player i (where we take $0/0$ to be 0, so if $c_i(\omega) = 0$, the payment to i is 0). Thus, we obtain a mechanism whose expected utility satisfies the robust-BIC condition, and IR holds with probability 1 for all $\overline{c}_i \in C_i, c_{-i} \in \mathcal{D}_{-i}$, A similar transformation can be applied to a (DSIC, IR)-in-expectation mechanism.

Optimization problems. Our main consideration is to minimize the expected total payment of the mechanism. It is natural to also incorporate the mechanism-designer's cost into the objective. Define the *disutility* of a mechanism $M = \big(f, \{p_i\}\big)$ under input v to be $\sum_i p_i(v) + \kappa \cdot pub\big(f(v)\big)$, where $\kappa \geq 0$ is a scaling factor. Our objective is to devise a polynomial-time robust (BIC-in-expectation, IR)-mechanism with minimum expected disutility. Since most problems we consider have $pub(\omega) = 0$ for all feasible allocations, in which case disutility equals the total payment, abusing terminology slightly, we refer to the above mechanism-design problem as the *payment-minimization* (PayM) problem. (An exception is metric *uncapacitated facility location* (UFL), where players provide facilities and the underlying metric is public knowledge; here, $pub(\omega)$ is the total client-assignment cost of the solution ω.) We always use O^* to denote the expected disutility of an optimal mechanism for the PayM problem under consideration.

The following technical lemma will prove quite useful, since it allows us to restrict the domain to a bounded set, which is essential to achieve IR with finite prices. (For example, in the single-dimensional setting, the payment depends on the integral going to ∞ of a certain quantity, one needs boundedness of the support to ensure that this is well defined.) Note that such complications do not arise for packing problems. We state the lemma here but defer its proof to the end of Section 3.5 as we need Theorem 3.5.2 in the proof. Let $\mathbf{1}_{\mathcal{T}_i}$ be the $|\mathcal{T}_i|$-dimensional all 1s vector. Let $\mathcal{I}$ denote the input size.

Lemma 3.2.3 *We can efficiently compute an estimate* $m_i > \max_{c_i \in \mathcal{D}_i, v \in \mathcal{T}_i} c_{i,v}$ *with* $\log m_i = \mathsf{poly}(\mathcal{I})$ *for all i such that there is an optimal robust-(BIC-in-expectation, IR) mechanism* $M^* = \big(\mathcal{A}^*, \{p_i^*\}\big)$ *where* $\mathcal{A}^*(m_i\mathbf{1}_{\mathcal{T}_i}, c_{-i}) = \emptyset$ *with probability* 1 *(over the random choices of* M^**) for all i and all* $c_{-i} \in \mathcal{D}_{-i}$.

It is easy to obtain the stated estimates if we consider only *deterministic* mechanisms, but it turns out to be tricky to obtain this when one allows randomized mechanisms due to the artifact that a randomized mechanism may choose arbitrarily high-cost solutions as long as they are chosen with small enough probability. In the sequel, we set $\overline{\mathcal{D}}_i := \mathcal{D}_i \cup \{m_i\mathbf{1}_{\mathcal{T}_i}\}$ for all $i \in [n]$, and $\overline{\mathcal{D}} := \bigcup_i (\overline{\mathcal{D}}_i \times \mathcal{D}_{-i})$. Note that $|\overline{\mathcal{D}}| = O(n|\mathcal{D}|^2)$.

3.3 Overview of Our Construction

The starting point for our construction is the observation that the problem of designing an optimal robust-(BIC, IR)-in-expectation mechanism can be encoded via an LP (P). This was also observed by [17] in the context of the revenue-maximization problem for CAs, but the covering nature of the problem renders various techniques utilized successfully in the context of packing problems inapplicable, and therefore from here on our techniques diverge.

We show that an optimal solution to (P) can be computed given an optimal algorithm $\mathcal{A}$ for the CM problem since $\mathcal{A}$ can be used to obtain a separation oracle for the dual LP. Next, we prove that a feasible solution to (P) yields a robust-(BIC-in-expectation, IR) mechanism with no larger objective value.

For single-dimensional problems, we show that even LP-relative ρ-approximation algorithms for the CM problem can be utilized, as follows. We move to a relaxation of (P), where we replace the set of allocations with the feasible region of the CM-LP. This can be solved efficiently, since the separation oracle for the dual can be obtained by optimizing over the feasible region of CM-LP, which can be done efficiently! But now we need to work harder to "round" an optimal solution (x, p) to the relaxation of (P) and obtain a robust-(BIC-in-expectation, IR) mechanism. Here, we exploit the Lavi-Swamy [32] convex-decomposition procedure, using which we can show (roughly speaking) that we can decompose ρx into a convex combination of allocations. This allows us to obtain a robust-(BIC-in-expectation, IR) mechanism while blowing up the payment by a ρ-factor.

3.4 Differences with Respect to Packing Problems

Note that [17] obtain (DSIC-in-expectation, IR)-mechanisms, which is a subtly stronger notion than the robust-(BIC-in-expectation, IR) solution concept that our mechanisms satisfy. This difference arises due to the different nature of covering and packing problems. [17] also first obtains a robust-(BIC, IR)-in-expectation mechanism. The key difference is that for combinatorial auctions, one can show that *any* robust-(BIC, IR)-in-expectation mechanism—in particular, the one obtained from the optimal LP solution—can be converted into a (DSIC-in-expectation, IR) mechanism without any loss in expected revenue (see Section A.1). Intuitively, this works because one can focus on a single player by allocating no items to the other players. Clearly, one cannot mimic this for covering problems: dropping players may render the problem infeasible, and it is not clear how to extend an LP-solution to a (DSIC-in-expectation, IR) mechanism for covering problems. We suspect that there is a gap between the optimal expected total payments of robust-(BIC-in-expectation, IR) and (DSIC, IR) mechanisms; we leave this as an open problem. Due to this complication, we sacrifice a modicum of the IC, IR properties in favor of obtaining polytime near-optimal mechanisms and settle for the weaker, but still quite robust notion of robust (BIC-in-expectation, IR). We consider this to be a reasonable starting point for exploring mechanism-design solutions for covering problems, which leads to various interesting research directions.

A more-stunning aspect where covering and packing problems diverge can be seen when one considers the idea of a k-lookahead auction [46, 17]. This was used by [17] to convert their results in the explicit model to the oracle model introduced by [46]. This however fails spectacularly in the covering setting. One can show that even for single-item procurement auctions, dropping even a single player can lead to an arbitrarily large payment compared to the optimum (see Section 3.9).

3.5 LP-relaxations for the Payment-Minimization Problem

The starting point for our results is the LP (P) that essentially encodes the payment-minimization problem. Throughout, we use i to index players, c to index type-profiles in $\overline{\mathcal{D}}$, and ω to index Ω. We use variables $x_{c,\omega}$ to denote the probability of choosing ω, and $p_{i,c}$ to denote the expected payment to player i, for input c. For $c \in \overline{\mathcal{D}}$, let $\Omega(c) = \Omega$ if $c \in \bigcup_i (\mathcal{D}_i \times \mathcal{D}_{-i})$, and otherwise if $c = (m_i \mathbf{1}_{\mathcal{T}_i}, c_{-i})$, let $\Omega(c) = \{\omega \in \Omega : \omega_i = \emptyset\}$ (which is non-empty since we are in a monopoly-free setting).

$$
\begin{aligned}
\min \quad & \sum_{c\in\mathcal{D}} \Pr_{\mathcal{D}}(c)\Big(\sum_i p_{i,c} + \kappa \sum_\omega x_{c,\omega} pub(\omega)\Big) && && \text{(P)}\\
\text{s.t.} \quad & \sum_\omega x_{c,\omega} = 1 && \forall c \in \overline{\mathcal{D}} && (3.1)\\
& p_{i,(c_i,c_{-i})} - \sum_\omega c_i(\omega) x_{(c_i,c_{-i}),\omega} \geq && && \\
& \qquad p_{i,(c'_i,c_{-i})} - \sum_\omega c_i(\omega) x_{(c'_i,c_{-i}),\omega} && \forall i, c_i, c'_i \in \overline{\mathcal{D}}_i, c_{-i} \in \mathcal{D}_{-i} && (3.2)\\
& \qquad p_{i,(c_i,c_{-i})} - \sum_\omega c_i(\omega) x_{(c_i,c_{-i}),\omega} \geq 0 && \forall i, c_i \in \overline{\mathcal{D}}_i, c_{-i} \in \mathcal{D}_{-i} && (3.3)\\
& \qquad p, x \geq 0, \quad x_{c,\omega} = 0 && \forall c, \omega \notin \Omega(c). && (3.4)
\end{aligned}
$$

(3.1) encodes that an allocation is chosen for every $c \in \overline{\mathcal{D}}$, and (3.2) and (3.3) encode the robust BIC and robust IR conditions respectively. Lemma 3.2.3 ensures that (P) correctly encodes PayM, so that $OPT := OPT_{\mathrm{P}}$ is a lower bound on the expected disutility of an optimal mechanism.

Our results are obtained by computing an optimal solution to (P), or a further relaxation of it, and translating this to a near-optimal robust (BIC-in-expectation, IR) mechanism. Both steps come with their own challenges. Except in very simple settings (such as single-item procurement auctions), $|\Omega|$ is typically exponential in the input size, and therefore it is not clear how to solve (P) efficiently. We therefore consider the dual LP (D), which has variables γ_c, $y_{i,(c_i,c_{-i}),c'_i}$ and $\beta_{i,(c_i,c_{-i})}$ corresponding to (3.1), (3.2) and (3.3) respectively.

$$
\begin{aligned}
\max \quad & \sum_c \gamma_c && && \text{(D)}\\
\text{s.t.} \quad & \sum_{i: c\in\overline{\mathcal{D}}_i\times\mathcal{D}_{-i}} \Big(\sum_{c'_i\in\overline{\mathcal{D}}_i} \big(c_i(\omega) y_{i,(c_i,c_{-i}),c'_i} - c'_i(\omega) y_{i,(c'_i,c_{-i}),c_i}\big) + c_i(\omega)\beta_{i,c}\Big) && && \\
& \qquad + \kappa\cdot \Pr_{\mathcal{D}}(c) pub(\omega) \ \geq \gamma_c && \forall c \in \overline{\mathcal{D}}, \omega \in \Omega(c) && (3.5)\\
& \sum_{c'_i\in\overline{\mathcal{D}}_i} \big(y_{i,(c_i,c_{-i}),c'_i} - y_{i,(c'_i,c_{-i}),c_i}\big) + \beta_{i,c_i,c_{-i}} \leq \Pr_{\mathcal{D}}(c) && \forall i, c_i \in \overline{\mathcal{D}}_i, c_{-i} \in \mathcal{D}_{-i} && (3.6)\\
& \qquad y, \beta \geq 0. && && (3.7)
\end{aligned}
$$

With additive types, the separation problem for constraints (3.5) amounts to determining if the optimal value of the CM problem defined by a certain input with possibly negative costs, is at least γ_c. Hence, an optimal algorithm for the CM problem can be used to solve (D), and hence, (P) efficiently.

Theorem 3.5.1 *With additive types, one can efficiently solve* (P) *given an optimal algorithm for the CM problem.*

Proof : Let $\mathcal{A}$ be an optimal algorithm for the CM problem. First, observe that we can use $\mathcal{A}$ to find a solution that minimizes $\sum_i c_i(\omega) + \kappa \cdot pub(\omega)$ for any $\kappa \geq 0$, even for an input $c = \{c_{i,v}\}_{i,v\in\mathcal{T}_i}$ where some of the $c_{i,v}$s are negative. Let $A_i = \{v \in \mathcal{T}_i : c_{i,v} < 0\}$. Clearly, if ω^* is an optimal solution, then $A_i \subseteq \omega^*_i$ (since $pub(.)$ does not increase upon adding covering objects). Define $c^+_{i,v} := \max(0, c_{i,v})$ and $c^+_i := \{c^+_{i,v}\}_{v\in\mathcal{T}_i}$.

Let $\Gamma = \frac{1}{\kappa}$ if $\kappa > 0$; otherwise let $\Gamma = NU$, where U is a strict upper bound on $\max_{\omega\in\Omega} pub(\omega)$ and N is an integer such that all the $c^+_{i,v}$s are integer multiples of $\frac{1}{N}$. Note that for any $\omega, \omega' \in \Omega$, if $\sum_i c^+_i(\omega) - \sum_i c^+_i(\omega')$ is non-zero, then its absolute value is at least $\frac{1}{N}$. Also, U and N may be efficiently computed (for rational data) and $\log(NU)$ is polynomially bounded. Let $(S_1, \dots, S_n)$ be the solution returned by $\mathcal{A}$ for the CM problem on the input where all the $c^+_{i,v}$s are scaled by Γ. The choice of Γ ensures that

$$\sum_i c^+_i(S_i) + \kappa \cdot pub\big(\bigcup_i S_i\big) \leq \sum_i c^+_i(\omega^*_i) + \kappa \cdot pub(\omega^*) = \sum_i \big(c_i(\omega^*_i) - c_i(A_i)\big) + \kappa \cdot pub(\omega^*).$$

So setting $\omega_i = A_i \cup S_i$ for every i yields a feasible solution such that $\sum_i c_i(\omega) + \kappa \cdot pub(\omega) \leq \sum_i c_i(\omega^*) + \kappa \cdot pub(\omega^*)$; hence ω is an optimal solution.

Given a dual solution (y, β, γ), we can easily check if (3.6), (3.7) hold. Fix $c \in \overline{\mathcal{D}}$ and player i. Notice that for every $\omega \in \Omega$, we have $\sum_{c'_i\in\overline{\mathcal{D}}_i}\big(c_i(\omega)y_{i,(c_i,c_{-i}),c'_i} - c'_i(\omega)y_{i,(c'_i,c_{-i}),c_i}\big) + c_i(\omega)\beta_{i,c} = \theta^c_i(\omega) := \sum_{v\in\omega_i}\theta^c_{i,v}$, where $\theta^c_{i,v} = \sum_{c'_i\in\overline{\mathcal{D}}_i}\big(c_{i,v}y_{i,(c_i,c_{-i}),c'_i} - c'_{i,v}y_{i,(c'_i,c_{-i}),c_i} + c_{i,v}\beta_{i,c}$.

Define $I := \{i : c \in \overline{\mathcal{D}}_i \times \mathcal{D}_{-i}\}$. Constraints (3.5) for c can then be written as

$\min_{\omega\in\Omega(c)}\left(\sum_{i\in I}\theta_i^c(\omega)+\kappa\cdot\Pr_{\mathcal{D}}(c)pub(\omega)\right)\geq\gamma_c$. Define $\tilde{c}$ as follows:

$$\tilde{c}_{i,v}=\begin{cases}\gamma_c+1 & \text{if } c_i=m_i\mathbf{1}_{\mathcal{T}_i},\\ \theta_{i,v}^c & \text{if } i\in I,\ c_i\in\mathcal{D}_i,\\ 0 & \text{otherwise.}\end{cases}$$

It is easy to see that (3.5) holds for c iff $\min_{\omega\in\Omega}\left(\sum_i\tilde{c}_i(\omega)+\kappa\cdot\Pr_{\mathcal{D}}(c)pub(\omega)\right)$, which can be computed using $\mathcal{A}$, is at least γ_c. Thus, we can use the ellipsoid method to solve (D). This also yields a compact dual consisting of constraints (3.6), (3.7) and the polynomially-many (3.5) constraints that were returned by the separation oracle during the execution of the ellipsoid method, whose optimal value is OPT_{D}. The dual of this compact dual is an LP of the same form as (P) but with polynomially many $x_{c,\omega}$-variables; solving this yields an optimal solution to (P). ■

Complementing Theorem 3.5.1, we argue that a feasible solution (x,p) to (P) can be "rounded" to a robust-(BIC-in-expectation, IR) mechanism having expected disutility at most the value of (x,p) (Theorem 3.5.2). Combining this with Theorem 3.5.1 yields the corollary that an optimal algorithm for the CM problem can be used to obtain an optimal mechanism for the PayM problem (Corollary 3.5.3).

Theorem 3.5.2 *We can extend a feasible solution* (x,p) *to* (P) *to a robust-(BIC-in-expectation, IR) mechanism with expected disutility* $\sum_c\Pr_{\mathcal{D}}(c)\left(\sum_i p_{i,c}+\kappa\sum_\omega x_{c,\omega}pub(\omega)\right)$.

Proof : Let $\Omega'=\{\omega : x_{c,\omega}>0 \text{ for some } c\in\overline{\mathcal{D}}\}$. We use x_c to denote the vector $\{x_{c,\omega}\}_{\omega\in\Omega'}$. Consider a player i, $c_{-i}\in\mathcal{D}_{-i}$, and $c_i,c_i'\in\overline{\mathcal{D}}_i$. Note that (3.2) implies that if $x_{(c_i,c_{-i})}=x_{(c_i',c_{-i})}$, then $p_{i,(c_i,c_{-i})}=p_{i,(c_i',c_{-i})}$. For $c_{-i}\in\mathcal{D}_{-i}$, define $R(i,c_{-i})=\left\{x_{(c_i,c_{-i})} : (c_i,c_{-i})\in\overline{\mathcal{D}}\right\}$, and for $y=x_{(c_i,c_{-i})}\in R(i,c_{-i})$ define $p_{i,y}$ to be $p_{i,(c_i,c_{-i})}$ (which is well defined by the above argument).

We now define the randomized mechanism $M=\left(\mathcal{A},\{q_i\}\right)$, where $\mathcal{A}(c)$ and $q_i(c)$ denote respectively the probability distribution over allocations and the expected payment to player i, on input c. We sometimes view $\mathcal{A}(c)$ equivalently as the random variable specifying the allocation chosen for input c. Fix an allocation $\omega_0\in\Omega$. Consider an input c. If

$c \in \overline{\mathcal{D}}$, we set $\mathcal{A}(c) = x(c)$, and $q_i(c) = p_{i,c}$ for all i. So consider $c \notin \overline{\mathcal{D}}$. If there is no i such that $c_{-i} \in \mathcal{D}_{-i}$, we simply set $\mathcal{A}(c) = \omega_0$, $q_i(c) = c_i(\omega_0)$ for all i; such a c does not figure in the robust (BIC, IR) conditions. Otherwise there is a unique i such that $c_{-i} \in \mathcal{D}_{-i}$, $c_i \in C_i \setminus \overline{\mathcal{D}}_i$. Set $\mathcal{A}(c) = \arg\max_{y \in R(i,c_{-i})}\big(p_{i,y} - \sum_{\omega\in\Omega'} c_i(\omega)y_\omega\big)$ and $q_j(c) = p_{j,\mathcal{A}(c)}$ for all players j. Note that (c_i, c_{-i}) figures in (3.2) *only* for player i. Crucially, note that since $y = x_{(m_i,c_{-i})} \in R(i, c_{-i})$ and $\sum_{\omega\in\Omega} c_i(\omega)y_\omega = 0$ by definition, we always have $q_i(c) - \mathrm{E}_{\mathcal{A}}[c_i(\mathcal{A}(c))] \geq 0$. Thus, by definition, and by (3.2), we have ensured that M is robust (BIC, IR)-in-expectation and its expected disutility is exactly the value of (x, p). This can be modified so that IR holds with probability 1.

This rounding is efficient if $\sum_{\omega\in\Omega'} c_i(\omega)x_{c,\omega}$ can be calculated efficiently. This is clearly true if $|\Omega'|$ is polynomially bounded, but it could hold under weaker conditions as well. ■

Corollary 3.5.3 *Given an optimal algorithm for the CM problem, we can obtain an optimal robust-(BIC-in-expectation, IR) mechanism for the PayM problem in multidimensional settings with additive types.*

As mentioned earlier, the CM problem is however often *NP*-hard (e.g., for vertex cover), and we would like to be able to exploit approximation algorithms for the CM problem to obtain near-optimal mechanisms. The usual approach is to use an approximation algorithm to "approximately" separate over constraints (3.5). However, this does not work here since the CM problem that one needs to solve in the separation problem involves negative costs, which renders the usual notion of approximation meaningless. Instead, if the CM problem admits a certain type of LP-relaxation (C-P), then we argue that one can solve a relaxation of (P) where the allocation-set is the set of extreme points of (C-P) (Theorem 3.5.4). For single-dimensional problems (Section 3.6), we leverage this to obtain strong and far-reaching results. We show that a ρ-approximation algorithm relative to (C-P) can be used to "round" the optimal solution to this relaxation to a robust-(BIC-in-expectation, IR)-mechanism losing a ρ-factor in the disutility (Theorem 3.6.2). Thus, we obtain near-optimal mechanisms for a variety of single-dimensional problems.

Suppose that the CM problem admits an LP-relaxation of the following form, where

$c = \{c_{i,v}\}_{i \in [n], v \in \mathcal{T}_i}$ is the input type-profile.

$$\min \quad c^T x + d^T z \qquad \text{s.t.} \qquad Ax + Bz \geq b, \quad x, z \geq 0. \tag{C-P}$$

Intuitively x encodes the allocation chosen, and $d^T z$ encodes $pub(.)$. For $x \geq 0$, define $z(x) := \arg\min\{d^T z : (x,z) \text{ is feasible to (C-P)}\}$; if there is no z such that (x,z) is feasible to (C-P), set $z(x) := \perp$. Define $\Omega_{\text{LP}} := \{x : z(x) \neq \perp, \quad 0 \leq x_{i,v} \leq 1 \ \forall i, v \in \mathcal{T}_i\}$. We require that: (a) a $\{0,1\}$-vector x is in Ω_{LP} iff it is the characteristic vector of an allocation $\omega \in \Omega$, and in this case, we have $d^T z(x) = pub(\omega)$; (b) $A \geq 0$; (c) for any input $c \geq 0$ to the covering problem, (C-P) is not unbounded, and if it has an optimal solution, it has one where $x \in \Omega_{\text{LP}}$; (d) for any c, we can efficiently find an optimal solution to (C-P) or detect that it is unbounded or infeasible.

We extend the type c_i of each player i and pub to assign values also to points in Ω_{LP}: define $c_i(x) = \sum_{v \in \mathcal{T}_i} c_{i,v} x_{i,v}$ and $pub(x) = d^T z(x)$ for $x \in \Omega_{\text{LP}}$. Let Ω_{ext} denote the finite set of extreme points of Ω_{LP}. Condition (a) ensures that Ω_{ext} contains the characteristic vectors of all feasible allocations. Let (P') denote the relaxation of (P), where we replace the set of feasible allocations Ω with Ω_{ext} (so ω indexes Ω_{ext} now), and for $c \in \overline{\mathcal{D}}$ with $c_i = m_i \mathbf{1}(\mathcal{T}_i)$, we now define $\Omega(c) := \{\omega \in \Omega_{\text{ext}} : \omega_{i,v} = 0 \ \forall v \in \mathcal{T}_i\}$. Since one can optimize efficiently over Ω_{LP}, and hence Ω_{ext}, even for negative type-profiles, we have the following.

Theorem 3.5.4 *We can efficiently compute an optimal solution to* (P')*.*

Now, given Theorem 3.5.2, we have the tool to prove Lemma 3.2.3.

Proof of Lemma 3.2.3 : Consider the following LP, which is the same as (P) except that we only consider $c \in \bigcup_i (\mathcal{D}_i \times \mathcal{D}_{-i})$.

$$
\begin{aligned}
\min \quad & \sum_{c \in \mathcal{D}} \Pr_{\mathcal{D}}(c) \Big(\sum_i q_{i,c} + \kappa \sum_\omega x_{c,\omega} pub(\omega) \Big) && && \text{(LP)} \\
\text{s.t.} \quad & \sum_\omega x_{c,\omega} = 1 && \forall c \in \bigcup_i (\mathcal{D}_i \times \mathcal{D}_{-i}) && (3.8) \\
& q_{i,(c_i,c_{-i})} - \sum_\omega c_i(\omega) x_{(c_i,c_{-i}),\omega} \geq \\
& \qquad q_{i,(c'_i,c_{-i})} - \sum_\omega c_i(\omega) x_{(c'_i,c_{-i}),\omega} && \forall i, c_i, c'_i \in \mathcal{D}_i, c_{-i} \in \mathcal{D}_{-i} && (3.9) \\
& \qquad q_{i,(c_i,c_{-i})} - \sum_\omega c_i(\omega) x_{(c_i,c_{-i}),\omega} \geq 0 && \forall i, c_i \in \overline{\mathcal{D}}_i, c_{-i} \in \mathcal{D}_{-i} && (3.10) \\
& \qquad q, x \geq 0. && && (3.11)
\end{aligned}
$$

Let $M = \big(\mathcal{A}, \{p_i\}\big)$ be an optimal mechanism. Define O^* to be the expected disutility of M. Then, M naturally yields a feasible solution (x, q) to (3.8)–(3.11) of objective value O^*, where $x_{c,\omega} = \Pr_M[\mathcal{A}(c) = \omega]$ and $q_{i,c} = \mathrm{E}_M[p_i(c)]$. Let $(\hat{x}, \hat{q})$ be an *optimal basic solution* to (LP). Then, for some N such that $\log N$ is polynomially bounded in the input size $\mathcal{I}$, we can say that the values of all variables are integer multiples of $\frac{1}{N}$, and $\log(N\hat{x}_{c,\omega}), \log(N\hat{q}_{i,c}) = \mathsf{poly}(\mathcal{I})$ for all i, $c \in \bigcup_i (\mathcal{D}_i \times \mathcal{D}_{-i})$, ω.

First, we claim that we may assume that for every $i, c_i \in \mathcal{D}_i, c_{-i} \in \mathcal{D}_{-i}$, if whenever $\hat{x}_{c,\omega} > 0$ we have $\omega_i = \emptyset$ (where $c = (c_i, c_{-i})$), then $\hat{q}_{i,c} = 0$. If not, then (3.9) implies that $\hat{q}_{i,(\tilde{c}_i,c_{-i})} - \sum_\omega \tilde{c}_i(\omega) x_{(\tilde{c}_i,c_{-i}),\omega} \geq \hat{q}_{i,c}$ for all $\tilde{c}_i \in \mathcal{D}_i$ and decreasing $\hat{q}_{i,(\tilde{c}_i,c_{-i})}$ by $\hat{q}_{i,c}$ for all $\tilde{c}_i \in \mathcal{D}_i$ continues to satisfy (3.9)–(3.11).

Set $m_i := \max\big(2 \sum_{i,v \in \mathcal{T}_i} \max_{c_i \in \mathcal{D}_i} c_{i,v}, N \sum_{i,c} \hat{q}_{i,c}\big)$ for all i. So $\log m_i = \mathsf{poly}(\mathcal{I})$. Recall that $\overline{\mathcal{D}}_i := \mathcal{D}_i \cup \{m_i \mathbf{1}_{\mathcal{T}_i}\}$ for all $i \in [n]$, and $\overline{\mathcal{D}} := \bigcup_i (\overline{\mathcal{D}}_i \times \mathcal{D}_{-i})$.

Now we extend $(\hat{x}, \hat{q})$ to $(\tilde{x}, \tilde{q})$ that assigns values also to type-profiles in $\overline{\mathcal{D}} \setminus \bigcup_i (\mathcal{D}_i \times \mathcal{D}_{-i})$ so that constraints (3.8)–(3.11) hold for all i, $c_i, c'_i \in \overline{\mathcal{D}}_i$, $c_{-i} \in \mathcal{D}_{-i}$. First set $\tilde{x}_{c,\omega} = \hat{x}_{c,\omega}$, $\tilde{q}_{i,c} = \hat{q}_{i,c}$ for all i, ω, $c \in \bigcup_i (\mathcal{D}_i \times \mathcal{D}_{-i})$. Consider $c \in \overline{\mathcal{D}} \setminus \bigcup_i (\mathcal{D}_i \times \mathcal{D}_{-i})$, and let i be such that $c_i = m_i \mathbf{1}_{\mathcal{T}_i}$ (there is exactly one such i). We "run" VCG on c considering only the cost incurred by the players. That is, we set $\tilde{x}_{c,\omega} = 1$ for $\omega = \omega(c) := \arg\min_{\omega \in \Omega} \sum_i c_i(\omega)$ and pay $\tilde{q}_{i,c} = \min_{\omega \in \Omega : \omega_i = \emptyset} \sum_j c_j(\omega) - \sum_{j \neq i} c_j\big(\omega(c)\big)$ to each player i. Note that the choice of m_i ensures that $\omega(c)_i = \emptyset$ and hence, $\tilde{q}_{i,c} = 0$.

We claim that this extension satisfies (3.8)–(3.11) for all i, $c_i, c'_i \in \overline{\mathcal{D}}_i$, $c_{-i} \in \mathcal{D}_{-i}$. Fix i, $c_i, c'_i \in \overline{\mathcal{D}}_i$, $c_{-i} \in \mathcal{D}_{-i}$. It is clear that (3.8), (3.11) hold. If $c_i \in \mathcal{D}_i$, then (3.10) clearly holds; if $c_i = m_i \mathbf{1}_{\mathcal{T}_i}$, then it again holds since $\tilde{x}_{c,\omega} = 1$ for $\omega = \omega(c)$ and $\omega(c)_i = \emptyset$. To verify (3.9), we consider four cases. If $c_i, c'_i \in \mathcal{D}_i$, then (3.9) holds since $(\tilde{x}, \tilde{q})$ extends $(\hat{x}, \hat{q})$. If $c_i = c'_i = m_i \mathbf{1}_{\mathcal{T}_i}$, then (3.9) trivially holds. If $c_i \in \mathcal{D}_i$, $c'_i = m_i \mathbf{1}_{\mathcal{T}_i}$, then (3.9) holds since the RHS of (3.9) is 0 (as $\tilde{x}_{(c'_i,c_{-i}),\omega(c'_i,c_{-i})} = 1$ and $\omega(c'_i, c_{-i})_i = \emptyset$). We are left with the case $c_i = m_i \mathbf{1}_{\mathcal{T}_i}$ and $c'_i \in \mathcal{D}_i$. If whenever $\tilde{x}_{(c'_i,c_{-i}),\omega} = \hat{x}_{(c'_i,c_{-i}),\omega} > 0$ we have $\omega_i = \emptyset$, then we also have $\tilde{q}_{i,(c'_i,c_{-i})} = \hat{q}_{i,(c'_i,c_{-i})} = 0$ by our earlier claim, so the RHS of (3.9) is 0, and (3.9) holds. Otherwise, we have $\sum_\omega c_i(\omega)\tilde{x}_{(c'_i,c_{-i}),\omega} \geq \frac{m_i}{N} \geq \hat{q}_{i,(c'_i,c_{-i})}$, so the RHS of (3.9) is at most 0, and (3.9) holds.

Thus, we have shown that $(\tilde{x}, \tilde{q})$ is a feasible solution to (P). Now we can apply Theorem 3.5.2 to extend $(\tilde{x}, \tilde{q})$ and obtain a robust-(BIC-in-expectation, IR) mechanism M^* whose expected disutility is at most $\sum_c \Pr_{\mathcal{D}}(c)\big(\sum_i \tilde{q}_{i,c} + \sum_\omega \tilde{x}_{c,\omega} pub(\omega)\big) \leq O^*$. Since $\tilde{x}_{(m_i \mathbf{1}_{\mathcal{T}_i}, c_{-i}),\omega} > 0$ implies that $\omega_i = \emptyset$ for all i, M^* satisfies the required conditions. ■

3.6 Single-Dimensional Problems

Corollary 3.5.3 immediately yields results for certain single-dimensional problems (see Table 3.1), most notably, *single-item procurement auctions*. We now substantially expand the scope of PayM problems for which one can obtain near-optimal mechanisms by showing how to leverage "LP-relative" approximation algorithms for the CM problem. Suppose that the CM problem can be encoded as (C-P). An *LP-relative ρ-approximation algorithm* for the CM problem is a polytime algorithm that for any input $c \geq 0$ to the covering problem, returns a $\{0,1\}$-vector $x \in \Omega_{\text{LP}}$ such that $c^T x + d^T z(x) \leq \rho OPT_{\text{C-P}}$. Using the convex-decomposition procedure in [32] (see Section 5.1 of [32]), one can show the following.

Lemma 3.6.1 *Let $x \in \Omega_{\text{LP}}$. Given an LP-relative ρ-approximation algorithm for the CM problem, $\mathcal{A}$, one can efficiently obtain $(\lambda^{(1)}, x^{(1)}), \ldots, (\lambda^{(k)}, x^{(k)})$, where $\sum_\ell \lambda^{(\ell)} = 1, \lambda \geq 0$, and $x^{(\ell)}$ is a $\{0,1\}$-vector in Ω_{LP} for all ℓ, such that $\sum_\ell \lambda^{(\ell)} x^{(\ell)}_{i,v} = \min(\rho x_{i,v}, 1)$ for all $i, v \in \mathcal{T}_i$, and $\sum_\ell \lambda^{(\ell)} d^T z(x^{(\ell)}) \leq \rho d^T z(x)$.*

Proof : It suffices to show that the LP (Q) can be solved in polytime and its optimal value is 1. Throughout, we use ℓ to index $\{0,1\}$ vectors in Ω_{LP}. (Recall that these correspond to feasible allocations.)

$$
\begin{aligned}
\max \quad & \sum_\ell \lambda^{(\ell)} && \text{(Q)}\\
\text{s.t.} \quad & \sum_\ell \lambda^{(\ell)} x^{(\ell)}_{i,v} = \min(\rho x_{i,v},1) \quad \forall i, v \in \mathcal{T}_i && (3.12)\\
& \sum_\ell \lambda^{(\ell)} d^T\big(z(x^{(\ell)})\big) \le \rho d^T z(x) && (3.13)\\
& \sum_\ell \lambda^{(\ell)} \le 1 && (3.14)\\
& \lambda \ge 0.
\end{aligned}
$$

$$
\begin{aligned}
\min \quad & \sum_{i,v\in\mathcal{T}_i} \min(\rho x_{i,v},1)\alpha_{i,v} + \rho d^T z(x)\cdot\beta + \theta && \text{(R)}\\
\text{s.t.} \quad & \sum_{i,v\in\mathcal{T}_i} x^{(\ell)}_{i,v}\alpha_{i,v} + d^T\big(z(x^{(\ell)})\big)\beta + \theta \ge 1 \quad \forall \ell && (3.15)\\
& \beta, \theta \ge 0.
\end{aligned}
$$

Here the α_ℓs, β and θ are the dual variables corresponding to constraints (3.12), (3.13), and (3.14) respectively. Clearly, $OPT_{(\text{D})} \le 1$ since $\theta = 1$, $\alpha_{i,v} = 0 = \beta$ for all i, v is a feasible dual solution.

Suppose $(\hat{\alpha}, \hat{\beta}, \hat{\theta})$ is a feasible dual solution of value less than 1. Set $\tilde{\alpha}_{i,v} = \hat{\alpha}_{i,v}$ if $\hat{\alpha}_{i,v} \ge 0$ and $\rho x_{i,v} \le 1$, and $\tilde{\alpha}_{i,v} = 0$ otherwise. Let $\Gamma = \frac{1}{\hat{\beta}}$ if $\hat{\beta} > 0$ and equal to $2Nd^T z$ otherwise, where N is is such that for all $\{0,1\}$-vectors $x^{(\ell)} \in \Omega_{\text{LP}}$, we have that $c^T x^{(\ell)} > c^T x$ implies $c^T x^{(\ell)} \ge c^T x + \frac{1}{N}$. Note that we can choose N so that its size is $\mathsf{poly}(\mathcal{I}, \text{size of } x)$. Consider the CM problem defined by the input $\Gamma\tilde{\alpha}$. Running $\mathcal{A}$ on this input, we obtain a $\{0,1\}$-vector $x^{(\ell)} \in \Omega_{\text{LP}}$ whose total cost is at most ρ times the cost of the fractional solution $\big(x, z(x)\big)$. This translates to

$$\sum_{i,v} x^{(\ell)}_{i,v}\tilde{\alpha}_{i,v} + d^T\big(z(x^{(\ell)})\big)\hat{\beta} \le \rho\Big(\sum_{i,v} x_{i,v}\tilde{\alpha}_{i,v} + d^T z(x)\cdot\hat{\beta}\Big). \qquad (3.16)$$

Now augment $x^{(\ell)}$ to the following $\{0,1\}$-vector $\hat{x}$: set $\hat{x}_{i,v} = 1$ if $\rho x_{i,v} > 1$ or $\hat{\alpha}_{i,v} < 0$, and $x^{(\ell)}_{i,v}$ otherwise. Then $\hat{x}$ is the characteristic vector of a feasible allocation, since we have only added covering objects to the allocation corresponding to $x^{(\ell)}$; hence $\hat{x} \in \Omega_{\text{LP}}$. We have $d^T z(\hat{x}) = pub(\hat{x}) \leq pub(x^{(\ell)}) = d^T\big(z(x^{\ell})\big)$ and

$$\sum_{i,v} \hat{x}_{i,v}\hat{\alpha}_{i,v} = \sum_{\substack{i,v:\rho x_{i,v}>1 \\ \text{or } \hat{\alpha}_{i,v}<0}} \hat{\alpha}_{i,v} + \sum_{i,v} x^{(\ell)}_{i,v}\tilde{\alpha}_{i,v} \leq \sum_{\substack{i,v:\rho x_{i,v}>1 \\ \text{or } \hat{\alpha}_{i,v}<0}} \min(\rho x_{i,v},1)\hat{\alpha}_{i,v} + \sum_{i,v} x^{(\ell)}_{i,v}\tilde{\alpha}_{i,v}.$$

Combined with (3.16), this shows that

$$\begin{aligned}\sum_{i,v} \hat{x}_{i,v}\hat{\alpha}_{i,v} + d^T z(\hat{x})\hat{\beta} &\leq \sum_{\substack{i,v:\rho x_{i,v}>1 \\ \text{or } \hat{\alpha}_{i,v}<0}} \min(\rho x_{i,v},1)\hat{\alpha}_{i,v} + \sum_{\substack{i,v:\tilde{\alpha}_{i,v}>0 \\ \text{or } \rho x_{i,v}\leq 1}} \rho x_{i,v}\tilde{\alpha}_{i,v} + \rho d^T z(x)\cdot\hat{\beta} \\ &= \sum_{i,v} \min(\rho x_{i,v},1)\hat{\alpha}_{i,v} + \rho d^T z(x)\cdot\hat{\beta} < 1 - \hat{\theta}\end{aligned}$$

which contradicts that $(\hat{\alpha}, \hat{\beta}, \hat{\theta})$ is feasible to (R). Hence, $OPT_{\text{(Q)}} = OPT_{\text{(R)}} = 1$.

Thus, we can add the constraint $\sum_{i,v\in\mathcal{T}_i} \min(\rho x_{i,v},1)\alpha_{i,v} + \rho d^T z(x)\cdot\beta + \theta \leq 1$ to (R) without altering anything. If we solve the resulting LP using the ellipsoid method, and take the inequalities corresponding to the violated inequalities (3.15) found by $\mathcal{A}$ during the ellipsoid method, then we obtain a compact LP with only a polynomial number of constraints that is equivalent to (R). The dual of this compact LP yields an LP equivalent to (Q) with a polynomial number of $\lambda^{(\ell)}$ variables which we can solve to obtain the desired convex decomposition. ■

Theorem 3.6.2 *Given an LP-relative ρ-approximation algorithm for the CM problem, one can obtain a polytime ρ-approximation robust-(BIC-in-expectation, IR) mechanism for the PayM problem.*

Proof : We solve (P') to obtain an optimal solution (x,p). Since $|\mathcal{T}_i| = 1$ for all i, it will be convenient to view $\omega \in \Omega_{\text{LP}}$ as a vector $\{\omega_i\}_{i\in[n]}$, where $w_i \equiv \omega_{i,v}$ for the single covering object $v \in \mathcal{T}_i$. Fix $c \in \overline{\mathcal{D}}$. Define $y_c = \sum_{\omega\in\Omega_{\text{ext}}} x_{c,\omega}\omega$ (which can be efficiently computed). Then, $\sum_{\omega\in\Omega_{\text{ext}}} c_i(\omega)x_{c,\omega} = c_i y_{c,i}$ and $d^T z(y) = \sum_{\omega\in\Omega_{\text{ext}}} pub(\omega)x_{c,\omega}$. By Lemma 3.6.1, we can efficiently find a point $\tilde{y}_c = \sum_{\omega\in\Omega} \tilde{x}_{c,\omega}\omega$, where $\tilde{x}_c \geq 0, \sum_{w\in\Omega} \tilde{x}_{c,\omega} = 1$, in the convex hull of the $\{0,1\}$-vectors in Ω_{LP} such that $\tilde{y}_{c,i} = \min(\rho y_{c,i}, 1)$ for all i, and $\sum_{w\in\Omega} \tilde{x}_{c,\omega} pub(\omega) \leq \rho d^T z(y)$.

We now argue that one can obtain payments $\{q_{i,c}\}$ such that $(\tilde{x}, q)$ is feasible to (P) and $q_{i,c} \leq \rho p_{i,c}$ for all $i, c \in \overline{\mathcal{D}}$. Thus, the value of $(\tilde{x}, q)$ is at most ρ times the value of (x, p). Applying Theorem 3.5.2 to $(\tilde{x}, q)$ yields the desired result.

Fix i and $c_{-i} \in \mathcal{D}_{-i}$. Constraints (3.4) and (3.2) ensure that $y_{(m_i,c_{-i}),i} = 0$, and $y_{(c_i,c_{-i}),i} \geq y_{(c'_i,c_{-i}),i}$ for all $c_i, c'_i \in \overline{\mathcal{D}}_i$ s.t. $c_i < c'_i$. Hence, $\tilde{y}_{(m_i,c_{-i})mi} = 0$, $\tilde{y}_{(c_i,c_{-i}),i} \geq \tilde{y}_{(c'_i,c_{-i}),i}$ for $c_i, c'_i \in \overline{\mathcal{D}}_i$, $c_i > c'_i$. Define $q_{i,(m_i,c_{-i})} = 0$. Let $0 \leq c_i^1 < c_i^2 < \ldots < c_i^{k_i}$ be the values in $\mathcal{D}_i$. For $c_i = c_i^\ell$, define

$$q_{i,(c_i,c_{-i})} = c_i \tilde{y}_{(c_i,c_{-i}),i} + \sum_{t=\ell+1}^{k_i} (c_i^t - c_i^{t-1}) \tilde{y}_{(c_i^t,c_{-i}),i}.$$

Since $\sum_{\omega \in \Omega} c_i(\omega) \tilde{x}_{(c_i,c_{-i}),\omega} = c_i \tilde{y}_{(c_i,c_{-i}),i}$, (3.3) holds. By construction, for consecutive values $c_i = c_i^\ell$, $c'_i = c_i^{\ell+1}$, we have

$$\begin{aligned} q_{i,(c_i,c_{-i})} - q_{i,(c'_i,c_{-i})} &= c_i \big(\tilde{y}_{(c_i,c_{-i}),i} - \tilde{y}_{(c'_i,c_{-i}),i} \big) \\ \leq \rho \cdot c_i \big(y_{(c_i,c_{-i}),i} - y_{(c'_i,c_{-i}),i} \big) &\leq \rho \big(p_{i,(c_i,c_{-i})} - p_{i,(c'_i,c_{-i})} \big). \end{aligned}$$

Since $q_{i,(m_i,c_{-i})} = 0 \leq \rho p_{i,(m_i,c_{-i})}$, this implies that $q_{i,(c_i,c_{-i})} \leq \rho p_{i,(c_i,c_{-i})}$. Finally, it is easy to verify that for any $c_i, c'_i \in \mathcal{D}_i$, we have $q_{i,(c_i,c_{-i})} - q_{i,(c'_i,c_{-i})} \geq c_i \big(\tilde{y}_{(c_i,c_{-i}),i} - \tilde{y}_{(c'_i,c_{-i}),i} \big)$, so $(\tilde{x}, q)$ satisfies (3.2). ■

Corollary 3.5.3 and Theorem 3.6.2 yield polytime near-optimal mechanisms for a host of single-dimensional PayM problems. Table 3.1 summarizes a few applications. Even for single-item procurement auctions, these are the *first* results for PayM problems with correlated players satisfying a notion stronger than (BIC, interim IR).

3.7 Multidimensional Problems

We obtain results for multidimensional PayM problems via two distinct approaches. One is by directly applying Corollary 3.5.3 (e.g., Theorem 3.7.1). The other approach is based on again moving to an LP-relaxation of the CM problem and utilizing Theorem 3.5.4 in conjunction with a *stronger* LP-rounding approach. This yields results for multidimensional (metric) UFL and its variants (Theorem 3.7.3).

Problem	Approximation	Due to
Single-item procurement auction: buy one item provided by n players	1	Corollary 3.5.3
Metric UFL: players are facilities, output should be a UFL solution	1.488 using [34]	Theorem 3.6.2
Vertex cover: players are nodes, output should be a vertex cover	2	Theorem 3.6.2
Set cover: players are sets, output should be a set cover	$O(\log n)$	Theorem 3.6.2
Steiner forest: players are edges, output should be a Steiner forest	2	Theorem 3.6.2
Multiway cut (a), *Multicut* (b): players are edges, output should be a multiway cut in (a), or a multicut in (b)	2 for (a) $O(\log n)$ for (b)	Theorem 3.6.2

Table 3.1: Results for some representative single-dimensional PayM problems.

Multi-item procurement auctions. Here, we have n sellers and k (heterogeneous) items. Each seller i has a *supply vector* $s_i \in \mathbb{Z}_+^k$ denoting his supply for the various items, and the buyer has a *demand vector* $d \in \mathbb{Z}_+^k$ specifying his demand for the various items. This is public knowledge. Each seller i has a *private* cost-vector $c_i \in \mathbb{R}_+^k$, where $c_{i,\ell}$ is the cost he incurs for supplying *one unit* of item ℓ. A feasible solution is an allocation specifying how many units of each item each seller supplies to the buyer such that for each item ℓ, each seller i provides at most $s_{i,\ell}$ units of ℓ and the buyer obtains d_ℓ total units of ℓ. The corresponding CM problem is a min-cost flow problem (in a bipartite graph), which can be efficiently solved optimally, thus we obtain a polytime optimal mechanism.

Theorem 3.7.1 *There is a polytime optimal robust-(BIC-in-expectation, IR) mechanism for multi-unit procurement auctions.*

Multidimensional budgeted (metric) uncapacitated facility location (UFL). In this problem, we have a set $\mathcal{E}$ of clients that need to be serviced by facilities, and a set $\mathcal{F}$ of locations where facilities may be opened. Each player i may provide facilities at the locations in $\mathcal{T}_i \subseteq \mathcal{F}$. We may assume that the $\mathcal{T}_i$s are disjoint. For each facility $\ell \in \mathcal{T}_i$ that is opened, i incurs a private opening cost $f_\ell \equiv f_{i,\ell}$, and assigning client j to an open facility ℓ incurs a publicly-known assignment cost $d_{\ell j}$, where the $d_{\ell j}$s form a metric. We are also given a public assignment-cost budget B. The goal in Budget-UFL is to open a subset $F \subseteq \mathcal{F}$ of facilities and assign each client j to an open facility $\sigma(j) \in F$ so as to minimize $\sum_{\ell \in F} f_\ell + \sum_{j \in \mathcal{E}} d_{\sigma(j)j}$ subject to $\sum_{j \in \mathcal{E}} d_{\sigma(j)j} \leq B$; UFL is the special case where $B = \infty$. We can define $pub(T)$ to be the total assignment cost if this is at most B, and ∞ otherwise.

Let O^* denote the expected disutility of an optimal mechanism for Budget-UFL. We obtain a mechanism with expected disutility at most $2O^*$ that always returns a solution with expected assignment cost at most $2B$. Consider the following LP-relaxation for Budget-UFL.

$$\min \sum_{\ell \in \mathcal{F}} f_\ell x_\ell + \sum_{j \in \mathcal{E}, \ell \in \mathcal{F}} d_{\ell j} z_{\ell j} \quad \text{s.t.} \tag{BFL-P}$$

$$\sum_{j \in \mathcal{E}, \ell \in \mathcal{F}} d_{\ell j} z_{\ell j} \leq B, \qquad \sum_{\ell \in \mathcal{F}} z_{\ell j} \geq 1 \;\; \forall j \in \mathcal{E}, \qquad 0 \leq z_{\ell j} \leq x_\ell \;\; \forall \ell \in \mathcal{F}, j \in \mathcal{E}. \tag{3.17}$$

Let (FL-P) denote (BFL-P) with $B = \infty$, and $OPT_{\text{(FL-P)}}$ denote its optimal value. Recall from Section 2.3 that an algorithm $\mathcal{A}$ is a *Lagrangian multiplier preserving* (LMP) ρ-approximation algorithm for UFL if for every instance, it returns a solution (F, σ) such that $\rho \sum_{\ell \in F} f_\ell + \sum_{j \in \mathcal{E}} d_{\sigma(j)j} \leq \rho \cdot OPT_{\text{(FL-P)}}$. Lemma 2.3.2 shows that given such an algorithm $\mathcal{A}$, one can take any solution (x, z) to (FL-P) and obtain a convex combination of *UFL solutions* $(\lambda^{(1)}; F^{(1)}, \sigma^{(1)}), \ldots, (\lambda^{(k)}; F^{(k)}, \sigma^{(k)})$, so $\lambda \geq 0$, $\sum_r \lambda^{(r)} = 1$, such that $\sum_{r: \ell \in F^{(r)}} \lambda^{(r)} = x_\ell$ for all ℓ and $\sum_r \lambda^{(r)} \big(\sum_j d_{\sigma^{(r)}(j)j}\big) \leq \rho \sum_{j,\ell} d_{\ell j} z_{\ell j}$. As mentioned in Chapter 2, an LMP 2-approximation algorithm for UFL is known [27].

Lemma 3.7.2 *Given an LMP ρ-approximation algorithm for UFL, one can design a poly-time robust-(BIC-in-expectation, IR) mechanism for Budget-UFL whose expected disutility is at most ρO^* while violating the budget by at most a ρ-factor.*

Proof : The LP-relaxation (BFL-P) for the CM problem is of the form (C-P) and satisfies the required properties. Recall that for $x \geq 0$, $z(x)$ denotes the min-cost completion of x to a feasible solution to (BFL-P) if one exists, and is $\perp$ if there is no such completion of x. Let $\Omega_{\text{LP}} := \{x : z(x) \neq \perp, \quad 0 \leq x_\ell \leq 1 \; \forall \ell\}$. For integral $\omega \in \Omega_{\text{LP}}$, $z(\omega)$ specifies the assignment where each client j is assigned to the nearest open facility. By Theorem 3.5.4, one can efficiently compute an optimal solution (X, p) to the relaxation of (P) where the set of feasible allocations is the set Ω_{ext} of extreme points of Ω_{LP}.

We proceed similarly to the proof of Theorem 3.6.2. Let Ω_{UFL} be the set of characteristic vectors of all integral UFL solutions. We use ℓ to index facilities in $\mathcal{F}$ and j to index clients in $\mathcal{E}$. Fix $c \in \overline{\mathcal{D}}$. Define $y_c = \sum_{\omega \in \Omega_{\text{ext}}} X_{c,\omega}\omega$, so $\sum_{w \in \Omega_{\text{ext}}} c_i(\omega) X_{c,\omega} = \sum_{\ell \in \mathcal{T}_i} f_\ell y_{c,\ell}$. Let $z_c = \sum_{\omega \in \Omega_{\text{ext}}} X_{c,\omega} z(\omega)$, so $\sum_{j,\ell} z_{c,\ell j} d_{\ell j} \leq B$. We use the LMP ρ-approximation algorithm to express y_c as a convex combination $\sum_{\omega \in \Omega_{\mathsf{UFL}}} \tilde{x}_{c,\omega}\omega$ of (integral) UFL-solutions such that the expected assignment cost $\sum_{\omega \in \Omega_{\mathsf{UFL}}} \tilde{x}_{c,\omega} \sum_{j,\ell} z(\omega)_{\ell j} d_{\ell j}$ is at most $\rho \sum_{j,\ell} d_{\ell j} z_{c,\ell j} \leq \rho B$. Hence, $(\tilde{x}, p)$ is a feasible solution to (P). Theorem 3.5.2 now yields the desired mechanism. ■

Theorem 3.7.3 *There is a polytime robust-(BIC-in-expectation, IR) mechanism for* $\mathsf{Budget\text{-}UFL}$ *with expected disutility at most* $2O^*$*, which always returns a solution with expected assignment cost at most* $2B$*.*

3.8 Extension: DSIC Mechanisms

We can strengthen our results from Section 3.6 to obtain (near-) optimal *dominant-strategy incentive compatible* (DSIC) mechanisms for single-dimensional problems in time exponential in n. Thus, if n is a constant, we obtain polytime mechanisms. We describe briefly the changes needed, and their implications.

The key change is in the LP (P) (or (P')), where we now enforce (3.1)—(3.4) for all type profiles in $\prod_i \overline{\mathcal{D}}_i$. The rounding procedure and arguments in Theorem 3.6.2 proceed essentially identically to yield a near-optimal solution to this LP. But we can now argue that in single-dimensional settings, a feasible solution to the LP can be rounded to a

(DSIC-in-expectation, IR) mechanism without increasing the expected disutility. Thus, we obtain the same guarantees as in Table 3.1, but under the stronger solution concept of DSIC-in-expectation and IR.

Analogous to Lemma 3.2.3, we can obtain estimates m_i such that there is an optimal mechanism M^* such that on any input $c \in \prod_i (\mathcal{D}_i \cup \{m_i\})$ where $c_i < m_i$ for at least one i, M^* only buys the item with non-zero probability from a player i with $c_i < m_i$ (the same proof approach works). Let $\tilde{\mathcal{D}} := \prod_i \overline{\mathcal{D}}_i$ and $\tilde{\mathcal{D}}_{-i} := \prod_{j \neq i} \overline{\mathcal{D}}_j$; also, let $\tilde{\mathcal{D}}_i := \overline{\mathcal{D}}_i := \mathcal{D}_i \cup \{m_i\}$ for uniformity of notation. The key change is in the LP (P) or (P') (where we consider an LP-relaxation of the CM problem and move to the allocation-set Ω_{ext}). For $c \in \tilde{\mathcal{D}}$, define $\Omega(c) = \{\omega \in \Omega : \omega_i = \emptyset \text{ for all } i \text{ s.t. } c_i = m_i\}$, if there is some i such that $c_i < m_i$, and Ω otherwise. In our LP, we now enforce (3.1)–(3.4) for all i, all $c_i, c'_i \in \tilde{\mathcal{D}}_i$ and all $c_{-i} \in \tilde{\mathcal{D}}_{-i}$. Let (K-P), (K-P') (with allocation-set Ω_{LP}) denote these new LPs. When n is a constant, both LPs have a polynomial number of constraints. So again by considering the dual, we can efficiently compute: (i) an optimal solution to (K-P) given an optimal algorithm for the CM problem; and (ii) an optimal solution to (K-P').

The rounding procedure and arguments in Theorem 3.6.2 proceed essentially identically to yield a near-optimal solution to (K-P). We argue that in single-dimensional settings, one can round a solution (x, p) to (K-P) to a (DSIC-in-expectation, IR) mechanism $M = \big(\mathcal{A}, \{q_i\}\big)$ without increasing the expected total payment. Here, $\mathcal{A}(c)$ and $q_i(c)$ denote as before the allocation-distribution and expected payment to i, on input c.

Define $y_c = \sum_\omega x_{c,\omega}\omega$, where we treat ω as a vector in $\{0,1\}^n$ with $\omega_i \equiv \omega_{i,v}$ for the single covering object $v \in \mathcal{T}_i$. Let $0 \le c_i^1 < c_i^2 < \ldots < c_i^{k_i} = c_i^{\max}$ be the values in $\mathcal{D}_i$, and set $c_i^{k_i+1} := m_i$. Define the mapping $H : C \to \tilde{\mathcal{D}}$ as follows: set $H(c) := \big(H_i(c_i)\big)_{i=1,\ldots,n}$, where $H_i(c_i)$ is c_i^{r+1} if $c_i \in (c_i^r, c_i^{r+1}]$, $r \le k_i$, and m_i if $c_i \ge m_i$. Define $H_{-i}(c_{-i}) := \big(H_j(c_j)\big)_{j \neq i}$.

Consider $c \in C$. If $c_i \le c_i^{\max}$ for at least one i, we set $\mathcal{A}(c) = y_{H(c)}$. If $c_i > c_i^{\max}$ for all i, we set $\mathcal{A}(c)$ as in the VCG mechanism. Since we are in the single-dimensional setting, generalizing Theorem 1.2.8 if we show that for all i, $c_{-i} \in C_{-i}$, $\mathcal{A}(c)_i$ is non-increasing in c_i and hits 0 at some point, then setting $q_i(c) = c_i\mathcal{A}(c)_i + \int_{c_i}^\infty \mathcal{A}(t, c_{-i})_i dt$ ensures such that $M = \big(\mathcal{A}, \{q_i\}\big)$ is (DSIC, IR)-in-expectation.

Consider some i, $c_{-i} \in C_{-i}$. If $c_j \le c_j^{\max}$ for some $j \neq i$, then $\mathcal{A}(c) = y_{H(c)}$. Since

H_i is non-decreasing in c_i and $y_{c,i}$ is non-increasing in c_i (which is easily verified), it follows that $\mathcal{A}(c)_i$ is non-increasing in c_i. Also, if $c_{-i} \in \tilde{\mathcal{D}}_{-i}$, then one can argue as in the proof of Theorem 3.6.2 that $q_i(c) \le p_{i,c}$. Hence, M has expected total payment at most $\sum_{c,i} \Pr_{\mathcal{D}}(c) p_{i,c}$. Suppose $c_j > c_j^{\max}$ for all $j \neq i$. Then, $H_j(c_j) = m_j$ for all $j \neq i$. So $\mathcal{A}(c) = y_{H(c)}$ for $c_i \le c_i^{\max}$, and is the VCG allocation for $c_i > c_i^{\max}$. Therefore, $\mathcal{A}(c)_i = 1$ for $c_i \le c_i^{\max}$, and the VCG allocation for $c_i > c_i^{\max}$, which is clearly non-increasing in c_i.

Theorem 3.8.1 *For single-dimensional problems with a constant number of players, we obtain the same guarantees as in Table 3.1, but under the stronger solution concept of DSIC-in-expectation and IR.*

3.9 Inferiority of k-lookahead Procurement Auctions

The following auction, called k-lookahead, was proposed by [17] for the single-item revenue-maximization problem generalizing the 1-lookahead auction considered by [46, 45]: on input $v = (v_1, \ldots, v_n)$, pick the set I of k players with highest values, and run the revenue-maximizing (DSIC, IR) mechanism for player-set I where the distribution we use for I is the conditional distribution of the values for I given the values $(v_i)_{i \notin I}$ for the other players. Dobzinski et al. [17] show that the k-lookahead auction achieves a constant-fraction of the revenue of the optimal (DSIC, IR) mechanism.

For any $k \ge 2$, we can consider an analogous definition of k-lookahead auction for the single-item procurement problem: on input c, we pick the set I of k players with smallest costs, and run the payment-minimizing robust (BIC-in-expectation, IR) mechanism for I for the conditional distribution of I's costs given $(c_i)_{i \notin I}$. We call this the *k-lookahead procurement auction*. The following example shows that the expected total payment of the k-lookahead procurement auction can be arbitrarily large, even when $k = n - 1$ (so we drop only 1 player), and compared to the optimal expected total payment of even a deterministic (DSIC, IR) mechanism.

Let $t = K + \epsilon$ where $\epsilon > 0$, and $\delta > 0$. The distribution $\mathcal{D}$ consists of n points: each c in $\bigcup_{i=1,\ldots,n-1} \{c : c_i = 0, \;\; c_j = K \; \forall j \neq i, n, \;\; c_n = t\}$ has probability $\Pr_{\mathcal{D}}(c) = \frac{1-\delta}{n-1}$, and the type-profile c where $c_i = K \; \forall i \neq n, \; c_n = t$ has probability $\Pr_{\mathcal{D}}(c) = \delta$.

Let $k = n - 1$. The k-lookahead procurement auction will always select the player-set $I = \{1, \ldots, n-1\}$, and the conditional distribution of values of players in I is simply $\mathcal{D}$. Let M' be the robust (BIC-in-expectation, IR) mechanism for the players in I under this conditional distribution $\mathcal{D}$. Suppose that on input $(K, K, \ldots, K, t)$, the k-lookahead auction (which runs M') buys the item from player i with probability at least $\frac{1}{n-1}$. Then, on the input $\tilde{c}$ where $\tilde{c}_i = 0$, $\tilde{c}_j = K$ for all $j \neq i, n$, $\tilde{c}_n = t$, the mechanism must also buy the item from player i with probability at least $\frac{1}{n-1}$ since M' is robust BIC. So since M' is robust (BIC-in-expectation, IR), the payment to player i under input $\tilde{c}$ is at least K, and therefore the expected total payment of the k-lookahead auction is at least $K\delta + \frac{K(1-\delta)}{n-1}$.

Now consider the following mechanism $M = (\mathcal{A}, p_i)$. Consider input c. If some player $i < n$ has $c_i = 0$, M buys the item from such a player i (breaking ties in some fixed way). Otherwise, if $c_n \leq t$, M buys from player n; else, M buys from the player i with smallest c_i. It is easy to verify that for every i and c_{-i}, this allocation rule $\mathcal{A}$ is monotonically decreasing in c_i. Let $p_i(c) = 0$ if M does not buy the item from i on input c, and $\max\{z : M \text{ buys the item from } i \text{ on input } (z, c_{-i})\}$ otherwise. By a well-known fact (see, e.g., Theorem 9.39 in [43]), $M = \big(\mathcal{A}, \{p_i\}\big)$ is DSIC and IR. Then, the total payment under any $c \in \mathcal{D}$ for which $c_i = 0$ for some i is 0, and the total payment under the input where $c_i = K$ for all $i \neq n$, $c_n = t$ is t. So the expected total payment of M is $t\delta$.

Thus the ratio of the expected total payments of M' and M is at least $\frac{K\delta + \frac{K(1-\delta)}{n-1}}{t\delta} = \frac{(n-2)K\delta + K}{t\delta}$, which can be made arbitrarily large by choosing δ and $\epsilon = t - K$ small enough.

Chapter 4

Frugal Mechanisms for VCP

4.1 Introduction

In this chapter, we consider the notion of *frugality*, wherein we again consider the payment of a mechanism (for a covering problem), but in a worst-case setting. That is, unlike Chapter 3, we do not assume that there is an underlying prior distribution over players' type profiles, and seek to prove *worst-case* performance guarantees.

An immediate difficulty that arises when working with worst-case guarantees is that there is no point-wise optimal (truthful) mechanism whose payment, for every type-profile c is minimum among the payment of all truthful, IR mechanisms for input c. Moreover (see Theorem 4.4.2), for every truthful, IR mechanism M and any factor α, one can tailor a truthful, IR mechanism M' and an input c, such that the payment of M on c is α times the payment of M' on c. Thus, one of the main issues when trying to measure frugality is in fact to come up with a reasonable benchmark for frugality relative to which one can obtain meaningful bounds.

The extant work on frugality has focused on the single-dimensional setting. The first proposed benchmark was for the (single-dimensional) shortest-path problem. In this problem each player owns an edge of a graph G and has a private cost for it. The goal of the mechanism is to buy the shortest (s, t)-path for a given pair $s, t \in V(G)$. The benchmark proposed was to compare the payment of the mechanism when the outcome is $\omega_1 \in \Omega$ (note that the algorithmic problem is polytime solvable and hence they assume that ω_1 is the shortest (s,t)-path, say $p(\omega_1)$), with the cost of the shortest outcome $\omega_2 \in \Omega$ where $\omega_1 \cap \omega_2 = \emptyset$ (see Archer et al. [2] and Talwar [48]). However, this benchmark has some limits and cannot be applied to several other problems or instances of problems in AMD, although it was generalized to "set-system" problems by [48]. (Set-systems are an alternative way of defining single-dimensional covering problems proposed by [48], which are captured by our general framework.)

Later, Karlin et al. [28] discuss the restrictions of the frugality notion in [2] and introduce a new frugality notion (Definition P_2) based on the cheapest Nash equilibrium in a monopoly-free set system. Elkind et al. [23] develop some other frugality notions that are in the same spirit as the rationale underlying the benchmark in [28]. Their proposed $\mathrm{NTU}_{\max}$ benchmark (see Definition 4.2.1) is used later in [29]. The problem of designing polytime truthful and IR mechanisms for single-dimensional VCP with bounded frugality ratios was initiated by [8] and followed by [23], [29].

4.1.1 Summary of Results

In Section 4.3, we consider the question of designing of *polytime* frugal truthful and IR mechanisms for single-dimensional *and* multi-dimensional VCP with respect to the frugality benchmarks proposed by [23] and [28]. To the best of our knowledge we are the first to consider and obtain polytime truthful and IR mechanisms for multi-dimensional settings. We show that some of our mechanisms defined in Section 2.4.1 for Multi-VCP enjoy good frugality properties. Thus, we obtain the *first* mechanisms for Multi-VCP that are polytime, truthful and *simultaneously* achieve bounded approximation ratio *and* bounded frugality ratio with respect to the benchmarks in [10, 29]. This nicely complements a result of [10], who devises such a mechanism for single-dimensional VCP.

We show that neither of the benchmarks $\mathrm{NTU}_{\max}$ and $\mathrm{NTU}_{\min}$ proposed in [28, 23] yield a pointwise lower-bound on the payment of a truthful, IR mechanism. Although, as noted earlier, no meaningful guarantees are possible when one compares against *all* truthful, IR mechanisms, the example that shows this involves a somewhat esoteric mechanism that may output a very sub-optimal vertex cover. Motivated by this, we ask whether meaningful guarantees are possible if we compare against the pointwise-optimal payment of a strict *subset* of truthful, IR mechanisms (e.g., truthful, IR approximation mechanisms). To study this question, we define a new frugality notion in Section 4.4, where we compare the payment of a mechanism with the pointwise-optimal payment of a given class of truthful, IR mechanisms. We obtain some preliminary results, both upper and lower bounds, that shed light on how the frugality ratio (under this new measure) is affected by restrictions on the approximation ratio or the type of optimality guarantee of the class of mechanisms that one compares against.

4.1.2 Related Work

Single-dimensional covering problems have been well studied from the perspective of frugality. Starting with the work of Archer and Tardos [2], various benchmarks for frugality have been proposed and investigated for various problems including VCP, k-edge-disjoint paths, spanning tree, s-t cut; see [28, 23, 29, 10] and the references therein. In [28], the authors propose a new benchmark for the shortest-path auctions problem and subsequently Elkind et al. [23] designed a frugal mechanism for single-dimensional VCP with respect to this benchmark. They also gave an alternative benchmark in their paper, which is used by Kempe et al. [29], to analyze the frugality ratio of their mechanism for single-dimensional VCP.

4.2 Problem Definition and Preliminaries

Recall from Chapter 2 that in the Multi-VCP mechanism design problem we have a graph $G = (V, E)$ with n nodes. Each player i provides a subset $\mathcal{T}_i$ of nodes. Given Remark 2.4.7, for simplicity, we assume that the $\mathcal{T}_i$s are disjoint, and given a cost-vector $\{c_{i,u}\}_{i\in[n],u\in\mathcal{T}_i}$, we use c_u to denote $c_{i,u}$ for $u \in \mathcal{T}_i$. The goal is to choose a minimum-cost *vertex cover*, i.e., a min-cost set $S \subseteq V$ such that every edge is incident to a node in S. Recall that we consider this problem in a monopoly-free setting.

As in Chapter 2 we use VCP to refer to the algorithmic vertex cover problem in this chapter. We use (G, c) to denote a Multi-VCP instance G with true players' cost vector c.

Frugality notion. Karlin et al. [28] and Elkind et al. [23] propose various benchmarks for measuring the *frugality ratio* of a mechanism, which is a measure of the (over-)payment of a mechanism where there is no underlying distribution over players' types. The mechanisms that we devised in Chapter 2 also enjoy good frugality ratios with respect to the following benchmark introduced by [23], which is denoted by $\mathrm{NTU}_{\max}$ in [23] (and $\nu^+(G, c)$ in [29]).

Definition 4.2.1 (Frugality benchmark $\mathbf{NTU_{max}}(G, c)$ [28, 23]) *Given an instance of VCP on a graph $G = (V, E)$ with node costs $\{c_u\}$, we define $NTU_{\max}(G, c)$ as follows. Fix an arbitrary min-cost vertex cover S (with respect to c).*

$$
\begin{aligned}
NTU_{\max}(G,c) \; := \; \max \quad & \sum_{v \in S} x_v && && (P_1) \\
\text{s.t.} \quad & x_v \geq c_v && \text{for all } v \in S \\
& \sum_{v \in S \setminus T} x_v \leq \sum_{v \in T \setminus S} c_v && \text{for all vertex covers } T.
\end{aligned}
$$

The above benchmark is defined along with the definition of $\mathrm{NTU}_{\min}$ which was also defined by Elkind et al. ([23]).

Definition 4.2.2 [28] *Given an instance of VCP on a graph $G = (V, E)$ with node costs $\{c_u\}$, we define $NTU_{\min}(G, c)$ as follows. Fix an arbitrary min-cost vertex cover S (with respect to c).*

$$
\begin{aligned}
NTU_{\min}(G,c) \; := \; \min \quad & \sum_{v \in S} x_v && && (P_2) \\
\text{s.t.} \quad & x_v \geq c_v && \text{for all } v \in S \\
& \sum_{v \in S \setminus T} x_v \leq \sum_{v \in T \setminus S} c_v && \text{for all vertex covers } T \\
& \forall v \in S, \; \exists \text{vertex cover } T \not\ni v \text{ s.t. } \sum_{u \in S \setminus T} x_u = \sum_{u \in T \setminus S} c_u
\end{aligned}
$$

Theorem 4.2.3 [23] *$NTU_{\min}(G, c)$ and $NTU_{\max}(G, c)$ are independent of the choice of S.*

The above theorem shows that $\mathrm{NTU}_{\min}(G, c)$ and $\mathrm{NTU}_{\max}(G, c)$ are well-defined even if (G, c) has more than one min-cost vertex cover. The following proposition regarding $\mathrm{NTU}_{\max}(G, c)$, $\mathrm{NTU}_{\min}(G, c)$ and c is used in some of our sequel arguments.

Lemma 4.2.4 [23] *Let S be a min-cost vertex cover of an instance (G, c) of VCP. We have $NTU_{\max} \geq NTU_{\min} \geq \max\{c(S), c(V \setminus S)\}$. Hence, $NTU_{\max} \geq NTU_{\min} \geq c(V)/2$.*

Definition 4.2.5 [28] *Let M be a truthful and IR mechanism for single-dimensional VCP with underlying graph G. Suppose c is the true players' cost vector. Let $p_M(G, c)$ denote the total payment of M to the players in (G, c) instance.*

With respect to the $NTU_{\min}$ benchmark, the frugality ratio of M on (G, c) is defined as $\phi_M(G, c) := \frac{p_M(G,c)}{NTU_{\min}(G,c)}$, and the frugality ratio of M is defined as $\phi_M(G) := \sup_c \phi_M(G, c)$.

Similarly, for the $NTU_{\max}$ benchmark, we define $\phi^+_M(G, c) := \frac{p_M(G,c)}{NTU_{\max}(G,c)}$ and $\phi^+_M(G) := \sup_c \phi^+_M(G, c)$.

Definition 4.2.6 [23] *An algorithm $\mathcal{A}$ for VCP, is called locally-optimal if for any vertex v in the outcome of $\mathcal{A}$ on any instance (G, c) we have $c_v \leq \sum_{uv \in E(G)} c_u$.*

For the ease of notation, we restated Definition 1.2.7 for single-dimensional VCP as follows.

Definition 4.2.7 *Let $\mathcal{A}$ be a MON algorithm for an instance G of single-dimensional VCP. Given players' cost vector c, for each vertex $v \in V(G)$, its critical value is defined as $b_v(\mathcal{A}, c_{-v}) := \sup\{c'_v | v \in \mathcal{A}(G, (c'_v, c_{-i}))\}$ whenever the RHS is finite, where i is the owner of vertex v.*

Note that, as mentioned in Section 1.2 if a mechanism for single-dimensional VCP has a bounded total cost, then by the monopoly-free nature of single-dimensional setting the critical values (recall Definition 1.2.7) are always defined. This fact is implicitly used by some sequel propositions such as Proposition 4.3.2.

4.3 Frugal Mechanisms for VCP

In this section, we consider the task of designing frugal truthful and IR mechanisms for single-dimensional VCP and Multi-VCP. Along with restating some needed notions and results we also mention our results on the frugality of the two threshold mechanisms defined in Section 2.4.1 with respect to the $\mathrm{NTU}_{\max}$ and $\mathrm{NTU}_{\min}$ benchmarks.

Theorem 4.3.1 *For any mechanism $M = (\mathcal{A}, p_1, \ldots, p_n)$, where $\mathcal{A}$ is a locally-optimal algorithm, we have $\phi_M(G, c) \leq 2\Delta$ for any instance G of single-dimensional VCP and any cost vector c, where Δ is the maximum degree of a vertex of G.*

Proof : By Lemma 4.2.4 we have $\mathrm{NTU}_{\min}(G, c) \geq c(V)/2$. On the other hand, by Theorem 1.2.8 we have

$$p_M(G, c) = \sum_{v \in \mathcal{A}(G,c)} b_v(\mathcal{A}, c_{-v}) \leq \sum_{v \in \mathcal{A}(G,c)} \Big(\sum_{uv \in E(G)} c_u \Big) \leq \sum_{u \in V(G)} \Delta c_u = \Delta c(V).$$

Hence, $\phi_M(G, c) = \frac{p_M(G,c)}{\mathrm{NTU}_{\min}(G,c)} \leq \frac{\Delta c(V)}{c(V)/2} = 2\Delta$. ■

Corollary 4.3.2 [23] *There exists a 2-approximation truthful mechanism M for single-dimensional VCP such that $\phi_M(G, c) \leq 2\Delta$ for any instance (G, c) where Δ is the max vertex degree of G.*

Proof : By fixing an order on the set of vertices and edges of G, the 2-approximation primal-dual algorithm for vertex cover problem mentioned in [49], where dual variables corresponding to edges are increased sequentially, has the MON property. Moreover, it is locally-optimal. Hence, by Theorem 1.2.8 we obtain the above mentioned 2-approximation truthful mechanism for any instance (G, c) of single-dimensional VCP. ■

As one might notice, the mechanism introduced in the previous proof is in fact a mechanism that returns a near-optimal VC and has bounded overpayment, i.e., it simultaneously has good approximation ratio and good frugality.

Recall from Section 2.4.1 that we define an x-*scaled* edge-threshold mechanism as follows: fix a vector $(x_u)_{u \in V}$, where $x_u > 0$ for all u, and set $t_u^{(uv)} := x_u c_v / x_v$ for every edge (u, v). We use $M_{\mathcal{A}}$ and $\mathcal{A}_x$ to denote the resulting edge-threshold mechanism and its allocation algorithm, respectively. Also, define $M_{\mathcal{B}}$ and $\mathcal{B}_x$ to be the neighbor-threshold mechanism and its allocation algorithm, respectively, where we set $t_u := \sum_{v \in N(u)} x_u c_v / x_v$. The proof of Lemma 2.4.3 is easily modified to show that the x-scaled mechanism $M_{\mathcal{A}}$ satisfies $\sum_i p_i(c) \le \sum_u t_u \le \beta(G; x) c(V)$, where $\beta(G; x) = \max_{u \in V} \frac{x(N(u))}{x_u}$. Since $\mathrm{NTU}_{\max}(G, c) \ge c(V)/2$ by Lemma 4.2.4, this implies that $\phi_{M_{\mathcal{A}}}(G) \le 2\beta(G; x)$. Also, if M is a decomposition mechanism constructed from threshold mechanisms $M_1, \ldots, M_k$, where each M_q satisfies $\sum_u t_u^q \le \phi_q \cdot c(V(G_q))$, then it is easy to see that $\phi_M(G) \le 2 \sum_{q=1}^k \phi_q$. Thus, we obtain the following results.

Theorem 4.3.3 *Let $G = (V, E)$ be a graph with n nodes. We can obtain a polytime, truthful, IR mechanism M with the following approximation $\rho = \rho_M(G)$ and frugality $\phi = \phi_M(G)$ ratios.*

(i) $\rho = (\alpha(G; x) + 1)$, $\phi \le 2\beta(G; x)$ *for* Multi-VCP *on* G*;*

(ii) $\rho = O(r^2 \log n))$, $\phi = O\big(r^2 \log n \cdot \Delta(G)\big)$ *for r-dimensional* VCP *on G (using a 2-approximation mechanism with frugality ratio* $2\Delta(G)$ *[23] for single-dimensional* VCP *in the construction of Theorem 2.4.9);*

(iii) $\rho = O(r\gamma \log n), \phi = O(r\gamma \log n \Delta(G))$ *for r-dimensional* VCP *on G when G is everywhere γ-sparse; hence, we achieve* $\rho = O(r \log n), \phi = O(r \log n \Delta(G))$ *for r-dimensional* VCP *on any proper minor-closed family.*

Kempe et al. [29], consider the $\mathrm{NTU}_{\max}(G, c)$ benchmark. Recall that with respect to this benchmark, the frugality ratio of M on (G, c) is denoted by $\phi_M^+(G, c) := \frac{p_M(G,c)}{\nu^+(G,c)}$ and the the frugality ratio of M is denoted by $\phi_M^+(G) := \sup_c \phi_M^+(G, c)$. They propose a computationally-inefficient mechanism for single-dimensional VCP that has frugality ratio at most λ_1' on any instance (G, c), where λ_1' is the largest eigenvalue of a scaled adjacency matrix of G.

A fundamental question in [29] concerns the motivation behind, and the validity of, using $\mathrm{NTU}_{\max}(G, c)$ as a benchmark. That is, is it the case that ϕ_M^+ is at least 1 for every truthful IR mechanism M, for general set system problems. Here we answer the question in the affirmative for single-dimensional VCP.

Proposition 4.3.4 *If M is a truthful and IR mechanism for single-dimensional VCP on a given graph G, then $\phi_M^+(G) \geq 1$ and $\phi_M(G) \geq 1$.*

Proof : By the definition of $\phi_M^+(G)$ and $\phi_M(G)$, we just need to show that there exists a cost vector c such that $\phi_M^+(G,c) \geq 1$ and $\phi_M(G,c) \geq 1$. Let $\mathcal{C} = \{v_1, \ldots, v_k\}$ be a maximal clique in G and let $a > 0$ be a real number. Suppose the cost vector c is as follows: $c_v = a$ for all $v \in \mathcal{C}$ and $c_v = 0$ for any $v \in V \setminus \mathcal{C}$. Now we claim that $\mathrm{NTU}_{\min}(G,c) = \mathrm{NTU}_{\max}(G,c) = c(S)$ where S is a min-cost vertex cover of (G,c). Since M is individually rational we have $p_M(G,c) \geq c(S)$. So if we prove the claim we are done.

First of all, any min-cost vertex cover S includes exactly $k-1$ (i.e.,$|\mathcal{C}|-1$) vertices of $\mathcal{C}$. Hence, $c(S) = a(k-1)$. By Theorem 4.2.3, WLOG, we can assume $S = V \setminus \{v_k\}$.

We need to show $\mathrm{NTU}_{\max}(G,c) = \mathrm{NTU}_{\min}(G,c) = a(k-1)(= c(S))$. As $\mathrm{NTU}_{\max}(G,c) \geq \mathrm{NTU}_{\min}(G,c) \geq c(S)$ we just need to show $\mathrm{NTU}_{\max}(G,c) = a(k-1)$. That means, by definition of $\mathrm{NTU}_{\max}(G,c)$, it suffices to show that if x^* is an optimal solution for the LP (P_1), then $x_v^* = c_v$ for every $v \in S$. This is equivalent to showing that for any $v \in S$ there exists a vertex cover T_v such that $c(T_v) = c(S)$ and $v \notin T_v$ (since then, the second constraint of (P_1) implies that $x^*(S \setminus T_v) = c(T_v \setminus S)$ thus $x_v^* = c_v$). We have the following two cases for any $v \in S$.

Case 1 $v \in \mathcal{C} \setminus \{v_k\}$: Let $T_v = V \setminus \{v\} \in O$.

Case 2 $v \in V \setminus \mathcal{C}$: Since $\mathcal{C}$ is maximal, there is $u \in \mathcal{C}$ such that $uv \notin E(G)$. Let $T_v = V \setminus \{u, v\} \in O$. ■

In the sequel, we may use mechanism M to represent the algorithm used in that mechanism, e.g. using $M(G,c)$ instead of $\mathcal{A}(G,c)$ or $b_v(M, c_{-v})$ instead of $b_v(\mathcal{A}, c_{-v})$ where $\mathcal{A}$ is the algorithm used in M.

Proposition 4.3.5 *There exist an instance (G,c) and a truthful mechanism M for it such that $p_M(G,c) < NTU_{\min}(G,c)$ [and consequently $p_M(G,c) < NTU_{\max}(G,c)$ since $NTU_{\min}(G,c) \leq NTU_{\max}(G,c)$].*

Proof : Let $G = (\{u_1, u_2, v_1, v_2\}, \{u_1v_1, u_2v_2\})$ and let $\mathcal{A}$ be an algorithm that on (G, c), for any cost vector $c \in \mathbb{R}^4_+$, outputs $\{u_1, u_2\}$ if $c_{u_1} + c_{u_2} \leq c_{v_1} + c_{v_2}$, $\{v_1, v_2\}$ if otherwise. M is MON, so by using the prices stated in Theorem 1.2.8 we obtain a truthful mechanism (more exactly, a truthful and IR mechanism).
Now if we set $c = (c_{u_1}, c_{u_2}, c_{v_1}, c_{v_2}) = (a, b, b, a)$ for some $a < b \in \mathbb{R}$, then $M(G, c)$ [i.e., $\mathcal{A}(G, c)$] is $\{u_1, u_2\}$ and $p_M(G, c) = a + b$. On the other hand, $\mathrm{NTU}_{\min}(G, c) = b + b = 2b$. Hence, $p_M(G, c) < \mathrm{NTU}_{\min}(G, c)$, as claimed.
Note that for this instance (G, c) we have $\mathrm{NTU}_{\min}(G, c) = \mathrm{NTU}_{\max}(G, c)$. ■

4.4 A New Frugality Benchmark

Proposition 4.3.5 shows that neither of $\mathrm{NTU}_{\min}(G, c)$ nor $\mathrm{NTU}_{\max}(G, c)$ is a lower bound on the payment of every truthful, IR mechanism for single-dimensional VCP on the instance (G, c). That is, they do not yield a pointwise lower bound on the payment of a truthful, IR mechanism. Moreover, an unsatisfactory aspect of these benchmarks is that while the benchmark is obtained by considering the payment of a mechanism that outputs a min-cost solution (i.e., min-cost vertex cover), there is no such (near-) optimality requirement on the mechanism whose frugality-factor is being evaluated. Taking a leaf from the area of approximation algorithms, we would ideally like to compare the payment of our mechanism on an instance (G, c) to the optimal payment among all truthful and IR mechanisms for single-dimensional VCP on the instance (G, c). This motivates the following notion of frugality.

Definition 4.4.1 *Let G be an instance of single-dimensional VCP where $|V(G)| = n$ and let M be a truthful and IR mechanism for G. Let $\mathcal{M}_1$ be a class of truthful and IR mechanisms for G. We define the payment ratio of M with respect to $\mathcal{M}_1$ to be*

$$\rho(G, M, \mathcal{M}_1) := \sup_{\substack{M_1 \in \mathcal{M}_1 \\ c \in \mathbb{R}^n_+}} \frac{p_M(G, c)}{p_{M_1}(G, c)}$$

and we define the payment ratio of class $\mathcal{M}_2$ of mechanisms with respect to $\mathcal{M}_1$ to be

$$\rho(G, \mathcal{M}_2, \mathcal{M}_1) := \inf_{M \in \mathcal{M}_2} \rho(G, M, \mathcal{M}_1).$$

As seen from the above definition, the above notion of payment ratio can be considered as the worst-case "frugality ratio" of a truthful mechanism. To the best of our knowledge, this notion of frugality is new and has not been considered in the literature.

Theorem 4.4.2 proves a basic impossibility result showing that if we compare against the class $\mathcal{M}$ of *all* truthful, IR mechanisms then no mechanism can achieve a bounded frugality ratio. This holds even for the simplest single-dimensional VCP instance where the graph $G = (\{u, v\}, \{uv\})$ consists of a single edge.

Consider the simplest non-trivial connected graph $G = (\{u, v\}, \{uv\})$. Let $\mathcal{A}$ be a vertex cover algorithm for G where $\mathcal{A}$ outputs $\{u\}$ if ($c_u = 0$ *and* $c_v \in [0, 1]$) or ($c_u \leq c_v$ *and* $c_v > 1$), and $\{v\}$ otherwise (see Figure 4.1). One can check that $\mathcal{A}$ is MON and hence by Theorem 1.2.8 there exist payments such that together with $\mathcal{A}$ form a truthful mechanism M^*.

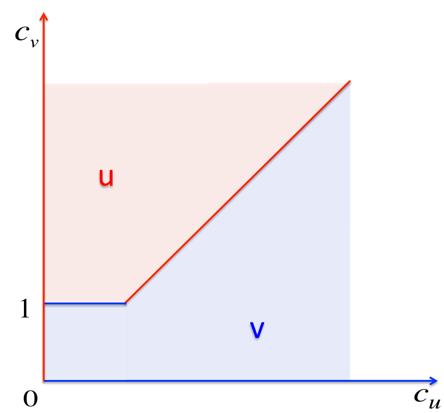

Figure 4.1: Output of $\mathcal{A}$ on graph of single-edge uv for all possible cost vectors (c_u, c_v); in the red areas u and in the blue areas v is the output.

Theorem 4.4.2 *Let $\mathcal{M}$ be the set of all truthful mechanisms for $G = K_2$, then we have $\rho(G, M, \mathcal{M}) = \infty$ for any $M \in \mathcal{M}$.*

Proof : Assume that $G = (\{u, v\}, \{uv\})$. For simplicity we break the argument into two cases.

Case 1 $M = M^*$: Let $c = (c_u, c_v) = (\epsilon, 0)$ for some $0 < \epsilon < 1$. Thus, $M^*(G, c) = \{v\}$ and $p_{M^*}(G, c) = b_v(M^*, c_{-v}) = 1$. Now let $M_1 \in \mathcal{M}$ be the VCG mechanism for G. So, $M_1(G, c) = \{v\}$ and $p_{M_1}(G, c) = \epsilon$. Hence

$$\frac{p_M(G, c)}{p_{M_1}(G, c)} = \frac{1}{\epsilon} \longrightarrow \infty \text{ as } \epsilon \longrightarrow 0$$

Case 2 $M \neq M^*$: Let $c = (0, a)$ for some $0 \leq a \leq 1$ and $M_1 = M^* \in \mathcal{M}$. We have $M_1(G, c) = \{u\}$ and $p_{M_1}(G, c) = 0$. Now if $M(G, c) = \{v\}$ by IR property of M, $p_M(G, c) \geq a$, hence $p_M(G, c)/p_{M_1}(G, c) = \infty$. Otherwise, $M(G, c)$ has to be $\{u\}$. But again, if $p_M(G, c) > 0$ then we have $p_M(G, c)/p_{M_1}(G, c) = \infty$. So, the only remaining case to be considered is when for every $0 \leq a \leq 1$ we have $M(G, c) = \{u\}$ and $p_M(G, c) = 0$ which by Theorem 1.2.8 implies that $b_u(M, c) = p_M(G, c) = 0$.
Hence, for every $a \in [0, 1]$, if $c' = (\epsilon, a)$ then $M(G, c') = \{v\}$ and $p_M(G, c') = b_v(M, c'_{-v}) \geq 1$. This behaviour of M is similar to M^* and hence the argument for Case 1 shows that comparing to VCG, the payment ratio of M tends to ∞. ■

Given the above theorem, in the sequel, we will focus on a large reasonable *subclass* of mechanisms — e.g. the class of γ-approximation truthful mechanisms, or the class of locally-optimal truthful and IR mechanisms — and compute the ρ-ratio of a mechanism (or class of mechanisms) with respect to this subclass. (Note that the mechanism M^*, used in Theorem 4.4.2, is neither a γ-approximation, for any $\gamma \geq 1$, nor locally-optimal, mechanism.)

Fix a single-dimensional VCP instance G and let $\mathcal{M}_l$ be the *set of all truthful and IR locally-optimal mechanisms* for G. Recall from Section 4.3 that we define *x-scaled* edge-threshold mechanism $M_{\mathcal{A}}$ as follows: fix a vector $(x_u)_{u \in V}$, where $x_u > 0$ for all u, and set $t_u^{(uv)} := x_u c_v / x_v$ for every edge (u, v). Also, we define $M_{\mathcal{B}}$ to be the neighbor-threshold mechanism where we set $t_u := \sum_{v \in N(u)} x_u c_v / x_v$. The following theorem shows that for $(x_u)_{u \in V} = (1, \ldots, 1)$, our proposed mechanisms $M_{\mathcal{A}}$ and $M_{\mathcal{B}}$, have bounded payment ratio with respect to the class $\mathcal{M}_l$.

Definition 4.4.3 *The local independence number, α_l, for a graph G is defined to be the maximum-size independent set of vertices that share a common neighbor:*

$$\alpha' = \max_{u \in V(G)} \max_{\substack{S \subseteq N(u) \\ S: independent}} |S|.$$

Theorem 4.4.4 *For $\mathcal{I} \in \{\mathcal{A}, \mathcal{B}\}$, we have*

$$\rho(G, M_{\mathcal{I}}, \mathcal{M}_l) \leq \alpha' + \Delta(2\alpha' + 1)$$

where α' is the local independence number and Δ is the maximum degree of graph G.

Proof :

Since for each vertex $v \in V(G)$ its threshold in $M_{\mathcal{B}}$ is not less than its threshold in $M_{\mathcal{A}}$, the output of $M_{\mathcal{A}}$ is a subset of the output of $M_{\mathcal{B}}$ and also the total payment of $M_{\mathcal{A}}$ is not more than the total payment of $M_{\mathcal{B}}$. Thus, it suffices to prove the lemma for $M_{\mathcal{B}}$.

Fix any cost vector $c \in \mathbb{R}^+$. Let $M_l \in \mathcal{M}_l$, $S_l = M_l(G, c)$, and $S = M_{\mathcal{B}}(G, c)$. We need to show $p_{M_{\mathcal{B}}}(G, c) \leq (\alpha' + \Delta(2\alpha' + 1))p_{M_l}(G, c)$.

For each vertex u let t_u denote its threshold in $M_{\mathcal{B}}$. By definition of $M_{\mathcal{B}}$ we know $p_{M_{\mathcal{B}}}(G, c) = \sum_{u \in S \setminus S_l} t_u + \sum_{u \in S \cap S_l} t_u$.

As S_l is a vertex cover we have $N(u) \subseteq S_l$ for all $u \in S \setminus S_l$ and since $S \setminus S_l$ is an independent set we obtain

$$\sum_{u \in S \setminus S_l} t_u \leq \sum_{v \in S_l} \alpha' c_v = \alpha' c(S_l) \leq \alpha' p_{M_l}(G, c)$$

where the last inequality follows since M_l is individually rational.

Hence it remains to show $\sum_{u \in S \cap S_l} t_u \leq \Delta(2\alpha' + 1)p_{M_l}(G, c)$. By t_u's definition we have $\sum_{u \in S \cap S_l} t_u = \sum_{u \in S \cap S_l} c(N(u)) = \Delta c(A) + \Delta c(B) + \Delta c(C) + \Delta c(D)$ for some subsets $A \subseteq S \setminus S_l$, $B \subseteq S \cap S_l$, $C \subseteq S_l \setminus S$, and $D \subseteq V \setminus (S \cup S_l)$.

Note that $c(A)$ can be upper bounded by $\sum_{u \in A} t_u \leq \alpha' p_{M_l}(G, c)$. Also, $c(B) + c(C) \leq c(S_l)$. So we focus on bounding $c(D)$. For each $v \in D$ since v is not in S we have

$c_v > c(N(v))$. In addition, as S and S_l are both vertex covers we have $N(v) \subseteq S \cap S_l$. If $c(N(v)) < c_v$ then $v \notin S_l$ (due to local optimality of M_l), and hence, $N(v) \subseteq S_l$. Hence, for every $u \in N(v)$, we have that if $c_u < c_v - c(N(v) \setminus \{u\})$, then $u \in S_l$. Hence, by Theorem 1.2.8, we have $p_u(M_l, c) \geq c_v - c(N(v) \setminus \{u\})$, which implies that $c_v \leq p_u(M_l, c) + c(N(v) \setminus \{u\})$, for every $u \in N(v)$. So $c(D) \leq \sum_{v \in D} \left(p_{\sigma(v)}(M_{\mathcal{B}}, c) + c(N(v) \setminus \{\sigma(v)\})\right)$ where $\sigma(v) \in N(v)$ denotes some arbitrary neighbour of v. Considering the maximum number of neighbours that each vertex in $S \cap S_l$ can have we obtain $c(D) \leq \sum_{u \in S_l}(r^u p_u(M_{\mathcal{B}}, c) + s^u c_u)$ where $r^u = |\{v \in D | \sigma(v) = u\}|$, and $s^u = |\{v \in D | u \in N(v), \sigma(v) \neq u\}|$. Note that $r^u + s^u \leq \alpha'$. As $c_u \leq p_u(M_{\mathcal{B}}, c)$ (by IR property of $M_{\mathcal{B}}$) we infer $c(D) \leq \alpha' p_{M_l}(G, c)$.

Therefore,

$$\sum_{u \in S \cap S_l} t_u \leq \Delta(\alpha' c(S_l) + c(S_l) + \alpha' c(S_l)) \leq \Delta(2\alpha' + 1) p_{M_l}(G, c)$$

as desired. ■

Next, we prove some lower bounds for the payment ratio when we compare against the class of approximation, truthful, IR mechanisms.

Proposition 4.4.5 *Let $G = K_{1,r}$ be the star with r leaves and $\gamma > 1$. If $\mathcal{M}$ is the set of all truthful and IR mechanisms and $\mathcal{M}(\gamma)$ is the set of all γ-approximation truthful and IR mechanisms, for single-dimensional* VCP *on G, then $\rho(G, \mathcal{M}, \mathcal{M}(\gamma)) \geq \gamma$.*

Proof : we need to show that for every $M \in \mathcal{M}$, there exist $c' \in \mathbb{R}_+^{r+1}$ and $M_1 \in \mathcal{M}(\gamma)$ such that $\frac{p_M(G,c')}{p_{M_1}(G,c')} \geq \gamma$.
Let $V(G) = \{u, v_1, \ldots, v_r\}$ where $\deg(u) = r$. Now set $c = (c_u, c_{v_1}, \ldots, c_{v_r}) = (a, a, 0, \ldots, 0)$ for some positive real number a.

Case 1 $u \in M(G, c)$: Let M_1 be the γ-approximation truthful mechanism that on an instance (G, c') of single-dimensional VCP outputs $\{u\}$ if $c'_u \leq \sum_{i=1}^{r} c'_{v_i}/\gamma$ and $\{v_1, \ldots, v_r\}$ otherwise. Note that M_1 is MON and hence by Theorem1.2.8 its payments are determined. So, M_1 is in $\mathcal{M}(\gamma)$.
Now set $c' = (a/\gamma, a, 0, \ldots, 0)$. $M_1(G, c') = \{u\}$ and $p_{M_1}(G, c') = b_u(M_1, c'_{-u}) = a/\gamma$. On

the other hand, as M is truthful, it is MON, thus $u \in M(G, c')$ (since $u \in M(G, c)$ and $c'_u \leq c_u$ while $c'_{-u} = c_{-u}$). So, $p_M(G, c') = b_u(M, c'_{-u}) \geq a$. Therefore, $p_M(G, c')/P_{M_1}(G, c') \geq a/(\frac{a}{\gamma}) = \gamma$.

Case 2 $u \notin M(G, c)$: Since $M(G, c)$ is a vertex cover, we must have $v_1 \in M(G, c)$. Let M_1 be the γ-approximation truthful mechansim that on an instance (G, c') of single-dimensional VCP, outputs $\{v_1, \ldots, v_r\}$ if $\sum_i c'_{v_i} \leq c'_u/\gamma$, and $\{u\}$ otherwise. One can easily check that $M_1 \in \mathcal{M}(\gamma)$. Now set $c' = (a, a/\gamma, 0, \ldots, 0)$. We have

$$v_1 \in M(G, c) \Rightarrow v_1 \in M(G, c') \Rightarrow p_M(G, c') = b_{v_1}(M, c'_{-v_1}) \geq a$$

and

$$M_1(G, c') = \{v_1, \ldots, v_r\} \Rightarrow p_{M_1}(G, c') = \sum_i b_{v_i}(M_1, c'_{-v_i}) = a/\gamma + 0 + \ldots + 0 = a/\gamma.$$

Therefore, $\frac{p_M(G,c')}{p_{M_1}(G,c')} \geq a/(\frac{a}{\gamma}) = \gamma$. ■

Proposition 4.4.6 *The above lower bound is tight for $r = 1$.*

Proof : Suppose $G = K_{1,1}$ and $V(G) = \{u, v\}$. Let M be the VCG mechansim. I.e., on an instance (G, c) of single-dimensional VCP, M outputs $\{u\}$ if $c_u \leq c_v$, and $\{v\}$ otherwise (and the payments are according to Theorem 1.2.8). It suffices to show that for every $c \in \mathbb{R}^2_+$ and every $M_1 \in \mathcal{M}(\gamma)$, we have

$$\frac{p_M(G, c)}{p_{M_1}(G, c)} \leq \gamma.$$

Note that $p_M(G, c) = \max(c_u, c_v)$ for any $c \in \mathbb{R}^2_+$.

Fix an arbitrary $c \in \mathbb{R}^2_+$ and $M_1 \in \mathcal{M}(\gamma)$. There are two cases.

Case 1 $M(G, c) = \{v\}$: So we have $c_u > c_v$ and $p_M(G, c) = c_u$.
If $p_{M_1}(G, c) < c_u/\gamma$, by the fact that M_1 is IR we conclude that $M_1(G, c) = \{v\}$. Thus, by Theorem 1.2.8, $p_{M_1}(G, c) = b_v(M_1, c_{-v})$.

Now choose $\epsilon > 0$ so that $\bar{c_v} := b_v(M_1, c_{-v}) + \epsilon < c_u/\gamma$ and let $\bar{c} = (c_u, \bar{c_v})$. The mechanism M_1 outputs $\{u\}$ on $(G, \bar{c})$ but $c_u/\bar{c_v} > \gamma$ which contradicts approximation ratio γ of M. Therefore, $p_{M_1}(G, c) \geq c_u/\gamma$, as we wanted to show.

Case 2 $M(G, c) = \{u\}$: The argument is along the lines of Case 1.

■

Definition 4.4.7 *A truthful mechanism M for single-dimensional* VCP *has the no-bossiness property if when a winner decreases his cost while others keep their costs the same, the outcome of M remains unchanged. In other words, none of the winners can change the outcome of the mechanism M by reducing his cost.*

Let $\mathcal{M}^{nb}$ be the class of all truthful mechanisms for K_r with the no-bossiness property. The following mechanism, which we call $\widetilde{M}$, will be used to show that $\rho(K_r, \mathcal{M}^{nb}, \mathcal{M}(\gamma)) \geq \gamma$.

Let $\gamma > 1$ and $G = K_r$. Assume that $V(G) = \{v_1, \ldots, v_r\}$ and let $\tilde{A}$ be the vertex cover algorithm for (G, c) such that for each $i = 1, \ldots, r-1$, $v_i \in \tilde{A}(G, c)$ iff $(c_i \leq c_r/\gamma)$ or $(c_i \leq c_j$ for some $1 \leq j \neq i \leq r-1)$, and $v_r \in \tilde{A}(G, c)$ iff $(c_r < \gamma c_j$ for some $1 \leq j \leq r-1)$. It is easy to check that $\tilde{A}$ is MON and hence, by Theorem 1.2.8, there exist unique payments with which it forms a truthful mechanism $\widetilde{M}$.

Lemma 4.4.8 $\widetilde{M} \in \mathcal{M}(\gamma)$ *for* $\gamma \geq 2$.

Proof Sketch : We only need to show that $\widetilde{M}$ has approximation ratio γ. Let, WLOG, assume that $c_1 \leq c_2 \leq \ldots \leq c_{r-1}$ and hence $\{v_1, \ldots, v_{r-2}\}$ is a min-cost vertex cover of (K_{r-1}, c_{-r}) induced by vertices $\{v_1, \ldots, v_{r-1}\}$.

We know that $\{v_1, \ldots, v_{r-2}\} \in \widetilde{M}(G, c)$ and also if $q = \sum_{i=1}^{r-2} c_i$ then the min-cost vertex cover of (G, c) has cost $q + \min\{c_{r-1}, c_r\}$. Now if $v_{r-1} \in \widetilde{M}(G, c)$ then either $c_{r-1} = c_{r-2}$ (which implies $q \geq c_{r-1}$) or $c_{r-1} \leq c_r/\gamma$. Moreover, if $v_r \in \widetilde{M}(G, c)$ then $c_r < \gamma c_{r-1}$. So, conditioning on whether v_{r-1} or v_r (or both) are in $\widetilde{M}(G,c)$, and whether $c_{r-1} \le c_r$ or the opposite, we can see that $\frac{c(\widetilde{M}(G,c))}{q+\min(c_{r-1},c_r)} \le \gamma$ for $\gamma \ge 2$.

■

Lemma 4.4.9 *We have*

$$\rho(K_r, \mathcal{M}^{nb}, \mathcal{M}(\gamma)) \ge \gamma$$

for all $\gamma \ge 2$.

Proof : Let $G = K_r$ and $M \in \mathcal{M}^{nb}$. We need to show there exist $M_1 \in \mathcal{M}(\gamma)$ and $c \in \mathbb{R}^r_+$ such that $p_M(G,c)/p_{M_1}(G,c) \ge \gamma$.

Let $V(G) = \{v_1, \ldots, v_r\}$ and $c = (c_1, \ldots, c_n)$. Suppose, WLOG, $\{v_1, \ldots, v_{r-1}\} \in M(G,(a,\ldots,a))$ for some $a > 0$. Since M has no-bossiness property, we have $M(G, c = (0,\ldots,0,a)) = M(G,(a,\ldots,a))$ and hence $p_M(G,c) = \sum_{i=1}^{r-1} b_{v_i}(M, c_{-v_i}) \ge \sum_{i=1}^{r-1} a = (r-1)a$.
Now let $M_1 = \widetilde{M}$. By Lemma 4.4.8 we have $M_1 \in \mathcal{M}(\gamma)$ for $\gamma \ge 2$. Furthermore, $M_1(G,c) = \{v_1, \ldots, v_{r-1}\}$ and $p_M(G,c) = \sum_{i=1}^{r-1} b_{v_i}(M, c_{-v_i}) = \sum_{i=1}^{r-1} a/\gamma = (r-1)a/\gamma$.

Therefore, $\frac{p_M(G,c)}{p_{M_1}(G,c)} \ge \frac{(r-1)a}{(r-1)a/\gamma} \ge \gamma$ for $\gamma \ge 2$.

■

Chapter 5

Future Directions

In this thesis, we initiated a systematic study of multidimensional covering mechanism design problems. We have only scratched the surface in this area, and a variety of open questions remain in this area. We briefly list some open questions regarding the various objectives considered in the thesis.

Cost Minimization (CM). A natural question is whether one can obtain efficient truthful, approximation mechanisms for other multidimensional covering problems (e.g., set cover, Steiner tree) where the approximation ratio is close to the approximation known for the underlying algorithmic problem.
For vertex cover, there is a large gap between the approximation known for the algorithmic problem and what we achieve for the multidimensional problem. It would be interesting to either obtain some lower bounds, or improve the upper bounds for Multi-VCP (i.e., the mechanism design problem). Even finding a lower bound on the approximation that can be achieved by threshold mechanisms for Multi-VCP would be illuminating.

Payment Minimization (PayM). The first question is if there is a gap between robust-BIC and DSIC notions in terms of optimal expected payment?
The running time of the robust-BIC mechanisms that we have obtained are polytime with respect to the number of items and players and the size of type-distribution, i.e. $|\mathcal{D}|$. The

question is if one can improve the running time to be dependent on marginal distribution, i.e. $\max_i |\mathcal{D}_i|$, instead of $|\mathcal{D}|$, parallel to the results in [5]?

Frugality. The frugality objective seems to be completely unexplored in multidimensional settings. It is worthwhile to investigate here whether there are other, potentially stronger, frugality benchmarks that are better suited to multidimensional settings. We have obtained some preliminary results for the frugality notion that we propose in Chapter 4, but it is clear that there are wide gaps in our understanding here. Improving the state-of-the-art here is an interesting and challenging research direction.

Finally, a general question that applies to both the CM and PayM objectives is whether there are other reasonable notions of truthfulness, stronger or weaker than the ones we have already considered, under which one can get interesting positive results?

APPENDICES

Appendix A

Black-box Reduction from Revenue- to Social-Welfare- Maximization for Packing Problems

We can leverage our ideas, in particular those in Theorem 3.6.2 and Lemma 3.7.2, to devise a powerful black-box reduction for packing problems showing that any "integrality-gap verifying" ρ-approximation algorithm $\mathcal{A}$ for the SWM problem can be used to obtain a DSIC-in-expectation mechanism whose revenue is at least a ρ-fraction of the optimum revenue. This is substantially stronger than the reduction in [17] in two respects: (a) it utilizes *approximation algorithms*, and (b) it applies also to *multidimensional* problems. In particular, we obtain, in the demand-oracle model, the first results for (multi-unit) combinatorial auctions with general valuations, and multi-unit auctions (Theorem A.1.1).

We briefly sketch the main ideas. we consider the prototypical problem of combinatorial auctions (CAs), where a feasible allocation ω consisting of allocating a disjoint set ω_i of items (which could be empty) to each player i, and player i's value under allocation ω is $v_i(\omega_i)$, where $v_i : 2^{[m]} \mapsto \mathbb{R}_+$ is player i's private valuation function. We use $v_i(\omega)$ to denote $v_i(\omega_i)$. Let V_i denote the set of all private types of player i, and $V_{-i} = \prod_{j \neq i} V_j$.

Dobzinski et al. [17] also consider an LP-relaxation for the revenue-maximization problem, but their variables encode the probability that a player i receives a set S of items.

Our LP (R-P) below is subtly, but as it turns out, crucially different from the LP in [17] in that we force the LP to also deliver a convex combination of integer/fractional allocations that will define the randomized mechanism. As Dobzinski et al. mention, there are two issues with their LP-relaxation: solving the LP, and extracting a mechanism from the LP solution. An integrality-gap verifying ρ-approximation algorithm for the SWM problem can be used to decompose a solution to the LP in [17] scaled by ρ into a convex combination of integer solutions, and thereby extract a DSIC-in-expectation mechanism; however, it is unclear how to solve their LP. This is precisely the power that our approach of moving *first* to an LP-relaxation of the SWM problem and then writing down the revenue-maximization LP in terms of *solutions to the SWM-LP* affords us: since one can efficiently optimize over the LP-relaxation of the SWM problem, we can give a separation oracle for the dual of (R-P), and hence solve (R-P).

Let Ω_{LP} denote the set of feasible solutions to the standard LP-relaxation for the CA problem. Note that Ω_{LP} is a polytope. let Ω_{ext} be the set of extreme points of Ω_{LP}. We consider the following LP. Since we are in a packing setting, we do not need the m_i estimates. We may assume that each $\mathcal{D}_i$ contains the valuation 0_i, where $0_i(\omega) = 0$ for all ω, since if not, we can just add this to $\mathcal{D}_i$, and set $\Pr_{\mathcal{D}}(0_i, v_{-i}) = 0$. Let $\overline{\mathcal{D}}' := \bigcup_i (\mathcal{D}_i \times \mathcal{D}_{-i})$. Throughout ω indexes Ω_{ext}.

$$\max \quad \sum_{v \in \mathcal{D}} \Pr_{\mathcal{D}}(v) \sum_i p_{i,v} \tag{R-P}$$

$$\text{s.t.} \quad \sum_\omega x_{v,\omega} = 1 \qquad \forall v \in \overline{\mathcal{D}}' \tag{A.1}$$

$$\sum_\omega v_i(\omega) x_{(v_i,v_{-i}),\omega} - p_{i,(v_i,v_{-i})} \geq \sum_\omega v_i(\omega) x_{(v'_i,v_{-i}),\omega} - p_{i,(v'_i,v_{-i})} \qquad \forall i, v_i, v'_i \in \mathcal{D}_i, v_{-i} \in \mathcal{D}_{-i} \tag{A.2}$$

$$\sum_\omega v_i(\omega) x_{(v_i,v_{-i}),\omega} - p_{i,(v_i,v_{-i})} \geq 0 \qquad \forall i, v \in \overline{\mathcal{D}}' \tag{A.3}$$

$$p, x \geq 0. \tag{A.4}$$

The dual is:

$$\min \quad \sum_{v} \gamma_v \tag{R-D}$$

$$\text{s.t.} \quad \sum_{i: v \in \mathcal{D}_i \times \mathcal{D}_{-i}} \Big(\sum_{v'_i \in \mathcal{D}_i} \big(v_i(\omega) y_{i,(v_i,v_{-i}),v'_i} - v'_i(\omega) y_{i,(v'_i,v_{-i}),v_i} \big) + v_i(\omega)\beta_{i,v} \Big)$$
$$\leq \gamma_c \qquad \forall v \in \overline{\mathcal{D}}', \omega \in \Omega \tag{A.5}$$

$$\sum_{v'_i \in \mathcal{D}_i} \big(y_{i,(v_i,v_{-i}),v'_i} - y_{i,(v'_i,v_{-i}),v_i} \big) \quad + \quad \beta_{i,v_i,v_{-i}} \geq Pr_D(v)$$
$$\forall i, v \in \overline{\mathcal{D}}' \tag{A.6}$$

$$y, \beta \geq 0. \tag{A.7}$$

As in the case of Theorem 3.5.4, the separation problem for (R-D) amounts to solving an SWM problem over the allocation space Ω_{ext} for an input that might involve negative valuations. Since we can efficiently optimize over Ω_{LP}, the dual, and hence the primal (R-P), can be solved efficiently. So let (x, p) be an optimal solution to (R-P) with polynomial-size support.

We argue that this can be converted to a (DSIC-in-expectation, IR) mechanism whose revenue is at least a ρ-fraction of the value of p. Let Ω be the set of all feasible (integer) allocations to the CA problem. Applying the convex-decomposition theorem in [32], we can use $\mathcal{A}$ to express x/ρ as a convex combination of integer allocations. Hence, $(\tilde{x} = x/\rho, \tilde{p} = p/\rho)$ yields a feasible solution to the LP (R-P') that is identical to (R-P), except that *the allocation-space is now the set Ω of all feasible integer allocations.*

A.1 Extending $(\tilde{x}, \tilde{p})$ to a (DSIC-in-expectation, IR) Mechanism

Finally, we show how to convert $(\tilde{x}, \tilde{p})$ to a (DSIC-in-expectation, IR) mechanism $M = \big(\mathcal{A}, \{q_i\}\big)$ with no smaller expected total revenue. Here, $\mathcal{A}(v)$ and $q_i(v)$ are the allocation-distribution and expected price of player i on input v. This construction works for *any* feasible solution to (R-P'); since any robust (BIC-in-expectation, IR) mechanism yields a feasible solution to (R-P'), this shows that in combinatorial-auction settings, *any robust (BIC-in-expectation, IR)-mechanism can be extended to a (DSIC-in-expectation, IR) mechanism without any loss in revenue.*

Our argument is similar to that in the proof of Theorem 3.5.2, but the packing nature of the problem simplifies things significantly. First, we set $\mathcal{A}(v) = \tilde{x}_v,\ q_i(v) = \tilde{p}_{i,v}$ for all $v \in \overline{\mathcal{D}}'$ and all i, so it is clear that the expected total revenue of M is the value of $(\tilde{x}, \tilde{p})$.

If $|\{i : v_i \notin \mathcal{D}_i\}| \geq 2$, then we give everyone the empty-set and charge everyone 0. Otherwise, suppose $v_i \notin \mathcal{D}_i,\ v_{-i} \in \mathcal{D}_{-i}$. Let $\bar{v}^{(i)} = \arg\max_{\tilde{v}_i \in \mathcal{D}_i} \big(\sum_\omega v_i(\omega)\tilde{x}_{(\tilde{v}_i, v_{-i}),\omega} - \tilde{p}_{i,(\tilde{v}_i, v_{-i})}\big)$ and $\bar{y}^{(i)} = \tilde{x}_{(\bar{v}_i^{(i)}, v_{-i})}$. For $\omega \in \Omega$, let $\mathsf{proj}_i(\omega)$ denote the allocation where player i receives $\omega_i \subseteq [m]$, and the other players receive $\emptyset$. Viewing $\mathcal{A}(v)$ as the random variable specifying the allocation selected, we set $\mathcal{A}(v) = \mathsf{proj}_i(\omega)$ with probability $\bar{y}^{(i)}_\omega$. We set $q_i(v) = \tilde{p}_{i,(\bar{v}^{(i)}, v_{-i})}$. Since $0_i \in \mathcal{D}_i$, we have $\mathrm{E}_{\mathcal{A}}[v_i(\mathcal{A}(v))] - q_i(v) \geq \sum_\omega 0_i(\omega)\tilde{x}_{(0_i, v_{-i}),\omega} - \tilde{p}_{i,(0_i, v_{-i})} \geq 0$, so M is IR-in-expectation.

To see that M is DSIC-in-expectation, consider some $i,\ v_i, v_i' \in V_i,\ v_{-i} \in V_{-i}$. If $v_{-i} \notin \mathcal{D}_{-i}$, then player i always receives the empty set and pays 0. Otherwise, we have ensured by definition that player i does not benefit by lying.

There are polytime integrality-gap verifying algorithms for multi-unit CAs with B copies per item, and multi-unit auctions (where all items are identical) with approximation ratios $O\big(m^{\frac{1}{B+1}}\big)$ and 2 respectively. Thus, we obtain the following.

Theorem A.1.1 *With demand oracles, we can obtain polytime (DSIC-in-expectation, IR) mechanisms with the following approximation ratios for the revenue-maximization problem: (i) $\sqrt{m}$ for CAs with general valuations; (ii) $O\big(m^{\frac{1}{B+1}}\big)$ for multi-unit CAs with B copies of each item and general valuations; (iii) 2-approximation for multi-unit auctions with general valuations.*

Permissions

With kind permission of Springer Science+Business Media for the following articles to be used in Chapter 2 and Chapter 3, respectively.

Hadi Minooei and Chaitanya Swamy. Truthful mechanism design for multidimensional covering problems. In *WINE*, pages 448–461, 2012.

Hadi Minooei and Chaitanya Swamy. Near-optimal and robust mechanism design for covering problems with correlated players. In *WINE*, 2013.

References

[1] Saeed Alaei, Hu Fu, Nima Haghpanah, Jason D. Hartline, and Azarakhsh Malekian. Bayesian optimal auctions via multi- to single-agent reduction. In *ACM Conference on Electronic Commerce*, page 17, 2012. 10, 38, 40

[2] Aaron Archer and Éva Tardos. Frugal path mechanisms. *ACM Trans. Algorithms*, 3:3:1–3:22, February 2007. 10, 64, 65

[3] Itai Ashlagi, Shahar Dobzinski, and Ron Lavi. An optimal lower bound for anonymous scheduling mechanisms. In *ACM Conference on Electronic Commerce*, pages 169–176, 2009. 10, 14

[4] Patrick Briest, Piotr Krysta, and Berthold Vöcking. Approximation techniques for utilitarian mechanism design. In *STOC*, pages 39–48, 2005. 8, 15, 17

[5] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *FOCS*, pages 130–139, 2012. 10, 38, 40, 41, 80

[6] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Reducing revenue to welfare maximization: Approximation algorithms and other generalizations. In *SODA*, pages 578–595, 2013. 10, 38, 40, 41

[7] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Understanding incentives: Mechanism design becomes algorithm design. *CoRR*, abs/1305.4002, 2013. 10, 38, 40, 41

[8] Gruia Calinescu. Bounding the payment of approximate truthful mechanisms. In *ISAAC*, pages 221–233, 2004. 64

[9] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *STOC*, pages 311–320, 2010. 10, 38, 40

[10] Ning Chen, Edith Elkind, Nick Gravin, and Fedor Petrov. Frugal mechanism design via spectral techniques. In *FOCS*, pages 755–764, 2010. 10, 24, 25, 64, 65

[11] Edward H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971. 13

[12] Jacques Cremer and Richard P McLean. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica*, 53(2):345–61, March 1985. 10, 38

[13] Jacques Cremer and Richard P McLean. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica*, 56(6):1247–57, November 1988. 10, 38

[14] Nikhil R. Devanur, Milena Mihail, and Vijay V. Vazirani. Strategyproof cost-sharing mechanisms for set cover and facility location games. *Decis. Support Syst.*, 39:11–22, March 2005. 15

[15] R. Diestel. *Graph Theory*. Electronic library of mathematics. Springer, 2005. 26

[16] Shahar Dobzinski. Two randomized mechanisms for combinatorial auctions. In *APPROX-RANDOM*, pages 89–103, 2007. 2, 14

[17] Shahar Dobzinski, Hu Fu, and Robert D. Kleinberg. Optimal auctions with correlated bidders are easy. In *STOC*, pages 129–138, 2011. 2, 10, 38, 40, 41, 44, 45, 46, 61, 83, 84

[18] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Truthful randomized mechanisms for combinatorial auctions. In *Proceedings of the thirty-eighth annual ACM*

symposium on Theory of computing, STOC '06, pages 644–652, New York, NY, USA, 2006. ACM. 2, 14

[19] Shahar Dobzinski and Jan Vondrák. The computational complexity of truthfulness in combinatorial auctions. In *ACM Conference on Electronic Commerce*, pages 405–422, 2012. 2, 14

[20] Shaddin Dughmi and Tim Roughgarden. Black-box randomized reductions in algorithmic mechanism design. In *FOCS*, pages 775–784, 2010. 2, 8, 10, 16, 17, 23

[21] Shaddin Dughmi, Tim Roughgarden, and Qiqi Yan. From convex optimization to randomized mechanisms: toward optimal combinatorial auctions. In *STOC*, pages 149–158, 2011. 2, 14

[22] Shaddin Dughmi and Jan Vondrák. Limitations of randomized mechanisms for combinatorial auctions. In *FOCS*, pages 502–511, 2011. 2, 14

[23] Edith Elkind, Leslie Ann Goldberg, and Paul W. Goldberg. Frugality ratios and improved truthful mechanisms for vertex cover. In *Proceedings of the 8th ACM conference on Electronic commerce*, EC '07, pages 336–345, New York, NY, USA, 2007. ACM. 10, 64, 65, 66, 67, 68, 69

[24] Theodore Groves. Incentives in teams. *Econometrica*, 41(4):617–31, July 1973. 13

[25] S.L. Hakimi. On the degrees of the vertices of a directed graph. *Journal of the Franklin Institute*, 279(4):290 – 308, 1965. 26

[26] Sergiu Hart and Noam Nisan. Approximate revenue maximization with multiple items. In *ACM Conference on Electronic Commerce*, page 656, 2012. 2, 10, 38, 40

[27] Kamal Jain, Mohammad Mahdian, Evangelos Markakis, Amin Saberi, and Vijay V. Vazirani. Greedy facility location algorithms analyzed using dual fitting with factor-revealing lp. *J. ACM*, 50(6):795–824, 2003. 23, 58

[28] Anna R. Karlin, David Kempe, and Tami Tamir. Beyond vcg: Frugality of truthful mechanisms. In *FOCS*, pages 615–626, 2005. 10, 64, 65, 66, 67

[29] David Kempe, Mahyar Salek, and Cristopher Moore. Frugal and truthful auctions for vertex covers, flows and cuts. In *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, FOCS '10, pages 745–754, Washington, DC, USA, 2010. IEEE Computer Society. 10, 25, 64, 65, 66, 69

[30] Jon Kleinberg and Eva Tardos. *Algorithm Design*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2005. 1

[31] Ron Lavi, Ahuva Mu'alem, and Noam Nisan. Towards a characterization of truthful combinatorial auctions. In *FOCS*, pages 574–583, 2003. 2, 14

[32] Ron Lavi and Chaitanya Swamy. Truthful and near-optimal mechanism design via linear programming. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '05, pages 595–604, Washington, DC, USA, 2005. IEEE Computer Society. 2, 8, 14, 15, 16, 17, 20, 23, 40, 45, 53, 85

[33] Ron Lavi and Chaitanya Swamy. Truthful mechanism design for multidimensional scheduling via cycle monotonicity. *Games and Economic Behavior*, 67(1):99–124, 2009. 10, 14

[34] Shi Li. A 1.488 approximation algorithm for the uncapacitated facility location problem. *Inf. Comput.*, 222:45–58, 2013. 57

[35] W. Mader. Homomorphieeigenschaften und mittlere kantendichte von graphen. *Mathematische Annalen*, 174(4):265–268, 1967. 26

[36] R Preston McAfee and Philip J Reny. Correlated information and mechanism design. *Econometrica*, 60(2):395–421, March 1992.

[37] Hadi Minooei and Chaitanya Swamy. Truthful mechanism design for multidimensional covering problems. In *WINE*, pages 448–461, 2012. 11

[38] Hadi Minooei and Chaitanya Swamy. Near-optimal and robust mechanism design for covering problems with correlated players. In *WINE*, 2013. 11

[39] Elchanan Mossel, Christos H. Papadimitriou, Michael Schapira, and Yaron Singer. Vc v. vcg: Inapproximability of combinatorial auctions via generalizations of the vc dimension. *CoRR*, abs/0905.1995, 2009. 2

[40] Ahuva Mu'alem and Noam Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. *Games and Economic Behavior*, 64(2):612–631, 2008. 17

[41] Roger B. Myerson. Optimal auction design. Discussion Papers 362, Northwestern University, Center for Mathematical Studies in Economics and Management Science, December 1978. 37

[42] Noam Nisan and Amir Ronen. Algorithmic mechanism design (extended abstract). In *STOC*, pages 129–140, 1999. 10, 14

[43] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007. 7, 14, 62

[44] Christos H. Papadimitriou and George Pierrakos. On optimal single-item auctions. In *Proceedings of the 43rd annual ACM symposium on Theory of computing*, STOC '11, pages 119–128, New York, NY, USA, 2011. ACM. 10, 38, 40

[45] A. Ronen and A. Saberi. On the hardness of optimal auctions. In *Foundations of Computer Science, 2002. Proceedings. The 43rd Annual IEEE Symposium on*, pages 396–405, 2002. 2, 61

[46] Amir Ronen. On approximating optimal auctions (extended abstract). In *The Third ACM Conference on Electronic Commerce (EC01)*, 2001. 40, 46, 61

[47] Tim Roughgarden and Inbal Talgam-Cohen. Optimal and near-optimal mechanism design with interdependent values. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, EC '13, pages 767–784, New York, NY, USA, 2013. ACM. 10, 38, 40

[48] Kunal Talwar. The price of truth: Frugality in truthful mechanisms. In *Proceedings of the 20th Annual Symposium on Theoretical Aspects of Computer Science*, STACS '03, pages 608–619, London, UK, UK, 2003. Springer-Verlag. 64

[49] Vijay V. Vazirani. *Approximation algorithms*. Springer, 2001. 1, 68

[50] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 03 1961. 13

[51] Robert J. Weber. Auction theory: By vijay krishna. academic press, 2002. *Games and Economic Behavior*, 45(2):488–497, November 2003. 10, 38