

On Sparsification of Stochastic Packing Problems

Shaddin Dughmi ✉ 🏠 

University of Southern California, Los Angeles, CA, USA

Yusuf Hakan Kalayci ✉ 🏠 

University of Southern California, Los Angeles, CA, USA

Neel Patel ✉ 

University of Southern California, Los Angeles, CA, USA

Abstract

Motivated by recent progress on stochastic matching with few queries, we embark on a systematic study of the sparsification of stochastic packing problems more generally. Specifically, we consider packing problems where elements are independently active with a given probability p , and ask whether one can (non-adaptively) compute a “sparse” set of elements guaranteed to contain an approximately optimal solution to the realized (active) subproblem. We seek structural and algorithmic results of broad applicability to such problems. Our focus is on computing sparse sets containing on the order of d feasible solutions to the packing problem, where d is linear or at most polynomial in $\frac{1}{p}$. Crucially, we require d to be independent of the number of elements, or any parameter related to the “size” of the packing problem. We refer to d as the “degree” of the sparsifier, as is consistent with graph theoretic degree in the special case of matching.

First, we exhibit a generic sparsifier of degree $\frac{1}{p}$ based on contention resolution. This sparsifier’s approximation ratio matches the best contention resolution scheme (CRS) for any packing problem for additive objectives, and approximately matches the best monotone CRS for submodular objectives. Second, we embark on outperforming this generic sparsifier for additive optimization over matroids and their intersections, as well as weighted matching. These improved sparsifiers feature different algorithmic and analytic approaches, and have degree linear in $\frac{1}{p}$. In the case of a single matroid, our sparsifier tends to the optimal solution. In the case of weighted matching, we combine our contention-resolution-based sparsifier with technical approaches of prior work to improve the state of the art ratio from 0.501 to 0.536. Third, we examine packing problems with submodular objectives. We show that even the simplest such problems do not admit sparsifiers approaching optimality. We then outperform our generic sparsifier for some special cases with submodular objectives.

2012 ACM Subject Classification Theory of computation → Packing and covering problems

Keywords and phrases Stochastic packing, sparsification, matroid

Digital Object Identifier 10.4230/LIPIcs.ICALP.2023.51

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version:* <https://arxiv.org/pdf/2211.07829.pdf>

Funding Supported by NSF grant CCF-2009060.

Acknowledgements We are grateful to the anonymous reviewers for their thoughtful feedback on the earlier version of this paper.

1 Introduction

Our starting point for this paper is the beautiful line of recent work on variants of the stochastic matching problem, seeking approximate solutions with limited query access to the (stochastic) data [9, 3, 2, 4, 8, 7, 6, 5]. Notably, many of the algorithms in these works are non-adaptive, and can therefore be interpreted as “sparsifiers” for the stochastic problem.



© Shaddin Dughmi, Yusuf Hakan Kalayci, and Neel Patel;

licensed under Creative Commons License CC-BY 4.0

50th International Colloquium on Automata, Languages, and Programming (ICALP 2023).

Editors: Kousha Etessami, Uriel Feige, and Gabriele Puppis; Article No. 51; pp. 51:1–51:17

Leibniz International Proceedings in Informatics



Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



These works feature powerful new algorithmic and analytic sparsification techniques of possibly more general interest, suggesting that effective sparsifiers might exist well beyond matching and closely related problems.

Our goal in this paper is to coalesce a broader agenda on the sparsification of combinatorial stochastic optimization problems more generally, beginning with the natural and broad class of packing problems. We ask, and make progress on, the fundamental questions: For which stochastic packing problems is effective sparsification possible? What are the algorithmic techniques and blueprints which are broadly applicable? What are the barriers to progress?

Concretely, we examine stochastic packing problems (SPPs) of the following (fairly general) form. We are given a set system (E, \mathcal{I}) , where E is a finite set of *elements* and $\mathcal{I} \subseteq 2^E$ is a downwards-closed family of *feasible sets* (often also referred to as *independent sets*, in particular for matroids). Also given is an objective function $f : 2^E \rightarrow \mathbb{R}_+$, which we assume to be either additive (a.k.a. modular) or submodular. The stochastic uncertainty is described by a given probability $p \in [0, 1]$: We assume that each element of E is *active*, i.e., viable for being selected, independently with probability p . The goal of the SPP is to select a feasible set of active elements maximizing the objective function.

When the set R of active elements is given, or can be queried without restriction, this reduces to non-stochastic optimization for the induced subproblem on R . We refer to the output of such an omniscient [approximation] algorithm as an [approximate] *stochastic optimum solution*. We are instead concerned with algorithms that approximate the stochastic optimum by querying the activity status of only a small, a.k.a. “sparse”, set of elements $Q \subseteq E$. In particular, as in much of the prior work we require the queried set Q to be chosen non-adaptively. Such algorithms can equivalently be thought of as factoring into two steps: First, a *sparsification algorithm* (or *sparsifier* for short) computes a (possibly random) set of elements Q . Second, we learn $R \cap Q$, and an [approximate] optimization algorithm is applied to the (now fully-specified) subproblem induced by $R \cap Q$. Since the second (optimization) step is familiar and well-studied, our focus is on the first step, namely sparsification.

We evaluate a sparsifier by two quantities. The first quantity is a familiar one, namely its *approximation ratio*. Specifically, a sparsifier is α -approximate if it guarantees an α -approximation to the stochastic optimum solution when combined with a suitable algorithm in the second (optimization) step. The second quantity is a measure of the “sparsity” of the set Q selected by the sparsifier. We say our sparsifier is of *degree* d if it guarantees $\mathbb{E}[|Q|] \leq d \cdot r$, where $r = \max\{|S| : S \in \mathcal{I}\}$ is the *rank* of the set system (E, \mathcal{I}) . Intuitively, the sparsification degree refers to the level of “contingency” or “redundancy” in the sparsified instance, relative to the size of maximal feasible solutions. Loosely speaking, the degree of a sparsifier roughly measures “how many” feasible solutions are maintained to account for uncertainty in the problem. Somewhat fortuitously, our definition of degree specializes to the (average) graph-theoretic degree in the special case of matching, lending consistency with prior work on stochastic matching with few queries.

We study sparsifiers whose degree admits an upperbound that is independent of the size of the system; The degree bound can not depend on the number of elements or the rank of the set system, for example. We focus especially on the “polynomial regime”, where the degree is restricted to be at most polynomial in $\frac{1}{p}$. We pursue sparsifiers which are constant-approximate, or in the best case $(1 - \epsilon)$ -approximate for arbitrarily small $\epsilon > 0$.

Results and Techniques

We begin with the observation that a degree of at least $\frac{1}{p}$ is necessary for constant-approximate sparsification, even for the simplest of packing problems: a rank one matroid and the unweighted additive objective. We then establish a “baseline” of possibility for all stochastic

packing problems, through a generic sparsifier with this same degree $\frac{1}{p}$. This sparsifier is simple: it computes (or estimates) the marginals $\{q_e\}_{e \in E}$ of the stochastic optimum solution, and outputs a set Q which includes each element e independently with probability $\frac{q_e}{p}$. For SPPs with an additive objective, we show that this sparsifier’s approximation ratio matches the balance ratio of the best contention resolution scheme (CRS)¹ for the set system. When the objective is submodular, we approximately match the balance ratio of the best *monotone* CRS up to a factor of $1 - \frac{1}{e}$. We note that contention resolution is only used as a proof tool to certify our sparsifier’s approximation guarantee, and is not invoked algorithmically. In settings where the marginals $\{q_e\}_{e \in E}$ are intractable to compute, this sparsifier can be made computationally efficient by resorting to approximation, in which case its approximation ratio degrades in the expected manner. This generic result yields constant-approximate sparsifiers of degree $\frac{1}{p}$ for a large variety of set systems for which contention resolution has been studied, including matroids and their intersections.

Next, we embark on “beating” this contention resolution baseline for natural SPPs. We succeed at doing so for additive (weighted) optimization over matroids, matroid intersections, and matchings. For a single matroid, we derive a simple greedy sparsifier which is $(1 - \epsilon)$ -approximate and has degree $\frac{1}{p} \cdot \log(1/\epsilon)$. This sparsifier repeatedly adds a maximum weight independent set of the matroid to the sparse set Q , and removes it from the matroid, until the desired degree is reached. Though our sparsifier is simple, its analysis is (we believe necessarily) less so.

For matroid intersections, we first argue that adaptations of our single-matroid sparsifier cannot succeed, due to feasible sets not “combining well” as they do in the case of a single matroid. Instead, our sparsifier for matroid intersections repeatedly samples the stochastic optimum solution and adds it to the sparse set Q , for a degree of $O(\frac{1}{p\epsilon} \cdot \log(1/\epsilon))$. The approximation ratio of our sparsifier for the intersection of k matroids is $\frac{1-\epsilon}{k+1/(k+1)}$, which beats the best known bound on the correlation gap of $1/(k+1)$ [1]. The analysis of this sparsifier is again nontrivial, and utilizes basis exchange maps.

For matroids and matroid intersections, we note that analysis techniques employed by prior work on matching do not appear to suffice. In particular, prior work on matching often employs concentration arguments on the active degree of matroid “flats” containing an element; this is sufficient in the case of matching, since each element is in at most two binding flats (one for each partition matroid). For general matroids, such concentration arguments fail to bound the degree in a manner independent of the number of elements, necessitating alternative proof approaches like ours.

For general (non-bipartite) matching, we augment our contention-resolution-baseline sparsifier with samples from the stochastic optimum solution, for a total degree of $O(1/p)$. We show that the samples from the stochastic optimum combine well with our baseline sparsifier. We obtain an approximation ratio which is a function of the (as yet not fully known) correlation gap of the matching polytope. This function exceeds the identity function everywhere, implying that our sparsifier strictly improves on the contention resolution baseline. Plugging in the best known lowerbound of 0.474 on the correlation gap from [17], we guarantee that our sparsifier is 0.536 approximate. This improves on the state of the art in the polynomial regime, 0.501-approximate sparsifier of degree $\text{poly}(1/p)$ due to [8]. In addition, assuming the conjecture from [20] which states the existence of 0.544 balanced CRS for general matching polytope implies that our sparsifier is 0.598 approximate.

¹ This is equal to the set system’s *correlation gap*, as shown by [12].

■ **Table 1** Summary of information theoretic sparsifiers for additive objectives. Here, n is the number of elements and W is the maximum element weight.

Constraint	Previous Results		This Work	
	Approx. Ratio	Sparsification Degree	Approx. Ratio	Sparsification Degree
Matroid	$1 - \epsilon$ [14]	$O\left(\frac{1}{p} \log \frac{\mathbf{Rank}}{\epsilon}\right)$	$1 - \epsilon$	$\frac{1}{p} \cdot \log\left(\frac{1}{\epsilon}\right)$
k -Matroid Intersection	$\frac{1-\epsilon}{2k}$ [18]	$O\left(\frac{W}{p} \log n \log\left(\frac{n}{\epsilon}\right) \frac{k}{\epsilon^3}\right)$	$\frac{1-\epsilon}{k+\frac{1}{k+1}}$	$\frac{1}{\epsilon \cdot p} \cdot \log\left(\frac{1}{\epsilon}\right)$
General Matching	$1 - \epsilon$ [6]	$O(\exp(\exp(\exp(1/p))))$	0.536	$O\left(\frac{1}{p}\right)$
General Matching	0.501 [8]	$O\left(\frac{1}{p}\right)$	0.536	$O\left(\frac{1}{p}\right)$

Finally, we further examine stochastic packing problems with submodular objectives. Our $(1 - \epsilon)$ -approximate sparsifier for weighted matroid optimization might tempt one to conjecture a similar result for submodular optimization over simple enough set systems. However, we show by way of an information-theoretic impossibility result that no sparsifier with degree bound independent of the number of elements can beat $(1 - 1/e)$, even for optimizing a coverage function subject to a uniform matroid constraint. We complement this impossibility result with algorithmic sparsification results for optimizing coverage functions over matroids, improving over the guarantees provided by our baseline generic sparsifier. Due to limited space, the results for submodular SPPs are detailed in the full version of this paper [13](Section 8).

Additional Discussion of Related Work

The exploration of sparsifying SPPs was initiated by [9], who focus on the unweighted stochastic matching problem. This problem has since been studied extensively in a series of works [4, 8, 7, 6] which attempt to beat the benchmark set by [9]. In the “polynomial-degree regime”, the state-of-the-art sparsifier for unweighted stochastic matching is a 0.66-approximation due to [2]. Recent work by [5] improves this approximation to $\frac{e}{e+1}$ for unweighted bipartite matching. For weighted stochastic matching in the polynomial-degree regime, the current best known sparsifier is a 0.501-approximation due to [8]. Going beyond polynomial degree, [7, 6] constructed a $(1 - \epsilon)$ -approximate sparsifier with degree $\exp(\exp(\exp(1/p)))$ for the weighted general matching problem. The sparsifiers designed for the stochastic matching problems rely heavily on structural properties particular to matching. Our techniques, on the other hand, are targeted at more general packing problems.

To the best of our knowledge, the work of [18, 19] stands alone in directly studying the sparsification of SPPs beyond matching. In [18], they proposed a general framework for solving stochastic packing integer programs. As a corollary of their techniques, they obtain non-adaptive sparsifiers for several additive SPPs. However, the degree of their sparsification algorithms intrinsically depends on the number of elements in settings where a single element may be in an exponential number of binding constraints (as is the case for matroids). Our work, in contrast, proposes several algorithmic techniques that yield approximate sparsifiers with degree independent of the number of elements.

Also related is the work of [14], which studies the covering analogue of our question for matroids. They show how to construct a set of size $O\left(\frac{\mathbf{Rank}}{p} \log \frac{\mathbf{Rank}}{\epsilon}\right)$ which is guaranteed to contain a minimum-weight base of the matroid with high probability. This implicitly

■ **Table 2** Summary of information theoretic sparsifiers for monotone submodular objectives. All mentioned results are shown in this paper.

Constraint	Approximation Ratio	Sparsification Degree	Note
r -Uniform Matroid	$(1 - \frac{1}{e}) \cdot (1 - \frac{1}{\sqrt{r+3}})$	$\frac{1}{p}$	$(1 - \frac{1}{e})$ upperbound, Optimal when $r \rightarrow \infty$
Matroid	$(1 - \frac{1}{e})^2$	$\frac{1}{p}$	$1 - \frac{1}{e}$ upperbound
k -Matroid Intersection	$(1 - \frac{1}{e}) \cdot \frac{1}{k+1}$	$\frac{1}{p}$	

implies an $O(\frac{1}{p} \log \frac{\mathbf{Rank}}{\epsilon})$ -degree sparsifier for weighted stochastic packing on matroids. Their analysis is tight for the covering setting, and it appears nontrivial to adapt their techniques for the packing setting in order to remove the degree's dependence on the rank. We compare our results for additive SPPs with prior work in Table 1.

The manuscript [19] proposes sparsifiers for SPPs with a monotone submodular objectives. However, their sparsification algorithms are intrinsically adaptive in nature. To the best of our knowledge, ours is the first work that analyzes SPPs with submodular objectives in the non-adaptive setting. We summarize our results for submodular SPPs in Table 2.

2 Problem Definition

We consider packing problems of the form $\langle E, \mathcal{I}, f \rangle$ where E is a ground set of *elements* with cardinality n , $f : 2^E \rightarrow \mathbb{R}_{\geq 0}$ is an objective function, and $\mathcal{I} \subseteq 2^E$ is a downwards-closed family of *independent sets* (a.k.a. *feasible sets*). We use $r = \operatorname{argmax}\{|I| : I \in \mathcal{I}\}$ to denote the *rank* of the set system \mathcal{I} . The aim of the packing problem is to select an independent set $O \in \mathcal{I}$ that maximizes $f(O)$.

In this paper, we study packing problems in a particular setting with uncertainty parameterized by $p \in [0, 1]$. In a *stochastic packing problem (SPP)* $\langle E, \mathcal{I}, f, p \rangle$, nature selects a random set $R \subseteq E$ of *active elements* such that $\Pr[e \in R] = p$ independently for all $e \in E$. We are then tasked with solving the induced (random) packing problem on the active elements, namely $\langle R, \mathcal{I}|R, f|R \rangle$ where $\mathcal{I}|R$ and $f|R$ denote the restriction of \mathcal{I} and f to subsets of R , respectively. We refer to an [approximately] optimum solution to $\langle R, \mathcal{I}|R, f|R \rangle$ as an [approximate] *stochastic optimum solution*. We use **OPT** to denote the expected value of a stochastic optimum solution, i.e.,

$$\mathbf{OPT}(E, \mathcal{I}, f, p) = \mathbb{E}_R \left[\max_{\substack{T \in \mathcal{I} \\ T \subseteq R}} f(T) \right],$$

where $R \subseteq E$ is the random set which each element of E independently with probability p .

We assume that the set R of active elements is a-priori unknown, and that we can *query* elements in E to check their membership in R . Motivated by settings in which queries are costly, we seek algorithms which query a small (we say “sparse”) subset of the elements, and moreover choose those queries non-adaptively. Such non-adaptive algorithms can be thought of as factoring into two steps: A *sparsification* step which selects the small set $Q \subseteq E$ of queries, and an *optimization* step which solves the packing problem $\langle R \cap Q, \mathcal{I}|R \cap Q, f|R \cap Q \rangle$ induced by the queried active elements. For the optimization step, we assume access to a traditional [approximation] algorithm. Our focus is on algorithms for the sparsification step, which we define formally next.

A *sparsification algorithm* (or *sparsifier* for short) \mathcal{A} takes as input an SPP $J = \langle E, \mathcal{I}, f, p \rangle$ from some family of SPPs, and outputs a (possibly random) set of elements $Q \subseteq E$. The twin goals here are for Q to be “sparse” in a quantified formal sense we describe shortly, while guaranteeing that optimally solving the “sparsified” SPP $J|Q = \langle Q, \mathcal{I}|Q, f|Q, p \rangle$ yields an approximate solution to the original SPP J . We say that the sparsification algorithm \mathcal{A} is α -*approximate* if it guarantees $\mathbf{OPT}(J|Q) \geq \alpha \mathbf{OPT}(J)$ – i.e., an optimal solution to the sparsified SPP is an α -approximate solution to the original SPP. We sometimes identify the sparsified SPP $J|Q$ with Q when J is clear from the context.

To quantify sparsity, we say that \mathcal{A} has *sparsification degree* d if it guarantees that $\frac{\mathbb{E}[|Q|]}{r} \leq d$, where r is the rank of the set system \mathcal{I} , and expectation is over the internal random coins of \mathcal{A} . Intuitively, the degree of sparsification refers to the level of “contingency” or “redundancy” in the sparsified instance, relative to the size of maximal feasible solutions. Loosely speaking, the degree of a sparsifier roughly measures “how many” feasible solutions it maintains to account for uncertainty in the problem.

In the absence of a bound on degree, an approximation factor of $\alpha = 1$ is trivially achievable. We aim to construct approximate sparsifiers of low degree for natural classes of SPPs. We begin by observing that a degree of $\Omega(1/p)$ is necessary for constant approximation, even for the simplest of constraints.

► **Example 1.** Consider the SPP with n elements, the unweighted additive objective function $f(S) = |S|$, a rank-one matroid constraint, and activation probability $p = 1/n$. There is at least one active element with probability $1 - (1 - p)^n \geq 1 - 1/e$, therefore $\mathbf{OPT} \geq 1 - 1/e$. On the other hand, a set of elements Q will contain no active elements with probability $(1 - p)^{|Q|} \geq 1 - |Q| \cdot p = 1 - \frac{|Q|}{n}$. When $|Q| = o(1/p) = o(n)$, there are no active elements in Q with probability $1 - o(1)$. Therefore, any constant-approximate sparsifier must have degree $\Omega(1/p)$.

We also show in the full version of this paper [13] that, unsurprisingly, there exist stochastic packing problems which do not admit constant approximate sparsifiers with degree $\text{poly}(1/p)$. Given these simple impossibility results, we ask a natural question:

► **Question 2.** Which stochastic packing problems admit constant approximate sparsifiers of degree $O\left(\frac{1}{p}\right)$, or more loosely $\text{poly}\left(\frac{1}{p}\right)$?

In this paper, we focus on designing sparsification algorithms for stochastic packing problems with additive or nonnegative monotone submodular objectives.

A Note on Input Representation

Many of our results are information theoretic, and therefore make no assumptions on how a stochastic packing problem is represented. Most of our algorithmic results, on the other hand, only require solving realized (non-stochastic) instances of the packing problem, possibly approximately. Specifically, for a stochastic packing problem $\langle E, \mathcal{I}, f, p \rangle$ we often assume access to a [β -approximate] *stochastic optimal oracle*. Such an oracle samples a [β -approximate] solution to the (random) packing problem $\langle R, \mathcal{I}|R, f|R \rangle$, where R includes each element of E independently with probability p , and $\mathcal{I}|R$ and $f|R$ denote the restriction of \mathcal{I} and f to R respectively. For our algorithmic results on matroids, we additionally assume access to an independence oracle, as is standard.

3 Sparsification from Contention Resolution

In this section, we show how to generically derive a sparsifier for a stochastic packing problem from bounds on contention resolution for the associated set system. First, we recall the relevant definition of contention resolution.

► **Definition 3** ([12]). Let (E, \mathcal{I}) be a set system, and let $P_{\mathcal{I}} = \text{convexhull}\{\mathbb{1}_I : I \in \mathcal{I}\}$ denote the associated polytope. A Contention Resolution Scheme (CRS) π for $P_{\mathcal{I}}$ is a (randomized) algorithm which takes as input a point $x \in P_{\mathcal{I}}$ and a set of active elements $R(x) \subseteq E$, including each element $i \in E$ independently with probability x_i , and outputs a feasible subset $\pi_x(R(x)) \subseteq R(x)$, $\pi_x(R(x)) \in \mathcal{I}$. For $b, c \in [0, 1]$, we say a CRS is (b, c) -balanced if for all $i \in E$ and $x \in b \cdot P_{\mathcal{I}}$, $\Pr[i \in \pi_x(R(x)) \mid i \in R(x)] \geq c$. A CRS π is monotone if for every $S \subseteq T \subseteq E$ we have that $\Pr[i \in \pi(S) \mid i \in S] \geq \Pr[i \in \pi(T) \mid i \in T]$.

Our generic sparsifier is randomized, has degree $\frac{1}{p}$, and is shown in Algorithm 1. Our sparsifier computes estimated marginals \mathbf{q} for the stochastic optimum solution. For an information-theoretic result, we can assume these to be exact. Then it samples each element $e \in E$ in a sparse set Q with probability $\frac{q_e}{p}$.

When the objective function f is additive, our sparsifier has an approximation factor that matches the balance ratio of the best CRS for $P_{\mathcal{I}}$.² For nonnegative monotone submodular functions, the approximation factor matches the balance ratio of the best monotone CRS for $P_{\mathcal{I}}$. This is due to the observation that each element $e \in E$ is included in the active subset of the sparse set Q with probability q_e and the fact that $\mathbf{q} \in P_{\mathcal{I}}$. The detailed proof for Theorem 4 can be found in the full version [13].

■ **Algorithm 1** Generic Sparsifier for a Stochastic Packing Problem $\langle E, \mathcal{I}, f, p \rangle$.

Input: Stochastic packing problem $\langle E, \mathcal{I}, f, p \rangle$

Compute the marginals \mathbf{q} of the stochastic optimum solution, or an approximation thereof.

$Q \leftarrow \emptyset$;

for all $e \in E$ **do**

 Add e to Q with probability $\frac{q_e}{p}$ (independently)

end for

Output: Sparse set Q .

► **Theorem 4.** Consider Algorithm 1, implemented with exact (possibly non-polynomial-time) computation of the marginals \mathbf{q} . When f is additive, and $P_{\mathcal{I}}$ admits a c -balanced CRS, the algorithm is a c -approximate sparsifier of degree $\frac{1}{p}$. When f is a nonnegative monotone submodular, and $P_{\mathcal{I}}$ admits a c -balanced monotone CRS, the algorithm is a $c(1 - \frac{1}{e})$ -approximate sparsifier of degree $\frac{1}{p}$.

To make our sparsifier algorithmically efficient, \mathbf{q} may be estimated by sampling from a (possibly approximate) stochastic optimum oracle, in which case our guarantees degrade in the expected manner due to sampling errors and/or the approximation inherent to the oracle. We present the detailed analysis with approximate stochastic optimal oracles in Appendix B in the full version [13]. Theorem 4 and Theorem 4.3 (In full version [13]) together with contention resolution schemes from prior work [1, 12, 17] and approximate stochastic optimal oracles that employ approximation algorithms from [11, 16], imply constant approximate sparsifier for a broad class of packing constraints summarized in Table 3.

² This balanced ratio is equal to the *correlation gap* of the set system \mathcal{I} , as per [12].

■ **Table 3** Approximation Ratio of Generic Sparsifier of degree $\frac{1}{p}$ for various packing constraint families with additive and non-negative monotone submodular function.

Constraint	Additive Objective		Submodular Objective	
	Information Theoretic	Poly-Time	Information Theoretic	Poly-Time
Matroid	$(1 - \frac{1}{e})$	$(1 - \epsilon) \cdot (1 - \frac{1}{e})$	$(1 - \frac{1}{e})^2$	$(1 - \epsilon) \cdot (1 - \frac{1}{e})^3$
k -matroid intersection	$\frac{1}{k+1}$	$(1 - \epsilon) \cdot \frac{1}{k^2-1}$	$\frac{1-1/e}{k+1}$	$(1 - \epsilon) \cdot \frac{1-1/e}{k^2-1}$
Matching	0.474	$(1 - \epsilon) \cdot 0.474$	$(1 - \frac{1}{e}) \cdot 0.43$	$(1 - \frac{1}{e})^2 \cdot 0.43$

The following proposition (whose proof is delegated to the full version [13]) shows that Algorithm 1 is optimal for matroids and additive objectives among sparsifiers of degree $1/p$. This strongly suggests that sparsification is intimately tied to contention resolution when the degree is restricted to $1/p$. In particular, exceeding degree $1/p$ appears necessary for outperforming the correlation gap of a set system in general.

► **Proposition 5.** *Consider the family of stochastic packing problems with matroid constraints and additive objectives. There is no degree $\frac{1}{p}$ sparsifier for this family that achieves an approximation ratio $1 - 1/e + \Omega(1)$.*

We note that $1 - 1/e$ is the best possible balance ratio for contention resolution on the rank one matroid, as shown in [12] through the correlation gap. Given the above discussion, it is natural to ask whether we can design sparsifiers of degree $O(1/p)$, or even $\text{poly}(1/p)$, whose approximation ratio α exceeds the best CRS balance ratio c , i.e., can we have $\alpha > c$ with degree linear or polynomial in $1/p$? Recent progress on this question for bipartite matching constraints came in a pair of recent works. Behnezhad et.al. [5] designed a $\frac{e}{e+1} \approx 0.731$ -approximate sparsifier with degree $\text{poly}(1/p)$ for unweighted bipartite matching. Their approximation factor is strictly better than a known upper bound of 0.544 on the correlation gap (and hence the best balance ratio) of bipartite matching, due to [15]. To our knowledge, this is the only sparsifier in the literature with degree polynomial in $\frac{1}{p}$ and approximation ratio provably exceeding the correlation gap of the set system. Another recent result due to Behnezhad et al [7] achieves a 0.501-approximate sparsification with degree polynomial in $1/p$ for *weighted* matching. This outperforms the best *known* contention resolution scheme for matching[10], though not clearly the best possible. Prior to our work, there was no known sparsifier for any *weighted* stochastic packing problem which provably outperforms the correlation gap using degree $\text{poly}(1/p)$.

In the following sections, we will construct degree $O(1/p)$ sparsifiers for matroids, matroid intersections, and matching which improve on the contention-resolution-based guarantees provided in this section. For matroids and matchings, our sparsifiers provably outperform contention resolution. For matroid intersections, we outperform the best *known* CRS.

4 Additive Optimization over a Matroid

In this section, we design an improved sparsifier for the stochastic packing problem $\langle E, \mathcal{I}, f, p \rangle$ when $\mathcal{M} = (E, \mathcal{I})$ is a matroid and f is additive. For an arbitrary $\epsilon > 0$, our sparsifier is $(1 - \epsilon)$ -approximate and has degree $\frac{1}{p} \log \frac{1}{\epsilon}$. Throughout, we use $\{w_e\}_{e \in E}$ to denote the weights associated with the additive function f , and use $R \subseteq E$ to denote the (random) set of active elements which includes each $e \in E$ independently with probability p . We also sometimes use r as shorthand for $\mathbf{Rank}(\mathcal{M})$. We present basic preliminaries of matroid theory in the full version [13]

► **Theorem 6.** *Let $\mathcal{M} = (E, \mathcal{I})$ be a matroid, f be an additive function and $p \in [0, 1]$. Algorithm 2 is a $(1 - \epsilon)$ -approximate polynomial time sparsifier for the stochastic packing problem $\langle E, \mathcal{I}, f, p \rangle$ with sparsification degree $\frac{1}{p} \cdot \log\left(\frac{1}{\epsilon}\right)$.*

Previously, the best known sparsifier for matroid was $(1 - \epsilon)$ -approximate with degree $O(1/p \log(\mathbf{Rank}/\epsilon))$ implicit in [14]. In contrast, the sparsification degree of our algorithm is independent of the “size” of the matroid. As we argued in introduction, such a size-independent guarantee appears to be beyond the techniques used in earlier works [14, 18].

■ **Algorithm 2** Sparsifier for (\mathcal{M}, f, p) , when \mathcal{M} is a matroid and f is additive.

Set $\mathcal{M}_0 = \mathcal{M}$ and $Q = \emptyset$.

for t in $\{1, \dots, \tau\}$ where $\tau = \frac{1}{p} \cdot \log\left(\frac{1}{\epsilon}\right)$

 Let $I_t \leftarrow \operatorname{argmax}_{I \in \mathcal{I}_{t-1}} f(I)$, where \mathcal{I}_{t-1} is the collection of independent sets in \mathcal{M}_{t-1} .

 Update $\mathcal{M}_t \leftarrow \mathcal{M}_{t-1} \setminus I_t$.

Output: $Q = \bigcup_{t=1}^{\tau} I_t$.

It is clear that the sparsifier in Algorithm 2 has degree $\tau = \frac{1}{p} \cdot \log\left(\frac{1}{\epsilon}\right)$, and can be implemented in polynomial time given an independence oracle for the matroid \mathcal{M} . The remainder of this section is devoted to proving that it is $(1 - \epsilon)$ -approximate, as needed to complete the proof of Theorem 6. Our proof will consist of two parts. First, we will analyze Algorithm 2 in the special case of unit weights (a.k.a. unweighted). Second, we reduce the analysis of the weighted problem to that of the unweighted problem.

4.1 Special Case: Unweighted Optimization

In this subsection, we assume that elements of the matroid \mathcal{M} all have unit weight. In this case, observe that Algorithm 2 repeatedly removes an arbitrary basis of the matroid and adds it to the sparse set Q . More precisely, in iteration t the set I_t is a basis of the remaining matroid $\mathcal{M}_{t-1} := \mathcal{M} \setminus \bigcup_{j=1}^{t-1} I_j$.

In this unweighted case, the stochastic optimal value is the expected rank of the active elements R , and our claimed approximation guarantee can be expressed as $\mathbb{E}[\mathbf{Rank}(Q \cap R)] \geq (1 - \epsilon) \mathbb{E}[\mathbf{Rank}(R)]$. To establish this, consider the following informal (but ultimately flawed) argument, starting with the observation that $I_t \cap R$ spans a p fraction of the rank of the remaining matroid \mathcal{M}_{t-1} in expectation. This observation suggests that the rank of elements not spanned by $Q \cap R$ should shrink by a factor of $(1 - p)$ with each iteration. Induction would then guarantee that after $\frac{1}{p} \cdot \log\left(\frac{1}{\epsilon}\right)$ iterations we have covered a $(1 - \epsilon)$ fraction of the rank of the matroid.

The above rough argument is a good starting point. Indeed, it succeeds when all (or many) of the bases I_1, \dots, I_τ are full-rank or close to it. These are precisely the scenarios in which $\mathbb{E}[\mathbf{Rank}(R)] \approx \mathbf{Rank}(\mathcal{M})$. However, in general $\mathbf{OPT} = \mathbb{E}[\mathbf{Rank}(R)]$ can be significantly smaller than $\mathbf{Rank}(\mathcal{M})$ – in the worst case up to a factor of p smaller – in which case the rank of I_t may drop precipitously with t and the above inductive analysis falls apart. Such scenarios are not simply outliers that we can assume away: they are unavoidable products of the weighted-to-unweighted reduction we present in the next subsection, and can account for a large fraction of the weighted stochastic optimal. This seems to necessitate a more nuanced proof approach in which we compare $\mathbb{E}[\mathbf{Rank}(Q \cap R)]$ with $\mathbb{E}[\mathbf{Rank}(R)]$. We present such a proof next, built upon the following definitions and structural properties.

51:10 On Sparsification of Stochastic Packing Problems

► **Definition 7.** A nested system of spanning sets (NSS) for a matroid \mathcal{M} is a sequence I_1, I_2, \dots, I_τ of sets such that for any $j \in [\tau]$, I_j is a full rank set of elements in $\mathcal{M} \setminus I_{1:j-1}$, where $I_{1:j-1} = \bigcup_{\ell=1}^{j-1} I_\ell$.

► **Observation 8.** The sets I_1, \dots, I_τ from Algorithm 2 are an NSS of \mathcal{M} .

The following lemma states that the property of being an NSS is preserved under contraction.

► **Lemma 9.** Let $\mathcal{M} = (E, \mathcal{I})$ be a matroid and let I_1, \dots, I_τ be an NSS of \mathcal{M} . For an arbitrary independent set S of \mathcal{M} , let $I'_j = I_j \setminus S$ for all j . Then, the sequence I'_1, \dots, I'_τ is an NSS of \mathcal{M}/S .

Proof. Fix an arbitrary $j \in \{1, \dots, \tau\}$. It is clear that I'_j is a subset of the elements of $\mathcal{M}/S \setminus I'_{1:j-1}$. It remains to show that I'_j is full rank in $\mathcal{M}/S \setminus I'_{1:j-1}$, as follows.

$$\begin{aligned}
 \mathbf{Rank}^{\mathcal{M}/S}(I'_j) &= \mathbf{Rank}^{\mathcal{M}}(I_j \cup S) - |S| && \text{(By (2) and definition of } I'_j) \\
 &= \mathbf{Rank}^{\mathcal{M}}((E \setminus I_{1:j-1}) \cup S) - |S| && (I_j \text{ is full rank in } \mathcal{M} \setminus I_{1:j-1}) \\
 &= \mathbf{Rank}^{\mathcal{M}}((E \setminus S \setminus I'_{1:j-1}) \cup S) - |S| && \text{(By definition of } I'_j) \\
 &= \mathbf{Rank}^{\mathcal{M}/S}(E \setminus S \setminus I'_{1:j-1}) && \text{(By (2))} \\
 &= \mathbf{Rank}(\mathcal{M}/S \setminus I'_{1:j-1}) && \blacktriangleleft
 \end{aligned}$$

► **Observation 10.** If I_1, \dots, I_τ is an NSS of \mathcal{M} , then I_2, \dots, I_τ is an NSS of $\mathcal{M} \setminus I_1$.

Now, we will prove the desired result for unweighted matroids.

► **Lemma 11.** Let \mathcal{M} be a matroid, and let I_1, \dots, I_τ be an NSS of \mathcal{M} . Then,

$$\mathbb{E}[\mathbf{Rank}(I_{1:\tau} \cap R)] \geq (1 - (1-p)^\tau) \cdot \mathbb{E}[\mathbf{Rank}(R)]$$

Proof. Let E denote the elements of \mathcal{M} . We will apply induction on τ to prove this result. The base case of $\tau = 0$ is trivial.

Consider $\tau \geq 1$. Let S be an arbitrary maximal independent subset of $R \cap I_1$, and let \mathbf{Rank}' denote the rank function of the (random) matroid $\mathcal{M}' = \mathcal{M}/S \setminus I_1$ with elements $E \setminus I_1$. Using (2) we can write

$$\mathbf{Rank}(R \cap I_{1:\tau}) = \mathbf{Rank}(R \cap I_1) + \mathbf{Rank}'(R \cap I_{2:\tau}) \quad (1)$$

The expectation of the first term is $\mathbb{E}[\mathbf{Rank}(R \cap I_1)] = r \cdot p$. To bound the expectation of the second term, we first condition on $R \cap I_1$, which also fixes S and \mathcal{M}' . It follows from Lemma 9 and Observation 10, as well as the fact that $S \subseteq I_1$ is disjoint from $I_{2:\tau}$, that I_2, \dots, I_τ is an NSS of \mathcal{M}' . This allows us to invoke the inductive hypothesis to obtain

$$\mathbb{E}[\mathbf{Rank}'(R \cap I_{2:\tau})] \geq (1 - (1-p)^{\tau-1}) \cdot \mathbb{E}[\mathbf{Rank}'(R \setminus I_1)].$$

We use a well-known fact about the rank function of the contracted matroid given by

$$\mathbf{Rank}^{\mathcal{M}/S}(T) = \mathbf{Rank}^{\mathcal{M}}(T \cup S) - \mathbf{Rank}^{\mathcal{M}}(S) = \mathbf{Rank}^{\mathcal{M}}(T \cup S) - |S|. \quad (2)$$

Equation (2) and the definition of S implies that $\mathbf{Rank}'(R \setminus I_1) = \mathbf{Rank}((R \setminus I_1) \cup S) - \mathbf{Rank}(S) = \mathbf{Rank}(R) - \mathbf{Rank}(R \cap I_1)$. Also using the fact $\mathbb{E}[\mathbf{Rank}(R \cap I_1)] = r \cdot p$, we obtain

$$\begin{aligned}
\mathbb{E}[\mathbf{Rank}'(R \cap I_{2:\tau})] &\geq (1 - (1 - p)^{\tau-1}) \cdot \mathbb{E}[\mathbf{Rank}'(R \setminus I_1)] \\
&= (1 - (1 - p)^{\tau-1})(\mathbb{E}[\mathbf{Rank}(R)] - \mathbb{E}[\mathbf{Rank}(R \cap I_1)]) \\
&= (1 - (1 - p)^{\tau-1})\mathbb{E}[\mathbf{Rank}(R)] - (1 - (1 - p)^{\tau-1}) \cdot r \cdot p
\end{aligned} \tag{3}$$

Finally, we combine (1), and (3) to conclude

$$\begin{aligned}
\mathbb{E}[\mathbf{Rank}(R \cap I_{1:\tau})] &\geq (1 - (1 - p)^{\tau-1})\mathbb{E}[\mathbf{Rank}(R)] + (1 - p)^{\tau-1} \cdot r \cdot p \\
&\geq (1 - (1 - p)^{\tau-1} + p(1 - p)^{\tau-1})\mathbb{E}[\mathbf{Rank}(R)] \\
&= (1 - (1 - p)^\tau)\mathbb{E}[\mathbf{Rank}(R)] \quad \blacktriangleleft
\end{aligned}$$

By observation 8, we get the following corollary of Lemma 11.

► **Corollary 12.** *Consider a stochastic matroid optimization problem $\langle E, \mathcal{I}, f, p \rangle$ for $p \in [0, 1]$ and $f(S) = |S|$ for all $S \subseteq E$. Algorithm 2 is a $(1 - \epsilon)$ -approximate sparsifier with degree $\frac{1}{p} \cdot \log\left(\frac{1}{\epsilon}\right)$.*

4.2 Proof of Theorem 6

In this section, we will complete the proof of Theorem 6 by reducing the analysis for a general (weighted) additive function to that of the unweighted case. We order the elements e_1, \dots, e_n in decreasing order of their weights $w_1 \geq \dots \geq w_n$. Without loss of generality we assume $w_n > 0$, and for notational convenience we define $w_{n+1} = 0$. The following lemma says that if a sparsifier is α -approximate for the unweighted problem on elements above any given weight threshold, then it is also α -approximate for the weighted problem.

► **Lemma 13.** *For all $j \in [n]$ with $w_j > w_{j+1}$, if a set $Q \subseteq E$ satisfies*

$$\mathbb{E}[\mathbf{Rank}(Q \cap R \cap \{e_1, \dots, e_j\})] \geq (1 - \epsilon)\mathbb{E}[\mathbf{Rank}(R \cap \{e_1, \dots, e_j\})], \tag{4}$$

then $\mathbb{E}[f(\text{opt}(Q \cap R))] \geq (1 - \epsilon)\mathbb{E}[f(\text{opt}(R))]$. Here, we denote $\text{opt}(S) \in \arg\max_{I \subseteq S} f(I)$, with ties broken arbitrarily.

The above lemma follows from the optimality of the greedy algorithm for weighted optimization over matroids. We relegate the (fairly standard) proof to the full version [13].

To conclude the proof of Theorem 6, we show in the following lemma that the output of Algorithm 2 satisfies condition (4).

► **Lemma 14.** *For all $j \in [n]$ with $w_j > w_{j+1}$, the output set Q of Algorithm 2 satisfies*

$$\mathbb{E}[\mathbf{Rank}(Q \cap R \cap \{e_1, \dots, e_j\})] \geq (1 - (1 - p)^\tau)\mathbb{E}[\mathbf{Rank}(R \cap \{e_1, \dots, e_j\})]$$

To provide more intuition, let I_1, \dots, I_τ be the sets defined in Algorithm 2, and $\bar{E} = \{e_1, \dots, e_j\}$ be the top weight j elements. It is sufficient to show that the sets $I_t \cap \bar{E}$ form a sequence of nested spanning sets for the restricted matroid $\bar{\mathcal{M}}$ on elements \bar{E} . The optimal choice of I_t in Algorithm 2, together with the matroid structure, implies that $I_t \cap \bar{E}$ has full rank in $\bar{\mathcal{M}} \setminus I_{1:t-1}$. We complete the proof in the full version [13]. Combining Lemmas 13 and 14 completes the proof of Theorem 6.

5 Improved Sparsifier for Stochastic Weighted Matching

In the instance of stochastic weighted matching $\langle E, \mathcal{I}, f, p \rangle$, the elements E are the edges of a known weighted graph $G := (V, E, w)$, \mathcal{I} is the set of all matchings in the graph G , and f is an additive function with element weights $\{w_e\}_{e \in E}$. For simplicity, we sometimes denote the stochastic matching instance $\langle E, \mathcal{I}, f, p \rangle$ by $\langle G, p \rangle$ when it is clear from the context.

The aim of a sparsifier for this problem is to query a poly $(1/p)$ -degree subgraph H of G such that the expected weight of the maximum matching on active edges of H approximates the optimum value of $\langle G, p \rangle$. The current state-of-the-art poly $(1/p)$ -degree sparsifier for the stochastic weighted matching problem achieves a 0.501 approximation ratio due to [8]³. In this section, we present a new poly $(1/p)$ -degree sparsifier for the stochastic weighted matching that improves the approximation ratio to 0.536.

Our sparsifier for the stochastic weighted matching problem consists of two phases. In the first phase, it samples a set of edges Q_{CRS} using the generic sparsifier described in Algorithm 1. In the second phase, we independently select T samples Q_1, \dots, Q_T from the stochastic optimum oracle \mathcal{D}_{opt} , which is similar to the method used in [8]. This second phase alone already provides a 0.501 approximation, but by incorporating the edges sampled in the first phase, we are able to improve the approximation ratio to 0.536. The main result of this section is presented in the following theorem.

■ **Algorithm 3** Sparsifier for Weighted Stochastic Matching Problem $\langle G, p \rangle$.

-
- 1: Compute the marginals \mathbf{q} of the stochastic optimum solution.
 - 2: Add each edge $e \in E$ to the set Q_{CRS} independently with probability $\frac{q_e}{p}$.
 - 3: Sample $Q_1, \dots, Q_T \sim \mathcal{D}_{\text{opt}}$ independently and add them to Q_{Greedy} for $T = 1/\epsilon^8 p$.
 - 4: **Output:** $Q = Q_{\text{CRS}} \cup Q_{\text{Greedy}}$.
-

► **Theorem 15.** *Let $G = (V, E, w)$ be a weighted graph and $p \in (0, 1)$. If the matching polytope of G admits an α -balanced contention resolution scheme, then Algorithm 3 is the $(1 - O(\epsilon)) \cdot \max\left\{\frac{1}{2}, \left(\frac{1+\alpha\epsilon^2}{1+\epsilon^2}\right)\right\}$ -approximate polynomial time sparsifier for the stochastic weighted matching problem $\langle G, p \rangle$ with sparsification degree $O(1/\epsilon^8 p)$.*

Our theorem combined with 0.474-balanced CRS for machining polytope from [17] implies 0.536-approximate sparsifier for stochastic weighted matching. Assuming the conjecture from [20] which states the existence of 0.544 balanced CRS for general matching polytope implies that Algorithm 3 is ~ 0.6 approximate.

The proof of Theorem 15 relies on p being small. So, before we prove the theorem, in Lemma 16, we show that for any $\epsilon > 0$ (constant), without loss of generality we can assume $p \leq \epsilon^4$. The proof of this part is rather technical and, we defer it to the full version [13] due to space constraints.

► **Lemma 16 (Reduction Lemma).** *If there exists an α -approximate sparsifier with degree d/p for the class of stochastic weighted matching with $p \leq \epsilon^4$ then there exists an α -approximate sparsifier for the same problem class and arbitrary $p \in (0, 1)$ with sparsification degree $\frac{d}{p \cdot \epsilon^4}$.*

For the rest of the section, we assume that $p \leq \epsilon^4$. We first define the set of crucial edges and non-crucial edges formally in the following definition.

³ Recent work by [6] constructs $(1-\epsilon)$ -approximate sparsifier with degree $\exp(\exp(\exp(1/\epsilon, 1/p)))$, however, in this work, we focus on sparsifiers with degree $\text{poly}(1/p)$

► **Definition 17.** Given $\langle G, p \rangle$, let q_e be the probability of an edge e being in the stochastic optimum solution. We define crucial edges as $\mathcal{C} := \{e \in E : q_e \geq \tau(\epsilon)\}$ and non-crucial edges as $\mathcal{NC} := \{e \in E : q_e < \tau(\epsilon)\}$ where $\tau(\epsilon) := \frac{\epsilon^3 p}{20 \cdot \log \frac{1}{\epsilon}}$ is the threshold.

Given $\langle G, p \rangle$ and set of crucial and non-crucial edges \mathcal{C} and \mathcal{NC} , we let $\mathbf{OPT}_{\mathcal{C}}$ and $\mathbf{OPT}_{\mathcal{NC}}$ be the contributions of crucial and non-crucial edges in the stochastic optimum, i.e. $\sum_{e \in \mathcal{C}} w_e \cdot q_e$ and $\sum_{e \in \mathcal{NC}} w_e \cdot q_e$. Note that

$$\mathbf{OPT} = \mathbf{OPT}_{\mathcal{C}} + \mathbf{OPT}_{\mathcal{NC}}.$$

In order to prove Theorem 15, we provide a procedure to construct a matching $M \subseteq Q \cap R$ such that $\mathbb{E}[\sum_{i \in M} w_e] \geq (1 - O(\epsilon)) \cdot \max\left\{\frac{1}{2}, \left(\frac{1 + \alpha \epsilon^2}{1 + \epsilon^2}\right)\right\} \cdot \mathbf{OPT}$. Our procedure constructs three matchings $M_{\mathcal{C}}, M_{\mathcal{NC}}, M_{\text{AUG}} \subseteq R \cap Q$ and then picks the matching with the maximum weight. We construct matchings $M_{\mathcal{C}}, M_{\mathcal{NC}}$ on the queried active crucial and non-crucial edges in Q_{Greedy} similar to the [8] which satisfies the desired properties described in Lemma 18 and Lemma 19. First, we state that each crucial edge $e \in \mathcal{C}$ appears in the Q_{Greedy} with probability $1 - \epsilon$ which shows the existence of matching $M_{\mathcal{C}} \subseteq Q \cap R \cap \mathcal{C}$ with expected weight at least $(1 - \epsilon) \cdot \mathbf{OPT}_{\mathcal{C}}$.

► **Lemma 18** (Crucial Edge Lemma [8]). *Given a stochastic weighted matching instance $\langle G, p \rangle$ and Q_{Greedy} is the set defined in Algorithm 3, let $M_{\mathcal{C}}$ be the maximum weight matching in the graph $Q_{\text{Greedy}} \cap \mathcal{C} \cap R$, then $\mathbb{E}[\sum_{e \in M_{\mathcal{C}}} w_e] \geq (1 - \epsilon) \cdot \mathbf{OPT}_{\mathcal{C}}$.*

Now, following the [8, Lemma 4.7], in Lemma 19, we construct a matching $M_{\mathcal{NC}} \subseteq R \cap Q_{\text{Greedy}} \cap \mathcal{NC}$ on active queried non-crucial edges, such that each $e \in \mathcal{NC}$ is present in $M_{\mathcal{NC}}$ with probability at least $(1 - O(\epsilon)) \cdot q_e$. We further prove an important property of $M_{\mathcal{NC}}$ that states that for any non-crucial edge $e \in \mathcal{NC}$, the probability of $e \in M_{\mathcal{NC}}$ can not decrease when we condition on the events that some of the neighbors of e are inactive⁴.

► **Lemma 19** (Non-Crucial Edges). *Given a stochastic weighted matching instance $\langle G, p \rangle$, let Q_{Greedy} be the set defined in Algorithm 3. There exists a matching $M_{\mathcal{NC}} \subseteq Q_{\text{Greedy}} \cap \mathcal{NC} \cap R$ such that for any non-crucial edge $e \in \mathcal{NC}$, $\Pr[e \in M_{\mathcal{NC}}] \geq (1 - 12\epsilon) \cdot q_e$. This implies that, $\mathbb{E}[\sum_{e \in M_{\mathcal{NC}}} w_e] \geq (1 - 12\epsilon) \cdot \mathbf{OPT}_{\mathcal{NC}}$. Moreover, for any subset $S \subseteq N(e)$ where $N(e)$ is the set of edges incident to e in graph G , we have*

$$\Pr[e \in M_{\mathcal{NC}} \mid S \cap R = \emptyset] \geq (1 - 12 \cdot \epsilon) \cdot q_e. \quad (5)$$

The proof of the lemma is technically involved and therefore it is delegated to the full version [13]. Lemma 18 and Lemma 19 together imply that our sparsifier is at least 1/2 approximate.

We note that Q_{CRS} is the output of generic sparsifier discussed in Algorithm 1 (Section 3). Let $M_{\text{CRS}} := \pi(Q_{\text{CRS}} \cap R)$ be the matching constructed by an α -balanced CRS π which ensures $\Pr[e \in M_{\text{CRS}}] \geq \alpha \cdot q_e$ for all $e \in E$. We refer M_{CRS} as **CRS-BaseMatching**. Crucially, M_{CRS} is independent of the edges sampled in Q_{Greedy} as well as $M_{\mathcal{NC}}$ and $M_{\mathcal{C}}$. Using Independence between M_{CRS} and $M_{\mathcal{NC}}$, we construct the third matching M_{AUG} on the set of edges $Q_{\text{CRS}} \cup (Q_{\text{Greedy}} \cap \mathcal{NC}) \cup R$. Our augmentation simply adds a non-crucial edge $e \in M_{\mathcal{NC}}$ to **CRS-BaseMatching** if both endpoints of the edge e are unmatched in **CRS-BaseMatching**. Algorithm 4 describes our augmentation procedure in detail.

⁴ We noticed a bug in the proof of a similar lemma presented in [8], further used in [7, 6]. In order to prove the lemma and the monotonicity property (5), we require slightly different proof techniques.

51:14 On Sparsification of Stochastic Packing Problems

Our key observation is that any non-crucial edge $e \in \text{NC}$ has a small probability of being sampled in the set Q_{CRS} . However, with some non-trivial probability, both endpoints of the edge e will be unmatched in **CRS-BaseMatching**. More formally, first we show that for any non-crucial edge $e := (u, v) \in \text{NC}$, both endpoints of e are unmatched in the **CRS-BaseMatching** with probability at least $1/e^2$. The intuition here is that as $p \leq \epsilon^4$, the number of incident edges on the endpoints of the edge e in the set Q_{CRS} are concentrated around $2/p$ with high probability. Such a property ensures that if all these incident edges are inactive, then both endpoints of e are unmatched in **CRS-BaseMatching**.

Later, we use the property (5) of M_{NC} from Lemma 19 to guarantee that when a non-crucial edge $e \notin Q_{\text{CRS}}$ and both endpoints of e are unmatched in **CRS-BaseMatching**, we can guarantee that $e \in M_{\text{NC}}$ with probability approximately q_e . Therefore, we can add such a non-crucial edge e to **CRS-BaseMatching** with probability approximately $\frac{q_e}{e^2}$. Combining this intuition, we prove the following key lemma whose proof is delegated to the full version [13].

► **Lemma 20.** *Let M_{AUG} be the output of the procedure described in Algorithm 4 then,*

$$\Pr[e \in M_{\text{AUG}}] \geq q_e \cdot \alpha \quad \forall e \in \mathcal{C} \quad \text{and} \quad \Pr[e \in M_{\text{AUG}}] \geq q_e \cdot \left(\alpha + \frac{1 - O(\epsilon)}{e^2} \right) \quad \forall e \in \text{NC}.$$

■ **Algorithm 4** Construction of the matching M_{AUG} on $Q \cap R$.

-
- 1: M_{NC} be the matching on $Q_{\text{Greedy}} \cap R \cap \text{NC}$ satisfying property of stated Lemma 19.
 - 2: $M_{\text{CRS}} \leftarrow \pi(Q_{\text{CRS}} \cap R)$ be the matching produced by α -balanced truncated CRS.
 - 3: $M_{\text{AUG}} \leftarrow M_{\text{CRS}}$.
 - 4: $\forall e \in M_{\text{NC}}$, add e to the matching M_{AUG} if both endpoints of e are unmatched in M_{AUG} .
-

Combining Lemma 18, Lemma 19 and Lemma 20, we show that the expected weight of the best matching among M_{CRS} , M_{NC} , and M_{AUG} exhibits the desired approximation ratio. We complete the proof of Theorem 15 in the full version [13].

6 Additive Optimization over the Intersection of k Matroids

Given our $(1 - \epsilon)$ -approximate sparsifier for additive optimization over a single matroid constraint, a natural question is whether the natural generalization of this algorithm to the intersection of matroids is $(1 - \epsilon)$ -approximate. This turns out to not be the case even for bipartite matching (the intersection of two partition matroids) due to [9]. The main challenge here is that, unlike for a single matroid, multiple solutions for matroid intersection do not always “combine” well. In this section, we prove a slightly weaker sparsification result for additive optimization over the intersection of k matroid constraints, which nevertheless beats the best known bound of $1/(k + 1)$ on the correlation gap of k -matroid intersection (see [1]), and therefore outperforms our generic sparsifier for this problem. The following theorem is the main result of this section.

► **Theorem 21.** *For each $\epsilon > 0$, there is a $\frac{(1-\epsilon)}{k + \frac{1}{k+1}}$ -approximate sparsifier of degree $O\left(\frac{1}{\epsilon \cdot p} \log \frac{1}{\epsilon}\right)$ for stochastic packing problem $\langle E, \mathcal{I}, f, p \rangle$ when (E, \mathcal{I}) is the intersection of k matroids and f is additive.*

Our sparsifier samples Q_1, \dots, Q_τ independently from stochastic optimum oracle \mathcal{D}_{opt} as a sparsifier. Similar algorithms with degree $\text{poly}(1/p)$ have been considered for the stochastic matching [8, 7], and were shown to be 0.6568-approximate for the unweighted and 0.501-approximate for a weighted matching with degree $\text{poly}(1/p)$.

■ **Algorithm 5** Sparsifier for additive optimization over the intersection of k matroid constraints.

Input: $\langle E, \mathcal{I}, f, p \rangle$ with the intersection of k matroids constraints and additive f ; D_{OPT}

Sample $Q_1, \dots, Q_\tau \sim D_{\text{OPT}}$ independently for $\tau \leftarrow \frac{2}{\epsilon p} \log \frac{2}{\epsilon}$;

Output: $Q = \cup_{i=1}^\tau Q_i$.

In order to prove Theorem 21, we provide a procedure for constructing a feasible solution $I \subseteq Q \cap R$ such that $\mathbb{E}[\sum_{e \in I} w_e] \geq \frac{(1-\epsilon)}{k+1/(k+1)} \text{OPT}$. The backbone of our analysis lies in Lemma 22. As a first step, let S_1 and S_2 be two independent sets of the same matroid and $R \subseteq E$ be the (random) set of active elements with parameter p . We propose a procedure (details in Algorithm 6 in the full version [13]) that swaps active elements from S_1 , i.e. $S_1 \cap R$, with elements of S_2 such that each element of S_2 is “protected” independently with probability $1 - p$. Hence, the expected value of updated set S_2 is $\geq \mathbb{E}[f(S_1 \cap R)] + (1 - p) \cdot f(S_2)$.

The key intuition here is that the exchange property of matroids allows us to swap any element $e \in S_1$ with a different element $f \in S_2$ without violating the feasibility of S_2 . Therefore, if e is inactive then e can not swap out f from S_2 and hence we “protect” f in S_2 with probability $1 - p$. However, the main challenge here is after a single swap between e and f , sets S_1 and S_2 get updated and f can potentially be swapped with some $f' \in S_2$. Our procedure overcomes this challenge by carefully choosing swaps of elements between S_1 and S_2 while maintaining feasibility.

We extend this idea to when S_1 and S_2 are two independent sets in the intersection of k matroids. We run the procedure described in Algorithm 6 in the full version [13] for each matroid and obtain sets T feasible in the intersection of all matroids such that each element of S_2 is added to T independently with probability $(1 - p)^k$. The details of procedure and proof of Lemma 22 is relegated to the full version of the paper [13]

► **Lemma 22.** *Let $\mathcal{M}_1, \dots, \mathcal{M}_k$ be matroids with $\mathcal{M}_\ell = (E, \mathcal{I}_\ell)$, and let $\mathcal{I} = \bigcap_{\ell=1}^k \mathcal{I}_\ell$ be their common independent sets. Let S_1 and S_2 be in \mathcal{I} . Let $R \subseteq E$ include each element of E independently with probability p . Let $T(\ell) \in \mathcal{I}_\ell$ be the output of Algorithm 6 in full version [13] for matroid \mathcal{M}_ℓ , for each $\ell \in [k]$. The set $T := \bigcap_{\ell=1}^k T(\ell)$ satisfies:*

1. $S_1 \cap S_2 \subseteq T$ with probability 1.
2. $T \in \mathcal{I}$ with probability 1.
3. $(S_1 \setminus S_2) \cap R \subseteq T$, i.e. $\Pr[e \in T] = p$ for all $e \in S_1 \setminus S_2$.
4. $\Pr[f \in T] \geq (1 - p)^k$ for all $f \in S_2 \setminus S_1$

We utilize the above lemma and propose a procedure to construct a feasible set $I \subseteq Q \cap R$. At a high level, our procedure iteratively observes active elements in the set Q_i and swaps elements in Q_{i+1}, \dots, Q_n by $Q_i \cap R$ using Lemma 22. To this end, Lemma 22 ensures that each element in Q_j for $j > i$ is not swapped (“protected”) with probability at least $(1 - p)$. Using this argument inductively, we prove the following lemma that lower bounds the probability of selecting each element $e \in Q$ whose proof is in the full version [13]. We then use the lemma and carefully analyze the probability of each element $e \in E$ in the constructed set $I \subseteq R \cap Q$ to conclude the proof of Theorem 21.

► **Lemma 23.** *Let $I^* = I(\tau)$ be the output of Algorithm 7 in full version [13]. For any $e \in E$, we have*

$$\Pr[e \in I^* \mid e \in Q_i \setminus \cup_{\ell=1}^{i-1} Q_\ell] \geq p \cdot (1 - p)^{k(i-1)}$$

7 Open Questions

- We believe that our results portend a deeper connection between the sparsification and contention resolution. The results of Section 3 show that contention resolution serves to lower-bound the sparsification ratio. We ask whether the connection goes both ways. In particular, does the existence of a c -sparsifier of degree $1/p$ imply a contention resolution scheme with balance c ? This is intimated by Proposition 5. Does the existence of a c -sparsifier of degree $\text{poly}(1/p)$ imply a contention resolution scheme with balance $\Omega(c)$ (or some other expression involving c and the degree)? Formalizing a tighter connection between sparsification and contention resolution (equivalently, the correlation gap) might lead to new structural and computational insights for the latter.
- In Section 4, we show that a greedy sparsifier $1 - \epsilon$ approximate with degree $O(1/p)$ for additive optimization subject to a matroid constraint. We conjecture that a similar greedy sparsifier exists for the intersection of k matroids, obtaining a $\frac{1-\epsilon}{k}$ -approximation with degree $O(1/p)$. A similar greedy sparsifier, albeit with degree $O(1/p^{1/\epsilon})$, was shown to be $1/2$ -approximate for the special case of unweighted bipartite matching in [9].
- Our results in Section 5 improve the state of the art sparsifier for weighted (non-bipartite) matching in the polynomial degree regime. Moreover, since our approximation guarantee is a function of the correlation gap, progress on the correlation gap of the matching polytope will lead to further improved sparsifiers. Finding the best possible sparsification ratio in the polynomial degree regime remains open, however, with $1 - \epsilon$ still on the table. Beyond polynomial degree, a $1 - \epsilon$ approximate sparsifier with degree $\exp(\exp(\exp(1/p)))$ was already shown by [6].

References

- 1 Marek Adamczyk and Michał Włodarczyk. Random order contention resolution schemes. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 790–801. IEEE, 2018.
- 2 Sepehr Assadi and Aaron Bernstein. Towards a unified theory of sparsification for matching problems. *arXiv preprint*, 2018. [arXiv:1811.02009](https://arxiv.org/abs/1811.02009).
- 3 Sepehr Assadi, Sanjeev Khanna, and Yang Li. The stochastic matching problem: Beating half with a non-adaptive algorithm. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 99–116, 2017.
- 4 Sepehr Assadi, Sanjeev Khanna, and Yang Li. The stochastic matching problem with (very) few queries. *ACM Transactions on Economics and Computation (TEAC)*, 7(3):1–19, 2019.
- 5 Soheil Behnezhad, Avrim Blum, and Mahsa Derakhshan. Stochastic vertex cover with few queries. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1808–1846. SIAM, 2022.
- 6 Soheil Behnezhad and Mahsa Derakhshan. Stochastic weighted matching: $(1-\epsilon)$ approximation. *arXiv preprint*, 2020. [arXiv:2004.08703](https://arxiv.org/abs/2004.08703).
- 7 Soheil Behnezhad, Mahsa Derakhshan, and MohammadTaghi Hajiaghayi. Stochastic matching with few queries: $(1-\epsilon)$ approximation. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1111–1124, 2020.
- 8 Soheil Behnezhad, Alireza Farhadi, MohammadTaghi Hajiaghayi, and Nima Reyhani. Stochastic matching with few queries: New algorithms and tools. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2855–2874. SIAM, 2019.
- 9 Avrim Blum, John P Dickerson, Nika Haghtalab, Ariel D Procaccia, Tuomas Sandholm, and Ankit Sharma. Ignorance is almost bliss: Near-optimal stochastic matching with few queries. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 325–342, 2015.

- 10 Simon Bruggmann and Rico Zenklusen. An optimal monotone contention resolution scheme for bipartite matchings via a polyhedral viewpoint. *Mathematical Programming*, pages 1–51, 2020.
- 11 Gruia Calinescu, Chandra Chekuri, Martin Pal, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM Journal on Computing*, 40(6):1740–1766, 2011.
- 12 Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM Journal on Computing*, 43(6):1831–1879, 2014.
- 13 Shaddin Dughmi, Yusuf Hakan Kalayci, and Neel Patel. On sparsification of stochastic packing problems. *arXiv preprint*, 2022. [arXiv:2211.07829](https://arxiv.org/abs/2211.07829).
- 14 Michel X. Goemans and Jan Vondrák. Covering minimum spanning trees of random subgraphs. In *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '04*, pages 934–941, USA, 2004. Society for Industrial and Applied Mathematics.
- 15 Richard M Karp and Michael Sipser. Maximum matching in sparse random graphs. In *22nd Annual Symposium on Foundations of Computer Science (sfcs 1981)*, pages 364–375. IEEE, 1981.
- 16 Jon Lee, Maxim Sviridenko, and Jan Vondrák. Submodular maximization over multiple matroids via generalized exchange properties. *Mathematics of Operations Research*, 35(4):795–806, 2010.
- 17 Calum MacRury, Will Ma, and Nathaniel Grammel. On (random-order) online contention resolution schemes for the matching polytope of (bipartite) graphs, 2022. [doi:10.48550/arXiv.2209.07520](https://doi.org/10.48550/arXiv.2209.07520).
- 18 Takanori Maehara and Yutaro Yamaguchi. Stochastic packing integer programs with few queries. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '18*, pages 293–310, USA, 2018. Society for Industrial and Applied Mathematics.
- 19 Takanori Maehara and Yutaro Yamaguchi. Stochastic monotone submodular maximization with queries. *arXiv preprint*, 2019. [arXiv:1907.04083](https://arxiv.org/abs/1907.04083).
- 20 Pranav Nuti and Jan Vondrák. Towards an optimal contention resolution scheme for matchings. *arXiv preprint*, 2022. [arXiv:2211.03599](https://arxiv.org/abs/2211.03599).