New Greedy Heuristics for Approximating Set Cover and Set Packing

by

David Kordalewski

A thesis submitted in conformity with the requirements for the degree of Master of Science Graduate Department of Computer Science University of Toronto

© Copyright 2013 by David Kordalewski

Abstract

New Greedy Heuristics for Approximating Set Cover and Set Packing

David Kordalewski Master of Science Graduate Department of Computer Science University of Toronto 2013

The Set Cover problem (SCP) and Set Packing problem (SPP) are standard NP-hard combinatorial optimization problems. Their decision problem versions are shown to be NP-Complete in Karp's 1972 paper. We specify a rough guide to constructing approximation heuristics that may have widespread applications and apply it to devise greedy approximation algorithms for SCP and SPP, where the selection heuristic is a variation of that in the standard greedy approximation algorithm. Our technique involves assigning to each input set a *valuation* and then selecting, in each round, the set whose valuation is highest. We prove that the technique we use for determining a valuation of the input sets yields a unique value for all Set Cover instances. For both SCP and SPP we give experimental evidence that the valuations we specify are unique and can be computed to high precision quickly by an iterative algorithm. Others have experimented with testing the observed approximation ratio of various algorithms over a variety of randomly generated instances, and we have extensive experimental evidence to show the quality of the new algorithm relative to greedy heuristics in common use. Our algorithms are somewhat more computationally intensive than the standard heuristics, though they are still practical for large instances. We discuss some ways to speed up our algorithms that do not significantly distort their effectiveness in practice on random instances.

Acknowledgements

I'm grateful for the support and patience of my supervisors, Allan Borodin and Ken Jackson. I would also like to thank Danny Ferriera for many excellent conversations relating to this work, and Stephanie Southmayd for her kind help in editing this document. A great deal of thanks is due to Ken Jackson, who was the first to prove both the existence and uniqueness results that form the body of chapter 5. I would also like to thank the OGS program for financial support.

Contents

1	Intr	roduction 1							
	1.1	Notation and Terminology	1						
	1.2	Organization of Thesis	2						
2	Pro	blem Definitions and Discussion	3						
	2.1	Problem Formulations	3						
		2.1.1 Set Cover	3						
		2.1.2 Hitting Set	4						
		2.1.3 (Hitting {Set) Cover}	4						
		2.1.4 Integer Programming Formulation	6						
		2.1.5 Set Packing and Formulations	6						
		2.1.6 Set Packing's Relationship to Maximum Independent Set	7						
	2.2	Greedy Algorithms for Set Cover/Packing Approximation	8						
		2.2.1 Set Cover/Packing Approximation	8						
		2.2.2 Set Cover Approximation	8						
		2.2.3 Set Packing Approximation	9						
		2.2.4 General Greedy Scheme	0						
	2.3	The Standard Greedy Cover Heuristic	0						
	2.4	Other Approximation Techniques	2						
3	Pre	processing 1	3						
	3.1	Basic Preprocessing 1	3						
	3.2	Subsumption Testing	4						

	3.3	Independent Subproblem Separation	15					
	3.4	Inferring Stronger Packing Constraints						
	3.5	Non-Minimal Covers	17					
4	Nev	w Greedy Heuristics for Set Cover and Set Packing	19					
	4.1	Motivation	19					
	4.2	The New Greedy Set Cover Heuristic	20					
		4.2.1 Relationship to the Standard Heuristic	20					
		4.2.2 Consistent Valuations	21					
		4.2.3 A Family of Heuristics	21					
		4.2.4 Relationship to Theory	22					
	4.3	The New Greedy Set Packing Heuristic	22					
5	Ma	thematical Results	24					
	5.1	Fixed Point Existence	24					
	5.2	Fixed Point Uniqueness	27					
6	Nu	merical Matters	28					
6	Nu 6.1	merical Matters Calculating Fixed Points In Practice	28 28					
6	Nu 6.1	merical Matters Calculating Fixed Points In Practice 6.1.1 An Alternate Iteration	282830					
6	Nu 6.1 6.2	merical Matters Calculating Fixed Points In Practice 6.1.1 An Alternate Iteration Additional Shortcuts	 28 30 32 					
6	Nu 6.1 6.2	merical Matters Calculating Fixed Points In Practice 6.1.1 An Alternate Iteration Additional Shortcuts 6.2.1 Using the Alternate Iteration for the Packing Heuristic	 28 28 30 32 32 					
6	Nu 6.1 6.2 6.3	merical Matters Calculating Fixed Points In Practice 6.1.1 An Alternate Iteration Additional Shortcuts 6.2.1 Using the Alternate Iteration for the Packing Heuristic Running Time of the New Heuristics	 28 28 30 32 32 33 					
6	Nu 6.1 6.2 6.3	merical Matters Calculating Fixed Points In Practice 6.1.1 An Alternate Iteration Additional Shortcuts 6.2.1 Using the Alternate Iteration for the Packing Heuristic Running Time of the New Heuristics 6.3.1 Running Time of the New Set Cover Heuristic	 28 28 30 32 32 33 33 					
6	Nu 6.1 6.2 6.3	merical Matters Calculating Fixed Points In Practice 6.1.1 An Alternate Iteration Additional Shortcuts Additional Shortcuts 6.2.1 Using the Alternate Iteration for the Packing Heuristic Running Time of the New Heuristics 6.3.1 Running Time of the New Set Cover Heuristic 6.3.2 Running Time of the New Set Packing Heuristic	 28 28 30 32 32 33 33 34 					
6	Nui 6.1 6.2 6.3	merical Matters Calculating Fixed Points In Practice 6.1.1 An Alternate Iteration Additional Shortcuts 6.2.1 Using the Alternate Iteration for the Packing Heuristic Running Time of the New Heuristics 6.3.1 Running Time of the New Set Cover Heuristic 6.3.2 Running Time of the New Set Packing Heuristic Exact Calculation of Fixed Points	 28 28 30 32 32 33 33 34 34 					
6	Nui 6.1 6.2 6.3 6.4 6.5	merical Matters Calculating Fixed Points In Practice 6.1.1 An Alternate Iteration Additional Shortcuts Additional Shortcuts 6.2.1 Using the Alternate Iteration for the Packing Heuristic Running Time of the New Heuristics 6.3.1 Running Time of the New Set Cover Heuristic 6.3.2 Running Time of the New Set Packing Heuristic Exact Calculation of Fixed Points Fixed Points For Broader Classes of Matrices	 28 28 30 32 32 33 34 34 35 					
6	Nui 6.1 6.2 6.3 6.4 6.5 Exp	merical Matters Calculating Fixed Points In Practice 6.1.1 An Alternate Iteration Additional Shortcuts Additional Shortcuts 6.2.1 Using the Alternate Iteration for the Packing Heuristic Running Time of the New Heuristics 6.3.1 Running Time of the New Set Cover Heuristic 6.3.2 Running Time of the New Set Packing Heuristic Exact Calculation of Fixed Points Fixed Points For Broader Classes of Matrices	 28 28 30 32 33 33 34 34 35 37 					
6	Nui 6.1 6.2 6.3 6.4 6.5 Exp 7.1	merical Matters Calculating Fixed Points In Practice 6.1.1 An Alternate Iteration Additional Shortcuts Additional Shortcuts 6.2.1 Using the Alternate Iteration for the Packing Heuristic Running Time of the New Heuristics 6.3.1 Running Time of the New Set Cover Heuristic 6.3.2 Running Time of the New Set Packing Heuristic Exact Calculation of Fixed Points Fixed Points For Broader Classes of Matrices perimental Results Random Instances	 28 28 30 32 33 33 34 34 35 37 37 					
6	Nui 6.1 6.2 6.3 6.4 6.4 6.5 Exp 7.1 7.2	merical Matters Calculating Fixed Points In Practice 6.1.1 An Alternate Iteration 6.1.1 An Alternate Iteration Additional Shortcuts 6.2.1 Using the Alternate Iteration for the Packing Heuristic Running Time of the New Heuristics 6.3.1 Running Time of the New Set Cover Heuristic 6.3.2 Running Time of the New Set Packing Heuristic Exact Calculation of Fixed Points Fixed Points For Broader Classes of Matrices perimental Results Random Instances Algorithms Used	 28 28 30 32 33 33 34 34 35 37 37 38 					

Bibliography							
8	Disc	cussion	ı	50			
		7.4.2	OR Library Instances	46			
		7.4.1	Comparison Between Packing Heuristics	46			
	7.4	Set Pa	cking Results	46			
		7.3.3	OR Library Instances	44			
		7.3.2	Comparison Between the Standard and New Heuristics	40			
		7.3.1	Varying γ for the New Heuristic	39			
	7.3	Set Co	over Results	39			
		7.2.2	Set Packing Algorithms	39			

Chapter 1

Introduction

The Set Cover Problem (SCP) and Set Packing Problem (SPP) are standard NP-hard combinatorial optimization problems. Their decision problem versions are shown to be NP-Complete in [13]. Because we cannot expect to solve all instances of these problems exactly in polynomial time, much effort has been expended on finding *approximation* algorithms for these problems. Many algorithms have been proven to obtain approximate solutions for SCP and SPP that are within some factor of the optimal solution. At the same time there are results showing that, assuming some complexity conjectures, no polynomial-time algorithm can approximate these problems to any constant ratio. Further results have placed even stronger constraints on what sort of approximation ratios are achievable by polynomial-time algorithms.

Our primary purpose here is to describe a technique for greedily obtaining high-quality approximate solutions for Set Cover and Set Packing problems. Our technique involves assigning to each input set a *valuation* and then selecting, in each round, the set whose valuation is highest. In order to specify the valuations, we define them recursively. For both SCP and SPP, we prove that a valuation satisfying our definition must exist, and in the case of Set Cover, that it is unique. We have not been able to show that our valuations result in a greedy algorithm that has some guaranteed approximation ratio, but do show experimentally that it performs somewhat better than the standard greedy algorithms for these problems on random instances.

We believe that the mathematics underlying our recursively defined valuations and the overall idea that motivates our particular definitions are of significant interest independent of our algorithms and their performance.

1.1 Notation and Terminology

The notation $\mathbf{d}^{\diamond -1}$ is borrowed from [6] to denote the *Hadamard inverse* of a matrix or vector with all non-zero entries. If **d** is a vector, then the components of $\mathbf{d}^{\diamond -1}$ are $(\mathbf{d}^{\diamond -1})_i = \frac{1}{\mathbf{d}_i}$.

We use the notation A - B to denote the set difference of A and B, $\{x \mid x \in A \land x \notin B\}$.

We use \mathbf{C} to denote a diagonal matrix having \mathbf{c} 's entries on the diagonal.

For a matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$, we use $diag(\mathbf{M})$ to denote the length *n* vector containing **M**'s diagonal entries.

A good proportion of the work on Set Cover uses n for the number of input sets and m for the size of the universe. In [5], however, Feige uses n for the universe size, as have some others since. By the rationale that n should denote the quantity most deeply involved in attainable approximation ratios, we use n for the universe size and m for the number of input sets.

The choice of using **A** to denote a $n \times m$ matrix (rather than its transpose) was made simply so that the constraints in the integer program formulation of Set Cover would not require a transpose operation to be written.

We frequently use inequalities with vector quantities on either side, by which we intend the conjunction of the elementwise inequalities. For instance, $\begin{pmatrix} a \\ b \end{pmatrix} \leq \begin{pmatrix} c \\ d \end{pmatrix}$ should be understood as stating that $a \leq c$ and $b \leq d$.

We sometimes use the notation 0 or 1 to represent the column vector (of whatever size is appropriate in the context) with a 0 or 1 in every component, respectively.

1.2 Organization of Thesis

In chapter 2 we define the Set Cover and Set Packing problems, discuss some of the relationships that they have to other problems, and describe previous work on approximation algorithms for these problems.

Chapter 3 outlines some preprocessing techniques that can be used to simplify SCP and SPP instances.

In chapter 4, we describe the overall idea that motivates our new heuristics and define the new heuristics themselves.

Chapter 5 contains our principal mathematical results, relating to vectors $\mathbf{v} \in \mathbb{R}^n_+$ for which $\mathbf{M}\mathbf{v} = \mathbf{v}^{\diamond -1}$ for certain $n \times n$ matrices \mathbf{M} .

Chapter 6 describes a few algorithms that we have used effectively to compute our new valuations and includes some discussion of cases for which fixed points can be calculated exactly.

Chapter 7 contains the results of experiments we have done, comparing the quality of approximate solutions to SCP and SPP instances obtained by a variety of simple greedy algorithms.

In a final brief chapter, we briefly summarize our results and suggest possible directions for future research.

Chapter 2

Problem Definitions and Discussion

2.1 Problem Formulations

2.1.1 Set Cover

An instance of the Set Cover Problem (SCP) has the following components:

- (a) There is a set *I*. This contains the *labels*, *names* or *indices* of the input sets and is only for notational convenience. This set can have arbitrary elements, but without loss of generality, we can take it to be {1,...,m}.
- (b) $m = |\mathcal{I}|$, the number of input sets.
- (c) For every $i \in \mathcal{I}$, there is an input set S_i . The S_i 's can, again, be over arbitrary elements but without loss of generality, we can take it that $S_i \subseteq \{1, \ldots, n\}$.
- (d) $\mathcal{U} = \bigcup_{i \in \mathcal{I}} S_i$. \mathcal{U} is the universe or basis set of the instance.
- (e) $n = |\mathcal{U}|$, the size of the universe.
- (f) Associated with every input $i \in \mathcal{I}$ is a cost $c_i \in \mathbb{R}$.

We say that $\mathcal{H} \subseteq \mathcal{I}$ covers \mathcal{U} or that \mathcal{H} is a cover for \mathcal{U} when $\bigcup_{i \in \mathcal{H}} S_i = \mathcal{U}$. That is, every basis element is included in at least one of \mathcal{H} 's sets. Define the cost of any $\mathcal{H} \subseteq \mathcal{I}$ to be $\sum_{i \in \mathcal{H}} c_i$, the sum of the costs of the sets included in \mathcal{H} .

The Set Cover Problem is to find, given \mathcal{I} , S_i and c_i for $i \in \mathcal{I}$, a set cover with minimal cost. The Unweighted Set Cover Problem describes instances for which $c_i = 1$ for all input sets i. The decision problem variant of Set Cover is to determine, given \mathcal{I} , S_i and c_i for $i \in \mathcal{I}$ and a cost threshold k, whether there is a cover $\mathcal{H} \subseteq \mathcal{I}$ with cost not exceeding k.

2.1.2 Hitting Set

Hitting Set is another common NP-hard combinatorial optimization problem. An instance of Hitting Set has the following components.

(a) There is a set \mathcal{U} . As in Set Cover, we call these the *labels* or *names* of the input sets.

(b) Define $n = |\mathcal{U}|$.

- (c) For every $i \in \mathcal{U}$, there is an input set S_i .
- (d) Define $\mathcal{I} = \bigcup_{i \in \mathcal{U}} S_i$.
- (e) Define $m = |\mathcal{I}|$.
- (f) Associated with every element $i \in \mathcal{I}$ is a cost $c_i \in \mathbb{R}$.

We will say that $\mathcal{H} \subseteq \mathcal{I}$ hits an input set S_i if $|\mathcal{H} \cap S_i| \ge 1$, and that $\mathcal{H} \subseteq \mathcal{I}$ is a hitting set for \mathcal{U} if for every $i \in \mathcal{U}$, \mathcal{H} hits S_i . Define the cost of any $\mathcal{H} \subseteq \mathcal{I}$ to be $\sum_{i \in \mathcal{H}} c_i$, the sum of the costs of the basis elements included in \mathcal{H} .

The Hitting Set problem is, then, to find a minimum cost hitting set $\mathcal{H} \subseteq \mathcal{I}$ for \mathcal{U} .

2.1.3 (Hitting $\{Set\}$ Cover $\}$

The construction given below is commonly used to show the equivalence of Hitting Set and Set Cover, but here we use it as a problem definition. We regard this problem as unifying Hitting Set and Set Cover into a uniform terminology that easily translates back to either problem. For lack of a standard name, we will call this problem (Hitting {Set) Cover} or HSC.

An HSC instance consists of the following objects.

- (a) $L = \{l_1, \ldots, l_m\}$ and $R = \{r_1, \ldots, r_n\}$ are disjoint sets.
- (b) Let $c: L \to \mathbb{R}$ give the costs of L's elements.
- (c) $G = \langle L \cup R, E \rangle$ is an undirected bipartite graph where $E \subseteq L \times R$, so every edge connects one element of L with one from R.

We then define a hitting set cover for the problem (G, c) to be any $\mathcal{H} \subseteq L$, the union of whose neighbours is R. That is, for every right element $r \in R$ there is some $l \in \mathcal{H}$ so that the edge $(l, r) \in E$ exists. For any subset $\mathcal{H} \subseteq L$, we define its cost $c(\mathcal{H})$ as the sum of its elements' costs $c(\mathcal{H}) = \sum_{l \in \mathcal{H}} c(l)$. A minimum hitting set cover, then, is any hitting set cover whose cost is least possible among all hitting set covers.

To see why this problem unifies hitting set and set cover, consider the adjacency matrix of G. Since G is bipartite, there are no edges between pairs of elements both in either of L or R, so we can write its

adjacency matrix as

$$\mathbf{M}_G = \left(\begin{array}{cc} 0 & \mathbf{A} \\ \mathbf{A}^T & 0 \end{array}\right)$$

where **A** is an $n \times m$ matrix with

$$\mathbf{A}_{i,j} = \begin{cases} 1 & \text{if } (l_j, r_i) \in E \\ 0 & o.w. \end{cases}$$

We will call **A** the fundamental matrix of a HSC problem. We would also like to define **c** as the vector of costs of L's elements, so that for $i \in \{1, ..., m\}$, we have $\mathbf{c}_i = c(l_i)$.

Now, \mathbf{c} and \mathbf{A} contain a complete description of our problem. When the columns of \mathbf{A} are set to the adjacency vectors of the input sets from a hitting set instance, we have a problem in this setting equivalent to the hitting set problem. Likewise, if we set \mathbf{A} 's rows to be the input sets from an instance of Set Cover, we have an equivalent problem. Because we can always transform instances of hitting set and set cover into problems of this sort, we can treat both problems uniformly by considering this problem instead.

As an example, consider the SCP instance $\mathcal{I} = \{1, 2, 3, 4, 5\}$, with $S_1 = \{1, 2, 3\}$, $S_2 = \{2, 4\}$, $S_3 = \{1, 3\}$, $S_4 = \{4\}$, and $S_5 = \{3, 4\}$. The costs are given by $c_1 = 3$, $c_2 = 1$, $c_3 = 2$, $c_4 = 1$, $c_5 = 2$. The figure below shows the HSC instance that this generates.



Figure 2.1: Graph structure of an HSC instance.

If we were to take the Hitting Set instance $\mathcal{U} = \{1, 2, 3, 4\}$, with $S_1 = \{1, 3\}$, $S_2 = \{1, 2\}$, $S_3 = \{1, 3, 5\}$, $S_4 = \{2, 4, 5\}$ and costs $c_1 = 3, c_2 = 1, c_3 = 2, c_4 = 1, c_5 = 2$, we obtain the same equivalent HSC instance.

This instance has fundamental matrix and cost vector given by the following:

$$\mathbf{A} = \begin{pmatrix} 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 1 & 1 \end{pmatrix} \quad \mathbf{c} = \begin{pmatrix} 3 \\ 1 \\ 2 \\ 1 \\ 2 \end{pmatrix}$$

2.1.4 Integer Programming Formulation

Set Cover can be written very simply as an equivalent integer program, in which:

(a) **A** is some particular $n \times m \ 0/1$ matrix.

- (b) **c** is a vector of costs: $\mathbf{c} \in \mathbb{R}^m$.
- (c) \mathbf{x} is a vector of binary variables.

Then the IP problem is

Minimize $\mathbf{c}^T \mathbf{x}$ subject to $\mathbf{A}\mathbf{x} \ge \mathbf{1}$ and $\mathbf{x}_i \in \{0, 1\}$ for all $i = 1, \dots, m$

This problem is equivalent to the SCP instance whose subsets are given by understanding the columns of \mathbf{A} as adjacency vectors that indicate which of the *n* basis elements are included in each input set. This is equivalent to the hitting set instance whose sets to hit are given by \mathbf{A} 's rows.

2.1.5 Set Packing and Formulations

Set Packing is also a well known NP-hard combinatorial optimization algorithm. As with SCP, a Set Packing instance consists of the following objects:

- (a) There is a set \mathcal{I} . This contains the *labels* or *names* of the input sets.
- (b) Define $m = |\mathcal{I}|$.
- (c) For every $i \in \mathcal{I}$, there is an input set S_i .
- (d) Define $\mathcal{U} = \bigcup_{i \in \mathcal{I}} S_i$. \mathcal{U} is the universe or basis set of the problem.
- (e) Define $n = |\mathcal{U}|$.
- (f) Associated with every input set $i \in \mathcal{I}$ is a cost $c_i \in \mathbb{R}$

Call $\mathcal{H} \subseteq \mathcal{I}$ a packing for \mathcal{U} if every basis element is included in at most one of \mathcal{H} 's sets. Equivalently, $\mathcal{H} \subseteq \mathcal{I}$ is a packing for \mathcal{U} if the input sets selected are pairwise disjoint. That is, $\forall i, j \in \mathcal{H}, i =$ *j* or $S_i \cap S_j = \emptyset$. Also define the weight of a packing $\mathcal{H} \subseteq \mathcal{I}$ to be the sum of the weights of the included sets, $c(\mathcal{H}) = \sum_{i \in \mathcal{H}} c_i$. The Set Packing problem is to find the maximum weight packing $\mathcal{H} \subset \mathcal{S}$ of \mathcal{U} .

It can be seen that the description of a Set Cover instance has precisely the same description as a Set Packing instance. Given \mathcal{I}, S_i, c_i for $i \in \mathcal{I}$, we can just as well ask what the maximum weight packing or the minimum cost cover is.

Much the same as with the relationship between Set Cover and Hitting Set, Set Packing has an analogous equivalent problem phrased in terms of elements with costs and constraints on subsets of these elements. This problem has not been given any particular attention that we are aware of, but we name and define it here for the sake of strengthening the analogy between Set Cover and Set Packing. We call this the *Jabbing Set Problem*.

A Jabbing Set instance consists of the following objects:

- (a) There is a set \mathcal{U} . This contains the *labels* or *names* of the input sets.
- (b) Define $n = |\mathcal{U}|$.
- (c) For every $i \in \mathcal{U}$, there is an input set S_i .
- (d) Define $\mathcal{I} = \bigcup_{i \in \mathcal{U}} S_i$. \mathcal{I} is the universe or basis set of the problem.
- (e) Define $m = |\mathcal{I}|$.
- (f) Associated with every input $i \in \mathcal{I}$ is a cost $c_i \in \mathbb{R}$

Call $\mathcal{H} \subseteq \mathcal{I}$ a jabbing set for \mathcal{U} if every subset S_i for $i \in \mathcal{U}$ is hit at most once, $\forall i \in \mathcal{I}, |S_i \cap \mathcal{H}| \leq 1$. The Jabbing Set problem is then to find the maximum weight jabbing set $\mathcal{H} \subseteq \mathcal{I}$ for \mathcal{U} .

Much as with Set Cover, we will often favor the Integer Program formulation of Set Packing. Set Packing can be written very simply as an integer program as follows, differing from the formulation for Set Cover only by the direction of an inequality and swapping the *minimize* for *maximize*:

- (a) **A** is some particular $n \times m \ 0/1$ matrix.
- (b) **c** is a vector of costs: $\mathbf{c} \in \mathbb{R}^m$.
- (c) \mathbf{x} is our vector of binary variables.

Then the IP problem is

Maximize $\mathbf{c}^T \mathbf{x}$ subject to $\mathbf{A}\mathbf{x} \leq \mathbf{1}$ and $\mathbf{x}_i \in \{0, 1\}$ for all $i = 1, \dots, m$

2.1.6 Set Packing's Relationship to Maximum Independent Set

The Weighted Maximum Independent Set (WMIS) problem (also called Vertex Packing) can be described as follows: For a graph $G = \langle V, E \rangle$ with vertex weights given by c_i for $i \in V$, what is the largest weight subset \mathcal{H} of V such that no pair of elements from \mathcal{H} are neighbours in G. This problem is also NP-hard. The reductions between Set Packing and WMIS are particularly clean. Given a WMIS problem defined by $G = \langle V, E \rangle$ and c, an equivalent Set Packing instance is $\mathcal{I} = V$, $S_i = \{e \in E \mid i \in e\}$. Any max weight independent set $\mathcal{H} \subseteq V$ for G, c is also a max weight packing for V, c and the S_i 's. To see this construction more clearly, consider the elements' neighbourhoods for the SPP instance. For every $e \in \mathcal{U}$ we can define $N_e = \{i \in \mathcal{I} \mid e \in S_i\}$. Then, for every $e \in E$, we make an element with a neighbourhood of size two.

Given instead some Set Packing instance \mathcal{I}, c and S_i for $i \in \mathcal{I}$, we can make an equivalent WMIS instance by defining $G = \langle \mathcal{I}, E \rangle$ where E has an edge (i, j) for every $i \neq j$ in \mathcal{I} iff $S_i \cap S_j \neq \emptyset$. Then the optimal solutions to these problems are the same. By considering the element neighbourhoods, we can construct the WMIS problem even more simply. For each $e \in \mathcal{U}$ we add a clique to G, yielding an edge between every $i, j \in N_e$.

Viewed this way, we can see SPP as an alternative way to specify WMIS instances where we can write a graph as a union of arbitrary size cliques, rather than just 2-cliques, which are simply edges. This will be relevant later in section 4.3, where we specify our new Set Packing heuristic.

To be more formal, consider a problem with input V some set, S_1, \ldots, S_n all subsets of V specifying cliques among V's elements and $c : V \to \mathbb{R}$ weights of the elements of V. The problem is to find the WMIS of $G = \langle V, E \rangle$ with weights given by c and $E = \bigcup_{i,j \in \{1,\ldots,n\}, i \neq j \text{ and } S_i \cap S_j \neq \emptyset} \{(i,j)\}$. If we restricted this to instances where $|S_i| = 2$ for all $i \in \{1,\ldots,n\}$, the problem is WMIS. Without that restriction, this problem is equivalent to Set Packing in precisely the same way that Set Cover and Hitting Set are equivalent.

2.2 Greedy Algorithms for Set Cover/Packing Approximation

2.2.1 Set Cover/Packing Approximation

Since the underlying problems are NP-hard, and are not expected to have polynomial time algorithms for exact solution, it has become an area of interest to find algorithms that run in polynomial time and perform well by some measure. Researchers have been particularly interested in algorithms that offer a guarantee on the ratio between the provided solution and the optimal solution. The standard greedy algorithm for SCP does just that, and provides a worst-case approximation ratio very close to what has been proven (under plausible assumptions) the best possible worst-case approximation ratio for any polynomial time algorithm.

2.2.2 Set Cover Approximation

The best studied technique for obtaining good approximate solutions to Set Cover problems was first discussed in [4]. Let the input be a \mathcal{I} , the sets S_i for all $i \in \mathcal{I}$ and costs $c_i \in \mathbb{R}_+$ for every $i \in \mathcal{I}$. We can describe the *Standard Greedy Algorithm for SCP* in the following pseudocode.

 $\mathcal{H} \leftarrow \emptyset$ while $\bigcup_{i \in \mathcal{I}} S_i \neq \emptyset$ do
$$\begin{split} b \leftarrow \operatorname{argmax}_{i \in \mathcal{I}} \frac{|S_i|}{c_i} \\ \mathcal{H} \leftarrow \mathcal{H} \cup \{b\} \\ \text{for all } i \in \mathcal{I} \text{ do} \\ S_i \leftarrow S_i - S_b \\ \text{end for} \\ \text{end while} \\ \text{return } \mathcal{H} \end{split}$$

Here, \mathcal{H} contains the indices of the sets that have been selected so far to form a cover of \mathcal{U} . In every run through the while loop, the algorithm selects the set maximizing the ratio between the number of uncovered elements and the cost of the set. It then modifies the input sets to remove the newly covered elements from them, effectively removing some elements of \mathcal{U} since they have already been covered and it is irrelevant whether they are covered again or not.

2.2.3 Set Packing Approximation

For Set Packing, we describe an analogous algorithm that greedily selects additional sets until no more can be added. For convenience, we describe the weights somewhat differently than in the above. Let the input be a set of input set names \mathcal{I} , each of the sets S_i for $i \in \mathcal{I}$, and weights given by $c_i \in \mathbb{R}_+$ for $i \in \mathcal{I}$

```
\begin{array}{l} \mathcal{H} \leftarrow \emptyset \\ \text{while } \mathcal{I} \neq \emptyset \text{ do} \\ b \leftarrow \operatorname{argmax}_{i \in \mathcal{I}} \frac{c_i}{\sqrt{|S_i|}} \\ \mathcal{H} \leftarrow \mathcal{H} \cup \{b\} \\ \text{for all } i \in \mathcal{I} \text{ do} \\ \text{ if } S_i \cap S_b \neq \emptyset \text{ then} \\ \mathcal{I} \leftarrow \mathcal{I} - i \\ \text{ end if} \\ \text{ end for} \\ \text{end while} \\ \text{return } \mathcal{H} \end{array} \right
```

In this algorithm, we select the remaining feasible set for which the cost per square root of the number of elements is largest. We then remove all sets that contain any element in common with the selected set, effectively removing some sets from S if our latest choice means that their selection would be infeasible. The particular choice of the set maximizing $c_i/\sqrt{|S_i|}$ is similar to the algorithm for Set Cover described above, though the square root of the set size may seem odd. In fact, this algorithm has the optimal worst-case approximation ratio for any greedy set packing algorithm in which the valuation of every set i is determined by $|S_i|$ and \mathbf{c}_i and so does not involve the detailed structure of the instance [8]. In section 2.4 we will discuss this and other interesting heuristics.

2.2.4 General Greedy Scheme

The most important line in each of the above algorithms is where the next set to use is chosen. We would like to describe a general greedy algorithm for approximating Cover and Packing problems, abstracting away the particular rationale for selection.

The previous two algorithms provide some valuation of the sets in the instance, assigning a real number to each of them indicating their relative desirability. Both also reduce the underlying instance to reflect the fact that some set has been selected. How this reduction is performed depends on whether we are dealing with a Cover or Packing problem. In addition to these, we can transform or preprocess the instance to a simpler one when the instance has some simple properties. One example suggesting the type of preprocessing we mean is the situation where (for SCP) some input set S_i is a subset of S_j and $c_i > c_j$, so that we immediately know that S_i will not be in any optimal cover since S_j is both larger and cheaper and we can remove S_i from the problem entirely. We will discuss this in greater detail in chapter 3.

The overall scheme can be described in pseudocode as follows, where P denotes an instance:

```
preprocess(P)
```

while a choice can or must be made **do**

 $\mathbf{v} \leftarrow valuation(P)$ $b \leftarrow \operatorname{argmax}_{i \in \mathcal{I}} \mathbf{v}_i$ $\mathcal{H} \leftarrow \mathcal{H} \cup b$ reduce(P, b) preprocess(P)end while return \mathcal{H}

The preprocessing step is not usually included in a description of the standard greedy algorithm for Set Cover, but it can make a difference in the performance of the algorithm. Usually, however, it is not seen as important whether preprocessing is done on SCP instances where the standard greedy algorithm is concerned, as it does not impact the worst-case approximation ratio. For the valuation technique we describe later, certain types of preprocessing are relatively cheap and have an impact on the valuations, whereas for the standard greedy algorithm they do not. This is discussed further in chapter 3.

2.3 The Standard Greedy Cover Heuristic

The standard greedy heuristic for Set Cover is to choose the set for which the number of uncovered elements per unit cost is largest. It was first described in [12] for unweighted instances. For a Set Cover instance given by the matrix **A** and element costs **c**, as in the IP formulation of the problem, the valuation \mathbf{v}_i of set S_i corresponding to column i of **A** will be

$$\mathbf{v}_i = \sum_{j=1}^n \mathbf{A}_{j,i} / \mathbf{c}_i = \frac{1}{\mathbf{c}_i} \sum_{j=1}^n \mathbf{A}_{i,j}^T$$

If we denote the $m \times m$ diagonal matrix with diagonal entries given by \mathbf{c}_i^{-1} for $1 \le i \le m$ by \mathbf{C}^{-1} , we can write the entire valuation vector simply as

$$\mathbf{v} = \mathbf{C}^{-1} \mathbf{A}^T \mathbf{1}$$

In [4], Chvàtal demonstrates that the approximate solutions for Set Cover problems obtained using this valuation must have have a cost of no more than H(n) times the optimal solution, where $H(n) = \sum_{i=1}^{n} \frac{1}{i}$ is the sum of harmonic series up to the *n*th term. In fact, Chvàtal showed that this approximation ratio cannot exceed H(k), where $k = \max_{i \in \mathcal{I}} |S_i|$, the size of the largest input set, which obviously cannot exceed n.

It is also known that for all n, there are instances for which an approximation arbitrarily close to H(n) is attained. We can give a class of such instances quite simply. Let $\mathcal{I} = \{1, \ldots, n+1\}$ and for every $1 \leq i \leq n$ let $S_i = \{r_i\}$ and $S_{n+1} = \{r_1, \ldots, r_n\}$. Figure 2.2 shows the HSC problem graph structure to which this corresponds. The costs of the input sets for $1 \leq i \leq n$ are $c_i = 1/i$, and the cost of the last set $c_{n+1} = 1 + \epsilon$ for any $\epsilon > 0$.

Now, for this instance, the valuation that the standard greedy heuristic provides is

$$\left(1, 2, 3, \dots, i, \dots, n, \frac{n}{1+\epsilon}\right)^T$$

Since the largest of these values is n, we will take the set corresponding to it, S_n , incurring cost 1/n and remove $S_n = \{r_n\}$ from every set. In the next iteration, we are given the valuation $(1, 2, 3, ..., i, ..., n - 1, \frac{n-1}{1+\epsilon})^T$ and select S_{n-1} . This will continue until the final set cover produced is $\mathcal{H} = \{1, ..., n\}$ with cost H(n). The optimal set cover, however, is $OPT = \{n+1\}$ so the attained approximation ratio is $\frac{c(\mathcal{H})}{c(OPT)} = \frac{H(n)}{1+\epsilon}$ and since epsilon can be arbitrarily small, these instances have approximation ratios arbitrarily close to H(n).



Figure 2.2: Graph structure of hard Instances for standard greedy set cover. The costs of l_i for $1 \le i \le n$ are $c_i = 1/i$ and the cost of l_{n+1} is $c_{n+1} = 1 + \epsilon$

The performance of this heuristic for solving Unweighted Set Cover problems, where all costs c_i are

fixed at 1, is also well understood. In [14], Slavík shows that the standard greedy heuristic provides solutions for unweighted instances whose cost is within a factor of $\ln(n) - \ln(\ln(n)) + 0.78$ of the cost of the optimal solution. He also shows that this ratio is tight. For every $n \ge 2$, there is an instance with $|\mathcal{U}| = n$ and the standard greedy heuristic yields a solution with cost more than $\ln(n) - \ln(\ln(n)) - 0.31$ times the optimal solution's cost. Note that this is better than the worst-case performance over weighted Set Cover instances, since $H(n) \approx \ln(n) + 0.58$.

2.4 Other Approximation Techniques

A variety of non-greedy poly-time algorithms have been proposed and used for approximating SCP. LP rounding, dual LP rounding, and primal/dual have been described and proven to have good worst-case approximation ratios. Williamson in [15], in particular, is a good resource for more information on the primal/dual technique.

In [10] Halldórsson describes an algorithm using local search that has worst-case approximation ratio better than the standard greedy algorithm by a constant. Beasley and Chu describe a genetic algorithm-based approach in [3]. In [9], Grossman and Wool describe some variations of the standard greedy algorithm in addition to a neural network-based algorithm. Caprara et al. give a Langrangianbased algorithm in [2] that won the 1994 Set Cover approximation competition *FASTER*.

Chapter 3

Preprocessing

Here we wish to describe some particular simplifications of Set Cover/Packing instances that can be done cheaply and result in problems that are equivalent to and no larger than the input problems. Any greedy algorithm utilizing the scheme presented in section 2.2.4 with some guaranteed approximation ratio that is non-decreasing in n and m (and perhaps $k = \max_{i \in \mathcal{I}} |S_i|$) cannot have a worse worst-case approximation ratio when these preprocessing steps are used.

We discuss some ways to take an instance I (in some convenient representation) and produce a new instance I' such that any optimal solution to I', possibly taken with some specified sets that must be included, is an optimal solution to I and $n' \leq n$, $m' \leq m$ and $d' \leq d$.

For transformations satisfying the above, it should generally be beneficial to perform the transformations and solve the transformed problems as opposed to solving the original problem with the same algorithm. However, it is certainly not the case that all algorithms for approximating SCP can solve the transformed problem with approximation ratio no worse than the original problem for all instances. Consider a transformation as unproblematic as reordering the basis elements, which could just as likely help as hinder an algorithm with some deterministic tie-breaking procedure. Nor can it be said that the approximation ratio must always be improved for all approximation algorithms. All we wish to claim for transformations of this sort is that they cannot worsen the worst-case approximation ratio of the greedy valuation technique used and they frequently improve the approximation ratio when applied.

3.1 Basic Preprocessing

First, the following are a few ways in which the resulting instances are very clearly equivalent.

- (a) Renaming the elements of the universe results in a problem that is semantically equivalent, though its representation can be somewhat different.
- (b) Renaming the input set labels.
- (c) Scaling all of the costs by some constant $\lambda > 0$. Even though the optimal cost after such a transfor-

mation will be different, it will be predictably scaled by λ . We wish to regard such scaled problems as equivalent.

Next, we describe some transformations that make important changes to instances. For every simplifying transformation that can be applied to a Set Cover problem there is an analogous transformation for Set Packing problems.

(d) For a Set Cover instance, if any set has a cost that is non-positive, we can take that set and remove all of its elements from the other input sets. Since it costs nothing or less than nothing to include, there must be some optimal cover that includes it. For a Set Packing instance, any set with non-positive weight can be rejected immediately, since it does not increase the total weight and its inclusion may hurt our ability to take other sets. There must be an optimal packing that excludes such a set, so we can safely exclude it.

In the literature, it is usually simply assumed that all costs/weights are positive. This transformation explains why that requirement does not diminish the generality of instances with only positive costs. Every one of the transformations discussed here will allow us to analogously trim the space of SCP and SPP instances with which we should be prepared to deal.

- (e) If some input set is empty (and assuming that it has positive cost, as the above permits us to), then it will not be in any optimal cover since it incurs a cost and does no work towards our goal, so we can remove it from the problem. It will, however, be included in every optimal packing, so we can require that it is taken.
- (f) Consider an instance where some universe element is not present in any of the input sets.¹ This means that no set cover exists for the instance. It is immediately infeasible and any attempt to find a cover will fail. For Set Packing, however, we can remove the uncovered elements from the universe and work with that equivalent simpler problem instead.
- (g) Consider an instance where some universe element is present in exactly 1 of the input sets. For SCP, the sole input set including this element must be included in every cover, so we can immediately include it and reduce the problem respecting that inclusion. For SPP, we can remove that element from the universe and the set that contains it. Since the constraint that every element must be included no more than once cannot fail to be satisfied for this element, we can leave it out of the problem entirely.

We call these 4 transformations *basic preprocessing*. When we talk about applying basic preprocessing to an instance, we mean that we apply all of these steps to an instance recursively until no more can be applied.

3.2 Subsumption Testing

There is an additional pair of preprocessing simplifications that we call *subsumption* preprocessing, due to the similarity they share with subsumption in the context of SAT instances. One of them operates

¹For the way that we have defined the problem, this is not possible, since the universe \mathcal{U} is defined as the union of the input sets. We describe this only for completeness.

from the perspective of the input sets, sometimes permitting the removal of an input set since in any cover/packing it can always be replaced by another input set with no loss of quality. The other operates from the perspective of the elements, enabling us to remove an element when some other element enforces a strictly stronger constraint on the problem.

(h) If there are 2 sets S_i, S_j such that $S_i \subseteq S_j$ and $c_i \ge c_j$, then we can see that *i* need not ever be included in a cover. In any cover using *i*, we can instead replace it with *j* to find a cover that is no more expensive. Therefore, we can remove set *i* from the problem.

For SPP, if there are 2 sets S_i, S_j such that $S_i \subseteq S_j$ and $c_i \ge c_j$, then we can remove set j from the instance. Any packing using j would run up against at least as many constraints as one using i, and it would do so for no more gain that i.

Basic preprocessing step (e) can be seen as a special case of this step.

(i) For any element e ∈ U, let N_e = {i ∈ I | e ∈ S_i} be the set of input sets of which e is a member.
For SCP, if there are two different elements e, e' ∈ U with N_e ⊆ N_{e'}, e is a stronger constraint, and every time one of its elements is selected, so will one of e's. Thus we can omit e', removing it from every set in N_{e'}, and have an equivalent problem.

For SPP, if there is a pair of different elements $e, e' \in \mathcal{U}$ where $N_e \subseteq N_{e'}$, e' is a strictly stronger constraint. If we remove e from all sets in N_e , we are left with a problem with precisely the same feasible packings.

Basic preprocessing step (f) can be seen as a special case of this step.

Running subsumption testing has high complexity relative to the basic preprocessing steps. With the most naïve approach we need to check $\mathcal{O}(m^2)$ different cases for set subsumption and $\mathcal{O}(n^2)$ cases for element subsumption. In the algorithm we use for our experiments, the only shortcut we use is checking sets/elements whose neighbourhoods have shrunk since we last checked whether they could be involved in some subsumption.

3.3 Independent Subproblem Separation

Consider the set cover instance $\mathcal{I} = \{1, 2, 3, 4, 5\}$ with $S_1 = \{1, 2\}$, $S_2 = \{2, 3\}$, $S_3 = \{1, 3\}$, $S_4 = \{4, 5\}$, $S_5 = \{4, 5\}$ and uniform costs. In figure 3.1 we see the HSC instance associated with this problem. We can easily observe that choosing input sets S_1 , S_2 and S_3 can never obtain elements 4 and 5. Likewise, sets S_4 and S_5 can never obtain elements 1, 2, or 3. We can solve instances with $\mathcal{I}' = \{1, 2, 3\}$ and $\mathcal{I}'' = \{4, 5\}$ independently and combine the covers for these to make a cover for the original instance. In general, we can solve each component of an HSC graph independently and combine them afterwards to obtain a cover.

Stated with reference only to the input of SCP or SPP instances, we can say that if any subset of the input sets $\mathcal{I}' \subset \mathcal{I}$ has the property that

$$\left(\bigcup_{i\in\mathcal{I}'}S_i\right)\cap\left(\bigcup_{i\in\mathcal{I}-\mathcal{I}'}S_i\right)=\emptyset$$



Figure 3.1: HSC graph structure for a problem with 2 independent subproblems.

then we can solve the instances with sets \mathcal{I}' and $\mathcal{I} - \mathcal{I}'$ independently.

The basic preprocessing step (e) can be viewed as a case of this preprocessing technique.

3.4 Inferring Stronger Packing Constraints

Recalling section 2.1.6, a Set Packing instance corresponds to an WMIS instance on the graph $\langle \mathcal{I}, E \rangle$ where E has a clique among the neighbourhood N_e of every element $e \in \mathcal{U}$. By finding larger cliques in this graph, we can rewrite the element neighbourhoods in the SCP instance, which may allow subsumption preprocessing to further simplify the instance.

In particular, if there is some $e \in \mathcal{U}$ and $i \in \mathcal{I} - N_e$ for which S_i intersects each of the sets in N_e , we can add e to S_i and the resulting problem will have logically equivalent constraints. We can see this process in action in figure 3.2.

We can use this observation to build maximal constraints, corresponding to maximal cliques in the WMIS graph, in polynomial time. There may be multiple descriptions of the graph as maximal cliques, and we may not find the *maximum* clique in our search, given the NP-hardness of the Max Clique problem. It is also not clear that this preprocessing step is always to our advantage in the context of k-Set Packing, since we allow ourselves to add elements to sets. Note, however, that in figure 3.2 we use constraint strengthening and subsumption to transform a 3-Set Packing instance into a 1-Set Packing instance, showing that this step is of use for at least some instances.



Figure 3.2: Strengthening packing constraints. In the first step, we can see that l_4 intersects l_1 on r_2 , l_2 on r_3 , and l_3 on r_4 . This permits us to add r_1 to l_4 , shown as a new green edge in the central diagram. In the rightmost diagram, we have removed the other 3 elements from the instance because they are subsumed by element r_1 .

3.5 Non-Minimal Covers

In some cases the standard greedy heuristic for SCP will yield solutions that are not minimal. Let \mathcal{H} be the returned solution having $\bigcup_{i \in \mathcal{H}} S_i = \mathcal{U}$. It sometimes happens that there is a strict subset of \mathcal{H} that is a set cover. Consider the following instance:

Example:

 $\mathcal{I} = \{1, 2, 3, 4\}$ $S_1 = \{1, 3\}$ with cost $c_1 = 2$ $S_2 = \{2, 4\}$ with cost $c_2 = 2$ $S_3 = \{3, 4\}$ with cost $c_3 = 1$ $S_4 = \{1, 2\}$ with cost $c_4 = 5$

For this instance, the initial valuations are $(1, 1, 2, 0.4)^T$, so S_3 is selected. In the next iteration, the valuations are $(0.5, 0.5, 0.4)^T$, so S_1 or S_2 is selected with the other one joining it in the following iteration. Thus the returned set cover is $\mathcal{H} = \{S_1, S_2, S_3\}$. $\{S_1, S_2\}$ is strictly smaller and is also the optimal cover.

In fact, it is quite common for solutions returned by the standard greedy algorithm to be redundant, so when we want to find approximate solutions to Set Cover instances, it is often advantageous to verify that the returned solution is minimal or reduce it to a minimal solution. We now describe some of the techniques that have been used to reduce non-minimal approximate solutions to minimal ones.

Wool and Grossman

In [9], the following technique is used to minimize approximate Set Cover solutions.

Let the Set Cover problem of interest be given by \mathbf{A}, \mathbf{c} as in the IP formulation, with the non-minimal solution given by the binary vector \mathbf{x} . Let $\mathbf{r} = \mathbf{A}\mathbf{x} - \mathbf{1}$. \mathbf{r}_i is the number of times that element *i* has been redundantly acquired by the chosen sets. Then, for each set $1 \le i \le m$, define the minimal redundancy $\mathbf{t}_i = \min_{j \in S_i} \mathbf{r}_j$. For any redundant input set *i*, $\mathbf{t}_i > 0$ and it can be eliminated from our cover. We choose the set *b* for which \mathbf{t}_b is largest (breaking ties by choosing the largest cost set), eliminate it from our cover by setting $\mathbf{x}_b = 0$ and continue until \mathbf{t} has 0's in every component.

Recursive Solution

It is possible to formulate the problem of minimizing a redundant Cover as an instance of Set Cover. Given a problem specified by \mathbf{A}, \mathbf{c} and a redundant solution given by the binary vector \mathbf{x} we can succinctly specify the IP with variable vector $\mathbf{z} \in \{0, 1\}^m$:

The constraint that $\mathbf{z} \leq \mathbf{x}$ represents the requirement that we are looking for a subset of the sets chosen in the first run through our problem. Once we remove the columns of \mathbf{A} for which $\mathbf{x}_i = 0$ and reduce the length of our variable vector \mathbf{z} appropriately, we have an IP with the precise form of a Set Cover problem. We can then solve this problem using the same greedy process with which we obtained \mathbf{x} and recursively minimize that solution. However, in order for us to guarantee that the recursive process terminates, we need to show that the resulting instance is strictly smaller than the original instance we set out to solve in the first place. This can only fail when $\mathbf{x} = \mathbf{1}$, when our greedy process selects every input set but the produced cover is redundant.

Since our algorithm stops after every element has been covered, there is some element that was covered for the first time in the last iteration of the greedy selection process. If we first reduce the instance **A**, **c** by requiring the inclusion of the set selected in the last iteration of the previous round, we will be solving a strictly smaller instance. For instance, in the example given earlier in this section, we run the greedy algorithm 3 times on subinstances with $(\mathcal{I}, \mathcal{U})$ given by, successively, $(\{1, 2, 3, 4\}, \{1, 2, 3, 4\})$, $(\{2, 3\}, \{2, 4\})$, $(\{3\}, \emptyset)$, resulting in the optimal cover being selected, since the final instance is trivial.

Alternatively, we could require *all* sets in the redundant solution that contain any element that was chosen exactly once to be included before forming the recursive subinstance or simply doing basic preprocessing between greedy iterations. Either would ensure that the recursive instance is smaller and that our process for minimization would terminate.

Chapter 4

New Greedy Heuristics for Set Cover and Set Packing

4.1 Motivation

In the general greedy scheme for Set Cover we have a variety of options available to us. To distinguish between them we generate a valuation for each alternative. We require that the valuations all be positive, corresponding to the fact that any choice makes *some* progress towards our goal. While constructing a cover, we know that every element must be obtained at some step. Every element forms a hard constraint on our eventual solution.

The standard greedy heuristic is to select the option which satisfies the most constraints for the least cost. If, instead, we consider each constraint not to have identical difficulty we might prefer to make the choice such that the sum of the difficulties overcome is largest per unit cost. How then, do we assign difficulties to the constraints? The standard heuristic can be seen to derive from the decision that all constraints are equally difficult to satisfy. We might instead try to evaluate the difficulty of the different constraints by referring recursively to the valuations of the choices that would satisfy them. A constraint that is satisfied by many available options with high valuations is going to be effectively less difficult to satisfy than one satisfied only by few with low valuations. Other choices are possible, but for our immediate purposes we assign constraint difficulties inversely proportional to the sum of the valuations of choices that would satisfy them.

If we write the vector of valuations as \mathbf{v} and the difficulties as \mathbf{d} , we have

$$\mathbf{v} = \mathbf{C}^{-1} \mathbf{A}^T \mathbf{d}$$
$$\mathbf{d} = (\mathbf{A}\mathbf{v})^{\diamond - 1}$$

This gives us a recursive definition of the valuation that we originally sought to find. In the following we show that this pair of definitions (and some variations) are satisfied only by a unique vector of valuations \mathbf{v} , can be found fairly quickly, and yield results that are generally preferable to those yielded by the

standard heuristic.

We are interested in this general idea, generating the valuations for our current options by defining the valuation recursively, and think that it can be effectively applied in a variety of circumstances. We also develop an intuitive heuristic for Set Packing using this idea and show that it performs better than other simple greedy heuristics. We are unable to prove anything about the quality of the covers and packings obtained from our heuristics, beyond our experimental results, but hope that future work may be able to whether they yield some worst case approximation guarantee.

More importantly, we believe that the particular heuristics we have defined can be of practical use immediately and that the general idea of calculating valuations recursively can be used effectively for a wide variety of tasks.

4.2 The New Greedy Set Cover Heuristic

4.2.1 Relationship to the Standard Heuristic

The standard greedy heuristic works by assigning a valuation for the set i of $\mathbf{v}_i = |S_i|/\mathbf{c}_i$ (that is, the number of input sets hit per unit cost) and irrevocably including the set maximizing this quantity to a set cover that we are building. We then reduce the instance in accordance with the newly obtained elements no longer constraining our future decisions, redo the preprocessing and continue on until we have hit every element.

Consider the standard greedy heuristic as our starting point

$$\mathbf{v}^{(1)} = \mathbf{C}^{-1} \mathbf{A}^T \mathbf{1}$$

Given these valuations, we can consider how *difficult* it might be to hit particular universe elements. Since we are less likely to select sets with low valuations, we might assign the inverse of the sum of the valuations of the sets of which it is a member. That is:

$$\mathbf{d}^{(1)} = \left(\mathbf{A}\mathbf{v}^{(1)}\right)^{\diamond - 1}$$

From this, we might wish to continue the process, letting the valuations of sets be the cost-weighted sum of the difficulties of their elements and recompute the difficulties, so we define:

$$\mathbf{v}^{(i)} = \mathbf{C}^{-1} \mathbf{A}^T \mathbf{d}^{(i-1)}$$
$$\mathbf{d}^{(i)} = \left(\mathbf{A} \mathbf{v}^{(i)}\right)^{\diamond - 1}$$

We can now calculate $\mathbf{v}^{(i)}$ for any *i* we like. In practice, as we use $\mathbf{v}^{(i)}$ for larger *i* as a valuation for greedy set cover approximation, the resulting covers tend to improve.

We can instead consider what valuation $\mathbf{v} \in \mathbb{R}^m_+$ and difficulty $\mathbf{d} \in \mathbb{R}^n_+$ could be selected so that

simultaneously $\mathbf{v} = \mathbf{C}^{-1} \mathbf{A}^T \mathbf{d}$ and $\mathbf{d} = (\mathbf{A} \mathbf{v})^{\diamond -1}$. For such a pair of vectors, we would have

$$\mathbf{d} = \left(\mathbf{A}\mathbf{C}^{-1}\mathbf{A}^T\mathbf{d}\right)^{\diamond - 1}$$

4.2.2 Consistent Valuations

In general, we can write a heuristic by recursively defining two vectors $\mathbf{v} \in \mathbb{R}^m_+$ and $\mathbf{d} \in \mathbb{R}^m_+$ and finding some pair of vectors that satisfy the definition. When there is some pair of vectors satisfying this recursive definition, we will call them a *consistent valuation* and may use \mathbf{v} as our greedy heuristic. We require that these vectors be strictly positive because we find the notion of a negative value or difficulty incoherent in this setting, since regardless of how poor a choice some set may be, it must make some positive progress towards the goal of collecting every basis element. The choice of

$$\mathbf{v} = \mathbf{C}^{-1} \mathbf{A}^T \mathbf{d}$$
$$\mathbf{d} = \mathbf{1}$$

yields the standard heuristic. For a recursive definition of this sort we would hope that a consistent valuation exists and would ideally be unique. For the definitions immediately above it is clear that, for any particular instance, there is a unique consistent valuation \mathbf{v}, \mathbf{d} satisfying it.

4.2.3 A Family of Heuristics

The form that we wish to propose is

$$\mathbf{v} = \mathbf{C}^{-1} \mathbf{A}^T \mathbf{d}$$
$$\mathbf{d} = (\mathbf{A} \mathbf{C}^{\gamma} \mathbf{v})^{\diamond -1}$$

where $\gamma \in \mathbb{R}$ is a free parameter that we will fix only later. For now we consider all valuations generated by γ ranging over the reals. We find that $\gamma = -3$ performs well. γ can be viewed as additionally penalizing sets with high cost.

If we can find a \mathbf{d} for which

$$\mathbf{d} = (\mathbf{A}\mathbf{C}^{\gamma}\mathbf{v})^{\diamond -1}$$
$$= (\mathbf{A}\mathbf{C}^{\gamma-1}\mathbf{A}^T\mathbf{d})^{\diamond -1}$$

then we have a consistent valuation, since we can immediately calculate \mathbf{v} given \mathbf{d} , and its value is uniquely determined by \mathbf{d} . If we have additionally determined that there is a unique \mathbf{d} for which the above holds, then the consistent valuation given by our recursive definition is also unique.

Let us write $\mathbf{M} = \mathbf{A}\mathbf{C}^{\gamma-1}\mathbf{A}^T$. **M** is symmetric positive semi-definite, since

$$\mathbf{M} = (\mathbf{C}^{\frac{\gamma-1}{2}})\mathbf{A}^T)^T ((\mathbf{C}^{\frac{\gamma-1}{2}})\mathbf{A}^T)$$

If we consider only instances with no empty input sets, which we can do by requiring that basic preprocessing be done on instances, every diagonal element of \mathbf{M} is strictly positive. Also, if we consider only instances for which every set's cost is strictly positive, which can be accomplished again by doing basic preprocessing, every component of \mathbf{M} is non-negative. From the results of chapter 5, these qualities ensure that there is a unique \mathbf{d} satisfying $\mathbf{d} = (\mathbf{M}\mathbf{d})^{\diamond -1}$, and thus that our recursive definitions specify a unique consistent valuation for every Set Cover instance that has undergone basic preprocessing.

4.2.4 Relationship to Theory

In our experiments, we will be most concerned with the heuristic obtained by setting the value $\gamma = -3$. It can be argued that $\gamma = 0$ is more natural, defining the difficulty of an element as the reciprocal of the sum of valuations of the sets it is in, instead of having the terms of this sum weighted by the inverse cube of the set's cost, but we have found that using $\gamma = -3$ performs better than other values of γ in an average case sense. An experiment justifying this choice can be seen in section 7.3.1. We have not been able to determine, for any value of γ , whether the new cover heuristic has some worst-case guaranteed approximation ratio.

Assuming that valuations given by the new cover heuristic can be calculated or approximated arbitrarily well in polynomial time (for which we provide evidence in section 6.3.1) and some plausible complexity assumptions, Feige's result [5] shows that there must be classes of Set Cover instances for which the approximation ratio obtained using the new heuristic exceeds $\ln(n) - c \ln(\ln(n))^2$ for some c > 0 for sufficiently large n. We have been unable to find any class of instances for which the new heuristic gives approximation ratios exceeding a constant as n grows arbitrarily large, though we expect that such classes of instances must exist. Feige's work could theoretically be used to build explicit instances, though it would be highly inconvenient to do so.

4.3 The New Greedy Set Packing Heuristic

A common greedy heuristic for WMIS is to select the vertex with largest weight divided by neighbourhood size. Considering the relationship between WMIS and SPP, this heuristic can be transferred to work on Set Packing instances. For every input set $i \in \mathcal{I}$, the valuation of i, $v_i = c_i/(|\{j \in \mathcal{I} \mid S_j \cap S_i \neq \emptyset\}| - 1)$ the set's weight per intersecting input set besides itself. The -1 is to exclude the input set from being counted among its neighbours. In the unweighted case, this heuristic behaves identically whether we include or exclude the set as a neighbour of itself. We could instead consider each set a neighbour of itself, but experimentally both valuations achieve good quality packings so we will focus on the WMIS heuristic, which does not consider a set to be in its own neighbourhood. We call this the valuation the *MIS heuristic* for Set Packing.

For the IP representation of SPP, we can write the valuation vector for the standard heuristic very simply. Define the *binarization* of a matrix $bin : \mathbb{R}^{n \times n} \to \{0,1\}^{n \times n}$ by letting $bin(\mathbf{M})$ be the matrix

having zeros where M has zeros, and 1's where M has non-zeros. Thus

$$bin(\mathbf{M})_{i,j} = \begin{cases} 0 & \text{if } \mathbf{M}_{i,j} = 0\\ 1 & \text{if } \mathbf{M}_{i,j} \neq 0 \end{cases}$$

With this function, the standard heuristic's valuation is written $\mathbf{v} = \mathbf{C}((bin(\mathbf{A}^T\mathbf{A}) - \mathbf{I})\mathbf{1})^{\diamond -1}$. Note that $bin(\mathbf{A}^T\mathbf{A})\mathbf{1} - \mathbf{I}$ is precisely the adjacency matrix of the WMIS instance equivalent to the SPP instance with constraints given by \mathbf{A} .

Our valuation comes from trying to generalize this idea. Instead of valuing a set as its weight divided by the quantity of neighbours it has, we could value it as its weight divided by the sum of its neighbours' weights, or its weight divided by the sum of its neighbours' ratios *weight divided by neighbourhood size* and so on. This leads us to the following recursive definition for our new valuation:

$$\mathbf{v} = \mathbf{C}(bin(\mathbf{A}^T\mathbf{A})\mathbf{v})^{\diamond -1}$$

Letting $\mathbf{M} = \mathbf{C}^{-1} bin(\mathbf{A}^T \mathbf{A})$, we have $\mathbf{v} = \mathbf{M} \mathbf{v}^{\diamond -1}$. If we insist on basic preprocessing, $bin(\mathbf{A}^T \mathbf{A})$ will have 1's on its diagonal and no negative entry. Left multiplying this by \mathbf{C}^{-1} , which is diagonal with strictly positive diagonal entries, results in a matrix \mathbf{M} with no negative entries and all diagonal entries strictly positive. From the result of section 5.1, we know that there must be some strictly positive vector \mathbf{v} satisfying our recursive definition. In general, it is not the case that \mathbf{M} will be positive semi-definite, so we cannot assert that it is unique. We will discuss the question of the uniqueness of our new Set Packing valuation in some detail in section 6.5. The results of our experiments with random instances, comparing the weight of packings produced by the new and 3 other heuristics, are presented in section 7.4.

We have tried to insert a parameter analogous to the new cover valuation's γ , but for all variations we have attempted, the straightforward definition above appears to have performance at least good as similar alternatives.

Chapter 5

Mathematical Results

The aim of this section is to investigate what conditions (on a $n \times n$ matrix **M**) are sufficient for the existence and uniqueness of *positive* solutions $\mathbf{v} \in \mathbb{R}^n_+$ for systems of equations of the form

$$(\mathbf{M}\mathbf{v})^{\diamond -1} = \mathbf{v}$$

As far as we are aware, it is not known how to characterize the \mathbf{M} for which there is a unique positive fixed point of this system. We have been particularly interested in \mathbf{M} that can be generated by our Set Cover and Packing heuristics, but there seem to be some \mathbf{M} that do not satisfy our assumptions which nevertheless have a unique positive solution.

For convenience, we define the function

$$F(\mathbf{v}) = (\mathbf{M}\mathbf{v})^{\diamond -1}$$

Any fixed point of F is a solution to our system.

We will use the following assumptions about our matrix \mathbf{M} here. The first two are sufficient to prove the existence of a positive fixed point of F, while existence and (c) enable us to show that the positive fixed point is unique. In what follows, we will refer to these from time to time.

(a) For all $1 \le i \le n$ and $1 \le j \le n$, $\mathbf{M}_{i,j} \ge 0$. Every matrix entry is non-negative.

- (b) For all $1 \le i \le n$ we have $\mathbf{M}_{i,i} > 0$. That is, every diagonal entry in the matrix is strictly positive.
- (c) **M** is positive semi-definite. For every vector $\mathbf{v} \in \mathbb{R}^n$, $\mathbf{v}^T \mathbf{M} \mathbf{v} \ge 0$.

5.1 Fixed Point Existence

To prove the existence of a fixed point, we will show that the function G has a zero.

$$G(\mathbf{d}) = \mathbf{M}(\mathbf{d}^{\diamond -1}) - \mathbf{d}$$

This will be accomplished by an application of the Poincaré-Miranda theorem (a generalization of the intermediate value theorem), as described in Idczak and Majewski[11]. A statement of this theorem is the following: Let $\mathcal{P} = [\alpha_1, \beta_1] \times \ldots \times [\alpha_n, \beta_n]$ with $\alpha_i \leq \beta_i$ set to fixed real constants for all $1 \leq i \leq n$, and $G : \mathbb{R}^n \to \mathbb{R}^n$ be a function that is continuous over \mathcal{P} . Then, if the following two statements hold,

 $G_i(\mathbf{d}) \ge 0$ for every $\mathbf{d} \in \mathcal{P}$ for which $\mathbf{d}_i = \alpha_i$ and $G_i(\mathbf{d}) \le 0$ for every $\mathbf{d} \in \mathcal{P}$ for which $\mathbf{d}_i = \beta_i$

there must be some $\mathbf{d} \in \mathcal{P}$ for which $G(\mathbf{d}) = \mathbf{0}$.

By definition, $G(\mathbf{d})$'s components for $1 \leq i \leq n$ are given by

$$G_i(\mathbf{d}) = \sum_{j=1}^n \frac{\mathbf{M}_{i,j}}{\mathbf{d}_j} - \mathbf{d}_i$$

Our function G's argument is inverted relative to F's. At the conclusion of the existence proof, we explain why a zero of G implies the existence of a fixed point of F.

We use G as defined above and now fix the α 's and β 's.

$$\alpha_i = \sqrt{\mathbf{M}_{i,i}}$$
$$\beta_i = \sum_{j=1}^n \frac{\mathbf{M}_{i,j}}{\alpha_j}$$

Note that fact (b) guarantees that the α 's are positive.

First, we demonstrate that $\alpha_i \leq \beta_i$, for all $1 \leq i \leq n$ as required.

$$\alpha_{i} = \sqrt{\mathbf{M}_{i,i}}$$

$$= \mathbf{M}_{i,i}\alpha_{i}^{-1}$$

$$\leq \mathbf{M}_{i,i}\alpha_{i}^{-1} + \sum_{j \neq i} \mathbf{M}_{i,j}\alpha_{j}^{-1}$$

$$= \sum_{j=1}^{n} \mathbf{M}_{i,j}\alpha_{j}^{-1}$$

$$= \beta_{i}$$

For the inequality, we use fact (a) and the fact that the α_i are strictly positive.

Also, we can see that G is continuous over \mathcal{P} . The only problem we might encounter is if \mathbf{d}_i is 0 for some *i*. Since $\alpha_i > 0$, \mathcal{P} has no such points and we are therefore quite safe.

There are two things left to do. First, assume, for arbitrary $1 \le i \le n$, that we have $\mathbf{d} \in \mathcal{P}$ with $\mathbf{d}_i = \alpha_i$. We must now show that this assumption guarantees that $G_i(\mathbf{d}) \ge 0$. Since $G_i(\mathbf{d}) = \sum_{j=1}^n \frac{\mathbf{M}_{i,j}}{\mathbf{d}_j} - \mathbf{d}_i$, it is sufficient to show that $\mathbf{d}_i \le \sum_{j=1}^n \frac{\mathbf{M}_{i,j}}{\mathbf{d}_j}$, which we derive below.

$$\begin{aligned} \mathbf{d}_{i} &= \alpha_{i} \\ &= \frac{\alpha_{i}^{2}}{\alpha_{i}} \\ &= \frac{\mathbf{M}_{i,i}}{\alpha_{i}} \\ &= \frac{\mathbf{M}_{i,i}}{\mathbf{d}_{i}} \\ &\leq \frac{\mathbf{M}_{i,i}}{\mathbf{d}_{i}} + \sum_{j \neq i} \frac{\mathbf{M}_{i,j}}{\mathbf{d}_{j}} \\ &= \sum_{j=1}^{n} \frac{\mathbf{M}_{i,j}}{\mathbf{d}_{j}} \end{aligned}$$

Second, assume, for arbitrary $1 \leq i \leq n$, that we have $\mathbf{d} \in \mathcal{P}$ with $\mathbf{d}_i = \beta_i$. We must now show that this assumption guarantees that $G_i(\mathbf{d}) \leq 0$. This is equivalent to $\mathbf{d}_i \geq \sum_{j=1}^n \frac{\mathbf{M}_{i,j}}{\mathbf{d}_j}$, which we will now show.

$$\mathbf{d}_{i} = \beta_{i}$$
$$= \sum_{j=1}^{n} \frac{\mathbf{M}_{i,j}}{\alpha_{j}}$$
$$\geq \sum_{j=1}^{n} \frac{\mathbf{M}_{i,j}}{\mathbf{d}_{j}}$$

These facts allow us to apply the Poincaré-Miranda theorem, establishing the existence of some vector $\mathbf{d} \in \mathcal{P}$ such that $G(\mathbf{d}) = \mathbf{0}$. For such a vector, we can define $\mathbf{v} = \mathbf{d}^{\diamond -1}$, and show that it is a fixed point of F.

$$F(\mathbf{v}) = (\mathbf{M}\mathbf{v})^{\diamond -1}$$

= $(\mathbf{M}\mathbf{d}^{\diamond -1})^{\diamond -1}$
= $(\mathbf{M}\mathbf{d}^{\diamond -1} - \mathbf{d} + \mathbf{d})^{\diamond -1}$
= $(G(\mathbf{d}) + \mathbf{d})^{\diamond -1}$
= $\mathbf{d}^{\diamond -1}$
= \mathbf{v}

So **v** is a fixed point of *F*. All of our assumptions were stated explicitly, so we now assert that for matrices **M** satisfying (a) and (b) there must be a positive vector **v** for which $(\mathbf{M}\mathbf{v})^{\diamond-1} = \mathbf{v}$.

Since there is a root of G in $[\alpha_1, \beta_1] \times \ldots \times [\alpha_n, \beta_n]$, there must be a fixed point of F in $[\beta_1^{-1}, \alpha_1^{-1}] \times \ldots \times [\beta_n^{-1}, \alpha_n^{-1}]$. Switching this to F-like vectors and simplifying, we have that there exists a fixed point **v** of F satisfying the following bounds:

$$F(diag(\mathbf{M})^{\diamond -1/2}) \le \mathbf{v} \le diag(\mathbf{M})^{\diamond -1/2}$$

5.2 Fixed Point Uniqueness

Assume that F has a fixed point. That is, there is some $\mathbf{v} \in \mathbb{R}^n_+$ with $F(\mathbf{v}) = (\mathbf{M}\mathbf{v})^{\diamond -1} = \mathbf{v}$. Because of this, $\mathbf{M}\mathbf{v} = \mathbf{v}^{\diamond -1}$, which we will use below. We will show that the additional requirement that \mathbf{M} is positive semi-definite is sufficient to guarantee that \mathbf{v} is the only positive fixed point of F.

Consider the function

$$K(\mathbf{u}) = (\mathbf{u} - \mathbf{v})^T (\mathbf{M}\mathbf{u} - \mathbf{u}^{\diamond -1})$$

This is a scalar-valued inner product of 2 vectors. We will show that for all $\mathbf{u} \in \mathbb{R}^n_+$ with $\mathbf{u} \neq \mathbf{v}$, $K(\mathbf{u}) > 0$. The existence of any fixed point $\mathbf{u} \in \mathbb{R}^n_+$ besides \mathbf{v} , then, is a contradiction, since any fixed point of F would have

$$K(\mathbf{u}) = (\mathbf{u} - \mathbf{v})^T (\mathbf{M}\mathbf{u} - \mathbf{u}^{\diamond -1})$$
$$= (\mathbf{u} - \mathbf{v})^T (\mathbf{u}^{\diamond -1} - \mathbf{u}^{\diamond -1})$$
$$= (\mathbf{u} - \mathbf{v})^T (\mathbf{0})$$
$$= 0$$

Let $\mathbf{u} \in \mathbb{R}^n_+$ with $\mathbf{u} \neq \mathbf{v}$. To establish that $K(\mathbf{u})$ must be positive, we will start by subtracting $(\mathbf{u} - \mathbf{v})^T \mathbf{M} (\mathbf{u} - \mathbf{v})$ and simplifying this. Note that this quantity is non-negative, since M is positive semi-definite.

$$\begin{split} K(\mathbf{u}) &= (\mathbf{u} - \mathbf{v})^T (\mathbf{M}\mathbf{u} - \mathbf{u}^{\diamond -1}) \\ &\geq (\mathbf{u} - \mathbf{v})^T (\mathbf{M}\mathbf{u} - \mathbf{u}^{\diamond -1}) - (\mathbf{u} - \mathbf{v})^T \mathbf{M} (\mathbf{u} - \mathbf{v}) \\ &= (\mathbf{u} - \mathbf{v})^T (\mathbf{M}\mathbf{u} - \mathbf{u}^{\diamond -1} - \mathbf{M} (\mathbf{u} - \mathbf{v})) \\ &= (\mathbf{u} - \mathbf{v})^T (\mathbf{M}\mathbf{u} - \mathbf{u}^{\diamond -1} - \mathbf{M}\mathbf{u} + \mathbf{M}\mathbf{v}) \\ &= (\mathbf{u} - \mathbf{v})^T (\mathbf{M}\mathbf{v} - \mathbf{u}^{\diamond -1}) \\ &= (\mathbf{u} - \mathbf{v})^T (\mathbf{v}^{\diamond -1} - \mathbf{u}^{\diamond -1}) \end{split}$$

Writing out this last expression as a sum, we have

$$\sum_{i=1}^{n} (\mathbf{u}_i - \mathbf{v}_i) \left(\frac{1}{\mathbf{v}_i} - \frac{1}{\mathbf{u}_i}\right) = \sum_{i=1}^{n} \frac{(\mathbf{u}_i - \mathbf{v}_i)^2}{\mathbf{u}_i \mathbf{v}_i}$$

Each term in this series is positive unless $\mathbf{u}_i = \mathbf{v}_i$, in which case it is 0. Since $\mathbf{u} \neq \mathbf{v}$, for some $i \mathbf{u}_i \neq \mathbf{v}_i$. Thus at least one term in the sum will be positive, and none are negative, so the entire sum must be positive, establishing our claim that $K(\mathbf{u})$ is positive.

Chapter 6

Numerical Matters

In order to calculate the consistent valuation for the new recursively defined valuations, we need to solve systems of the form $\mathbf{M}\mathbf{v} = \mathbf{v}^{\diamond -1}$ for some $\mathbf{v} \in \mathbb{R}^n_+$ given a matrix $\mathbf{M} \in \mathbb{R}^{n \times n}_{\geq 0}$. In this section we describe a technique that has worked well for us in practice, discuss the rate of convergence of our technique, address the time complexity of computing our new heuristic valuations and explore the prospects for solving these sorts of systems exactly.

6.1 Calculating Fixed Points In Practice

Let us consider a particular matrix for which we wish to calculate a vector \mathbf{v} for which $\mathbf{M}\mathbf{v} = \mathbf{v}^{\diamond -1}$,

A naïve way of attempting to calculate such a \mathbf{v} would be to start at some initial vector, possibly $v^{(0)} = (1, 1, 1, 1)^T$, and then iteratively calculate $\mathbf{v}^{(k+1)} = (\mathbf{M}\mathbf{v}^{(k)})^{\diamond -1}$ until $||\mathbf{v}^{(k-1)} - \mathbf{v}^{(k)}||$ is sufficiently small.

In practice, this does not work. Some values of $\mathbf{v}^{(k)}$ for different k are shown below:

$$\mathbf{v}^{(0)} = (1, 1, 1, 1)^{T}$$
$$\mathbf{v}^{(1)} = (0.125, 0.142857, 0.2, 0.166667)^{T}$$
$$\mathbf{v}^{(2)} = (0.837488, 0.95672, 1.19829, 1.07969)^{T}$$
$$\mathbf{v}^{(999)} = (0.128746, 0.146926, 0.188802, 0.167022)^{T}$$
$$\mathbf{v}^{(1000)} = (0.831298, 0.948681, 1.21907, 1.07844)^{T}$$
$$\mathbf{v}^{(9999)} = (0.128746, 0.146926, 0.188802, 0.167022)^{T}$$
$$\mathbf{v}^{(10000)} = (0.831298, 0.948681, 1.21907, 1.07844)^{T}$$

The vector settles into a repeating cycle of length 2, getting us no closer to a fixed point.

We have found the following iteration to be very effective in finding fixed points. Let $\mathbf{v}^{(0)} \in \mathbb{R}^n_+$ be some positive starting vector. For $k \ge 1$, define

$$\mathbf{v}^{(k+1)} = \frac{\mathbf{v}^{(k)} + (\mathbf{M}\mathbf{v}^{(k)})^{\diamond - 1}}{2}$$

Every iteration is the arithmetic mean of the previous point and the naïve iteration's next point. This iteration appears to have straightforward linear convergence, as can be seen in figure 6.1, showing the distance of successive values of $\mathbf{v}^{(k)}$ from the true fixed point for this problem, $\mathbf{v} \approx (0.327149, 0.373344, 0.479752, 0.424410)^T$. Note that at the right side, convergence halts because we have obtained the fixed point to floating-point double-precision. The value that we are using for the *true fixed point* is a vector computed to agree with the fixed point up to 80 decimal places for each component.



Figure 6.1: Convergence of $\mathbf{v}^{(k)}$ towards the true fixed point \mathbf{v} starting with $\mathbf{v}^{(k)} = \mathbf{1}$ and using a weighted iteration.

We have found that this technique consistently achieves a linear rate of convergence, regardless of the initial point used, so long as it is strictly positive. We have found no non-negative \mathbf{M} with strictly positive diagonals for which this iteration does not converge, though we have no proof that it must converge in all such cases.

Some simple variations of this iteration are possible and can have a significant impact on the rate of convergence. Choosing $\mathbf{v}^{(k+1)}$ to be any weighted mean of $\mathbf{v}^{(k)}$ and $(\mathbf{M}\mathbf{v}^{(k)})^{\diamond-1}$ appears to be satisfactory for convergence to occur. In practice, we have found that using the arithmetic mean and putting half as much weight on $\mathbf{v}^{(k)}$ as the new term causes the series to converge fairly quickly relative to alternatives. We have also found that using the starting point with components $\mathbf{v}_i^{(0)} = (\sum_{j=1}^n \mathbf{M}_{i,j})^{-1/2}$, the inverse square roof of the row sum, works particularly well. This can also be written $\mathbf{v}^{(0)} = (\mathbf{M}\mathbf{1})^{\diamond-1/2}$.

Thus, our final recommendation for calculating solutions to these systems is the following iteration:

$$\mathbf{v}^{(0)} = (\mathbf{M}\mathbf{1})^{\diamond -1/2}$$
$$\mathbf{v}^{(k+1)} = \frac{\mathbf{v}^{(k)} + 2(\mathbf{M}\mathbf{v}^{(k)})^{\diamond -1}}{3}$$

Using this scheme we obtain the fixed point to floating-point double-precision in nearly 20 fewer iterations, as indicated by figure 6.2.



Figure 6.2: Convergence of $\mathbf{v}^{(k)}$ for our final proposed iteration.

This convergence behavior appears to be universal and is typical even for very large matrices \mathbf{M} , as suggested by figure 6.3. For floating-point double-precision, 30 iterations appear to suffice. It seems we can obtain the fixed point to any required precision in a number of iterations that is constant with respect to the instance size. The time required to perform one iteration is dominated by the time taken to perform the matrix multiplication. This fact makes our valuation technique competitive with the standard greedy algorithm, as we discuss further in 6.3.1.

6.1.1 An Alternate Iteration

We have found an alternative iteration scheme that is significantly different from the one presented above. It proceeds from the simple idea of asking how should \mathbf{v}_i be set if all other components of \mathbf{v} are



Figure 6.3: Convergence of $\mathbf{v}^{(k)}$ for our final proposed iteration for a matrix **M** of size 10000×10000 with entries selected uniformly from the reals in the interval (0, 5)

correct. We want $\mathbf{v}_i = \frac{1}{\sum_{j=1}^n \mathbf{M}_{i,j} \mathbf{v}_j}$. Writing this as a 2nd degree polynomial in \mathbf{v}_i , we have

$$\mathbf{v}_i^2 + \frac{\left(\sum_{j \neq i} \mathbf{M}_{i,j} \mathbf{v}_j\right)}{\mathbf{M}_{i,i}} \mathbf{v}_i - \frac{1}{\mathbf{M}_{i,i}} = 0$$

Defining $\mathbf{s}_i = \frac{1}{\mathbf{M}_{i,i}} \sum_{j \neq i} \mathbf{M}_{i,j} \mathbf{v}_j$, the unique positive solution to this equation is

$$\mathbf{v}_i = \frac{1}{2} \left(\sqrt{\mathbf{s}_i^2 + 4/\mathbf{M}_{i,i}} - \mathbf{s}_i \right)$$

This leads us to propose the following iteration scheme, starting at any positive $\mathbf{v}^{(0)}$. Let $\mathbf{s}^{(k+1)}$ be the vector with components $\mathbf{s}_i^{(k+1)} = \frac{1}{\mathbf{M}_{i,i}} \sum_{j \neq i} \mathbf{M}_{i,j} \mathbf{v}_j^{(k)}$. This can also be found by computing $\mathbf{M}\mathbf{v}^{(k)}$, dividing each component by the appropriate diagonal entry of \mathbf{M} and then subtracting $\mathbf{v}^{(k)}$.

Then define the new iteration $\mathbf{v}^{(k+1)}$ to be the vector with components

$$\mathbf{v}_{i}^{(k+1)} = \frac{1}{2} \left(\sqrt{\left(\mathbf{s}_{i}^{(k)}\right)^{2} + 4/\mathbf{M}_{i,i}} - \mathbf{s}_{i}^{(k)} \right)$$

We have found that this iteration converges to a fixed point regardless of the initial point $\mathbf{v}^{(0)}$ used, although only slowly for large instances. When combined with the previous iteration in a weighted average, as is done with our previous iteration, we find that this converges to a fixed point similarly quickly. Additionally, we have had good results with updating the valuation vector one component at a time, as we describe in section 6.2.1.

6.2 Additional Shortcuts

We have found that we can speed this process up even more with a few optimizations. Since we are only calculating the full valuation in order to see which of its components is largest, very high precision is often not required and we can terminate the iteration well before reaching machine precision. At the fixed point for the random instances we have examined, we found that the difference between the valuations of the highest valuation set and the second-highest valuation set is typically around 1%. Assuming that the iteration converges as straightforwardly as suggested in figure 6.3, we should be safe halting the iteration when $||\mathbf{v}(k) - \mathbf{v}(k-1)||_1$ is less that the difference between the largest and second-largest components of $\mathbf{v}^{(k)}$.

Note that for the new cover heuristic, this is less straightforward. The valuations are given by $\mathbf{v} = \mathbf{C}^{-1} \mathbf{A}^T \mathbf{d}$ where \mathbf{d} is the fixed point we seek. Terminating early here should be done with the knowledge that the maximum of \mathbf{v} will not change in later iterations. Conveniently, it is possible to straightforwardly examine what the valuation would be at each iteration as an intermediate result of the matrix product needed to perform the next iteration. That is, given $\mathbf{d}^{(k)}$ we can consider $\mathbf{v}^{(k)} = \mathbf{C}^{-1} \mathbf{A}^T \mathbf{d}^{(k)}$ and continue on to calculate $\mathbf{d}^{(k+1)} = (\mathbf{A} \mathbf{C}^{\gamma} \mathbf{v}^{(k)})^{\diamond -1}$.

Another shortcut we have found is reusing the valuations or difficulties found in the previous greedy iteration as a starting point for the next fixed point iteration. After the greedy scheme selects one set to include in the growing cover or packing, the instance is modified and preprocessed again, but the new system is still fairly similar to the previous one, making the use of the previous fixed point vector a better starting point than the one defined in section 6.1.

Additionally, we find that the performance of the overall greedy algorithm does not suffer significantly when the maximum number of iterations is fixed at even very small numbers. In random instances we have found that, after an average of around two iterations the same set is selected as would be after any number of subsequent iterations. The figure of two is slightly misleading, however, because the variance of the number of iterations before the correct choice is made is fairly large. Despite this, fixing a maximum number of iterations can be an effective way of controlling the running time of the overall algorithm without hurting its performance too badly.

6.2.1 Using the Alternate Iteration for the Packing Heuristic

We have found the alternate iteration technique to be particularly useful in obtaining fixed points for the packing heuristic. Updating each component independently, we find that we reach fixed-point double precision in around 15 iterations. The form of the iterations is particularly clean for this heuristic, and we present pseudocode for the entire iteration procedure below.

 $\begin{array}{ll} \text{for all } i \in \mathcal{I} \text{ do} \\ M_i \leftarrow \{j \in \mathcal{I} \mid S_i \cap S_j \neq \emptyset\} - \{i\} \\ \mathbf{v}_i \leftarrow (|M_i|/\mathbf{c}_i)^{-1/2} \\ \text{end for} \\ \text{for each iteration do} \\ \text{for all } i \in \mathcal{I} \text{ do} \end{array} \qquad \triangleright \text{ Generate sparse version of } \mathbf{M} \\ \end{array}$

$$\begin{array}{l} s \leftarrow \sum_{j \in M_i} \mathbf{v}_j \\ \mathbf{v}_i \leftarrow \frac{1}{2} (\sqrt{s^2 + 4\mathbf{c}_i} - s) \\ \text{end for} \\ \text{end for} \\ \text{return } \mathbf{v} \end{array} \triangleright \text{Update the valuation for set } i \end{array}$$

The convergence of this approach can be seen in figure 6.4.



Figure 6.4: Convergence of $\mathbf{v}^{(k)}$ towards the true fixed point for a random 5000 × 5000 matrix.

6.3 Running Time of the New Heuristics

6.3.1 Running Time of the New Set Cover Heuristic

Let **A** be the fundamental matrix and **C** be the diagonal matrix of costs for a Set Cover instance. For many problems of practical interest and for the random instances we explore in our experiments, **A** is often quite sparse. Let ρ be the *density* of a Set Cover instance calculated as $\rho = \frac{\sum_{i,j} \mathbf{A}_{i,j}}{mn}$. In practice we use a sparse representation of **A**, equivalent to keeping track of every input set and the neighbourhood of every element. This enables us to perform matrix products $\mathbf{A}\mathbf{v}$ and $\mathbf{A}^T\mathbf{d}$ for vectors **v** and **d** in time $\mathcal{O}(mn\rho)$. Similar products using **C** only take time $\mathcal{O}(m)$ since **C** is diagonal.

Let $\mathbf{M} = \mathbf{A}\mathbf{C}^{\gamma-1}\mathbf{A}^T$. In order to calculate the new cover heuristic, we need to find $\mathbf{d} \in \mathbb{R}^n_+$ for which $\mathbf{d} = (\mathbf{M}\mathbf{d})^{\diamond-1}$. and then calculate the valuations $\mathbf{v} = \mathbf{C}^{-1}\mathbf{A}^T\mathbf{d}$. The iterative fixed point approximation technique described in section 6.1 enables us to rapidly approximate the required \mathbf{d} , obtaining it to floating-point double-precision in around 30 iterations. The running time of each iteration is dominated by the time taken to perform a matrix multiplication by \mathbf{M} . With our sparse representations of \mathbf{A} and \mathbf{C} , this is again time $\mathcal{O}(mn\rho)$. This means that the time taken for one iteration of our greedy Set Cover algorithm with the new heuristic only takes a constant multiple of the time taken by the standard greedy algorithm.

6.3.2 Running Time of the New Set Packing Heuristic

The cost of approximating fixed points for our new packing heuristic is significantly worse than it is for cover. For a Set Packing instance given by **A** and **C**, we are interested in the matrix $\mathbf{M} = \mathbf{C}bin(\mathbf{A}^T\mathbf{A})$. In order to perform matrix multiplications by **M**, we are not helped substantially by the sparse data structures we maintain and our ability to do matrix multiplication by **A** and **A**^T efficiently.

In practice, to multiply a vector \mathbf{v} by \mathbf{M} , we find each row of \mathbf{M} independently and use each one to calculate a single component of $\mathbf{M}\mathbf{v}$. Using the notation of input sets S_i and element neighbourhoods N_e , we obtain $(\mathbf{M}\mathbf{v})_i$ by first finding $R_i = \bigcup_{e \in S_i} N_e$ and then calculating $(\mathbf{M}\mathbf{v})_i = \mathbf{c}_i \left(\sum_{j \in R_i} \mathbf{v}_j\right)^{-1}$. To estimate the running time of this technique, note that we need to compute the union of around $m\rho$ neighbourhoods, each of size around $n\rho$ for each of the m input sets. Thus we expect the running time for one multiplication by \mathbf{M} to be in $\mathcal{O}(m^2n\rho^2)$, a factor of $m\rho$ larger than the time required for matrix multiplication in the case of the cover heuristic. This can make finding approximate packings with the new heuristic less practical than finding approximate covers. Note that the MIS heuristic has a similar issue, needing $\mathcal{O}(m^2n\rho^2)$ time to calculate, and it runs in around 1/30th of the time needed for the new heuristic.

6.4 Exact Calculation of Fixed Points

The principal reason that we have so few solid mathematical results here is that we find it difficult to characterize the solutions to $\mathbf{M}\mathbf{v} = \mathbf{v}^{\diamond -1}$. In a few cases, however, we can write out the solution explicitly:

- 1. When **M** is a diagonal matrix with all diagonal entries positive, there is a unique positive fixed point **v** whose components are given by $\mathbf{v}_i = \mathbf{M}_{i,i}^{-\frac{1}{2}}$. In other words, $\mathbf{v} = diag(\mathbf{M})^{\diamond -1/2}$.
- 2. When **M** is non-negative with all non-zero diagonal entries and **M**'s rows all have the same sum, there is a positive fixed point with all components equal to $s^{-\frac{1}{2}}$ where s is the sum of any row of **M**. Equivalently, there is a fixed point at $(\mathbf{M1})^{\diamond-1/2}$.
- 3. When **M** is positive with all entries on the same row being equal, every column of **M** is the same. Let **u** be one column of **M**, and let $s = \sqrt{\sum_{j=1}^{n} \frac{1}{\mathbf{u}_j}}$, the square root of the sum of the reciprocals of each row's entry. There is a positive fixed point at $\mathbf{v} = (s\mathbf{u})^{\diamond -1}$
- 4. In some other situations we can find fixed points exactly by calculating them numerically and then using an inverse symbolic calculator to find a representation of the number. For instance, when

$$\mathbf{M} = \left(\begin{array}{rrr} 1 & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{array} \right)$$

there must be a unique $\mathbf{v} \in \mathbb{R}^3_+$ satisfying $\mathbf{M}\mathbf{v} = \mathbf{v}^{\diamond -1}$. Numerically, it is given by $(0.48587, 0.78615, 0.78615)^T$.

It turns out that the exact value of \mathbf{v} is

$$\left(\begin{array}{c} \sqrt{\sqrt{5}-2} \\ \sqrt{\frac{1}{2}(\sqrt{5}-1)} \\ \sqrt{\frac{1}{2}(\sqrt{5}-1)} \end{array}\right)$$

In the general case we do not expect that fixed points can be described by any simple closed form. Without a deeper understanding of these fixed points, we have had to rely on experimental results to provide the force of our overall argument for the value of the new heuristics. There is much room for progress in understanding this problem.

6.5**Fixed Points For Broader Classes of Matrices**

The assumptions we make about matrices \mathbf{M} in order to prove the existence and uniqueness of fixed points are these:

- (a) No entries of **M** are negative.
- (b) All diagonal entries of **M** are positive.
- (c) **M** is positive semi-definite.

If we relax constraint (a), we can still find matrices with positive fixed points. For instance

$$\left(\begin{array}{rrrr} 3 & 1 & -1 \\ 1 & 3 & 1 \\ 0 & 1 & 3 \end{array}\right)$$

appears to have a fixed point at approximately $(0.59242, 0.42200, 0.51129)^T$.

For a similar matrix $\begin{pmatrix} 3 & -1 & -1 \\ -1 & 3 & -1 \\ 0 & -1 & 3 \end{pmatrix}$, however, we are unable to find any fixed points with all

positive components.

Relaxing constraint (b), we can consider the matrix $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. This has fixed points at $(n, \frac{1}{n})^T$ for all $n \in \mathbb{R}_+$. There is a positive fixed point, but it is not unique. If we look at a matrix that is only a slight perturbation of this one, $\begin{pmatrix} 0 & 1+\epsilon\\ 1 & 0 \end{pmatrix}$ for any small $\epsilon \in \mathbb{R}_+$, we find it has no positive fixed points at all.

This observation is what has motivated us, in the new packing heuristic's definition, to regard each set as a neighbour of itself. Without this, it is possible to formulate Set Packing instances for which the new packing heuristic has no consistent valuations. We have also considered a class of heuristics where we obtain valuations for the definition

$$\mathbf{v} = \mathbf{C}((bin(\mathbf{A}^T\mathbf{A}) - (1-\epsilon)\mathbf{I})\mathbf{v})^{\diamond - 1}$$

for arbitrarily small positive epsilon. Matrices of this form satisfy the preconditions for our fixed point existence proof, and the iteration described in section 6.1 does tend to find a fixed point. Using these fixed points for valuations in the greedy scheme appears to generate packings with high quality, but they are not obviously superior to the packings found by the new packing heuristic as presented in section 4.3. This looks to us like a potentially valuable direction for further investigation.

We have not found any explicit matrix \mathbf{M} that satisfies (a) and (b) but not (c) and is known to have more than one positive fixed point. The matrix \mathbf{M} used in our new packing heuristic are generally of this type. We believe it is possible that our new packing heuristic generates unique valuations, but do not have much confidence in either possible resolution. In the event that our packing heuristic produces unique valuations, we would regard it as more natural, but in either case the results of section 7.4 indicate that our heuristic produces high quality packings relative to alternative greedy packing heuristics.

Chapter 7

Experimental Results

7.1 Random Instances

In order to test the heuristics we have described, we need to run the greedy algorithm on particular instances. To this end, we define a distribution of random instances. Let $\mathcal{D}(m, n, \rho, C)$ where $m \in \mathbb{N}$ represents the number of sets, $n \in \mathbb{N}$ the number of elements, $\rho \in (0, 1)$ the *density* of the instance and C the distribution of the set costs. We use $\mathcal{D}(m, n, \rho, C)$ to represent the distribution of instances made by the following process:

- 1. Fix $\mathcal{I} = \{1, \dots, m\}$
- 2. For each $i \in \mathcal{I}$, set c_i to a sample drawn uniformly at random from C.
- 3. For each $i \in \mathcal{I}$ set S_i , to a subset of $\{1, \ldots, n\}$ with each element selected independently with probability ρ .
- 4. For each element $e \in \{1, ..., n\}$, if e is in fewer than 2 input sets, add e to S_i for $i \in \{1, ..., m\}$ selected uniformly at random until e is in 2 input sets.

The cost distributions that we will consider are the following:

- 1. *unweighted*. All costs are set to 1. In order to save space in tables of results, we will sometimes write *u* instead of *unweighted*.
- 2. discrete(a, b) for $a, b \in \mathbb{N}$ with 0 < a < b. All costs are selected with uniformly probability from the set $\{a, \ldots, b\}$. We will sometimes write d(a, b) for this distribution.
- 3. continuous(a, b). All costs are selected with uniformly probability from the real interval (a, b). We occasionally use c(a, b) to denote this.

We are interested in instances with these different distributions mainly because the standard heuristic is highly prone to ties in the unweighted case and only somewhat less so in the discrete cost setting. When the costs are random real numbers, the standard greedy algorithm is deterministic for all intents and purposes. When the standard valuation gives ties for some elements, we select the set to include at random. We do the same for the new heuristics, but they obtain ties far less frequently, even for unweighted instances. In order to provide a fair comparison between standard greedy and the new algorithm, in many cases we give both algorithms approximately equal time by doing multiple independent runs of the standard greedy algorithm and using the best cover/packing that it finds. The ability to break ties in different ways is the only advantage obtained by running the standard algorithms multiple times. This permits them to effectively sample from the space of solutions that they could potentially return. The new algorithms would not usually obtain any benefit from multiple runs, since they produce ties so rarely.

7.2 Algorithms Used

In all of the experiments below, we use only basic preprocessing. We have found that the more involved preprocessing techniques are not clearly of value, so we do not consider them here. Further work is needed to evaluate their usefulness.

For all of the valuation techniques discussed here, we consider it a tie if the valuations of input sets differ by less than 10^{-7} . In the event of a tie, we select an input set at random from the sets with maximum valuation.

7.2.1 Set Cover Algorithms

For SCP, the algorithms we are comparing all use the general greedy scheme as described in 2.2.4 so they differ only in how they compute the valuations of the sets at each step. The following Cover valuations are considered:

- 1. The standard heuristic (STD). $\mathbf{v} = \mathbf{C}^{-1}\mathbf{A}^T\mathbf{1}$. For some tests, we run the standard heuristic many times. We denote the best cover found with k independent runs of the standard algorithm by STD_k .
- 2. The new heuristic $(NEWC(\gamma))$, with parameter γ .
 - $\mathbf{v} = \mathbf{C}^{-1} \mathbf{A}^T \mathbf{d}$ for \mathbf{d} such that $\mathbf{d} = (\mathbf{A} \mathbf{C}^{\gamma 1} \mathbf{A}^T \mathbf{d})^{\diamond 1}$

Multiple input sets obtaining the same valuation is very rare, relative to the standard heuristic, so we only ever use a single run of the new algorithm.

We have omitted many other possible algorithms. The main reason for this is that other simple algorithms (e.g. Primal/Dual, LP Rounding) do not appear to be competitive with the standard greedy algorithm, as can be seen from Gomes et al.'s experimental work in [7].

Except when otherwise noted, we always minimize returned covers by the Wool and Grossman technique described in section 3.5. In practice, this is far more beneficial for the standard algorithm than for the new algorithm.

7.2.2 Set Packing Algorithms

For Set Packing, we consider the following valuation techniques:

- 1. A variation of the standard Set Cover heuristic (STDP), where we pick the set with greatest weight per element. The valuations are given by $\mathbf{v} = \mathbf{C}(\mathbf{A}^T \mathbf{1})^{\diamond -1}$. When run multiple times, we indicate this by $STDP_k$ for k independent runs.
- 2. The heuristic valuing sets by their weight divided by the square root of their size, which we call the (*ROOT*). $\mathbf{v} = \mathbf{C}(\mathbf{A}^T \mathbf{1})^{\diamond -\frac{1}{2}}$. Multiple runs are denoted by $ROOT_k$.
- 3. The standard MIS heuristic (*MIS*). $\mathbf{v} = \mathbf{C}((bin(\mathbf{A}^T\mathbf{A}) \mathbf{I})\mathbf{1})^{\diamond -1}$. Multiple runs are denoted by $STDP_k$.
- 4. The new heuristic (*NEWP*). We choose $\mathbf{v} \in \mathbb{R}^m_+$ such that $\mathbf{v} = \mathbf{C}(bin(\mathbf{A}^T\mathbf{A})\mathbf{v})^{\diamond-1}$.

7.3 Set Cover Results

7.3.1 Varying γ for the New Heuristic

In order to justify the choice of $\gamma = -3$ for the remaining tests, we have tested a variety of different values of γ on a variety of different random problem distributions. For each distribution, we generate 100 instances and solve them with each heuristic. We have compared the new algorithm for γ ranging between 0 and -4. The results are presented in tables 7.1 and 7.2. Unweighted distributions are not tested, because all values of γ yield the same valuations when all costs are identical.

In table 7.1 we show the proportion of the 100 instances that the new heuristic with the stated choice of γ performed best on. Table 7.2 shows the *quality* of the obtained solutions. It is calculated as the average, over the 100 instances, of the cost obtained by each algorithm divided by the cost of the best solution found by any of the algorithm runs. It is effectively a proxy for approximation ratio, which we cannot compute because the instances are too large to be solved exactly.

It can be readily seen that $\gamma = -3$ is best among the algorithms both in terms of finding the smallest cover found most frequently and in terms of quality for a majority of distributions. It is interesting, however, that $\gamma = -3$ does not dominate any of the alternatives. It seems reasonable that running the new heuristic with different values of γ can be used to obtain different solutions, enabling us to utilize multiple runs of the new algorithm in the same way that its tendency to tie makes multiple runs of the standard heuristic valuable.

It is interesting that for most of the distributions, there are few instances for which the best cover is found by more than one of these algorithms. With $\mathcal{D}(5000, 500, 0.02, continuous(1, 50))$, for instance, the best cover was found by exactly one of our settings for γ . The fairly wide spread between heuristics finding the best solution also suggests that multiple runs with different γ can be valuable in practice.

The setting $\gamma = -3$ appears to be the best overall, but it should be noted that there is no reason that γ must be an integer. Arbitrary real values of γ yield alternative versions of the new cover heuristic, and

Distribution	$\gamma = 0$	$\gamma = -1$	$\gamma = -2$	$\gamma = -3$	$\gamma = -4$
$\mathcal{D}(500, 50, 0.05, discrete(1, 50))$	22%	27%	75%	66%	57%
$\mathcal{D}(1000, 200, 0.02, discrete(1, 50))$	6%	9%	43%	50%	27%
$\mathcal{D}(1000, 200, 0.05, discrete(1, 50))$	9%	9%	37%	57%	49%
$\mathcal{D}(2000, 200, 0.02, discrete(1, 50))$	2%	12%	56%	47%	29%
$\mathcal{D}(3000, 300, 0.02, discrete(1, 50))$	4%	7%	42%	59%	35%
$\mathcal{D}(5000, 500, 0.02, discrete(1, 50))$	3%	5%	35%	52%	33%
$\mathcal{D}(5000, 500, 0.05, discrete(1, 50))$	10%	21%	29%	52%	55%
$\mathcal{D}(5000, 500, 0.1, discrete(1, 50))$	29%	54%	57%	63%	69%
$\mathcal{D}(10000, 1000, 0.02, discrete(1, 50))$	3%	6%	27%	50%	49%
$\mathcal{D}(500, 50, 0.05, continuous(1, 50))$	10%	17%	45%	42%	31%
$\mathcal{D}(1000, 200, 0.02, continuous(1, 50))$	1%	9%	43%	31%	19%
$\mathcal{D}(1000, 200, 0.05, continuous(1, 50))$	8%	14%	28%	26%	26%
$\mathcal{D}(2000, 200, 0.02, continuous(1, 50))$	2%	6%	36%	40%	18%
$\mathcal{D}(3000, 300, 0.02, continuous(1, 50))$	1%	3%	36%	40%	20%
$\mathcal{D}(5000, 500, 0.02, continuous(1, 50))$	5%	6%	29%	39%	21%
$\mathcal{D}(5000, 500, 0.05, continuous(1, 50))$	6%	8%	16%	39%	31%
$\mathcal{D}(5000, 500, 0.1, continuous(1, 50))$	10%	16%	15%	36%	34%
$\mathcal{D}(10000, 1000, 0.02, continuous(1, 50))$	1%	6%	21%	40%	32%
Average over all instances	7.33%	13.06%	37.22%	46.06%	35.28%

Table 7.1: Comparing the effectiveness of different values for γ for the new cover heuristic. Each cell indicated what proportion of 100 instances drawn from that row's distribution that column's heuristic performed best on.

it may well be that some other value between -2 and -4 performs better than -3. Further experiments in this direction may be valuable.

7.3.2 Comparison Between the Standard and New Heuristics

In order to compare the standard and new heuristics as fairly as possible, we have to examine a few different situations. We run both algorithms on a variety of different distributions, generating 100 problems for every row of the tables below. For each problem, we run the standard heuristic 50 times, taking the best solution it obtains, and the new heuristic once. We track the cost of the best solution obtained before minimization and also the cost of the best solution after minimization.

Table 7.3 summarizes our results for the 2 algorithms, $STDC_{50}$ and NEWC(-3), considering the quality of the solutions obtained for the random problems. Both algorithms are run using only the basic preprocessing steps described in section 3.1. The columns labelled $STDC_{50}$ and NEWC(-3) show the percentage of the 100 problems for which each algorithm obtained the best solution. It is possible for the sum of these values to exceed 100% if for some of the instances, both algorithms return a set with the same cost. Under the columns labelled *Quality*, we calculate the average ratio relative to the best solution found (by either algorithm *after* minimization) which we use as a proxy for approximation ratio, since many of these problems cannot be solved exactly within a reasonable period of time. Let s_i be the cost of the best of 50 runs of STDC on instance number i and t_i be the cost of the set returned by NEWC(-3). Then the " $STDC_{50}$ Quality" column contains the value $\frac{1}{100} \sum_{i=1}^{100} \frac{s_i}{\min(s_i,t_i)}$. In every row the cell for the algorithm that

nces drawn from that row's distribution.									
Distribution	$\gamma = 0 \ Q$	$\gamma = -1 \ Q$	$\gamma = -2 \text{ Q}$	$\gamma = -3 \ \mathrm{Q}$	$\gamma = -4 \text{ Q}$				
$\mathcal{D}(500, 50, 0.05, d(1, 50))$	1.0383	1.0282	1.0078	1.0082	1.0171				
$\mathcal{D}(1000, 200, 0.02, d(1, 50))$	1.0337	1.0253	1.0081	1.0068	1.0123				
$\mathcal{D}(1000, 200, 0.05, d(1, 50))$	1.0447	1.0367	1.0178	1.0116	1.0152				
$\mathcal{D}(2000, 200, 0.02, d(1, 50))$	1.0384	1.0252	1.0079	1.0077	1.0149				
$\mathcal{D}(3000, 300, 0.02, d(1, 50))$	1.0374	1.0287	1.0107	1.0056	1.0117				
$\mathcal{D}(5000, 500, 0.02, d(1, 50))$	1.0407	1.0298	1.0132	1.0079	1.0119				
$\mathcal{D}(5000, 500, 0.05, d(1, 50))$	1.0489	1.0337	1.0243	1.0144	1.0133				
$\mathcal{D}(5000, 500, 0.1, d(1, 50))$	1.0436	1.0251	1.0222	1.0164	1.0133				
$\mathcal{D}(10000, 1000, 0.02, d(1, 50))$	1.0463	1.0309	1.0173	1.0091	1.0082				
$\mathcal{D}(500, 50, 0.05, c(1, 50))$	1.0428	1.0286	1.0132	1.0098	1.0142				
$\mathcal{D}(1000, 200, 0.02, c(1, 50))$	1.0337	1.0237	1.0057	1.0080	1.0149				
$\mathcal{D}(1000, 200, 0.05, c(1, 50))$	1.0442	1.0288	1.0186	1.0158	1.0182				
$\mathcal{D}(2000, 200, 0.02, c(1, 50))$	1.0420	1.0319	1.0099	1.0092	1.0154				
$\mathcal{D}(3000, 300, 0.02, c(1, 50))$	1.0432	1.0317	1.0088	1.0094	1.0153				
$\mathcal{D}(5000, 500, 0.02, c(1, 50))$	1.0420	1.0327	1.0141	1.0077	1.0136				
$\mathcal{D}(5000, 500, 0.05, c(1, 50))$	1.0522	1.0372	1.0217	1.0148	1.0167				
$\mathcal{D}(5000, 500, 0.1, c(1, 50))$	1.0472	1.0347	1.0308	1.0200	1.0173				
$\mathcal{D}(10000, 1000, 0.02, c(1, 50))$	1.0430	1.0290	1.0145	1.0083	1.0086				
Average over all instances	1.0424	1.0301	1.0148	1.0106	1.0140				

Table 7.2: Comparing the effectiveness of different values for γ for the new cover heuristic. Each cell indicated the average performance of the new heuristic for the stated column's value of γ over 100 instances drawn from that row's distribution.

performs best on the highest proportion of instances and the cell for the algorithm with best quality are highlighted.

It can be seen that on the majority of weighted problems the new heuristic generally performs better regardless of instance size, though the standard heuristic performs better on some instances. The reason the standard algorithm works well on unweighted instances is that they have more situations where the standard valuation gives ties, allowing the 50 runs allocated to the standard algorithm to explore a variety of the possible solutions accessible to it. For the problems with continuous cost distributions, the standard heuristic is effectively deterministic, since the likelihood of ties occurring is very low.

In table 7.4 we show the results after minimizing all returned covers. The new heuristic fares only a little more poorly, though uniformly so. In table 7.5 we show the difference in quality between the pre-minimized and the minimized solutions. Minimizing the covers returned by the new heuristic does not substantially reduce their sizes, improving them by around 0.1% on average. It is effectively built into the new heuristic to avoid making selections that will later be made wholly redundant. This is not so for the standard heuristic. It is very common that the best minimized solution returned by the standard heuristic is significantly better than the best non-minimized solution. For this reason, we believe that the new heuristic can be valuable for situations in which the minimization step is not possible, as in some formulations of online Set Cover or Hitting Set problems. One scheme that we suspect the new heuristic is particularly well suited for is model M2 described in [1].

Overall, we can see that the new heuristic is usually superior for instances where the set costs are drawn from a continuous distribution, but the results are mixed for distributions where the costs are

Table 7.3: Comparison between the standard and new set cover heuristics for a range of random instance distributions. All results reported before the returned covers are minimized.

Distribution	$STDC_{50}$	$STDC_{50}$ Q	NEWC(-3)	NEWC(-3) Q
$\mathcal{D}(1000, 200, 0.02, unweighted)$	74%	1.0085	62%	1.0138
$\mathcal{D}(1000, 200, 0.05, unweighted)$	99%	1.0015	53%	1.0244
$\mathcal{D}(2000, 200, 0.02, unweighted)$	91%	1.0033	52%	1.0170
$\mathcal{D}(3000, 300, 0.02, unweighted)$	97%	1.0013	42%	1.0186
$\mathcal{D}(5000, 500, 0.02, unweighted)$	95%	1.0010	36%	1.0157
$\mathcal{D}(5000, 500, 0.05, unweighted)$	99%	1.0007	42%	1.0239
$\mathcal{D}(5000, 500, 0.1, unweighted)$	100%	1.0000	46%	1.0318
$\mathcal{D}(10000, 1000, 0.02, unweighted)$	99%	1.0002	28%	1.0153
$\mathcal{D}(1000, 200, 0.02, discrete(1, 50))$	0%	1.0749	100%	1.0025
$\mathcal{D}(1000, 200, 0.05, discrete(1, 50))$	2%	1.0734	98%	1.0082
$\mathcal{D}(2000, 200, 0.02, discrete(1, 50))$	0%	1.0711	100%	1.0047
$\mathcal{D}(3000, 300, 0.02, discrete(1, 50))$	0%	1.0708	100%	1.0042
$\mathcal{D}(5000, 500, 0.02, discrete(1, 50))$	1%	1.0646	100%	1.0044
$\mathcal{D}(5000, 500, 0.05, discrete(1, 50))$	39%	1.0363	82%	1.0186
$\mathcal{D}(5000, 500, 0.1, discrete(1, 50))$	83%	1.0147	70%	1.0196
$\mathcal{D}(10000, 1000, 0.02, discrete(1, 50))$	1%	1.0481	100%	1.0066
$\mathcal{D}(1000, 200, 0.02, continuous(1, 50))$	0%	1.1000	100%	1.0015
$\mathcal{D}(1000, 200, 0.05, continuous(1, 50))$	0%	1.1036	100%	1.0026
$\mathcal{D}(2000, 200, 0.02, continuous(1, 50))$	0%	1.0868	100%	1.0003
$\mathcal{D}(3000, 300, 0.02, continuous(1, 50))$	0%	1.0911	100%	1.0012
$\mathcal{D}(5000, 500, 0.02, continuous(1, 50))$	0%	1.0833	100%	1.0024
$\mathcal{D}(5000, 500, 0.05, continuous(1, 50))$	2%	1.0652	98%	1.0048
$\mathcal{D}(5000, 500, 0.1, continuous(1, 50))$	15%	1.0448	85%	1.0065
$\mathcal{D}(10000, 1000, 0.02, continuous(1, 50))$	1%	1.0679	99%	1.0030
Average over all instances	37.4%	1.0464	78.9%	1.0105

Table 7.4: Comparison between the standard and new set cover heuristics for a range of random instance distributions. All results reported after the returned covers are minimized.

Distribution	$STDC_{50}$	$STDC_{50}$ Q	NEWC(-3)	NEWC(-3) Q
$\mathcal{D}(1000, 200, 0.02, unweighted)$	74%	1.0080	62%	1.0130
$\mathcal{D}(1000, 200, 0.05, unweighted)$	99%	1.0005	52%	1.0244
$\mathcal{D}(2000, 200, 0.02, unweighted)$	92%	1.0027	51%	1.0170
$\mathcal{D}(3000, 300, 0.02, unweighted)$	97%	1.0010	42%	1.0186
$\mathcal{D}(5000, 500, 0.02, unweighted)$	95%	1.0010	36%	1.0157
$\mathcal{D}(5000, 500, 0.05, unweighted)$	99%	1.0004	41%	1.0239
$\mathcal{D}(5000, 500, 0.1, unweighted)$	100%	1.0000	46%	1.0318
$\mathcal{D}(10000, 1000, 0.02, unweighted)$	99%	1.0002	28%	1.0153
$\mathcal{D}(1000, 200, 0.02, discrete(1, 50))$	27%	1.0169	77%	1.0025
$\mathcal{D}(1000, 200, 0.05, discrete(1, 50))$	55%	1.0117	71%	1.0078
$\mathcal{D}(2000, 200, 0.02, discrete(1, 50))$	40%	1.0117	73%	1.0044
$\mathcal{D}(3000, 300, 0.02, discrete(1, 50))$	39%	1.0109	77%	1.0037
$\mathcal{D}(5000, 500, 0.02, discrete(1, 50))$	38%	1.0131	80%	1.0029
$\mathcal{D}(5000, 500, 0.05, discrete(1, 50))$	63%	1.0099	57%	1.0140
$\mathcal{D}(5000, 500, 0.1, discrete(1, 50))$	89%	1.0049	58%	1.0180
$\mathcal{D}(10000, 1000, 0.02, discrete(1, 50))$	20%	1.0167	90%	1.0015
$\mathcal{D}(1000, 200, 0.02, continuous(1, 50))$	9%	1.0357	91%	1.0012
$\mathcal{D}(1000, 200, 0.05, continuous(1, 50))$	10%	1.0435	90%	1.0017
$\mathcal{D}(2000, 200, 0.02, continuous(1, 50))$	2%	1.0381	98%	1.0002
$\mathcal{D}(3000, 300, 0.02, continuous(1, 50))$	7%	1.0420	93%	1.0007
$\mathcal{D}(5000, 500, 0.02, continuous(1, 50))$	3%	1.0443	97%	1.0005
$\mathcal{D}(5000, 500, 0.05, continuous(1, 50))$	7%	1.0438	93%	1.0019
$\mathcal{D}(5000, 500, 0.1, continuous(1, 50))$	19%	1.0362	81%	1.0035
D(10000, 1000, 0.02, continuous(1, 50))	2%	1.0446	98%	1.0004
Average over all instances	49.4%	1.0182	70.1%	1.0094

uniform or discrete.

Distribution	$STDC_{50} \Delta$ Quality	$NEWC(-3) \Delta$ Quality
$\mathcal{D}(1000, 200, 0.02, unweighted)$	0.0005	0.0008
$\mathcal{D}(1000, 200, 0.05, unweighted)$	0.0010	0.0000
$\mathcal{D}(2000, 200, 0.02, unweighted)$	0.0006	0.0000
$\mathcal{D}(3000, 300, 0.02, unweighted)$	0.0003	0.0000
$\mathcal{D}(5000, 500, 0.02, unweighted)$	0.0000	0.0000
$\mathcal{D}(5000, 500, 0.05, unweighted)$	0.0003	0.0000
$\mathcal{D}(5000, 500, 0.1, unweighted)$	0.0000	0.0000
$\mathcal{D}(10000, 1000, 0.02, unweighted)$	0.0000	0.0000
$\mathcal{D}(1000, 200, 0.02, discrete(1, 50))$	0.0580	0.0000
$\mathcal{D}(1000, 200, 0.05, discrete(1, 50))$	0.0617	0.0004
$\mathcal{D}(2000, 200, 0.02, discrete(1, 50))$	0.0594	0.0003
$\mathcal{D}(3000, 300, 0.02, discrete(1, 50))$	0.0599	0.0005
$\mathcal{D}(5000, 500, 0.02, discrete(1, 50))$	0.0515	0.0015
$\mathcal{D}(5000, 500, 0.05, discrete(1, 50))$	0.0264	0.0046
$\mathcal{D}(5000, 500, 0.1, discrete(1, 50))$	0.0098	0.0016
$\mathcal{D}(10000, 1000, 0.02, discrete(1, 50))$	0.0314	0.0051
$\mathcal{D}(1000, 200, 0.02, continuous(1, 50))$	0.0643	0.0003
$\mathcal{D}(1000, 200, 0.05, continuous(1, 50))$	0.0601	0.0009
$\mathcal{D}(2000, 200, 0.02, continuous(1, 50))$	0.0487	0.0001
$\mathcal{D}(3000, 300, 0.02, continuous(1, 50))$	0.0491	0.0005
$\mathcal{D}(5000, 500, 0.02, continuous(1, 50))$	0.0390	0.0019
$\mathcal{D}(5000, 500, 0.05, continuous(1, 50))$	0.0214	0.0029
$\mathcal{D}(5000, 500, 0.1, continuous(1, 50))$	0.0086	0.0030
$\mathcal{D}(10000, 1000, 0.02, continuous(1, 50))$	0.0233	0.0026
Average over all instances	0.0281	0.0011

Table 7.5: Quality difference between cover solutions before and after minimization.

7.3.3 OR Library Instances

The OR Library is a collection of optimization problems maintained by J.E. Beasley at http://people. brunel.ac.uk/~mastjjb/jeb/info.html. The Set Cover problems that we will be approximating are described at http://people.brunel.ac.uk/~mastjjb/jeb/orlib/scpinfo.html. They have perviously been used in experiments with SCP approximation in [9], and [7]. All of these problems have set costs chosen from *discrete*(1, 100).

The results are shown in table 7.6. The new heuristic has performance modestly better than 50 runs of the standard heuristic over these instances. We find the new heuristic to obtain the best solution in 60.3% of the instances, and the standard heuristic only 54%.

Table 7.6: A comparison of the approximate solutions obtained by the standard and new Set Cover heuristics for the OR Library instances.

Instance Name	m	n	ρ	$STDC_{50}$	NEWC(-3)
scp41	1000	200	2.00%	434	436
scp42	1000	200	1.99%	529	513

Instance Name	m	n	ρ	$STDC_{50}$	NEWC(-3)
scp43	1000	200	1.99%	537	526
scp44	1000	200	2.00%	504	512
scp45	1000	200	1.97%	518	514
scp46	1000	200	2.04%	585	565
scp47	1000	200	1.96%	447	438
scp48	1000	200	2.01%	502	493
scp49	1000	200	1.98%	663	659
scp410	1000	200	1.95%	521	516
scp51	2000	200	2.00%	269	259
scp52	2000	200	2.00%	323	318
scp53	2000	200	2.00%	230	230
scp54	2000	200	1.98%	247	245
scp55	2000	200	1.96%	212	212
scp56	2000	200	2.00%	225	218
scp57	2000	200	2.01%	301	299
scp58	2000	200	1.98%	300	294
scp59	2000	200	1.97%	290	281
scp510	2000	200	2.00%	273	272
scp61	1000	200	4.92%	142	143
scp62	1000	200	5.00%	153	150
scp63	1000	200	4.96%	148	149
scp64	1000	200	4 93%	135	134
scp65	1000	200	1.007%	178	169
sepol	3000	300	2.01%	250	258
scpa1	3000	300	2.01%	253	256
sepa2	2000	200	2.0170	204	201
scpa3	3000	300	2.01%	239	240
scpa4	2000	200	2.01%	240	237
scpa5	3000	300	2.0170	240	242
sepb1	2000	200	4.9970	70	73
scpb2	3000	300	4.99%	01	10
scpb5	3000	300	4.99%	01	82
scpb4	3000	300	4.99%	83	82
scpb5	3000	300	4.99%	12	13
scpc1	4000	400	2.00%	236	234
scpc2	4000	400	2.00%	224	224
scpc3	4000	400	2.00%	248	251
scpc4	4000	400	2.00%	231	225
scpca	4000	400	2.00%	220	219
scpd1	4000	400	5.01%	62	63
scpd2	4000	400	5.01%	68	68
scpd3	4000	400	5.01%	73	75
scpd4	4000	400	5.00%	63	63
scpd5	4000	400	5.00%	62	63
scpnre1	5000	500	9.98%	29	30
scpnre2	5000	500	9.97%	31	33
scpnre3	5000	500	9.97%	28	28
scpnre4	5000	500	9.97%	30	31
scpnre5	5000	500	9.98%	30	29
scpnrf1	5000	500	19.97%	15	15
scpnrf2	5000	500	19.97%	15	16
scpnrf3	5000	500	19.97%	15	16
scpnrf4	5000	500	19.97%	15	15
scpnrf5	5000	500	19.97%	14	14
scpnrg1	10000	1000	1.99%	184	186
scpnrg2	10000	1000	1.99%	161	162
scpnrg3	10000	1000	1.99%	175	177
scpnrg4	10000	1000	1.99%	178	179
scpnrg5	10000	1000	1.99%	179	173
scpnrh1	10000	1000	4.99%	67	70

Table 7.6: (continued)

Instance Name	m	n	ρ	$STDC_{50}$	NEWC(-3)
scpnrh2	10000	1000	4.99%	68	66
scpnrh3	10000	1000	4.99%	63	64
Best overall				54.0%	60.3%

Table 7.6: (continued)

7.4 Set Packing Results

7.4.1 Comparison Between Packing Heuristics

For this experiment, we run all 4 Set Packing heuristics on 100 instances from each of a variety of different random problem distributions. We use basic preprocessing for all algorithms and run all but the new heuristic 50 times, taking the best packing found in any of those runs as the result. Note that the running time of MIS_{50} and NEWP are comparable, but the running time of $STDP_{50}$ and $ROOT_{50}$ are significantly less than the running time of NEWP. The results of this experiment are shown in tables 7.7 and 7.8. Table 7.7 shows what proportion of the instances for that row's distribution were solved with the largest weight packing among the 4 algorithms. Table 7.8 shows the quality of the packings produced. The quality is computed as the average, over all instances for that row, of the ratio between the best packing found and that algorithm's packing. This means that values nearer 1 are better. In both tables, we have highlighted the best achievement in each row.

Here, the new heuristic generally performs better than the alternatives. Much as with the new cover heuristic, it performs relatively poorly on the unweighted instances. This is again because the other algorithms see a larger range of possible solutions because of the multiple runs they are given. The results for the new packing heuristic are stronger than those for the cover heuristic. Here over two thirds of the time the new heuristic finds the best packing amongst all of the heuristics studied.

7.4.2 OR Library Instances

scp45

scp46

scp47

As in section 7.3.3, we have run the packing heuristics on the OR Library instances. The results are shown in table 7.9.

The performance of the new heuristic over these instances is significantly better than the other heuristics, obtaining the best solution found over three quarters of the time.

the four set packing heuristics for the Oft Library histalices.								
Instance Name	m	n	ρ	STD_{50}	$ROOT_{50}$	MIS_{50}	NEWI	
scp41	1000	200	2.00%	5695	5639	5749	5887	
scp410	1000	200	1.95%	5940	6131	6066	6251	
scp42	1000	200	1.99%	5611	5747	5835	6044	
scp43	1000	200	1.99%	5996	6054	5850	6035	
scp44	1000	200	2.00%	5701	5808	5669	5905	

1.97%

2.04%

1.96%

5746

6026

5949

5897

6136

6095

5949

6196

6040

6025

6234

6272

1000

1000

1000

200

200

200

Table 7.9: A comparison of the approximate solutions obtained by the four set packing heuristics for the OR Library instances.

Average over all instances

Instance Name	m	n	ρ	STD_{50}	$ROOT_{50}$	MIS_{50}	NEWP
scp48	1000	200	2.01%	6154	6152	6147	6378
scp49	1000	200	1.98%	6393	6345	6360	6539
scp51	2000	200	2.00%	8206	8197	8333	8459
scp510	2000	200	2.00%	7716	7791	7724	7951
scp52	2000	200	2.00%	7845	7964	7981	8155
scp53	2000	200	2.00%	7750	7655	7874	7990
scp54	2000	200	1.98%	8073	8091	8176	8212
scp55	2000	200	1.96%	8030	8067	8038	8276
scp56	2000	200	2.00%	7991	7956	8024	8144
scp57	2000	200	2.01%	7841	7837	7748	7964
scp58	2000	200	1.98%	7928	7901	7968	8139
scp59	2000	200	1.97%	7918	8095	7993	8110
scp61	1000	200	4.92%	1597	1571	1597	1625
scp62	1000	200	5.00%	1693	1693	1815	1774
scn63	1000	200	4.96%	1559	1597	1581	1797
scp64	1000	200	4 93%	1910	1753	1812	1910
sep65	1000	200	4.97%	1627	1457	1592	1618
sepal	3000	300	2.01%	7410	7380	7580	7605
separ	3000	300	2.01%	8000	7844	7880	8105
scpa2	3000	300	2.01%	7483	7446	7499	7555
scpas	2000	200	2.0170	7400	7440	7919	7000
scpa4	3000	300	2.01%	7699	7690	7569	7006
scpa5	3000	300	2.0170	1459	1080	1308	1990
scpb1	3000	300	4.99%	1453	1459	1425	1440
scpb2	3000	300	4.99%	1419	1513	1571	1545
scpb3	3000	300	4.99%	1360	1458	1447	1503
scpb4	3000	300	4.99%	1314	1398	1413	1491
scpb5	3000	300	4.99%	1437	1486	1519	1617
scpc1	4000	400	2.00%	6684	6438	6669	6786
scpc2	4000	400	2.00%	6076	6091	6242	6224
scpc3	4000	400	2.00%	6246	6233	6260	6567
scpc4	4000	400	2.00%	6652	6716	6579	6718
scpc5	4000	400	2.00%	6794	6646	6565	6949
scpd1	4000	400	5.01%	1141	1036	1052	1131
scpd2	4000	400	5.01%	1084	1088	1144	1262
scpd3	4000	400	5.01%	1183	1216	1131	1200
scpd4	4000	400	5.00%	1195	1224	1142	1224
scpd5	4000	400	5.00%	1128	1145	1159	1201
scpnre1	5000	500	9.98%	228	363	296	296
scpnre2	5000	500	9.97%	224	359	275	293
scpnre3	5000	500	9.97%	356	241	274	274
scpnre4	5000	500	9.97%	221	239	295	295
scpnre5	5000	500	9.98%	354	238	294	294
scpnrf1	5000	500	19.97%	95	95	100	100
scpnrf2	5000	500	19.97%	99	99	100	100
scpnrf3	5000	500	19.97%	99	99	100	100
scpnrf4	5000	500	19.97%	99	99	100	100
scpnrf5	5000	500	19.97%	99	99	100	100
scpnrg1	10000	1000	1.99%	2922	2821	3097	3035
scpnrg2	10000	1000	1.99%	2993	3006	2993	2953
scpnrg3	10000	1000	1.99%	3100	3069	2872	3122
scpnrg4	10000	1000	1.99%	2946	2812	2987	3159
scpnrg5	10000	1000	1.99%	2942	2787	2980	3125
scpnrh1	10000	1000	4.99%	532	532	532	532
scpnrh2	10000	1000	4.99%	531	531	531	531
scpnrh3	10000	1000	4.99%	531	531	531	531
			//	~ ~ -		~ ~ -	

12.7%

17.5%

20.6%

76.2%

Table 7.9: (continued)

Distribution	STDP ₅₀	$ROOT_{50}$	MIS ₅₀	NEWP
$\mathcal{D}(1000, 100, 0.02, unweighted)$	100%	100%	100%	100%
$\mathcal{D}(1000, 200, 0.02, unweighted)$	44%	44%	61%	52%
$\mathcal{D}(1000, 200, 0.05, unweighted)$	69%	72%	48%	36%
$\mathcal{D}(2000, 200, 0.02, unweighted)$	82%	79%	66%	31%
$\mathcal{D}(3000, 300, 0.02, unweighted)$	8%	5%	62%	66%
$\mathcal{D}(5000, 500, 0.02, unweighted)$	17%	27%	64%	45%
$\mathcal{D}(5000, 500, 0.05, unweighted)$	78%	77%	68%	67%
$\mathcal{D}(5000, 500, 0.1, unweighted)$	62%	62%	92%	90%
$\mathcal{D}(10000, 1000, 0.02, unweighted)$	62%	55%	43%	27%
$\mathcal{D}(1000, 100, 0.02, discrete(1, 50))$	17%	7%	16%	70%
$\mathcal{D}(1000, 200, 0.02, discrete(1, 50))$	0%	0%	5%	95%
$\mathcal{D}(1000, 200, 0.05, discrete(1, 50))$	11%	9%	15%	70%
$\mathcal{D}(2000, 200, 0.02, discrete(1, 50))$	2%	3%	6%	91%
$\mathcal{D}(3000, 300, 0.02, discrete(1, 50))$	2%	5%	5%	88%
$\mathcal{D}(5000, 500, 0.02, discrete(1, 50))$	7%	6%	10%	77%
$\mathcal{D}(5000, 500, 0.05, discrete(1, 50))$	17%	25%	32%	49%
$\mathcal{D}(5000, 500, 0.1, discrete(1, 50))$	37%	35%	51%	55%
$\mathcal{D}(10000, 1000, 0.02, discrete(1, 50))$	13%	14%	28%	47%
$\mathcal{D}(1000, 100, 0.02, continuous(1, 50))$	8%	10%	19%	64%
$\mathcal{D}(1000, 200, 0.02, continuous(1, 50))$	4%	2%	1%	93%
$\mathcal{D}(1000, 200, 0.05, continuous(1, 50))$	4%	16%	19%	61%
$\mathcal{D}(2000, 200, 0.02, continuous(1, 50))$	1%	0%	2%	97%
$\mathcal{D}(3000, 300, 0.02, continuous(1, 50))$	2%	0%	1%	97%
$\mathcal{D}(5000, 500, 0.02, continuous(1, 50))$	2%	4%	5%	89%
$\mathcal{D}(5000, 500, 0.05, continuous(1, 50))$	15%	20%	19%	58%
$\mathcal{D}(5000, 500, 0.1, continuous(1, 50))$	43%	43%	36%	45%
$\mathcal{D}(10000, 1000, 0.02, continuous(1, 50))$	5%	12%	17%	66%
Average over all instances	26.4%	27.1%	33.0%	67.6%

Table 7.7: A comparison between the performance of 4 different heuristics for the Set Packing problem. The proportion of the 100 problems that each heuristic performed best on is reported.

Distribution	$STDP_{50} Q$	$ROOT_{50}$ Q	MIS_{50} Q	NEWP Q
$\mathcal{D}(1000, 100, 0.02, unweighted)$	1.0000	1.0000	1.0000	1.0000
$\mathcal{D}(1000, 200, 0.02, unweighted)$	1.0067	1.0064	1.0041	1.0061
$\mathcal{D}(1000, 200, 0.05, unweighted)$	1.0122	1.0111	1.0246	1.0341
$\mathcal{D}(2000, 200, 0.02, unweighted)$	1.0013	1.0015	1.0029	1.0074
$\mathcal{D}(3000, 300, 0.02, unweighted)$	1.0176	1.0184	1.0041	1.0045
$\mathcal{D}(5000, 500, 0.02, unweighted)$	1.0156	1.0140	1.0062	1.0112
$\mathcal{D}(5000, 500, 0.05, unweighted)$	1.0193	1.0200	1.0293	1.0319
$\mathcal{D}(5000, 500, 0.1, unweighted)$	1.1267	1.1267	1.0267	1.0333
$\mathcal{D}(10000, 1000, 0.02, unweighted)$	1.0094	1.0118	1.0189	1.0295
$\mathcal{D}(1000, 100, 0.02, discrete(1, 50))$	1.0039	1.0060	1.0038	1.0005
$\mathcal{D}(1000, 200, 0.02, discrete(1, 50))$	1.0334	1.0254	1.0250	1.0003
$\mathcal{D}(1000, 200, 0.05, discrete(1, 50))$	1.0535	1.0589	1.0378	1.0088
$\mathcal{D}(2000, 200, 0.02, discrete(1, 50))$	1.0223	1.0208	1.0164	1.0003
$\mathcal{D}(3000, 300, 0.02, discrete(1, 50))$	1.0299	1.0274	1.0236	1.0007
$\mathcal{D}(5000, 500, 0.02, discrete(1, 50))$	1.0344	1.0394	1.0284	1.0025
$\mathcal{D}(5000, 500, 0.05, discrete(1, 50))$	1.0713	1.0781	1.0464	1.0235
$\mathcal{D}(5000, 500, 0.1, discrete(1, 50))$	1.0970	1.1065	1.0883	1.0548
$\mathcal{D}(10000, 1000, 0.02, discrete(1, 50))$	1.0435	1.0458	1.0269	1.0086
$\mathcal{D}(1000, 100, 0.02, continuous(1, 50))$	1.0053	1.0065	1.0039	1.0005
$\mathcal{D}(1000, 200, 0.02, continuous(1, 50))$	1.0290	1.0247	1.0253	1.0003
$\mathcal{D}(1000, 200, 0.05, continuous(1, 50))$	1.0633	1.0553	1.0415	1.0106
$\mathcal{D}(2000, 200, 0.02, continuous(1, 50))$	1.0265	1.0246	1.0198	1.0001
$\mathcal{D}(3000, 300, 0.02, continuous(1, 50))$	1.0356	1.0352	1.0294	1.0002
$\mathcal{D}(5000, 500, 0.02, continuous(1, 50))$	1.0479	1.0435	1.0335	1.0013
$\mathcal{D}(5000, 500, 0.05, continuous(1, 50))$	1.0660	1.0587	1.0471	1.0188
$\mathcal{D}(5000, 500, 0.1, continuous(1, 50))$	1.0623	1.0795	1.1354	1.1002
$\mathcal{D}(10000, 1000, 0.02, continuous(1, 50))$	1.0560	1.0618	1.0361	1.0062
Average over all instances	1.0367	1.0373	1.0291	1.0147

Table 7.8: A comparison between the performance of 4 different heuristics for the Set Packing problem. The average quality of each algorithm over 100 instances is reported.

Chapter 8

Discussion

We have demonstrated a novel method for devising greedy approximation heuristics. For two problems, Set Cover and Set Packing, we have constructed new heuristics and shown that their performance is better than alternative greedy algorithms. For the Set Cover heuristic, we have demonstrated that the valuation we define exists and is unique, while we are unsure whether the Set Packing heuristic guarantees a unique valuation. This leaves a variety of open questions. Is the valuation determined by the Packing heuristic unique? Is there some simple way to characterize matrices **M** for which there is a unique solution **v** to $\mathbf{Mv} = \mathbf{v}^{\diamond -1}$? Do the iterations we have described always converge? Is there a more efficient way to find these fixed points? Do either of the new heuristics guarantee some approximation ratio? Most importantly, can the overall technique of defining valuations recursively be used to construct high quality approximation algorithms for other problems?

In our experimental results, it is striking that none of the algorithms considered appears to dominate the others. Although the new heuristics are generally preferable to the alternatives on the sorts of random instances we have used, it is still beneficial to use a variety of heuristics when looking for good covers and packings. We can straightforwardly recommend that any real-world software in which any of the standard cover or packing heuristics are used exclusively to obtain approximate solutions could substantially benefit from also considering the solutions given by our new heuristics. Although their runtimes are longer, they are not substantially so, and is should be possible to engineer them to run in only 5 or 10 times what's required for the standard heuristics.

It is both a blessing and a curse that the new heuristics are less prone to ties than the standard heuristics. It is convenient that a single run deterministically generates a particular solution, but at the same time, this means that multiple runs will not allow us to sample from the space of possible solutions in the same way as the standard algorithms do. In order to generate a wider variety of solutions using the new heuristics there are a few different approaches that might be made. For the cover heuristic, varying the parameter γ permits us to obtain different covers. Additionally, for any valuation-producing heuristic, we could randomize the general greedy scheme to select sets with probabilities determined by the valuations produced. We leave it to others to determine how valuable are the solutions built with this approach to obtaining variety.

Bibliography

- Giorgio Ausiello, Nicolas Bourgeois, Telis Giannakos, and Vangelis Th. Paschos. Greedy algorithms for on-line set-covering. *Algorithmic Operations Research*, 4(1):36–48, 2009.
- [2] Alberto Caprara, Matteo Fischetti, and Paolo Toth. A heuristic method for the set covering problem. Operations Research, 47(5):730-743, 1999.
- [3] Paul C Chu and John E Beasley. A genetic algorithm for the multidimensional knapsack problem. Journal of Heuristics, 4(1):63–86, 1998.
- [4] V. Chvatal. A greedy heuristic for the set-covering problem. Mathematics of Operations Research, 4(3):233-235, 1979.
- [5] Uriel Feige. A threshold of ln n for approximating set cover. Journal of the ACM (JACM), 45(4):634– 652, 1998.
- [6] C.H. Fitzgerald and R.A. Horn. On fractional hadamard powers of positive definite matrices. Journal of Mathematical Analysis and Applications, 61(3):633-642, 1977.
- [7] Fernando C Gomes, Cláudio N Meneses, Panos M Pardalos, and Gerardo Valdisio R Viana. Experimental analysis of approximation algorithms for the vertex cover and set covering problems. *Computers & Operations Research*, 33(12):3520–3534, 2006.
- [8] Rica Gonen and Daniel Lehmann. Optimal solutions for multi-unit combinatorial auctions: Branch and bound heuristics. In *Proceedings of the 2nd ACM conference on Electronic commerce*, pages 13–20. ACM, 2000.
- [9] Tal Grossman and Avishai Wool. Computational experience with approximation algorithms for the set covering problem. *European Journal of Operational Research*, 101(1):81–92, 1997.
- [10] Magnús M Halldórsson. Approximating discrete collections via local improvements. In Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, pages 160–169. Society for Industrial and Applied Mathematics, 1995.
- [11] D. Idczak and M. Majewski. A generalization of the Poincaré-Miranda theorem with an application to the controllability of nonlinear repetitive processes. In *International Workshop on Multidimen*sional (nD) Systems, 2009.
- [12] David S Johnson. Approximation algorithms for combinatorial problems. Journal of Computer and System Sciences, 9(3):256–278, 1974.

- [13] Richard M. Karp. Reducibility among combinatorial problems. In Complexity of Computer Computations, pages 85–103, 1972.
- [14] Petr Slavík. A tight analysis of the greedy algorithm for set cover. In STOC, pages 435–441, 1996.
- [15] David P Williamson. The primal-dual method for approximation algorithms. Mathematical Programming, 91(3):447–478, 2002.