University of California
Santa Barbara

The Impact of Information in Cooperative and Noncooperative Systems

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Electrical and Computer Engineering

by

David Randall Grimsman

Committee in charge:

Professor Jason R. Marden, Chair
Professor João P. Hespanha
Professor Francesco Bullo
Professor Ambuj K. Singh

June 2021

The Dissertation of David Randall Grimsman is approved.

Professor João P. Hespanha

Professor Francesco Bullo

Professor Ambuj K. Singh

Professor Jason R. Marden, Committee Chair

May 2021

The Impact of Information in Cooperative and Noncooperative Systems

by

David Randall Grimsman

To Kalisha, Rose, Harvey, Elaine, and Walter

Acknowledgements

I have felt incredibly supported in my journey as a graduate student, on so many levels. First and foremost, thank you to my wife and best friend Kalisha. Without your love and patience, this work doesn't happen. We made a conscious decision to somewhat abruptly change our family's lifestyle so that I could pursue a PhD. I told you many times that if balancing finances and family in this way ever became too much of a burden, that I would quit and we would go do something else. While I'm sure this has crossed your mind frequently, I have never felt anything but strength and support from you throughout; truly you are the ultimate partner and teammate.

Thank you to my children. Entering the program, we had two twin babies: Rose and Harvey. During this journey we have added two more children: Elaine and Walter. No matter the frustrations of research or classes, I have always been able to count on fun, laughter, and unconditional love as soon as I walk through the door. I hope you kids remember this time with fondness, and that it serves as a reminder of the importance of becoming the people that you're meant to be.

Thank you to my advisor, Jason. You have truly been the epitome of patience with me. There have been several times in research when we have disagreed, you have convinced me that you're right, then I go back and convince myself that you're wrong, only to have you convince me once again that you're right. I appreciate your ability to simplify complex ideas and ask insightful questions. I have also learned a great deal in writing in such a way that is approachable and intuitive – I hope that this dissertation lives up to that standard. I also appreciate your emphasis on building relationships within our academic world, especially in reaching out to students. You've created a great culture of comraderie and rigorous study in our lab. I mean, how many PhD students can say that they've had a meaningful discussion with their advisor while surfing? Thank

you for all you've done to mentor and lift me!

Thank you also to the rest of my committee. João, thank you for always taking the time to carefully consider my work and give insightful feedback. I'm constantly amazed at how you seem to be an expert on so many topics; I've often found myself lost in the details of a presentation during a seminar, only to have you ask a very pointed question about an assumption from three slides ago. I admire your work ethic and humility and hold you as an example of scholarship. Ambuj, thank you for your efforts in creating the IGERT Network Science program, which has been a highly influential part of my experience at UCSB. Thank you for your encouragement throughout the program and helpful insights into new research directions. Francesco, thank you for being kind and gracious in your feedback, and thank you for example of service at UCSB and in the controls community.

Finally, thank you to my fellow students in the lab. You have made my time at UCSB enjoyable, and I look up to each of you for your dedication to your work. I'm excited to see the various important contributions that I know you all will make!

I hope that this work is a commensurate reflection of all the wonderful support that I have received during my time as a graduate student.

Curriculum Vitæ
David Randall Grimsman

Education

2021	Ph.D. in Electrical and Computer Engineering (Expected), University of California, Santa Barbara.
2016	M.S. in Computer Science, Brigham Young University.
2006	B.S in Electrical Engineering, Brigham Young University.

Publications

Journal Publications

1. P. E. Paré, D. Grimsman, A. T. Wilson, M. K. Transtrum and S. Warnick, "Model boundary approximation method as a unifying framework for balanced truncation and singular perturbation approximation," *IEEE Transactions on Automatic Control*, 2019.
2. D. Grimsman, M. S. Ali, J. P. Hespanha, and J. R. Marden, "The impact of information in greedy submodular maximization," *IEEE Transactions on Control of Network Systems*, 2018.
3. L. D. R. Beal, D. Peterson, D. Grimsman, S. Warnick and J. D. Hedengren, "Integrated scheduling and control in discrete-time with dynamic parameters and constraints," *Computers and Chemical Engineering*, 2018.

Proceedngs and Refereed Conferences

1. D. Grimsman, M. R. Kirchner, J. P. Hespanha and J. R. Marden, "The impact of message passing in agent-based submodular maximization," *IEEE Conference on Decision and Control*, 2020.
2. D. Grimsman, J. H. Seaton, J. R. Marden and P. N. Brown, "The cost of denied observation in multiagent submodular optimization," *IEEE Conference on Decision and Control*, 2020.
3. H. Sun, D. Grimsman and J. R. Marden, "Distributed submodular maximization with parallel execution," *American Control Conference*, 2020.
4. D. Grimsman, J. P. Hespanha and J. R. Marden, "Stackelberg equilibria for two-player network routing games on parallel networks," *American Control Conference*, 2020.
5. D. Grimsman, J. P. Hespanha and J. R. Marden, "Strategic information sharing in greedy submodular maximization," *IEEE Conference on Decision and Control*, 2018.

6. D. Grimsman, M. S. Ali, J. P. Hespanha, and J. R. Marden, "Impact of information in greedy submodular maximization," *IEEE Conference on Decision and Control*, 2017.
7. D. Grimsman and S. Warnick, "Deadbeat-like approximations for sequencing non-rigid heaps," *IEEE Conference on Decision and Control*, 2016.
8. D. Grimsman, V. Chetty, N. Woodbury, E. Vaziripour, S. Roy, D. Zappala and S. Warnick, "A case study of a systematic attack design method for critical infrastructure cyber-physical systems," *American Control Conference*, 2016.

Awards/Honors

1. UCSB ECE Outstanding TA Award, 2020
2. UC Santa Barbara Grad Slam Semifinalist , 2018
3. NSF IGERT Network Science Trainee Fellowship, 2016
4. BYU 3-Minute Thesis CS Department Winner, 2016
5. BYU Student Research Conference Best Session Presentation 2014, 2016
6. BYU BYU Heritage Scholarship Award Winner 2000

Abstract

The Impact of Information in Cooperative and Noncooperative Systems

by

David Randall Grimsman

Large-scale autonomous systems are systems comprised of many components, each acting according to its own preferences, local information and capabilities. Such systems are ubiquitous in our world and include automated warehouses, UAV swarms, traffic systems, sensor networks, the internet of things, auctions, ridesharing systems, and social networks. We focus on autonomous systems which are engineered: each component is human-designed.

While such systems are attractive because they can process a high amount of data and are generally robust against single points of failure, there are often many challenges in their design. System designers must take into account that each component has its own capabilities, model of the environment, data set, and local objective. Furthermore, the system designer often cannot make decisions for each component at each time step, rather, decision-making rules are assigned to allow components to react autonomously. Small adjustments to such rules can often have cascading effects throughout the system. Finally, the interconnected nature of the system opens doors to new kinds of system-wide attacks and vulnerabilities.

In this work, we focus on the challenge of information sharing constraints: each agent does not have access to all of the system information at every given time step. These constraints often arise naturally (i.e., no router can access all available data on the internet before making a routing decision), but they can also arise from privacy, trust or political issues. Thus it is imperative for the system designer to understand the relevant

information constraints on the system and their effects on the emergent behavior. In this work we endeavor to answer the two following questions:

1. How do a set of information sharing constraints impact the resulting emergent system-wide behavior?

2. How can a system designer strategically set decision-making rules for the components to offset any negative effects caused by information sharing constraints?

We answer the first question by assuming that the emergent behavior is a system equilibrium, and then comparing the value of the worst-case equilibrium to the value of the optimal decision set, where value is based on the system designer objective. Different types of information sharing constraints among the components are evaluated on this basis. We answer the second question in certain settings by showing that small deviations from standard decision-making rules can improve the system's performance guarantees.

These questions are addressed in two settings: first in a cooperative setting, where the system designer can design the decision-making rules for each agent. The system designer objective function is assumed to be submodular, and this property is leveraged to show closeness of equilibrium to optimal. The second is a noncooperative setting, where the system design must operate in the presence of an attacker. Here, the constraints on information sharing are related to how much knowledge about the attacker the other players have.

Contents

Curriculum Vitae **vii**

Abstract **ix**

1 Introduction **1**
1.1 Model 3
1.2 Illustrative Example 6
1.3 Summary of Contributions 9

2 Technical Preliminaries **13**
2.1 Submodular Objective Functions 13
2.2 The Price of Anarchy 15

3 Blind and Isolated Agents **17**
3.1 Chapter Model 19
3.2 Effects of Compromised Agents 21
3.3 Simulation 27

4 Graph Constraints **30**
4.1 Review of Graph Theory Terms 32
4.2 Valid Utility Games with Graph Constraints 35
4.3 A Bound on Optimal Utilities 39

5 The Greedy Algorithm **42**
5.1 Introduction 42
5.2 Examples 45
5.3 Price of Anarchy Bounds 46
5.4 Optimal Structures 48
5.5 Strategic Information Sharing 57

6 Augmenting Action Sets **68**
6.1 Chapter Model 70
6.2 Comparison to Greedy 74
6.3 Suboptimal Selections 83
6.4 Numerical Example 87

7 Network Security **91**
7.1 Model 94
7.2 Problem Hardness 98
7.3 Equilibria 102
7.4 The Value of Information 105
7.5 A Robust Policy 110

8 Conclusions **113**
8.1 Future Work and Open Questions 114

A Proofs for Selected Results **116**
A.1 Proof for Theorem 3.2 116
A.2 Proof for Theorem 4.1 120
A.3 Proof for Lower Bound in Example 5.3 127
A.4 Proof for Lemma 5.3 127
A.5 Proof for Lemma 5.2 128
A.6 Proof for Lemma 6.1 129
A.7 Proof for Lemma 6.2 131
A.8 Proof for Lemma 7.2 134
A.9 Proof for Lemma 7.3 134
A.10 Proof for Lemma 7.4 135
A.11 Proof for Theorem 7.3 137
A.12 Proof for Theorem 7.4 140

Bibliography **144**

Chapter 1

Introduction

Large-scale autonomous systems are systems comprised of many components, each acting according to its own preferences, local information and capabilities. Historically, such systems have been purely social: each component is a human. These include political systems, economic markets, and other systems of competition, culture, and customs. Advancements in automation have caused the nature of these systems to change; some or all components of an autonomous system are human-built and driven by data and algorithms [1, 2, 3]. For instance, the modern business must understand how to blend software and human decision-making in order to drive growth [4]. In many cases, such as automated warehouses [5] or robot swarms [6], all components are technological. We refer to such systems as *engineered autonomous systems*.

Engineered autonomous systems are often advantageous when solving problems that involve a massive amount of data and cannot be solved centrally [7]. They can also offer a robustness against failure of a single component in the face of an uncertain environment or an attacker. However, the design of such systems is often challenging for the following reasons:

1. Each component has its own set of capabilities, model of the environment, data

set, local objective, and time scale.

2. The system designer cannot (either for tractability or privacy reasons) directly make decisions for each component. Rather, the designer can only design a set of decision-making rules for each component.

3. Small adjustments to local behavior can have cascading effects.

4. The connected nature of the system introduces new vulnerabilities to system-wide attacks [8].

The system designer, with its own objective, must account for these challenges. Thus the overall goal of this research is to design systems that are distributed and subject to time and information constraints to provide high performance guarantees.

In this work, our focus will be on understanding how information sharing constraints within the system affect the overall system behavior. Information sharing constraints arise from several causes. For instance, a common constraint in sensor networks is that individual components may have low computational power, especially in cases where a component may need to operate without access to a power source [9]. Components may also be connected via networks of limited bandwidth, such as in disaster recovery, where sharing all system information would take too long for the system to be responsive [10]. It could also be that the time scale on which decisions need to be made is such that components can only receive or process a small amount of information, such as in a UAV swarm [11]. Information sharing constraints can also arise for nontechnical reasons, for instance, if two components do not fully trust one another, as in privacy-preserved learning scenarios [12]. Lastly, the information sharing constraints could be the result of an outside attacker seeking to disrupt the system [13].

Information sharing constraints in engineered systems give rise to two important

questions:

1. How do a set of information sharing constraints impact the resulting emergent system-wide behavior?

2. How can a system designer strategically set decision-making rules for the components to offset any negative effects caused by information sharing constraints?

We address these questions in two settings: first in a cooperative setting, where the system designer can design the decision-making rules for each agent, and second, in a noncooperative setting, where the system designer must operate in the presence of an attacker.

1.1 Model

Consider a set of $N = \{1, \ldots, n\}$ agents, where each agent i is endowed with a set of actions X_i. Each action $x_i \in X_i$ is evaluated according to a utility function $U_i : X_1 \times \cdots \times X_n \to \mathbb{R}$, which is dependent on the actions chosen by the other agents. To highlight this dependence, we often use the notation $U_i(x_i, x_{-i})$, where $x_{-i} := (x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n)$. In addition to the agents, we assume (following [14]) the presence of a system designer, who is able to assess the quality of a set of actions using a global objective function $f : X_1 \times \cdots \times X_n \to \mathbb{R}$. The goal of the system designer is for agents to choose an action profile $x^{\text{opt}} \in X = X_1 \times \cdots \times X_n$ which satisfies

$$x^{\text{opt}} \in \arg\max_{x \in X} f(x) \tag{1.1}$$

One challenge in finding x^{opt} is that (1.1) is generally intractable, therefore even if the system designer had the ability to explicitly assign actions to agents, it could not guarantee to find x^{opt} in polynomial time. In most cases the system designer does not have

this ability; therefore, assume that the system designer can only *incentivize* agents to choose desirable actions by adjusting utility functions.

Under these constraints, one can define the emergent action profile x^{eq} to be an equilibrium, which satisfies

$$U_i(x_i^{\text{eq}}, x_{-i}^{\text{eq}}) \geq U_i(x_i, x_{-i}^{\text{eq}}), \tag{1.2}$$

for all $x_i \in X_i$ and for all $i \in N$, i.e., given the actions of other agents, agent i has no unilateral incentive to deviate from is choice x_i^{eq}. Studying equilibria allows the system designer abstract away the dynamics or algorithms of the system, and instead focus on designing utility functions that are guaranteed to result in high-quality equilibria, regardless of the choice of dynamics.

In this standard game theoretic model, it is generally assumed by the definition of U_i that each agent can access the full actions of all other agents in the system. As discussed, this may not be a realistic assumption in for many applications. Therefore, one can consider two types of information sharing constraints, both of which are addressed in this work

1. Type 1: Agents can only observe the actions of a subset of agents in the system.
2. Type 2: Agents can only observe the messages (which may or may not contain the actions) of a subset of agents in the system.

We will precisely describe each.

1.1.1 Type 1 Information Sharing Constraints

The Type 1 information sharing constraint assumes that agent i only has access to the actions of a subset $\mathcal{N}_i$ of other agents in the system. This is enforced by imposing

that for any two action profiles $x^1, x^2 \in X$, all U_i have the following property:

$$x^1_j = x^2_j \ \forall j \in \mathcal{N}_i \implies U_i(x_i, x^1_{-i}) = U_i(x_i, x^2_{-i}) \ \forall x_i \in X_i. \tag{1.3}$$

We use the notation $U_i(x_i, x_{\mathcal{N}_i})$, where $x_{\mathcal{N}_i}$ is the tuple of actions of agents in $\mathcal{N}_i$. Thus the Type 1 information sharing constraints are completely defined by the sets $\mathcal{N}_i$. If $\mathcal{N}_i = \emptyset$ for all i, then each agent acts independently. On the other hand, if $\mathcal{N}_i = N \setminus \{i\}$ for all i, then we recover the standard setting.

1.1.2 Type 2 Information Sharing Constraints

The Type 2 information sharing constraint assumes further that agent i does not necessarily have access to the actions of agents in $\mathcal{N}_i$, rather if $j \in \mathcal{N}_i$, agent i has access to a *message* m_j. One way to model this is to consider a meta-action $a_i = (x_i, m_i)$, where m_i belongs to some message set M_i, and corresponding meta-utility function $V_i(a_i, a_{-i})$. Again one can impose that for two meta-action profiles a^1, a^2, that V_i satisfies

$$m^1_j = m^2_j \ \forall j \in \mathcal{N}_i \implies V_i(a_i, a^1_{\mathcal{N}_i}) = V_i(a_i, a^2_{\mathcal{N}_i}) \ \forall a_i \in X_i \times M_i, \tag{1.4}$$

i.e., V_i is not dependent on the action of any agent in $\mathcal{N}_i$, only the message. Using meta-actions allows one to therefore separate what each agent communicates to the other agents from its contribution to the global objective function, which is still only dependent on x. The Type 2 information sharing constraints are imposed by limiting the sets M_i. For instance, if $M_i = \{0, 1\}$ then agent i can share 1 bit with other agents, and they must choose their actions without knowing x_i. On the other hand, one could allow more information by letting $M_i = 2^{X_i}$; here agent i could share its entire action set as a message.

For a system with Type 2 information sharing constraints, $a^{\text{eq}} = (x^{\text{eq}}, m^{\text{eq}})$ is considered an equilibrium if

$$V_i(a_i^{\text{eq}}, a_{-i}^{\text{eq}}) \geq V_i(a_i, a_{-i}^{\text{eq}}), \tag{1.5}$$

for all $a_i \in X_i \times M_i$ and for all $i \in N$.

1.2 Illustrative Example

In this section, and several places throughout this work, we will leverage the Weighted Set Cover (WSC) problem.

Example 1.1 (Weighted Set Cover Problem [15]) *Consider a set of n agents and a base set of resources S, where every $s \in S$ has a value v_s. Each agent's action set is a family of subsets of S, i.e., $x_i \subseteq S$ for all $x_i \in X_i$, where X_i is not dependent on the actions of other agents. The system designer objective is for the agents to select the most valuable set of resources, i.e.,*

$$f(x) = \sum_{s \in \cup_i x_i} v_s. \tag{1.6}$$

We first assume the presence of Type 1 information sharing constraints, in that agent i has access to the actions of the agents in $\mathcal{N}_i$. Consider the problem instance in Figure 1.1a with 3 agents and 4 resources. The black circles represent the agents, and the white boxes represent the resources. The number next to each resource represents its value, for instance resource s_0 has value $1+\varepsilon$ for some small $\varepsilon > 0$. The blue lines (dashed and solid) represent the action set X_i of each agent: $X_i = \{\{s_0\}, \{s_i\}\}$ for all i. The solid black lines between agents represent the fact that $\mathcal{N}_i = N \setminus \{i\}$ for all i. We will assume first that the system designer endows each agent with a utility function $U_i(x_i, x_{\mathcal{N}_i}) = f(x_i, x_{\mathcal{N}_i})$. Thus the value of an action compared to others in the action set is its added value relative to the actions of the agents in $\mathcal{N}_i$ according to f.

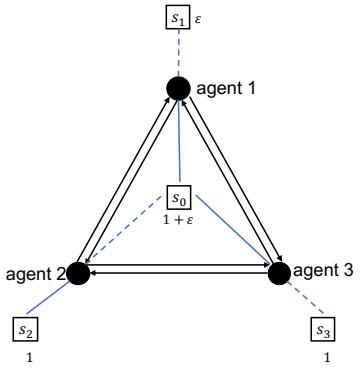

(a) In this system, all agents share information with one another. The equilibrium represented has value $f(x^{\text{eq}}) = 2 + 2\varepsilon$, or roughly $2/3$ of the optimal value.

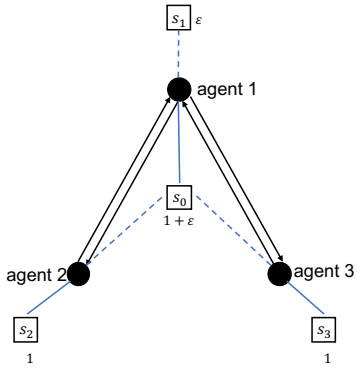

(b) In this system, agents 2 and 3 share information with each other. The equilibrium represented has value $f(x^{\text{eq}}) = 1 + 2\varepsilon$, or roughly $1/3$ of the optimal value.

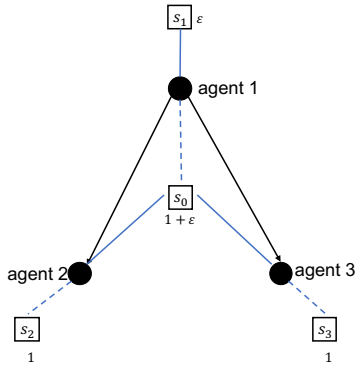

(c) In this system, agent 1 shares information with agents 2 and 3. The equilibrium represented is the optimal allocation, despite stricter information sharing constraints.

Figure 1.1: An example WSC problem with 3 different Type 1 information sharing constraints. The black circles are the agents and the white squares are the resources. Each action is a single resource, thus the blue lines between agents and resources represent the various action spaces, for instance, $X_1 = \{\{s_0\}, \{s_1\}\}$. The optimal decision set is $x^{\text{opt}} = (\{s_0\}, \{s_2\}, \{s_3\})$, and $f(x^{\text{opt}}) = 3 + \varepsilon$. The solid black lines between agents represent the information sharing constraints on $\mathcal{N}_i$: an edge exists between i and j if and only if $j \in \mathcal{N}_i$. Under these conditions, the dashed lines represent a possible equilibrium of the system.

In this setting, there are 3 possible equilibria: $(\{s_1\}, \{s_0\}, \{s_3\})$, $(\{s_1\}, \{s_2\}, \{s_0\})$, and $(\{s_0\}, \{s_2\}, \{s_3\})$ – the last one is also the optimal allocation, with value $f(x^{\text{eq}}) = 3 + \varepsilon$. Figure 1.1a also shows the second listed equilibrium with the dashed blue lines. In either non-optimal equilibrium, $f(x^{\text{eq}}) = 2 + 2\varepsilon$, roughly 2/3 of the optimal, since ε is small. In the literature it is well-known that with these utility functions and information sharing constraints, the equilibrium for any resource allocation problem is guaranteed to be within 1/2 of the optimal [16].

Figure 1.1b shows how different information sharing constraints can affect the possible equilibria. The agents, action sets, utility functions, and objective function are the same, however, here we impose that $2 \notin \mathcal{N}_3$ and $3 \notin \mathcal{N}_2$, i.e., agents 2 and 3 do not share information with each other. In Figure 1.1b this is illustrated by the fact that there are no solid black lines between agents 2 and 3. This new set of information sharing constraints yields a different set of equilibria: $(\{s_1\}, \{s_0\}, \{s_0\})$ and the optimal allocation. Since agents 2 and 3 no longer consider each other's actions in equilibrium, they both choose s_0 in the non-optimal equilibrium x^{eq}, where $f(x^{\text{eq}}) = 1 + 2\varepsilon$, roughly 1/3 of the optimal. Therefore, we see in this instance that constraining information sharing has allowed for a possibly-worse equilibrium – and a lower equilibrium guarantee than when all agents share information with each other.

Figure 1.1c shows a problem instance where information sharing has been further constrained: $\mathcal{N}_1 = \emptyset$ and $\mathcal{N}_2 = \mathcal{N}_3 = \{1\}$. Essentially, in equilibrium, agent 1 chooses independently of the other agents, and agents 2 and 3 only consider the choice of agent 1. Here, although these are tighter sharing constraints, *the only equilibrium is the optimal allocation.* A more constrained setting has yielded a higher performance. While achieving the optimal in this particular instance is an artifact of how the action sets intersect, it will be shown in Chapter 5 that, like the full information setting, *these information sharing constraints recover the guarantee that any equilibrium will be within 1/2 of the optimal.*

Returning to the example in Figure 1.1b, one can find the optimal allocation if the information sharing constraints are relaxed in a different way: suppose that we have Type 2 information sharing constraints where $M_i = 2^{X_i}$, i.e., each message can be some subset of the action set. A possible meta-utility function is

$$V_i(a_i, a_{\mathcal{N}_i}) = |m_i| + \max_{x_{\mathcal{N}_i} \in \Pi_{j \in \mathcal{N}_i} m_j} f(x_i, x_{\mathcal{N}_i}), \tag{1.7}$$

recalling that $a_i = (x_i, m_i)$. In equilibrium, each agent is incentivized to choose $m_i = X_i$ (the largest subset of X_i), and to choose the x_i that maximizes f over all possible actions of agents in $\mathcal{N}_i$. For the example in Figure 1.1b, this causes agents 2 and 3 to recognize that the best choice is $\{s_2\}$ and $\{s_3\}$, respectively. Likewise, since agent 1 can see all action sets, it knows that the best choice is $\{s_0\}$, thus the optimal allocation is the only equilibrium. Of course it will not hold that this choice of meta-utility will always incentivize agents to choose an optimal decision set. The guarantees of such rules across the set of WSC problems remains an open question.

1.3 Summary of Contributions

The example in the previous section illustrates that the study of how information sharing constraints affect equilibria is nontrivial, and in some cases unintuitive. A rigorous understanding of the "value of information" is thus imperative to the efficiency of these systems. Chapters 3–6 cover a cooperative setting, and Chapter 7 covers a noncooperative setting.

Chapters 3–6 consider a class of problems where the system objective function is submodular: it exhibits a "diminishing returns" property that will be defined precisely in Chapter 2. Under these conditions, we are interested in the guarantees of any equilibrium,

as compared to the optimal set of choices. Chapters 3–5 focus on Type 1 information sharing constraints, whereas Chapter 6 focuses on Type 2.

Chapter 3 assumes the presence of an attacker, which has the capability to compromise k agents in the system by either making them "blind" (they do not share information) or isolated (they act independently of the other agents). We assume that the system designer has chosen from a set of valid utilities that are related to the objective function f. Under these conditions, we show how the worst-case equilibrium gets incrementally worse as k increases. It is also shown that choosing a specific utility from within that set can achieve a slightly higher guarantee.

Chapter 4 explores the slightly more general setting, where any Type 1 constraint is allowed. The sets $\mathcal{N}_i$ can be modeled as a directed graph $G = (V, E)$, where the vertices are the agents and $(i, j) \in E$ if and only if $i \in \mathcal{N}_j$. It is well-known that if the graph is complete, then the system designer can choose utilities within a set of *valid utilities* with the guarantee that the resulting equilibrium will be valued within $1/2$ of the optimal. However, we show that this guarantee can degrade quickly when edges are removed from the graph - in fact, for many graphs the guarantee is arbitrarily bad for large systems. We propose constraining the set of utility functions further to those that additionally have a consistency property, which mitigates this degradation.

Chapter 5 operates under many of the same assumptions as Chapter 4, however, it is assumed that the graph G which represents the information sharing constraints, is a directed acyclic graph (DAG). This particular set of graphs is important, because they admit a simple greedy algorithm for arriving at an equilibrium. Given any DAG G, one can sequence the nodes such that $(i, j) \in E$ only if $i < j$. The algorithm proceeds as follows: each agent sequentially chooses an action which maximizes its utility function based on the actions of those agents which it can observe. Once each agent has chosen, the resulting decision set is an equilibrium. Under these assumptions, the literature

shows that a fully-connected graph will yield an equilibrium guaranteed to be within 1/2 of the optimal. We show that the performance guarantees degrade as the independence number (the largest set of nodes among which there are no edges) of the DAG increases. We then show, given a fixed number of edges, the graphs that provide the best efficiency in this regard. Section 5.5 leverages Type 2 information sharing constraints, where the message m_i can either be x_i, or the message m_j from $j \in \mathcal{N}_i$. For a particular set of DAGs, we show the optimal set of meta-utility functions.

Type 2 constraints are also used in Chapter 6, where we assume G is a fully-connected DAG. Here, agent i, in addition to sharing its chosen action, can augment the action sets of future agents in the sequence by including some elements of its own action sets. This could effectively offset a poor performance due to an agent not having access to any valuable actions. We show bounds on how much this type of information increase could boost performance and present a set of utilities and information sharing rules that are near-optimal in this sense. Finally, we describe how performance is affected when these near-optimal policies are intractable and can only be approximated.

Chapter 7 covers a noncooperative environment with 2 agents: a router (agent 1) and an attacker (agent 2). The system designer cannot affect the utility function of the attacker, only the router. The router is able to route some amount of traffic through a network of links of limited capacity - each possible route is an action. The attacker has a budget of traffic r^a that it can use to flood traffic on some subset of links - each such policy is an action. The information sharing constraints are Type 2 in that the router is only aware of r^a, i.e. $m_2 = r^a$ and the attacker is aware of the router's chosen action, i.e. $m_1 = x_1$. Therefore, the equilibrium takes the form of the Stackelberg Equilibrium. The system designer's objective is to minimize the amount of traffic blocked, and it endows the router with its utility function accordingly. We show that in this scenario, finding the equilibrium policy is NP-Hard for both agents. However, for small networks, we give

an analytical expression for the equilibrium and show how it is affected when the router only knows that r^a is on some interval, rather than its exact value. Finally, we show that if $\mathcal{N}_1 = \emptyset$, i.e., the router knows nothing of the attacker's capabilities, there exists a tractable routing policy that offers robust guarantees against any value of r^a.

Chapter 2

Technical Preliminaries

In this chapter, we will further discuss technical background that will be relevant to many of the chapters in this work. Where specific sections deviate from the model presented in Section 1.1 or the details presented in this chapter, it will be noted at the beginning of the section.

2.1 Submodular Objective Functions

The focus of Chapters 3–6 is on system designer objective functions that are *submodular*. The optimization of such functions is a well-studied topic due to its wide application space. Examples include sensor placement [17], maximizing and inferring influence in a social network [18, 19], image segmentation in image processing [20], multiple object detection in computer vision [21], document summarization [22], path planning of multiple robots [23], sensor placement [24], outbreak detection in networks [26], clustering [27], assigning satellites to targets [28], path planning for multiple robots [23], and leader selection and resource allocation in multiagent systems [29, 25]. The key thread in these problems is that each exhibits some form of a "diminishing returns" property, e.g., adding

more sensors to a sensor placement problem improves performance, but every additional sensor marginally contributes less to the overall performance as the number of sensors increases. Any problem exhibiting such behavior can likely be formulated as a submodular optimization problem.

While polynomial algorithms exist to solve submodular minimization problems, [30, 31, 32], maximization has been shown to be NP-Hard for important subclasses of submodular functions [33]. Therefore, even if the system designer could centrally assign actions to agents, it would still not be computationally feasible to guarantee the optimal set of actions.

Consider a system designer objective function of the form $f(x) = g(\cup_i x_i)$, where $g : 2^S \to \mathbb{R}$ is:

1. *submodular*: $g(A \cup \{s\}) - g(A) \geq g(B \cup \{s\}) - g(B)$ for all $A \subseteq B \subseteq S$ and $s \in S \setminus B$,
2. *monotone*: $g(A) \leq g(B)$ for all $A \subseteq B \subseteq S$,
3. *normalized*: $g(\emptyset) = 0$.

For simplicity, we refer to f as being submodular, monotone, and normalized, without referencing g. In doing so, we also abuse notation so that the input to f need not be an action profile: rather, it can be the actions of a subset of agents, or it could be multiple action profiles. For instance, $f(x_i, x_j, x_k) = g(x_i \cup x_j \cup x_k)$ and $f(x^{\text{opt}}, x^{\text{eq}}) = g\left(x_1^{\text{opt}} \cdots \cup x_n^{\text{opt}} \cup x_1^{\text{eq}} \cup \cdots \cup x_n^{\text{eq}}\right)$.

Since the submodularity property is a statement about the marginal contribution of an element to a set, and since this property is leveraged heavily thoughout this work, it is convenient to define the marginal contribution of $x_A \in \Pi_{i \in A} X_i$ given $x_B \in \Pi_{i \in B} X_i$, according to f:

$$\Delta(x_A | x_B) := f(x_A, x_B) - f(x_B). \tag{2.1}$$

Note that the following holds for two action profiles $x, x' \in X$:

$$f(x, x') = f(x') + \sum_{i \in N} \Delta(x_i | x_1, \ldots, x_{i-1}, x'). \tag{2.2}$$

2.2 The Price of Anarchy

Define a system to be the tuple $H = (N, X, \{\mathcal{U}_i\}_i, \mathcal{I}, f)$, where $\mathcal{U}_i$ is the utility function U_i for Type 1 constraints or the meta-utility function V_i for Type 2, and where $\mathcal{I}$ represents the information sharing constraints. The set $\mathcal{I} = \{\mathcal{N}_i\}_i$ for Type 1 constraints and $\mathcal{I} = (\{\mathcal{N}_i\}_i, \{M_i\}_i)$ for Type 2. Denote $\mathcal{H}$ to be the set of all such systems. Here we consider equilibria either of the form (1.2) for Type 1 constraints or (1.5) for Type 2 constraints. As mentioned, we assume in this work that such equilibria exist, and in most of our contexts, we can prove that they do. One cannot assume, however, that an equilibrium is unique. Since the system designer in our model cannot control which equilibrium is reached, one way to evaluate a system H is by the value of its worst-case equilibrium action profile compared to the value of the optimal action profile. Likewise, when considering a set of possible utility functions or information sharing constraints, one can evaluate the entire set $\mathcal{H}' \subseteq \mathcal{H}$ by finding the worst ratio between the worst-case equilibrium and the optimal action profile. More formally, we define the *price of anarchy* for a set of systems $\mathcal{H}'$ as

$$\text{PoA}(\mathcal{H}') = \min_{\substack{H \in \mathcal{H}' \\ x^{\text{eq}} \in \text{EQ}(H)}} \frac{f(x^{\text{eq}})}{f(x^{\text{opt}}(H))} \in [0, 1], \tag{2.3}$$

where $\text{EQ}(H)$ is the set of all possible equilibria for system H, and $x^{\text{opt}}(H)$ is the optimal action profile for system H. Note that the closer the price of anarchy is to 1, the more desirable the resulting equilibrium performance.

The price of anarchy has been well-studied in the literature [34, 35, 36], and has been used in many applications, including job scheduling [37], congestion games [38], and auctions [39]. We remark that many of our results endeavor to describe the price of anarchy under various circumstances. The general approach to such proofs is to find a lower bound on the PoA by leveraging the properties of submodular monotone functions and other properties. Then one shows some form of tightness by designing a system or set of systems such that $f(x^{\text{eq}})/f(x^{\text{opt}})$ is close or meets that lower bound.

Chapter 3

Blind and Isolated Agents

In this chapter we begin with the assumption that f is submodular, and that information sharing constraints are Type 1. This class of models offers a multitude of attractive theoretical guarantees. For instance the submodularity of the objective function can be leveraged in combination with a wide variety of utility function designs (yielding the class of so-called *valid utility games*) to ensure that all equilibria are within a factor of 2 of the optimal; i.e., the *price of anarchy* is 1/2 [16]. These types of results have a great deal of synergy with the broader literature on submodular maximization.

Following from the initial successes of this game-theoretic model, recent work has begun to critically investigate the robustness properties of this approach. For instance, it has been shown that in general settings, slight changes to agent utility functions can lead to dramatic changes in the quality of emergent behavior [40], and that faulty or misbehaving agents can easily lead stochastic learning dynamics astray [41]. While a comprehensive measure of robustness for such systems remains elusive, positive results exist as well. In particular, for submodular maximization, it is known that performance guarantees can be quite robust to discrepancies of information availability among the agents when the agents are endowed with the specific *marginal-contribution* utility func-

tion [42, 43]. Specifically, these papers show that the price of anarchy associated with marginal-cost utility functions degrades gracefully as information is denied to agents — this will be the topic of Chapter 5. While attractive, these preliminary positive robustness results are limited in scope as they consider only the specific marginal-contribution agent utility function, despite the fact that this is only one possible choice of utility design and is not optimal in all settings [15, 44].

Accordingly, this chapter initiates a study on the robustness of performance guarantees for the broad class of valid utility games - i.e., the system designer chooses utility functions from within the set of valid utilities. We first study the effects of information sharing constraints of a particular structure: a set of k agents is compromised either by becoming *blind* (unable to observe the action choices of any other agent, but still observable by others) or becoming *isolated* (unable to observe other agents or be observed by other agents). Theorem 3.1 states that the price of anarchy when k agents are compromised is $1/(2+k)$, and that this bound is tight for any combination of blind or isolated agents. This result is significant in at least two dimensions: first, it shows for general valid utility games that, in line with the narrower characterization of earlier work [42, 43], performance guarantees degrade gracefully as information is denied from agents. Second, and perhaps more surprisingly, Theorem 3.1 illustrates that isolation is no worse than blindness. Intuitively, this suggests that if an agent is blind, it might as well be invisible also.

This raises the question: are blindness and isolation equivalent for all forms of utility functions for the agents? Theorem 3.2 answers this in the negative, showing that *if* the non-compromised agents are endowed with the marginal-contribution utility function, isolation has a cost: the price of anarchy resulting when k agents are compromised improves if some of the compromised agents are not isolated. Specifically, if at least 1 of the k compromised agents is blind (but not isolated), then the price of anarchy improves

to $1/(1+k)$. Thus, Theorem 3.2 also demonstrates graceful degradation of performance guarantees, and in addition shows that for some utility function designs, blindness can indeed be strictly better than isolation.

3.1 Chapter Model

As stated, we assume that f is submodular and that the information sharing constraints are Type 1, i.e., they are defined by the sets $\mathcal{N}_i$. One utility function that is of note for this chapter and subsequent chapters is marginal contribution (MC), wherein each agent maximizes its marginal contribution to the objective function f, with respect to the remaining agents. More formally stated:

$$\mathrm{MC}_i(x_i, x_{\mathcal{N}_i}) := f(x_i, x_{\mathcal{N}_i}) - f(x_{\mathcal{N}_i}). \tag{3.1}$$

Of course, this class of problems admits many other utility functions as well. In this work we consider those which satisfy the *valid utility game* assumptions of [16]:

Definition 3.1 *A Valid Utility Game (VUG) is a system with no information sharing constraints that satisfies the following three conditions:*

1. *f is submodular, nondecreasing, and normalized,*[1]
2. *$U_i(x_i, x_{-i}) \geq f(x_i, x_{-i}) - f(x_{-i})$*
3. *$\sum_i U_i(x_i, x_{-i}) \leq f(x_i, x_{-i})$*

Note that when f satisfies 1), MC is one possible choice of utility function that satisfies 2) and 3).

[1] The original defintion used in [16] did not require monotonicity or normalization, yet we impose that as it suits our purposes here.

3.1.1 Compromised Agents

We now describe additional information sharing constraints as compromised agents of various forms. We begin with the assumption that H satisfies Definition 3.1, and then some subset of agents is compromised in a way that fixes the sets $\mathcal{N}_i$. We consider three ways in which an agent can be compromised:

1. *Blind agents*: a blind agent does not know the actions of any other agents, i.e., if agent i is blind, then $\mathcal{N}_i = \emptyset$. We denote the set of blind agents as $B \subseteq N$.

2. *Isolated agents*: an isolated agent is blind, and the other agents are also blind to it. In other words, if agent i is isolated, then $\mathcal{N}_i = \emptyset$, and there is no j such that $i \in \mathcal{N}_j$ We denote the set of isolated agents as $I \subseteq N$.

3. *Disabled agents*: a disabled agent i cannot contribute to system welfare and always selects $x_i = \emptyset$. The remaining agents are unaffected: $U_j = U_j$ for $j \neq i$. We denote the set of disabled agents as $D \subseteq N$.[2]

One way to model $j \notin \mathcal{N}_i$ is to suppose that agent i assumes $x_j = \emptyset$. We denote $K = B \cup I \cup D$ as the set of all compromised agents; i.e., K may contain any combination of blind, isolated, and disabled agents. Note that a blind or isolated agent can still select among its usual actions and its action choice still contributes to the system objective f despite the denied observations. For an example of a VUG with compromised agents, see Figure 3.1.

In this context, the equilibrium in (1.2) takes the form

$$U(x_i^{\text{eq}}, x_{\mathcal{N}_i}^{\text{eq}}) \geq U(x_i, x_{\mathcal{N}_i}^{\text{eq}}),\ \forall\ x_i \in X_i, i \in N \tag{3.2}$$

[2]Note that for the case of blind or isolated agents, if agent i cannot "see" the actions of agent j, agent i effectively assumes that agent j is disabled. It should be noted that assuming disability is merely one possibility and that optimal modeling of unobservable agents is an open area of research [43].

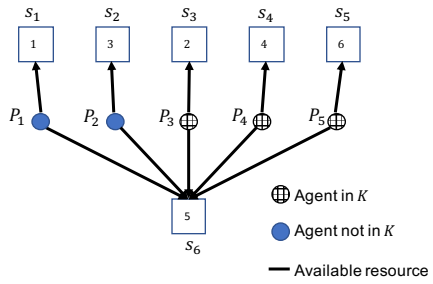

(a) A valid utility game (WSC problem) where some agents have been compromised. The agents are represented by circles, with the black cross-hatch agents being compromised. The action set for agent i is represented by the black lines to boxes, which are resources, i.e., the action set for agent i is $\{\{s_i\}, \{s_6\}\}$. Each agent is endowed with the marginal contribution utility MC_i.

Scenario	x_1^{eq}	x_2^{eq}	x_3^{eq}	x_4^{eq}	x_5^{eq}	$f(x)$
Optimal	s_6	s_2	s_3	s_4	s_5	20
No compromised agents	s_1	s_2	s_3	s_6	s_5	17
P_3, P_4, P_5 are blind	s_1	s_2	s_6	s_6	s_5	15
P_3, P_4, P_5 are isolated	s_1	s_6	s_6	s_6	s_5	12
P_3, P_4, P_5 are disabled	s_1	s_6	$\emptyset$	$\emptyset$	$\emptyset$	6

(b) Worst-case equilibria (and corresponding evaluations) for five scenarios. The first two assume that no agents have been compromised - note that the optimal allocation is also an equilibrium. The final three assume that agents $3, 4, 5$ have been compromised in the same way.

Figure 3.1: A valid utility game (WSC problem) where some agents have been compromised. We see that making agents $3, 4, 5$ blind causes agents 3 and 4 to both choose s_6. When $3, 4, 5$ are blind and $1, 2$ cannot see their actions (i.e., $3, 4, 5$ are isolated), this additionally causes agent 2 to choose s_6. Finally, when $3, 4, 5$ are disabled, we see that they no longer contribute to the welfare of the system. While not every problem instance would degrade in this manner, we shall see in the results of this paper that this example is indicative of worst-case behavior.

In any VUG, such an equilibrium is not guaranteed to exist [16], although a mixed equilibrium will [45]. Here we operate under the assumption that the equilibrium exists, and leave the precise characterization of when this is true for future work. Again, we focus on the solution concept of equilibrium so as to abstract away the mechanics of specific learning rules and algorithms. Accordingly, we measure the effectiveness of a given utility design by the price of anarchy.

3.2 Effects of Compromised Agents

In this section, we demonstrate guarantees about the price of anarchy in a complex system when one or more agents in the system have been compromised. The effect this has on the price of anarchy depends somewhat on whether the agents in the compromised

set K are blind, isolated, or disabled.

3.2.1 Results

Theorem 3.1 *Let $\mathcal{H}_{\text{VUG}}(k) \subseteq \mathcal{H}$ be the set of valid utility games satisfying Definition 3.1, where agents in $K \subseteq N$ are subsequently compromised with $|K| \leq k$. If at least one agent is disabled, then* $\text{PoA}(\mathcal{H}_{\text{VUG}}(k)) = 0$. *Otherwise,*

$$\text{PoA}(\mathcal{H}_{\text{VUG}}(k)) = \max(1/(2+k), 1/n) \tag{3.3}$$

Before giving the formal proof, we give a brief overview and some discussion of the significance of this result. It should be clear that having a disabled agent can be arbitrarily bad, thus the PoA of 0 should not be surprising. In order to show the remaining cases, we leverage the properties in Definition 3.1 and the definitions of blind and isolated agents to give a lower bound on $\text{PoA}(\mathcal{H}_{\text{VUG}}(k))$. We then consider a subclass of VUGs where agents are endowed with a Shapley value utility function [46] as an example to show that the lower bound is tight.

Perhaps unintuitively, blind agents and isolated agents affect the PoA in the same way; the information provided to the uncompromised agents by the actions of the blind agents has no effect. The key deterrent to the PoA is that the compromised agents do not consider the actions of others, not that others cannot see the actions of the compromised agents.

Another way to look at this is to think about a directed graph (V, E), where each node represents an agent and an edge (i, j) in the graph means that agent j's utility function depends on the action of agent i. In a general sense, one might expect that the more "connected" the graph, the better the resulting PoA. Under the nominal setting, where

no agent is compromised, the graph is complete, and we have $\text{PoA}(\mathcal{H}_{\text{VUG}}(0)) = 1/2$. When a single agent i becomes blind, every edge (j, i) for $j \neq i$ is removed from the graph, but all edges (i, j) remain. According to Theorem 3.1, this results in a decrease in the PoA to $\text{PoA}(\mathcal{H}_{\text{VUG}}(1)) = 1/3$. If agent i becomes isolated, this further removes all edges (i, j) from the graph. However, Theorem 3.1 shows that the price of anarchy is unchanged at $\text{PoA}(\mathcal{H}_{\text{VUG}}(1)) = 1/3$.

Proof: First we establish the case where agent $i \in K$ is disabled. Then one could construct an example with f and $x_i \in X_i$ such that $f(x_i, x_{-i})$ is arbitrarily large and $f(x_{-i}) = 0$ for any x_{-i}. Since agent i is forced to choose $\emptyset$, we see that $\text{PoA}(\mathcal{H}_{\text{VUG}}(k)) = 0$.

For the remainder of the proof we assume that all agents in K are either blind or isolated. In the first case we assume that $|K| < n - 1$. We show through the properties of Definition 3.1 and the definition of blind and isolated agents that $1/(2+|K|)$ is a lower bound on $\text{PoA}(\mathcal{H}_{\text{VUG}}(k))$. Then we show that the bound is tight by choosing a particular welfare function f and utility profile U such that $f(x^{\text{eq}})/f(x^{\text{opt}}) = 1/(2 + |K|)$.

Denote $x_{j<i}$ to mean $x_1, \ldots x_{i-1}$, and denote $P_i = N \setminus (I \cup D \cup \{i\})$, i.e., P_i is the set

of agents whose actions agent $i \notin K$ "observes". Then we see that

$$f(x^{\text{opt}}) \leq f(x^{\text{opt}}, x^{\text{eq}}) \tag{3.4}$$

$$\leq f(x^{\text{eq}}) + \sum_i f(x_i^{\text{opt}}, x_{j<i}^{\text{opt}}, x^{\text{eq}}) - f(x_{j<i}^{\text{opt}}, x^{\text{eq}}) \tag{3.5}$$

$$\leq f(x^{\text{eq}}) + \sum_i f(x_i^{\text{opt}}, x_{P_i}^{\text{eq}}) - f(x_{P_i}^{\text{eq}}) \tag{3.6}$$

$$\leq f(x^{\text{eq}}) + \sum_{i \notin K} U_i(x_i^{\text{opt}}, x_{P_i}^{\text{eq}}) + \sum_{i \in K} f(x_i^{\text{opt}}) \tag{3.7}$$

$$\leq f(x^{\text{eq}}) + \sum_{i \notin K} U_i(x_i^{\text{eq}}, x_{P_i}^{\text{eq}}) + \sum_{i \in K} U_i(x_i^{\text{opt}}) \tag{3.8}$$

$$\leq f(x^{\text{eq}}) + f(x^{\text{eq}}) + \sum_{i \in K} U_i(x_i^{\text{eq}}) \tag{3.9}$$

$$\leq f(x^{\text{eq}}) + f(x^{\text{eq}}) + \sum_{i \in K} f(x_i^{\text{eq}}) \tag{3.10}$$

$$\leq (2 + |K|) f(x^{\text{eq}}), \tag{3.11}$$

where (3.4) is true since f is nondecreasing; (3.5) is true via (2.2); (3.6) is true by submodularity of f; (3.7) holds since the original U_i satisfy 2) in Definition 3.1 (2nd term), and by submodularity of f (3rd term); (3.8) is true by defintion of equilibrium (2nd term) and by the utlities of the blind and isolated agents (3rd term); (3.9) is true since all U_i satisfy 3) in Definition 3.1 (2nd term) and by defintion of equilibrium (3rd term); (3.10) is true by the defintion of U_i for agents in K; and (3.11) is true since f is nondecreasing.

To see the upper bound, consider a scenario where f is of the form

$$f(x) = \sum_{s \in S} f_s(|x|_s), \tag{3.12}$$

where $|x|_s$ denotes the number of agents which have selected element s under action profile

x. The functions $f_s : \{1, \ldots, N\} \to \mathbb{R}$ are nonnegative, i.e., $f_s(i) \geq 0$; nondecreasing, i.e., $f_s(i+1) \geq f_s(i)$; and have decreasing marginal returns, i.e., $f_s(i+1) - f_s(i) \geq f_s(i+2) - f_s(i+1)$. When f has this form, this represents a well-studied set of games called distributed resource allocation games (see, for instance [47]).

We also assume that, before agents are compromised, all are endowed with the equal share (ES) utility, wherein each agent chooses an action according to the following:

$$\mathrm{ES}_i(x_i, x_{-i}) = \sum_{s \in S(x_i)} \frac{1}{|x|_s} f_s(|x|_s), \tag{3.13}$$

where $S(x_i)$ is the set of elements in action x_i. Essentially, when multiple agents choose the same resource, the agents divide the utility $f_s(|x|_s)$ equally. We note that the ES utility is an instance of the more general Shapley value utility, also a subject of much study within the literature.

Based on the construction of f_s and therefore f, it should be clear that f satisfies 1) in Definition 3.1. Likewise, it should be immediately clear that when ES is employed, 3) in Definition 3.1 is satisfied with equality. We can also see that 2) is satisfied since

$$f(x) - f(\emptyset, x_{-i}) = \sum_{s \in S(x_i)} \frac{f_s(|x|_s)}{|x|_s} - \frac{f_s(|x_{-i}|_s)}{|x_{-i}|_s}, \tag{3.14}$$

$$\leq \sum_{s \in S(x_i)} \frac{f_s(|x|_s)}{|x|_s} = \mathrm{ES}_i(x). \tag{3.15}$$

Therefore, since f satisfies (3.12), the system is a VUG.

Assume that the example shown in Figure 3.2 is such a game, where now a subset of agents K are all blind. The blind agents are represented by the black cross-hatch circles, and the other agents as blue circles. Each agent has access to its own resource (the box closest to it) and a central resource. The value v_s of resource r is the value listed in the

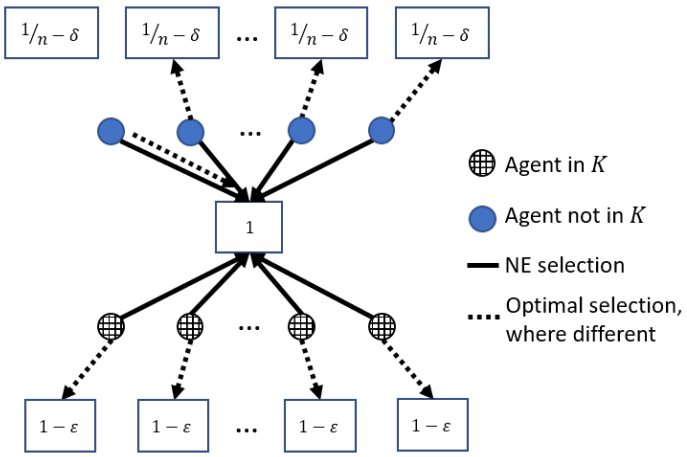

Figure 3.2: An example game used in the proof for Theorem 3.1, a WSC problem where now a subset of agents K are all blind. Each agent has access to its own resource (the box closest to it) and a central resource. The value v_s of resource r is the value listed in the box, where $\varepsilon, \delta > 0$ are small, and $f(x) = \sum_{s \in S(x)} v_s$. We see that the agents in K will all choose the central resource, since they all act independently. Agents not in K are endowed with ES, and therefore are also incentivized to choose the central resource. This implies $f(x^{\text{eq}}) = 1$ and as $\varepsilon, \delta \to 0$ and $n \to \infty$, $f(x^{\text{opt}}) \to 2 + |K|$.

box, where $\varepsilon, \delta > 0$ are small. Then $f_s = v_s$, i.e., $f(x) = \sum_{s \in S(x)} v_s$. We see that the agents in K will all choose the central resource, since they all act independently. We also see that in any equilibrium, all agents not in K are also incentivized to choose the central resource, since they are endowed with ES_i. Therefore, $f(x^{\text{eq}}) = 1$. The optimal allocation is for one agent not in K to choose the central resource and the remaining agents to choose their alternates, implying that $f(x^{\text{opt}}) = 1 + (n - |K| - 1)(1/n - \delta) + |K|(1 - \varepsilon)$. As $\varepsilon, \delta \to 0$, and $n \to \infty$, we see that

$$\frac{f(x^{\text{eq}})}{f(x^{\text{opt}})} \to \frac{1}{2 + |K|}. \tag{3.16}$$

■

3.2.2 Marginal Contribution Utility

In this section, we consider the use of the marginal contribution utility and whether using this specific utility function design can offset the decrease in PoA that we saw in Theorem 3.1. We show that marginal contribution utility can give a higher PoA in the presence of blind agents.

Theorem 3.2 *Let* $\mathcal{H}_{\mathrm{MC}}(k) \subseteq \mathcal{H}_{\mathrm{VUG}}(k)$ *be the of systems which where* f *is submodular and which leverage* MC_i *for all agents, where agents* $K \subseteq N$ *are compromised such that* $|K| \leq k$*. If one agent is disabled, then* $\mathrm{PoA}(\mathcal{H}_{\mathrm{MC}}(k)) = 0$*. Otherwise*

$$\mathrm{PoA}(\mathcal{H}_{\mathrm{MC}}(k)) = \begin{cases} \frac{1}{1+|K|}, & \text{if } |B| > 0, \\ \max\left(\frac{1}{2+|K|}, \frac{1}{n}\right), & \text{if } |B| = 0. \end{cases} \tag{3.17}$$

The proof is found in Appendix A.1.

3.3 Simulation

In this section we present empirical findings from the results of running a simulation of stochastic learning dynamics applied to a VUG in which agents are endowed with a marginal contribution utility function. We simulate the popular *log-linear learning* dynamics [48, 49] to validate the results and explore the effect of "noisy" behavior on these low-quality equilibria. Log-linear learning operates in discrete steps at times $t = 0, 1, \ldots,$ producing a sequence of joint actions $x(0), x(1), \ldots$. We assume agents begin with an arbitrary joint action $x(0) \in X$, and let $x(t) = (x_i, x_{-i}) \in X$. At time $t \in \mathbb{N}$, agent $i \in N$ is selected uniformly at random to update its action for time $t + 1$; all other agents' actions will remain fixed for time $t + 1$. At time $t + 1$, agent i chooses action $x_i \in X_i$

with probability

$$p_i^{x_i}(t+1) = \frac{e^{U_i(x_i, x_{-i}(t))/T}}{\sum_{\tilde{a}_i \in X_i} e^{U_i(\tilde{a}_i, x_{-i}(t))/T}}. \tag{3.18}$$

"Temperature" parameter $T > 0$ dictates an updating agent's degree of rationality and is identical for all agents $i \in N$. As $T \to 0$, agents are increasingly likely to select utility-maximizing actions, and as $T \to \infty$, agents tend to choose their next actions uniformly at random. The joint action at time $t+1$ is $x(t+1) = (x_i(t+1), x_{-i}(t))$.

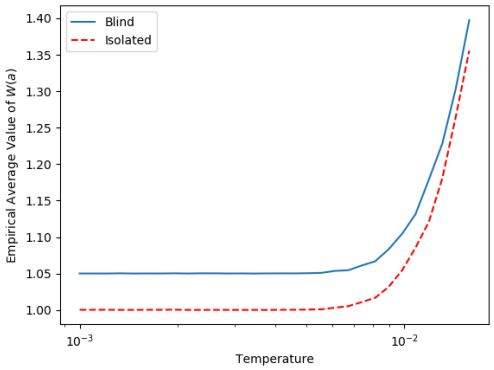

(a) Simulations with temperatures ranging from 0.001 to 0.016.

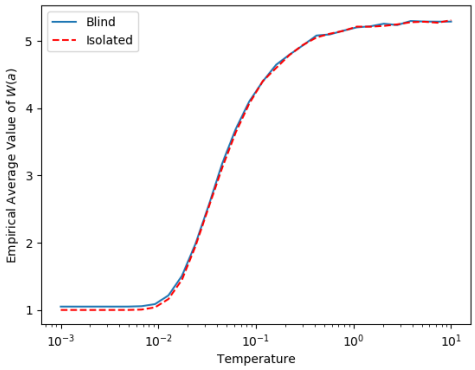

(b) Simulations with temperatures ranging from 0.001 to 10.

Figure 3.3: Plots representing simulation results of the game shown in Figure 3.2, with $n = 10$, $|K| = 9$, with the exceptions that the agent not in K has the MC utility and its alternate resource has value $\varepsilon = 0.05$. For each trace, log-linear learning is run for 200,000 iterations for each temperature value. The solid blue trace corresponds to all agents in K being blind, and the dashed red trace corresponds to all agents in K being isolated. Note that for very low temperatures (effectively the agents are playing asynchronous best-response dynamics), blindness has a slight advantage over isolation in accordance with Theorem 3.2.

We run log-linear learning on the game depicted in Figure 3.2 with the following exceptions: there is only 1 agent not in K, it is endowed with the MC utility, and its "alternate" resource option has value ε. We use $n = 10$, $|K| = 9$, and $\varepsilon = 0.05$, and for each value of T, we report the average value of $f(x)$ for 200,000 trials. The difference in

a game in which all the compromised agents are isolated versus one in which at least one of these agents is blind can be seen in Figure 3.3a. The compromised agents in K are all blind in the first simulation (solid blue), and all are isolated in the second (dashed red).

For this game, since $\varepsilon = 0.05$, the optimal selection of resources yields a value of the welfare function of 9.55. The equilibrium yields a value of 1 when all agents in K are isolated, and 1.05 when at least one agent in K is blind. In the simulation, the game in which all the agents were blind had a minimum average value of the welfare function of 1.050015, hence the price of anarchy is 0.10995. The game in which all the agents were isolated had an minimum average value of welfare function of 1.000034, giving a price of anarchy of 0.1047. These values of the price of anarchy differ slightly from those given by Theorem 3.2 since $\varepsilon \neq 0$. As temperature increases, the instability of the equilibrium becomes apparent, and the average of the welfare function increases with an increase in temperature until this value is indistinguishable from that produced by a purely random strategy as seen in Figure 3.3b.

An intriguing aspect of this example is that the equilibrium, representing a worst-possible equilibrium in the class of games $\mathrm{PoA}(\mathcal{H}_{\mathrm{MC}}(9))$, is actually among the worst action profiles in the game. Hence, a large value of T (i.e., agents selecting actions uniformly at random) results in play that is of far higher quality than the equilibrium.

Chapter 4

Graph Constraints

This chapter operates under much the same setting as Chapter 3: we consider the set of systems which are valid utility games and are subsequently attacked or otherwise constrained on information sharing. The set of constraints, however, is slightly more general. Rather than an agent being blind to all other agents, this chapter examines the scenario where an agent can be blind to some subset of other agents. In other words, agent i has some subset of agents $\mathcal{N}_i$ whose messages it observes in equilibrium.

The sets $\mathcal{N}_i$ define a directed graph $G = (V, E)$, where $V = N$ and $(j, i) \in E$ if $j \in \mathcal{N}_i$. We refer to such a graph as an *information sharing constraint graph*, as it effectively represents the information sharing constraints on the system. Given a graph G, two natural questions arise:

1. How does the structure of G affect the price of anarchy of the class of VUGs subject to G?

2. What utility functions can provide a better price of anarchy?

Theorem 4.1 addresses the first question: essentially, the price of anarchy is tied to the number of groups of agents who have access to the same information - for instance,

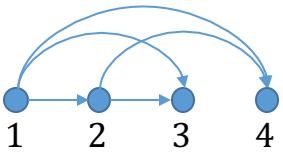

(a) In this graph, there are 4 cliques of size 1 (one for each node), 5 cliques of size 2 (one representing each edge), and 2 cliques of size 3 (the sets $\{1,2,3\}$ and $\{1,2,4\}$). Thus $\omega(G) = 3$. A minimum clique cover is $\{1,3\},\{2,4\}$, so $k(G) = 2$. The maximum independent set is $\{3,4\}$, thus $\alpha(G) = 2$. Since $\alpha(G) = k(G) = 2$, we know that $\alpha^*(G) = k^*(G) = 2$. Appendix A.3 shows that the price of anarchy is $1/2$, making it a graph that meets the upper bound for Theorem 5.1 (see Section 5.3.1). Lastly, it is also an example of a graph without the Sibling Property (see Section 5.4.3), since no such w exists from Definition 5.1.

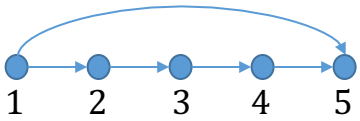

(b) A graph where $\alpha(G) = 2$, $k(G) = 3$, $\alpha^*(G) = k^*(G) = 2.5$, and $z = [1/2, 1/2, 1/2, 1/2, 1/2]^T$ maximizes (4.3). As a note, this is the graph with the fewest number of nodes and edges such that $\alpha(G) \neq k(G)$. This is also a graph with the Sibling Property (see Section 5.4.3), since for maximum independent set $\{2,4\}$, $2 \in \mathcal{N}_3$, so $w = 3$ from Definition 5.1.

Figure 4.1: Two example graphs showcasing the graph properties defined in Section 4.1. These graphs will be referred to throughout this to illustrate the tightness of bounds in Theorem 5.1 and to illustrate the Sibling Property (see Section 5.4.3).

in the complete graph, this number is 1. As edges are removed, this number increases, and it is shown that strategically removing $n-1$ edges from the complete graph degrades the price of anarchy to $1/(n+1)$, arbitrarily bad for large systems.

Theorem 4.2 addresses the second question by introducing a subset of VUGs, where the utility functions must all satisfy a certain consistency property. Among this smaller set of systems, a lower bound on the price of anarchy is instead tied to the largest set of nodes in G, among whom there are no reciprocal edges. This lower bound is at least as high as the value of the price of anarchy for the general VUG setting, and is strictly better for most graphs. Additionally, Proposition 4.1 provides an upper bound on the price of anarchy for any fixed utility function. This bound is also given in terms of the largest set of agents in G among whom there are not any edges.

4.1 Review of Graph Theory Terms

For all definitions in this work, we assume that $G = (V, E)$ is a simple directed graph. We begin with cliques:[1] a *clique* is a set of nodes $C \subseteq V$ such that for every $i, j \in C$, either $(i, j) \in E$ or $(j, i) \in E$. The *clique number* $\omega(G)$ is the number of nodes in the largest clique in G. We denote by $K(G)$ the set of all cliques in G. A *clique cover* is a partition on V such that the nodes in each set of the partition form a clique. The *clique cover number* $k(G)$ is the minimum number of sets needed to form a clique cover of G. For an example, see Figure 4.1a.

Another important notion in graph theory is that of independence. An *independent set* $J \subseteq V$ is a set of vertices such that $v_1, v_2 \in J$ implies $(v_1, v_2), (v_2, v_1) \notin E$. A *maximum independent set* is an independent set of G such that no other independent set has more vertices. The *independence number* $\alpha(G)$ is the number of nodes in the largest independent set in G. For an example, see Figure 4.1a.

The work in [50] equivalently characterizes the independence number as the solution to an integer linear program [2]. Let $Q \in \mathbb{R}^{|K(G)| \times n}$ be the binary matrix whose rows are indicator vectors for the cliques in G. In other words, $Q_{ij} = 1$ if node j belongs to clique i in G, and 0 otherwise. Note that Q also includes cliques of size 1 (the individual nodes). Then $\alpha(G)$ is given by

$$\begin{aligned} \max_{z} \quad & z^T \mathbf{1} \\ \text{subject to} \quad & Qz \leq \mathbf{1} \\ & z \in \mathbb{Z}^n \geq \mathbf{0}. \end{aligned} \tag{4.1}$$

It is similarly shown that $k(G)$ is characterized by the dual to this problem, implying

[1] The terms *clique* and *independence set* are traditionally defined only for undirected graphs, however, we adapt those terms for our purposes here.

[2] It is actually the chromatic number and clique number that are defined this way in [50]. However, using graph complementarity, it is an easy extension to show that the solution to the linear program in (4.1) yields a maximum independent set.

that $\alpha(G) \leq k(G)$. As an example, for the graph in Figure 4.1a,

$$Q = \begin{array}{c} \begin{matrix} \text{Node 1} & \text{Node 2} & \text{Node 3} & \text{Node 4} \end{matrix} \\ \left[\begin{array}{cccc} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \\ 1 & 1 & 0 & 1 \end{array}\right] \begin{matrix} \{1\} \\ \{2\} \\ \{3\} \\ \{4\} \\ \{1,2\} \\ \{1,3\} \\ \{1,4\} \\ \{2,3\} \\ \{2,4\} \\ \{1,2,3\} \\ \{1,2,4\} \end{matrix} \end{array} \tag{4.2}$$

Using this Q in (4.1), it is straightforward to show that the optimal solution is $z = [0, 0, 1, 1]^T$, i.e., $\alpha(G) = 2$, and the maximum independent set is $\{3, 4\}$.

Note by defintion that $\alpha(G)$ and $k(G)$ are always positive integers. However, in many applications, it is helpful to consider a real-valued relaxation on these notions: this is the motivation for fractional graph theory [50]. Here we leverage the *fractional independence*

number $\alpha^*(G)$, which we define as the real-valued relaxation to (4.1): [3]

$$\begin{aligned} \alpha^*(G) := \max_{z} \quad & z^T \mathbf{1} \\ \text{subject to} \quad & Qz \leq \mathbf{1} \\ & z \geq \mathbf{0}. \end{aligned} \tag{4.3}$$

Likewise, $k^*(G)$, the *fractional clique cover number* of G, can be defined by its dual

$$\begin{aligned} k^*(G) := \min_{y} \quad & y^T \mathbf{1} \\ \text{subject to} \quad & Q^T y \geq \mathbf{1} \\ & y \geq \mathbf{0}. \end{aligned} \tag{4.4}$$

In accordance with the Strong Duality of Linear Programming [52], it follows that:

$$\alpha(G) \leq \alpha^*(G) = k^*(G) \leq k(G). \tag{4.5}$$

An example of a graph where the independence number differs from the fractional independence number is found in Figure 4.1b.

In this work, we introduce the notion of an *information group*: a set of nodes which is fully connected, and which have the same incoming neighbors. Formally stated, $T \subseteq N$ is an information group of G if for all $i, j \in T$, $\mathcal{N}_i \cup \{i\} = \mathcal{N}_j \cup \{j\}$. An alternate definition is that if $A(G)$ is the adjacency matrix of G, then all rows of $A(G) + I$ associated with the nodes in T are the same. A *maximal information group* is an information group that is not a subset of any other information group. We denote the set of all maximal information groups of G as $\mathcal{T}(G)$, which is both unique and a partition on the nodes of

[3] Another defintion of fractional independence exists in the literature (see [51]), which was created to preserve certain properties of graph independence (such as nested maximality), but has not been shown to preserve $\alpha^*(G) = \omega^*(\bar{G})$, where $\bar{G}$ is the complement graph of G and $\omega^*(G)$ is the fractional clique number of G.

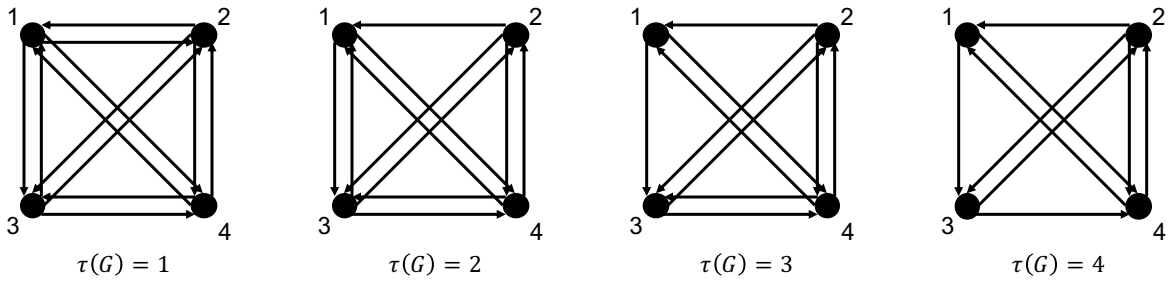

Figure 4.2: An illustration of information groups. The first graph is a complete graph, thus all nodes are in the only information group. The second graph has edge (1, 2) removed, thus $T(G) = (\{1, 3, 4\}, \{2\})$, since 2 is the only node without an incoming edge from 1. The third graph further has edge (3, 1) removed, and $T(G) = \{3, 4\}, \{1\}, \{2\}$. Finally the last graph has edge (4, 3) removed, and each node is its own information group.

G. Denote $\tau(G) = |\mathcal{T}(G)|$. See Figure 4.2 for an example.

4.2 Valid Utility Games with Graph Constraints

In this section, we endeavor to show how the structural properties of information sharing constraint graph G are related to the set of possible resulting equilibria. As in the previous chapter, we make this comparison in terms of the price of anarchy. Assume that the system designer uses a set of utility functions that satisfy Definition 3.1, i.e., the system is a valid utility game. Then we define $\mathcal{H}_{\text{VUG}}(G) \subseteq \mathcal{H}$ to be the set of all systems which are valid utility games that are subsequently subjected to the information sharing constraints graph G.

Theorem 4.1 *For any G where G has at least one edge,*

$$\text{PoA}(\mathcal{H}_{\text{VUG}}(G)) = \frac{1}{1 + \tau(G)} \geq \frac{1}{n+1}. \tag{4.6}$$

The full proof is given in Appendix A.2, but here we give an overview. Submodularity and the properties from Definition 3.1 are used to establish that $\text{PoA}(\mathcal{H}_{\text{VUG}}(G)) \geq \frac{1}{1+\tau(G)}$.

Tightness is shown by example, where one carefully constructs a problem instance and set of utility functions that satisfy Definition 3.1 so that agents are incentivized into making poor decisions.

This result effectively generalizes the results from Chapter 3, since each compromised agent is its own information group, and the remaining agents form another. While these results were somewhat positive: compromising one agent only increases the denominator for the PoA by one, Theorem 4.1 effectively shows that the general class of valid utility games is not robust against these types of information constraints. For instance, consider the example set forth in Figure 4.2. In the leftmost graph, which is a complete graph, we see that $\text{PoA}(\mathcal{H}_{\text{VUG}}(G)) = 1/2$, recovering the well-known result from [16]. However, the rightmost graph only has 3 edges removed, yet $\tau(G) = 4$, and $\text{PoA}(\mathcal{H}_{\text{VUG}}(G)) = 1/5$. In fact, for any number of agents this example is instructive: there exist $n-1$ edges that can be removed form the complete graph such that $\text{PoA}(\mathcal{H}_{\text{VUG}}(G))$ moves from $1/2$ to $1/(n+1)$ — arbitrarily bad. For large systems, this implies that the system designer cannot be content to simply choose utilities that satisfy Definition 3.1.

To this end, we introduce an additional utility function property called *consistency*. A utility function is consistent if

$$U_i(x_i, x_A) \geq U_i(x_i, x_B), \tag{4.7}$$

for all $A \subseteq B \subseteq N \setminus \{i\}$, for all $x_i \in X_i, x_A \in \Pi_{j\in A} x_i, x_B \in \Pi_{j\in B} x_j$, and for all $i \in N$. Here the sets A and B represent possible choices of incoming neighbors; the consistency property simply states that an agent's preference for any action decreases as the set of agents that it can observe grows. Many common choices of utility functions, including the marginal contribution utility function MC_i, satisfy this property. Considering now valid utility functions that are also consistent, we have the following result.

Theorem 4.2 *For any graph G,*

$$\mathrm{PoA}(\mathcal{H}_c(G)) \geq \frac{1}{1+\alpha^*(\bar{G})}, \tag{4.8}$$

where $\mathcal{H}_c(G) \subseteq \mathcal{H}_{\mathrm{VUG}}(G)$ is the set of all consistent valid utility games which are subsequently subjected to G, and $\bar{G}$ is the subgraph of G such that any "non-reciprocal" edges from G are removed, i.e. if $\bar{G} = (\bar{V}, \bar{E})$, then $\bar{V} = V$, and $(i,j) \in \bar{E}$ iff $(i,j),(j,i) \in E$.

Proof: Begin with

$$\begin{aligned}
f(x^{\mathrm{opt}}) \leq & f(x^{\mathrm{eq}}) + \Delta(x^{\mathrm{opt}}|x^{\mathrm{eq}}), && (4.9)\\
= & f(x^{\mathrm{eq}}) + \sum_i \Delta(x_i^{\mathrm{opt}}|x_{1:i-1}^{\mathrm{opt}}, x^{\mathrm{eq}}), && (4.10)\\
\leq & f(x^{\mathrm{eq}}) + \sum_i \Delta(x_i^{\mathrm{opt}}|x_{\mathcal{N}_i}^{\mathrm{eq}}), && (4.11)\\
\leq & f(x^{\mathrm{eq}}) + \sum_i U_i(x_i^{\mathrm{opt}}, x_{\mathcal{N}_i}^{\mathrm{eq}}), && (4.12)\\
\leq & f(x^{\mathrm{eq}}) + \sum_i U_i(x_i^{\mathrm{eq}}, x_{\mathcal{N}_i}^{\mathrm{eq}}), && (4.13)
\end{aligned}$$

where (4.9) and (4.11) are true by submodularity of f, (4.12) is true from 2) of Definition 3.1, and (4.13) is true by definition of equilibrium. Now suppose that we have a set

of scalars $\{y_k\}_{k\in K(\bar{G})}$, such that $y_k \geq 0$ and $\sum_{k:i\in k} y_k \geq 1$ for all i. Then

$$\begin{aligned}
\sum_i U_i(x_i^{\mathrm{eq}}, x_{\mathcal{N}_i}^{\mathrm{eq}}) \leq & \sum_i U_i(x_i^{\mathrm{eq}}, x_{\mathcal{N}_i}^{\mathrm{eq}}) \left(\sum_{k\in K(\bar{G}):i\in k} y_k \right) && (4.14)\\
= & \sum_i \sum_{k\in K(\bar{G}):i\in k} y_k U_i(x_i^{\mathrm{eq}}, x_{\mathcal{N}_i}^{\mathrm{eq}}) && (4.15)\\
= & \sum_{k\in K(\bar{G})} \sum_{i\in k} y_k U_i(x_i^{\mathrm{eq}}, x_{\mathcal{N}_i}^{\mathrm{eq}}) && (4.16)\\
\leq & \sum_{k\in K(\bar{G})} \sum_{i\in k} y_k U_i(x_i^{\mathrm{eq}}, x_{k\setminus\{i\}}^{\mathrm{eq}}) && (4.17)\\
\leq & \sum_{k\in K(\bar{G})} y_k \sum_{i\in k} U_i(x_i^{\mathrm{eq}}, x_{k\setminus\{i\}}^{\mathrm{eq}}) && (4.18)\\
\leq & \sum_{k\in K(\bar{G})} y_k f(x_k^{\mathrm{eq}}) && (4.19)\\
\leq & f(x^{\mathrm{eq}}) \sum_{k\in K(\bar{G})} y_k, && (4.20)
\end{aligned}$$

where (4.17) is true by the consistency property, (4.18) is true by 3) of Definition 3.1, and (4.19) is true by the monotonicity of f. Combining this with (4.13) yields

$$\frac{f(x^{\mathrm{eq}})}{f(x^{\mathrm{opt}})} \geq \frac{1}{1+\sum_{k\in K(\bar{G})} y_k}. \tag{4.21}$$

The choice of $\{y_k\}_{k\in K(\bar{G})}$ that minimizes $\sum_{k\in K(\bar{G})} y_k$ will therefore give the highest lower bound. One can formulate this as

$$\begin{aligned}
\min_{\{y_k\}_{k\in K(\bar{G})}} \quad & \textstyle\sum_{k\in K(\bar{G})} y_k \\
\text{subject to} \quad & \textstyle\sum_{k:i\in k} y_k \geq 1, \text{ for all } i \\
& y_k \geq 0, \text{ for all } k.
\end{aligned} \tag{4.22}$$

This is equivalent to the formulation of $k^*(\bar{G})$ in (4.4). Since $k^*(\bar{G}) = \alpha^*(\bar{G})$, this

completes the proof. ■

The consistency property allows one to make a stronger guarantee about the set of resulting equilibria. For instance, consider again the example in Figure 4.2. Of course, the complete graph on the left is such that $\alpha^*(\bar{G}) = \alpha^*(G) = \alpha(G) = 1$. Therefore, Theorem 4.2 gives the same bound as the more general case: that $\text{PoA}(\mathcal{H}_c(G)) \geq 1/2$. The rightmost graph G is such that $\bar{G}$ is a line graph: edges (2, 1), (1, 3), and (3, 4) are removed since they have no reciprocal. Here $\alpha^*(\bar{G}) = \alpha(\bar{(}G)) = 2$, ensuring that $\text{PoA}(\mathcal{H}_c(G)) \geq 1/3$, compared to 1/5 for the more general case. In fact, its trivial to show that $\tau(G) \geq \alpha(\bar{G})$ for any G, since if 2 nodes are in the same information group in G, they cannot be independent in $\bar{G}$. Therefore, the system designer is better off (in terms of equilibrium guarantees) implementing consistent utilities within the valid utility framework.

4.3 A Bound on Optimal Utilities

In this section, we relax the assumption that the system is a valid utility game. Instead, we consider the class of all admissible utility functions, and we show an upper bound on the price of anarchy given the information sharing constraint graph G.

Proposition 4.1 *For any admissible utility function profile $U = (U_1, \ldots, U_n)$ and any graph G,*

$$\text{PoA}(\mathcal{H}_U(G)) \leq \frac{1}{\alpha(G)}, \tag{4.23}$$

where $\mathcal{H}_U(G) \subseteq \mathcal{H}$ is the set of systems that employ U and are subject to the information sharing constraint graph G.

Proof: This proof is given by example. For ease of notation, denote α to mean $\alpha(G)$ and let $J \subseteq N$ to be a fixed maximum independent set. Consider a WSC problem

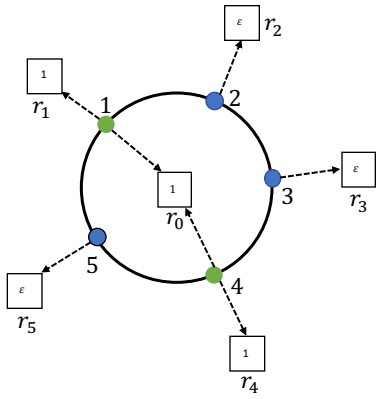

Figure 4.3: An example for the proof of Proposition 4.1. The agents are labeled $\{1, \ldots, 5\}$, the solid lines represent the graph G, and the dashed lines represent the actions available to each agent. G is a ring graph with 5 agents, and $J = \{1, 4\}$ (the green nodes). In a worst-case equilibrium, the green agents choose s_0 and the rest choose s_i. The optimal choices are for the green agents to choose s_i and one of the blue agents choose s_0.

with base set of resources $S = \{s_0, \ldots, s_n\}$. Let $f(s_0) = 1$, let $f(s_i) = 1$ for $i \in J$, and let $f(s_i) = \varepsilon$ for $i \notin J$ and for some small ε. For every agent $i \in J$, the action set is $X_i = \{\{s_0\}, \{s_i\}\}$. For every agent $i \notin J$, the action set is $X_i = \{\{s_i\}\}$, in other words these agents have only a single action to choose. See Figure 4.3 for an example.

Based on G agents in J cannot have a utility which directly accounts for the action of any other agent in J at equilibrium. One can assume without loss of generality that for $i \in J$, $U_i(s_i) \leq U_i(s_0)$, since the two elements are indistinguishable except by indexing, which could easily be switched. Therefore, a worst-case equilibrium decision set x^{eq} would be all agents in J choose $\{s_0\}$ and all other agents choose $\{s_i\}$. In this case $f(x^{\text{eq}}) = 1 + (N - \alpha)\varepsilon$. On the other hand, the optimal action profile x^{opt} is where all agents choose s_i, implying that $f(x^{\text{opt}}) = \alpha + 1 + \varepsilon + (N - \alpha - 1)\varepsilon$. Then

$$\lim_{\varepsilon \to 0} \frac{f(x^{\text{eq}})}{f(x^{\text{opt}})} = \lim_{\varepsilon \to 0} \frac{1 + (N - \alpha)\varepsilon}{\alpha + (N - \alpha)\varepsilon} = \frac{1}{\alpha}. \tag{4.24}$$

By definition, this is then an upper bound on PoA. ■

For many graphs, there is still a large gap between the upper bound on PoA shown in Proposition 4.1 and the lower bound for consistent valid utilities shown in Theorem 4.2. For instance, if G is a fully-connected directed acyclic graph, then $\text{PoA}(\mathcal{H}_c(G)) \geq 1/(n+1)$. However, we will see in the next chapter that deploying the utility MC can guarantee a PoA of $1/2$ for this graph constraint. Thus utilities which are optimal in this sense are a study of future work.

Chapter 5

The Greedy Algorithm

5.1 Introduction

The scope of this chapter is to consider Type 1 information sharing constraints that take the form of a directed acyclic graph (DAG). DAGs are particularly noteworthy, because they imply a sequencing of the agents, such that $(i, j) \in E$, then i comes before j in the sequence; assume without loss of generality that the agents are numbered according to such a sequence. This implies that at equilibrium, agent 1 has chosen independently it's best action without regard to the actions or messages of any other agents. Once it has made such a choice, it has no reason to deviate. Likewise, once agent 1 has made its choice and sent its messages, agent 2 makes its choice based (possibly) off that information. After agent 1 has chosen, agent 2 has no incentive to deviate from its choice, and so on through the sequence. This consequential behavior of an information sharing constraint graph in the form of a DAG has two benefits:

1. an equilibrium always exists.

2. a simple greedy algorithm always exists for fniding such an equilibrium.

Because of this, even if the problem doesn't inherently have an information sharing constraint graph G that is a DAG, one method of system design is to endow the agents with utility functions U_i that artificially create a DAG by "ignoring" the information that come from future agents in the sequence. Whether artifically imposed or not, all graphs in this chapter are assumed to be DAGs. Therefore, we shift focus from equilibria to greedy algorithms, keeping in mind that the result of such algorithms will also be an equilibrium.

As discussed in Chapter 2, submodular maximization is an NP-Hard problem in general. Thus a tremendous effort has been placed on developing fast algorithms that approximate the solution to the submodular maximization problem [53, 54, 55, 56, 57, 58, 28]. A resounding message from this extensive research is that very simple algorithms can provide strong guarantees on the quality of the approximation.

The seminal work in [54] demonstrates that a centralized greedy algorithm provides a solution that is within $1/2$ of the quality of the optimal solution. In fact, more sophisticated algorithms can often be derived for certain classes of submodular maximization problems that push these guarantees from $1/2$ to $1 - 1/e$ [53, 59, 60]. Progress beyond this level of suboptimality is not possible in general, because it was also shown that no polynomial-time algorithm can achieve a higher guarantee than $(1-1/e)$, unless $P = NP$ [61].

One appealing trait of the greedy algorithm is that it can be implemented in a distributed way while still maintaining the $1/2$ performance guarantee: this is the case where G is a fully connected DAG and $U_i = \text{MC}_i$ for all i. Other research has begun to explore how limited information can impact the performance of this distributed algorithm. For example, [25] focuses on the submodular resource allocation problem, modeled as a game played among agents. The resulting Nash equilibria have the familiar $1/2$ performance guarantee, however it is shown that when information is limited to be local instead of

global, the performance guarantee degrades to $1/n$, where n is the number of agents. The work in [62] formulates the problem of selecting representative data points from a large corpus as a submodular maximization problem. In order to perform the optimization in a distributed way, agents are partitioned into sets, where the full greedy algorithm is performed among agents within a set, while no information is transferred between sets. In this setting, the paper shows that the algorithm performance is worse than $1/2$, even when a preprocessing algorithm is used to intelligently assign decision sets to each agent. Other work in [28] discusses the role of information in the task assignment problem. It is shown that the distributed greedy algorithm can be implemented asynchronously, with convergence in a finite number of steps. Additionally, when agent action sets are based on spatial proximity, agents need only consider local information to achieve the $1/2$ bound. Finally, the work in [42] studies the performance of the distributed greedy algorithm when an agent can only observe a local subset of its predecessors. It is shown that localizing information, particularly when agents are partitioned from each other, leads to a degradation in performance. For instance, in the case where agents are partitioned into sets, performing the full greedy algorithm within the set and obtaining no information outside the set, the performance degrades proportionally to the number of sets in the partition.

This chapter studies system design under a few different scenarios. First, we assume that $U_i = \text{MC}_i$ for all i, and show in Theorem 5.1, similar to the results in Chapter 4, that the price of anarchy guarantees are related to the fractional independence number $\alpha^*(G)$. Theorem 5.2 addresses the scenario where the system designer can design the DAG, given the number of agents and the number of edges. We show that disconnected cliques maximize the use of edges in terms of price of anarchy. Finally, we explore Type 2 information sharing constraints, where the message m_i can either be x_i or some other message m_j for $j \in \mathcal{N}_i$. Under this scenario, we present the optimal utility functions U_i

for a particular class of DAGs.

5.2 Examples

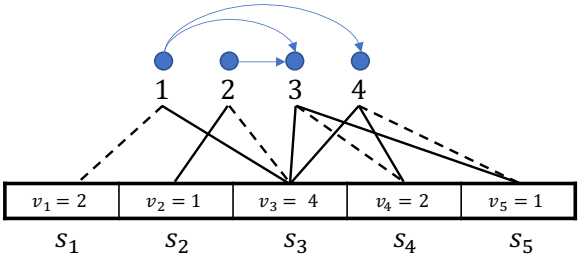

(a) The setup of a WSC problem. The resources are $S = \{s_1, ..., s_5\}$, each represented by a box, and each with a corresponding value. The available choices to each agent are represented by the black lines (both dotted and solid) - for instance $X_1 = \{\{s_1\}, \{s_3\}\}$, $X_2 = \{\{s_2\}, \{s_3\}\}$, etc. The dashed lines represent an optimal set of choices. The goal for the agents is to maximize $f(x)$ in (1.6). Using the generalized distributed algorithm (i.e., agents choose according to (5.1)), agent 1 chooses s_3, since $v_3 > v_1$. Then, agent 2, who (according to the graph) does not know that agent 1 has chosen s_3, also chooses s_3, since $s_3 > v_2$. Agent 3 observes that agents 1 and 2 have both chosen s_3, so it chooses s_4, since $v_4 > v_5$. Finally, agent 4, observing that agent 1 has chosen s_3 (but not that agent 3 has chosen s_4), chooses s_4, since $v_4 > v_5$. These results are summarized in the table below.

Algorithm	x_1^{sol}	x_2^{sol}	x_3^{sol}	x_4^{sol}	$f(x^{\text{sol}})$
Optimal	$\{s_1\}$	$\{s_3\}$	$\{s_4\}$	$\{s_5\}$	9
Distributed Greedy	$\{s_3\}$	$\{s_2\}$	$\{s_4\}$	$\{s_5\}$	8
Generalized Distributed Greedy	$\{s_3\}$	$\{s_3\}$	$\{s_4\}$	$\{s_4\}$	6

(b) For the WSC problem outlined to the left, this table shows the agents' decisions in an optimal case, the case where the distributed greedy algorithm is used (G is a fully connected DAG)and the case where the generalized distributed algorithm is used (agents choose according to (5.1), constrained to the graph shown above). The difference between the distributed greedy algorithm and the generalized version can be seen in the choices of agents 2 and 4. Agent 2 chooses s_3 when it can observe that s_4 has already been chosen by agent 1, otherwise it chooses s_4. Likewise, agent 4 chooses s_5 only when it knows that s_4 has already been selected. Therefore, as the informational constraints grow, the solution quality decreases. As a note, in this case we see that $f(x^{\text{sol}})/f(x^{\text{opt}}) = 6/9$.

Figure 5.1: An instance of the weighted set cover problem and the performance of the greedy algorithm in solving it.

To start we assume that $U_i = \text{MC}_i$ for all i. Then the greedy algorithm proceeds as follows: each agent sequentially chooses an action x_i^{sol} that satisfies the following rule:

$$x_i^{\text{sol}} \in \arg\max_{x_i \in X_i} \text{MC}_i\left(x_i, x_{\mathcal{N}_i}^{\text{sol}}\right) = \arg\max_{x_i \in X_i} f(x_i, x_{\mathcal{N}_i}^{\text{sol}}) - f(x_{\mathcal{N}_i}^{\text{sol}}). \tag{5.1}$$

In other words, each agent chooses the best action, according to f based on some subset of previoius agents in the sequence $\mathcal{N}_i$ whose actions it can observe. See Figure 5.1 for

an example.

We next present another relevant problem that can be modeled accordingly. This serves to give a scope and relevance to the model, as well as provide an example that will be leveraged later.

Example 5.1 (Vehicles target assignment problem [63]) *Consider the classic vehicles target assignment problem where there are a collection of targets $\mathcal{T}$ and each target $t \in \mathcal{T}$ has an associated value $v_t \geq 0$. Further, there exists a collection of n agents, and each agent i is associated with a success probability $p_i \in [0, 1]$ and a set of possible assignments $X_i \subseteq 2^{\mathcal{T}}$. The agents make decisions to reach a feasible allocation of agents to targets $x = (x_1, \ldots, x_n) \in X_1 \times \cdots \times X_n$ that optimizes a system-level performance metric of the form:*

$$f(x) = \sum\nolimits_{t \in \cup_i x_i} v_t \left(1 - \prod\nolimits_{i:t \in x_i} (1 - p_i)\right). \tag{5.2}$$

Note that the objective function given in (5.2) is submodular, as f can be expressed as a function of the form $f : 2^S \to \mathbb{R}_{\geq 0}$ for an appropriate choice of the domain set S, i.e., $S = N \times 2^{\mathcal{T}}$ and the action sets can be expressed as disjoint sets in S.

5.3 Price of Anarchy Bounds

In this section we present bounds on the price of anarchy when G is a DAG and when $U_i = \mathrm{MC}_i$ for all i. We show that the performance degrades proportionally to the fractional independence number of G [1].

[1] In comparison to the bounds shown in [42], the bounds shown in our work are tighter in all cases. In fact, except in certain corner cases (for example, both bounds are the same on a full clique), our results are strictly tighter.

Theorem 5.1 *For any graph DAG G,*

$$\frac{1}{\alpha(G)} \geq \text{PoA}(\mathcal{H}_{\text{MC}}(G)) \geq \frac{1}{\alpha^*(G)+1}. \tag{5.3}$$

The upper bound shows that it is impossible to construct a graph G such that the greedy algorithm's performance is better than $1/\alpha(G)$ for all possible $H \in \mathcal{H}_{\text{MC}}(G)$. Likewise, the lower bound means that no $H \in \mathcal{H}_{\text{MC}}(G)$ can result in a performance lower than $1/(\alpha^*(G)+1)$.

The formal proof for this theorem is given in Appendix A.2.1, but here we give a brief outline of the argument. The upper bound is simply a consequence of Proposition 4.1. The lower bound is found by leveraging the properties of submodularity and monotonicity, similar to the proof in Theorem 4.2.

We note that the lower bound in Theorem 4.2 leverages the fractional independence number of the graph $\bar{G}$, which is the graph G where all "non-reciprocal" edges are removed. This means for a DAG G that $\bar{G}$ is the empty graph: the lower bound on price of anarchy is $1/(n+1)$. Theorem 5.1 shows that the marginal contribution utility specifically is well-suited to DAGs, as opposed to other utilities that are valid and consistent.

5.3.1 Examples

Theorem 5.1 shows lower and upper bounds on price of anarchy, but we have not shown whether either of these bounds is tight. There exist graphs for which H can be chosen appropriately to meet the lower bound, and there also exist graphs whose lower bound can be proven to meet the upper bound. In this section, we provide an example of each.

Example 5.2 *The weighted set coverage problem presented in Figure 5.2 is an example*

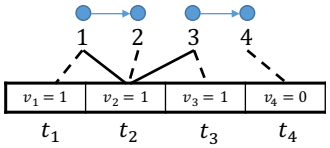

Figure 5.2: An example of a graph G where $\text{PoA}(\mathcal{H}_{\text{MC}}(G)) = 1/(\alpha^*(G)+1)$, and an instance of a weighted set cover problem using the same notation as in Figure 5.1. Here $\alpha(G) = \alpha^*(G) = 2$, and we can see that $f(x^{\text{opt}}) = 3$. The worst-case results from the generalized distributed greedy algorithm occur when $x_1^{\text{sol}} = x_2^{\text{sol}} = x_3^{\text{sol}} = \{t_2\}$, and therefore $f(x^{\text{sol}}) = 1$. This means $f(x^{\text{sol}})/f(x^{\text{opt}}) = 1/(\alpha^*(G)+1) = 1/3$, so the lower bound in Theorem 5.1 is tight for this graph.

showing that the lower bound from Theorem 5.1 is tight. For this graph G, $\alpha(G) = \alpha^*(G) = 2$. As shown, $f(x^{\text{sol}})/f(x^{\text{opt}}) = 1/(\alpha(G)+1)$, so the bound is tight for this system.

Example 5.3 *The graph G in Figure 4.1a is an example where the upper bound from Theorem 5.1 is tight. Here $\alpha^*(G) = 2$, and it is shown in Appendix A.3 for this graph that no $H \in \mathcal{H}_{\text{MC}}(G)$ can be constructed to give a worse efficiency than $1/2$.*

5.4 Optimal Structures

In this section, we describe how to build a graph G that yields the highest price of anarchy subject to a constraint on the number of edges.

5.4.1 Preliminaries

We denote $\mathcal{G}_{m,n} := \{G = (V,E) : |V| = n, |E| \leq m, G \text{ is a DAG}\}$ and $G^*_{m,n} \in \arg\max_{G \in \mathcal{G}_{m,n}} \text{PoA}(\mathcal{H}_{\text{MC}}(G))$, i.e., $G^*_{m,n}$ is a graph in $\mathcal{G}_{m,n}$ that maximizes efficiency. The complement $\bar{G} = (\bar{V}, \bar{E})$ of graph $G = (V,E)$ is such that $\bar{V} = V$ and $(i,j) \in \bar{E}$ if and only if $(i,j) \notin E$. It is straightforward to show that $\alpha(G) = \omega(\bar{G})$.

In graph theory a *Turán graph* $T(n,r)$ is a graph with n vertices created with the following algorithm:

(a) The Turán graph $T(8,3)$, where $\omega(G) = 3$. No other graph with 8 vertices can have more edges without also having a clique of size 4 or higher.

(b) The complement Turán graph $\overline{T(8,3)}$, where $\alpha(G) = 3$. No other graph with 8 vertices can have fewer edges without also having an independent set of size 4 or higher.

Figure 5.3: A Turán graph and its complement.

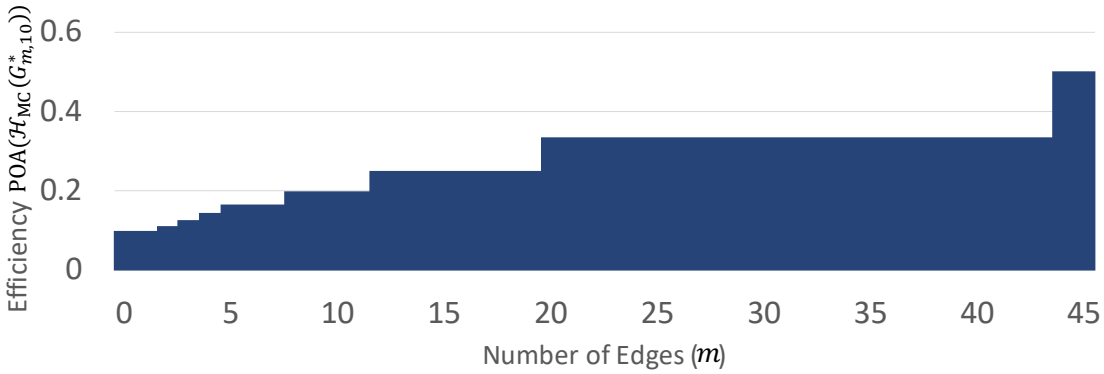

Figure 5.4: The efficiency of $G^*_{m,10}$ for all values of m, with example graphs for a few values of m. Notice the "dead zones", where adding more edges does not lead to any higher efficiency guarantees.

1. Partition the vertices into r disjoint sets $C_1, ..., C_r$ such that $|C_i| - |C_j| \leq 1$ for all $i, j \in \{1, ..., r\}$.

2. Create edges between all nodes not within the same set.

A result known as Turán's theorem states that $T(n, r)$ is an n-node graph with the highest number of edges that has clique number r or smaller [64]. Alternatively stated,

$$T(n,r) \in \underset{G=(V,E):\omega(G)\leq r}{\arg\max} |E|. \tag{5.4}$$

The complement of a Turán graph, denoted $\overline{T(n,r)}$, is created with the same procedure as a Turán graph, except that in Step 2, edges are created among all nodes *within* the same set. An example of a Turán graph and its complement is found in Figure 5.3.

Thus we can also state, similar to (5.4), that

$$\overline{T(n,r)} \in \underset{G=(V,E):\alpha(G)\leq r}{\arg\min} |E|. \tag{5.5}$$

In words $\overline{T(n,r)}$ is a graph with the fewest edges that has independence number r. It should also be clear that

$$\alpha\left(\overline{T(n,r)}\right) = \alpha^*\left(\overline{T(n,r)}\right) = k\left(\overline{T(n,r)}\right) = r. \tag{5.6}$$

Lastly, we define the graph

$$\hat{T}(n,m) := \underset{\left\{\overline{T(n,r)}:|E|\leq m\right\}}{\arg\min} r, \tag{5.7}$$

which is the complement n-node Turán graph with the lowest independence number among all graphs with the number of edges less than or equal to m [2].

5.4.2 Result

The main result of this section regarding efficient graph structures is stated below and later proved in Section 5.4.4.

Theorem 5.2 *Consider two nonnegative integers n and m such that $m \leq \frac{1}{2}n(n-1)$. If $m \neq \frac{1}{2}n(n-1)-1$, then $G^*_{m,n} = \hat{T}(n,m)$. If $m = \frac{1}{2}n(n-1)-1$, then $G^*_{m,n}$ is the full clique on n nodes, minus the edge $(n-1,n)$.*

An illustration of $\mathrm{PoA}(\mathcal{H}_{\mathrm{MC}}(G^*_{m,n}))$ as a function of the number of edges m is given

[2] Searching over the space of complement Turán graphs can be done simply. Adapting part of Turán's theorem, we see that $r \geq \lceil n^2/(2m+n)\rceil$. Therefore, one can start by setting r to this minimum value, and then determining whether $m \geq M(n,r)$, see Lemma 5.2 below. If the statement is not true, r can be incremented until it is.

in Figure 5.4. One item to note is that there may be extra edges not used in our design of $G^*_{m,n}$. For instance, in Figure 5.4, the efficiency is the same when $12 \leq m \leq 19$. This implies that $G^*_{12,10}$ and $G^*_{19,10}$ can be the same graph, and for any value of m in between. Hence, there are "dead zones" seen in the graph in Figure 5.4.

5.4.3 The Sibling Property

Here we present a graph property, along with a corollary to Theorem 5.1. These results are key to the proof for Theorem 5.2.

Definition 5.1 *Let $G \in \mathcal{G}$. Then G has the* Sibling Property *if for some maximum independent set J, there exist $w \in V \setminus J$ and $i \in J$ such that $i \in \mathcal{N}_w$ (see Figure 4.1).*

Lemma 5.1 *If a graph G lacks the Sibling Property, then*

1. *There is a unique maximum independent set J.*
2. *The set J must include nodes n and $n-1$.*
3. *The induced subgraph G' created by removing the set J from G must be such that $\alpha(G) > \alpha(G')$.*
4. *Every node outside the set J must have outgoing edges to at least 2 nodes in J* [3].

Proof: We prove each Propoerty separately:

Property 1: Suppose there are 2 maximum independent sets J and J'. Let $i \in J' \setminus J$ and $j \in J \setminus J'$. By definition, all nodes in either set cannot have any outgoing edges. This implies that $(i,j),(j,i) \notin E$: in other words, i and j are independent from each other, and neither J nor J' are maximum, a contradiction.

[3] In the literature, such a J is called a *perfect independent set*. We present a proof here that suits the needs of this work, but it is also shown in [65] that every unique maximum independent set is perfect.

Property 2: First, suppose that n is not included in J. Then there exists an edge (i, n) for some $i \in J$. By defintion, this means G has the Sibling Property, a contradiction. Now suppose that $n-1$ is not in J. Since G does not have the Sibling Property, then by definition $(j, n-1) \notin E$ for all $j \in J \setminus n$. This means that another maximum independent set is $\{n-1\} \cup J \setminus n$, which is a contradiction to statement 1.

Property 3: If this were not true, then J would not be a unique maximum independent set.

Property 4: Let $i \notin J$. By definition, i cannot have any incoming edges from J and if there are no edges between i and J, then i must be part of J, a contradiction. Therefore, we consider the case where $(i, j) \in E$ for some $j \in J$, but no outgoing edges from i to $J \setminus j$ exist. This means that another maximum independent set is $i \cup J \setminus j$, and J is not unique. By Property 1, this is a contradiction. ■

Corollary 5.1 *For a DAG G with the Sibling Property,*

$$\mathrm{PoA}(\mathcal{H}_{\mathrm{MC}}(G)) \leq \frac{1}{1+\alpha(G)}, \tag{5.8}$$

with equality when $\alpha(G) = k(G)$.

Proof: We provide an example which gives us the upper bound using a weighted set cover problem. Let J be a maximum independent set of G and let w be defined as in Definition 5.1. Then $S = \{s_1, ..., s_n\}$, where $v_i = 1$ if $i \in J$ or $i = w$, and $v_i = 0$ otherwise. The action sets are

$$X_i = \begin{cases} \{\{s_w\}, \{s_i\}\} & \text{if } i \in J, \\ \{\{s_w\}\} & \text{if } i = w \\ \{\{s_i\}\} & \text{otherwise.} \end{cases} \tag{5.9}$$

Each agent in J is equally incentivized to choose either option, since none of them can access to the choice of the others. Therefore, the worst case in the greedy algorithm is for every agent in J to choose s_w, implying $f(x^{\text{sol}}) = v_w = 1$. Each agent makes the other choice in the optimal, so $f(x^{\text{opt}}) = v_w + \sum_{i \in J} v_i = 1 + \alpha(G)$. Therefore $f(x^{\text{sol}})/f(x^{\text{opt}}) = 1/(1 + \alpha(G))$ is an upper bound on $\mathcal{H}_{\text{MC}}(G)$.

In the case where $\alpha(G) = k(G)$, (4.5) shows that $\alpha^*(G) = \alpha(G)$, which implies by Theorem 5.1 that $\mathcal{H}_{\text{MC}}(G) \geq 1/(1 + \alpha(G))$. ■

5.4.4 Proof for Theorem 5.2

In this section we present the proof for Theorem 5.2, beginning with two lemmas. The first characterizes the number of edges in a complement Turán graph (proof in Appendix A.5), and the second characterizes the fewest number of edges in a graph without the Sibling Property (proof in Appendix A.4).

Lemma 5.2 *Let $G = \overline{T(n, r)}$. Then the number of edges in G is*

$$M(n, r) := \frac{1}{2}(n \bmod r)\left\lceil \frac{n}{r} \right\rceil \left(\left\lceil \frac{n}{r} \right\rceil - 1 \right) + \frac{1}{2}(r - n \bmod r)\left\lfloor \frac{n}{r} \right\rfloor \left(\left\lfloor \frac{n}{r} \right\rfloor - 1 \right). \tag{5.10}$$

Lemma 5.3 *Let $G \in \mathcal{G}$ have n nodes, be without the Sibling Property, and such that $\alpha(G) = r$. Then the number of edges m in G satisfies*

$$m \geq M(n - r, r - 1) + 2(n - r). \tag{5.11}$$

Furthermore, for any values of n and r, such a G can be constructed so that (5.11) is at equality.

We now commence with the proof for Theorem 5.2. The case $m = 0$ trivially holds, so we assume that $m > 0$. Recall that the graph $\hat{T}(n, m)$ is a set of disconnected cliques,

which implies that any maximum independent set has one node from each clique, and that no maximum independent set is unique. Therefore, by Lemma 5.1, Property 1, $\hat{T}(n,m)$ has the Sibling Property. In light of (5.6), It follows from Corollary 5.1 that $\text{PoA}(\mathcal{H}_{\text{MC}}(\hat{T}(n,m))) = 1/(1+\alpha(\hat{T}(n,m))$. The statement in (5.5) also shows that no other graph with $\leq m$ edges can have a smaller independence number. Combining this with Corollary 5.1 implies that no other graph with the Sibling Property (and same number of nodes and edges) can have a higher efficiency.

It remains to confirm that any graph without the Sibling Property cannot have a higher efficiency than $\hat{T}(n,m)$, given n nodes and m edges – with the exception when $m = \frac{1}{2}n(n-1)$. Let G be a graph with n nodes, m edges, without the Sibling Property, and with independence number $r+1$. We assume that G has the fewest number of edges (as dictated by Lemma 5.3), with the highest possible efficiency $\text{PoA}(\mathcal{H}_{\text{MC}}(G)) = 1/(r+1)$. By Corollary 5.1 and (5.6), this is the same efficiency as $\hat{T}(n,m)$, thus we seek to characterize when the number of edges in G is greater than or equal to that of $\hat{T}(n,m)$. In other words, $\mathcal{G}^*_{m,n} = \hat{T}(n,m)$ only if

$$M(n,r) \leq m = M(n-r-1,r) + 2(n-r-1). \tag{5.12}$$

In order to show when this condition holds, we divide the remainder of the proof into four cases, the union of which covers all possible values of n and r. In the first case, when $r = 1$, we prove (5.12) is false for all values of n (which corresponds to the case in the theorem statement when $m = \frac{1}{2}n(n-1)-1$). In the other cases, we show that (5.12) is true, justifying that $G^*_{m,n} = \hat{T}(n,m)$.

Case 1: $r = 1$. Here, $\hat{T}(n,m)$ is a clique, and has $\frac{1}{2}n(n-1)$ edges. The graph G is

such that

$$m = M(n-2,1) + 2(n-2) = \frac{1}{2}n(n-1) - 1, \tag{5.13}$$

which is one less edge than $\hat{T}(n,m)$. Thus, for any value of n, there exists a G where (5.12) is false. Such a G is shown for $n = 4$ in Figure 4.1a, and a trivial extension to the proof in Appendix A.3 shows that $\text{PoA}(\mathcal{H}_{\text{MC}}(G)) = \text{PoA}(\mathcal{H}_{\text{MC}}(\hat{T}(n,m))) = 1/2$ for any value of n. Since G is created with the fewest number of edges, it follows that (5.12) is false *only* when $m = \frac{1}{2}n(n-1) - 1$. By the construction in the proof of Lemma 5.3, such a G is the full clique minus the edge $(n-1, n)$.

Case 2: $r = n - 1 \geq 2$. In this case, $\hat{T}(n,m)$ is the graph with no edges and efficiency $1/n$. Any graph with 1 or 0 edges must have this same efficiency, so (5.12) is true in this case.

In the remaining cases, we assume that $2 \leq r \leq n-2$, which also implies that $n \geq 4$. In both cases, we show that (5.12) holds.

Case 3: $n \bmod r \geq 1$. This condition implies the following:

- $(n - r - 1) \bmod r = n \bmod r - 1$
- $\lceil n/r \rceil = \lfloor n/r \rfloor + 1$
- $\lfloor (n - r - 1)/r \rfloor = \lfloor r/n \rfloor - 1$
- $\lceil (n - r - 1)/r \rceil = \lfloor r/n \rfloor$

Leveraging the above statements, $M(n,r)$ and $M(n-r-1,r)$ become:

$$M(n,r) = \frac{1}{2}\left\lfloor \frac{n}{r} \right\rfloor \left(2(n \bmod r) + r\left\lfloor \frac{n}{r} \right\rfloor - r \right) \tag{5.14}$$

$$M(n-r-1,r) = (n \bmod r)\left(\left\lfloor \frac{n}{r} \right\rfloor - 1 \right) + \frac{r}{2}\left\lfloor \frac{n}{r} \right\rfloor^2 - 3\left\lfloor \frac{n}{r} \right\rfloor \tag{5.15}$$

Using these expressions to evaluate (5.12) yields

$$\left\lfloor \frac{n}{r} \right\rfloor (r+1) + n \bmod r \leq 2n - r - 1. \tag{5.16}$$

We can now use the identity $n \bmod r = n - r\lfloor n/r \rfloor$ to change the requirement in (5.16) to

$$\left\lfloor \frac{n}{r} \right\rfloor \leq n - r - 1. \tag{5.17}$$

Since $\lfloor n/r \rfloor \leq n/r$, a sufficient statement for (5.16) to hold can be found by replacing $\lfloor n/r \rfloor$ with n/r, which can be simplified to

$$\frac{r^2 + 1}{r - 1} \leq n. \tag{5.18}$$

The expression on the left side of the inequality is nondecreasing in r. Since $r \leq n - 2$, if the inequality is true for $r = n - 2$, then it is true for all relevant values of n, r. If we let $r = n - 2$ in (5.18) and simplify, we conclude that (5.18) holds for all $n \geq 5$. By the premise that $n \geq 4$, the only values of n, r that could make this false are $n = 4, r = 2$, however, this would imply that $n \bmod r = 0$, not allowable by the condition for this case. Thus, (5.18) is true and by extension so is (5.12).

Case 4: $n \bmod r = 0$. The condition implies the following:

- $(n - r - 1) \bmod r = r - 1$
- $\lfloor n/r \rfloor = \lceil n/r \rceil = n/r$
- $\lfloor (n - r - 1)/r \rfloor = \lfloor r/n \rfloor - 2$
- $\lceil (n - r - 1)/r \rceil = \lfloor r/n \rfloor - 1$

Leveraging the above statements, $M(n,r)$ and $M(n-r-1,r)$ become:

$$M(n,r) = \frac{n}{2}\left(\frac{n}{r} - 1\right), \tag{5.19}$$

$$M(n-r-1,r) = \frac{n^2}{2r} - \frac{3}{2}n + r - \frac{n}{r} + 2. \tag{5.20}$$

Using these expressions to evaluate (5.12) yields

$$\frac{r^2}{r-1} \leq n. \tag{5.21}$$

It is straightforward to see that if (5.18) holds, then so does (5.21). Case 3 shows that (5.18) is true unless $r = 2, n = 4$. However, (5.21) is true for these values, which implies (5.21) holds for all relevant values of n, r. Thus, in this case, (5.12) also holds. ■

5.5 Strategic Information Sharing

Thus far a graph constraint G has meant that when $(i,j) \in E$, agent i shares its action with agent j. However, this need not be the case, as in Type 2 information sharing constraints. In this section we address the question of whether endowing agents with a more strategic information sharing policy can increase performance guarantees via the price of anarchy. We show what such a policy would look like for a certain class of graphs.

5.5.1 Section Model

The previous sections show that when G is a DAG, the marginal contribution utility can provide strong guarantees on the price of anarchy. Theorem 5.2 gives insight that if the system designer can choose G given a number of edges, creating a set of disconnected

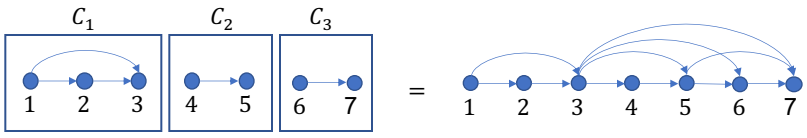

Figure 5.5: An example of the model and notation used in this section. We assume that each clique in $\mathcal{C}$ passes a single action from among its agents to all agents in future cliques. This is equivalent to the graph on the right, where again we assume that the last agent in the clique is charged with sharing the clique's information. It should be clear that for the graph on the right, $\alpha(G) = k(G) = |\mathcal{C}|$.

cliques is optimal. In light of this, assume that the agents are partitioned into a series of disconnected cliques $\mathcal{C} = \{C_1, ..., C_m\}$. We denote $C_i > C_j$ to mean that the agents of C_i come before the agents in C_j in the sequence. As an example, for Figure 5.5, $C_1 = \{1, 2, 3\}, C_2 = \{4, 5\}, C_3 = \{7, 8\}$. We will henceforth in this section refer to a graph as $\mathcal{C}$ when appropriate to avoid confusion. For such graph structures and utilities, Theorem 5.1 and Corollary 5.1 demonstrate that the resulting efficiency guarantees associated with the greedy algorithm are precisely

$$\text{PoA}(\mathcal{H}_{\text{MC}}(\mathcal{C})) = \frac{1}{1 + |\mathcal{C}|}, \tag{5.22}$$

since $\alpha^*(\mathcal{C}) = \alpha(\mathcal{C}) = |\mathcal{C}|$, and all $\mathcal{C}$ has the Sibling Property.

In this section we begin to explore Type 2 information sharing constraints, which will relax the assumption that agents communicate their own actions to other agents. Recall that a meta-action is the tuple (x_i, m_i), where m_i is a message communicated to other agents. Since G is a DAG, each agent sequentially chooses a meta-action choosing x_i and m_i, which can be done independently - x_i leveraging the utility function U_i and m_i similarly leveraging a function W_i. We will restrict information sharing by imposing that $M_i = \{x_i\} \cup \{m_j\}_{j \in \mathcal{N}_i}$, i.e., agent i can share its chosen action *or* the message that it has received from a previous agent in the sequence. Thus, while each agent is still only sharing a single action (either its own or one of its neighbors), what is different from

previous sections in this chapter is that we are allowing agents to be more strategic with what information is being shared.

For agent i which is not last in its clique, the optimal message choice is $m_i = x_i^{\text{sol}}$, since future agent j in the same clique will already have access to the same messages as agent i – see Figure 5.5 for reference. Thus the only nontrivial message passing decisions are made by the last agent in each clique. In this section, we refer to the message sent by the last agent in clique C as m_C, and will also refer to this as the message sent by clique C.

Here we relax the requirement that $U_i = \text{MC}_i$, rather we allow the system designer to assign any feasible utility function to each agent. However, since MC_i is part of our analysis in this chapter, note that we adapt its definition to be $\text{MC}_i(x_i, m_{\mathcal{N}_i}) = f(x_i, m_{\mathcal{N}_i}) - f(m_{\mathcal{N}_i})$, which is valid since $m_j \subseteq S$. Clique C shares message m_C with agents in future cliques using the following rule

$$m_C = W_C(X_C, \{m_{C'}\}_{C'<C}) \in \{x_i^{\text{sol}}\}_{i \in C}, \tag{5.23}$$

where $X_C = \bigcup_{i \in C} X_i$. Note that unlike U_i, which provides a value to an action given the messages of others, W_C is a policy which deterministically chooses the message, a property allowable due to the sequential nature of the system. We refer to W_C as the information sharing strategy for clique C and $W = \{W_C\}_{C \in \mathcal{C}}$ as the information strategy profile for $\mathcal{C}$. We likewise refer to $\mathcal{H}_{U,W}(\mathcal{C})$ as the set of systems with information sharing constraints represented by $\mathcal{C}$ which use utility profile $U = (U_1, \dots U_n)$ and information sharing strategy profile $W = (W_{C_1}, \dots, W_{C_{|\mathcal{C}|}})$. We now give two examples of information sharing strategies and show how they are utilized in conjunction with the graph constraint. These will both be used heavily in our results. The first is a pre-committed

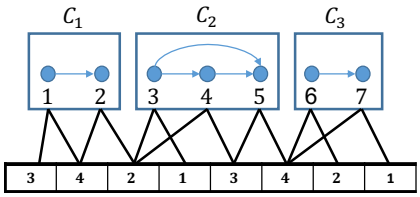

$\{I, \pi\}$	Agent 1 choice	Agent 2 choice	Agent 3 choice	Agent 4 choice	Agent 5 choice	Agent 6 choice	Agent 7 choice	z_{C_1}	z_{C_2}	Value
Optimal	L	L	R	L	L	R	L	-	-	19
$\{I^P, \pi^M\}$	R	R	L	R	R	L	L	x_1^{sol}	x_3^{sol}	13
$\{I^M, \pi^M\}$	R	R	L	R	R	R	R	x_1^{sol}	x_5^{sol}	16

Figure 5.6: An example WSC problem for the greedy algorithm including strategy. The boxes along the bottom represent S, and f is simply the sum of the values in a set of boxes. The blue circles are the agents, and the black lines represent X: each agent can choose between two boxes. In the table, these choices are denoted L for left and R for right. The partition $\mathcal{C}$ has 3 sets: C_1, C_2, C_3, each forming clique among the agents (blue edges) in its set. Results for two strategies are shown in the table.

strategy:

$$m_C = x_i^{\text{sol}}, \tag{5.24}$$

where $i \in C$ is determined a priori. The second is

$$m_C = \arg\max_{i \in C} \Delta(x_i^{\text{sol}} | \{x_j^{\text{sol}}\}_{j \in \mathcal{N}_i}, \{m_{C'}\}_{C' < C}), \tag{5.25}$$

Essentially, this strategy shares the decision of the agent with the highest marginal contribution. As we will be referring to these strategies often, denote W^P as the profile when W_C is (5.24) and W^M as the profile when W_C is (5.25) for all $C \in \mathcal{C}$. Figure 5.6 gives an example to show how these strategies work.

5.5.2 The Benefit of Strategy

In this section, we characterize the benefit of using strategic information sharing. To begin, we show the efficiency of the greedy algorithm when a pre-committed information strategy is used, again an application of Theorem 5.1 and Corollary 5.1.

Corollary 5.2 *Let $\mathcal{C}$ be a partition of agents such that $|C| > 1$ for all C, and assume*

$W = W^P$ and $U = \text{MC}$. Then

$$\text{PoA}(\mathcal{H}_{U,W}(\mathcal{C})) = \frac{1}{1 + |\mathcal{C}|}. \tag{5.26}$$

We omit a formal proof here, but essentially this statement follows from the observation that utilizing W^P and MC is equivalent to adding edges to the graph $\mathcal{C}$ between the partitions. Since no clique is sharing the information of every agent, the clique cover number remains the same, and the efficiency is the same as in (5.22).

Comparing Corollary 5.2 to the baseline in (5.22), we can glean that if W pre-commits to sharing the decision of a set of agents, then this strategy offers no benefit to efficiency. By pre-comitting, one is essentially removing the dependence of W on f from (5.23). Even in the case that one relaxes the restriction in (5.23) that only one piece of information is shared, there is no benefit to price of anarchy, assuming that the decision of at least one agent in every clique is not shared. Therefore, any strategy that increases PoA must utilize f in some way.

Theorem 5.3 *Let $\mathcal{C} = \{C_1, .., C_p\}$ be a partition such that $|C_j| > 1$ for all j. Then for any admissible U, W,*

$$\text{PoA}(\mathcal{H}_{U,W}(\mathcal{C})) \leq \frac{1}{2 + \sum_{i=1}^{p-1} \prod_{j=1}^{i} (1 - 1/|C_j|)}, \tag{5.27}$$

with equality when $W = W^M$ and $U = \text{MC}$.

The proof for this theorem will be shown at the end of the section, in favor of some discussion up front. First, given the current graph structure, Theorem 5.3 states that an optimal choice for U_i for every agent is MC_i. Not only is this convenient given the simplicity of the rule, but it also allows one to leverage the intuition and insights from the previous sections in a new setting. Given this, it should not be surprising that the action

with the highest marginal contribution to its predecessors is the best information to send to future agents in the sequence. If $\mathcal{C}$ is comprised of a single clique (i.e. $\mathcal{C} = \{C_1\}$), then the sum in the denominator is 0, and the expression simplifies to the familiar 1/2 guarantee. This is also true for the expression in Corollary 5.2. For the rest of the discussion we assume this is not the case.

Any term in the sum in (5.27) is strictly less than 1. The fact that there are $|\mathcal{C}| - 1$ terms in the sum then confirms that $\text{PoA}(\mathcal{H}_{W^M,\text{MC}}(\mathcal{C})) > \text{PoA}(\mathcal{H}_{W^P,\text{MC}}(\mathcal{C}))$. This is significant, because as mentioned above it holds even when allowing many actions to be shared from within the clique rather than just one. To further this point, if one restricted information sharing to only be between two cliques C and C' (where $C < C'$) then (5.27) becomes

$$\text{PoA}(\mathcal{H}_{U,W}(\mathcal{C})) = \frac{1}{(1 - 1/|C|) + |\mathcal{C}|}, \tag{5.28}$$

which still has a strictly higher efficiency than (5.26). In essence, one always sees a benefit to strategic information sharing.

Another observation from Theorem 5.3 is that the order of the sizes of the different cliques matters. For instance, consider the graph in Figure 5.5. According to Theorem 5.3,

$$\text{PoA}(\mathcal{H}_{W^M,\text{MC}}(\mathcal{C})) = \frac{1}{2 + (2/3) + (2/3)(1/2) + (2/3)(1/2)(1/2)} \tag{5.29}$$

$$\approx 0.3158. \tag{5.30}$$

However, if one reordered the cliques so that the the clique of size 3 were last, then

$$\text{PoA}(\mathcal{H}_{W^M,\text{MC}}(\mathcal{C})) = \frac{1}{2 + (1/2) + (1/2)(1/2) + (1/2)(1/2)(2/3)} \tag{5.31}$$

$$\approx 0.3429. \tag{5.32}$$

Intuitively, this makes sense since moving the larger clique to the end will require more edges in the equivalent graph - see again Figure 5.5. Therefore, sequencing the smaller cliques first is advantageous. We now proceed with the proof.

Proof: First we give an example which shows a universal upper bound on $\text{PoA}(\mathcal{H}_{U,W}(\mathcal{C}))$. Then we show that if $W = W^M$ and $U = \text{MC}$, there is a lower bound on $\text{PoA}(\mathcal{H}_{U,W}(\mathcal{C}))$ which matches the upper bound. Thus the bound is tight and we know that $\{W^M, \text{MC}\}$ is the optimal strategy.

Upper Bound

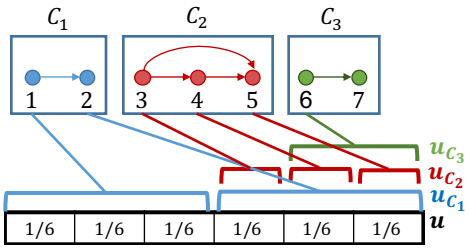

(a) The set of worst-case greedy choices as described in the upper bound proof for Theorem 5.3. Notice that u_{C_1} and u_{C_2} are divided up into 2 and 3 parts, respectively, for each agent in the cliques. Only the first agent in C_3 chooses u_{C_3}.

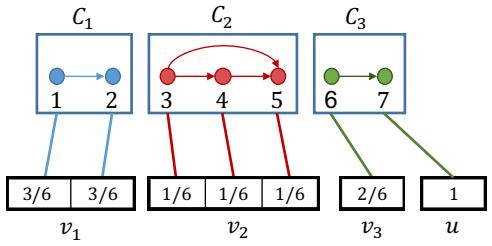

(b) The set of optimal choices as described in the upper bound for the Theorem 5.3 proof. Notice that agent 7 chooses the full set of boxes chosen in the greedy algorithm.

Figure 5.7: An example of the upper bound f, X for Theorem 5.3. As argued in the proof, any I and π must value every decision and agent, respectively, the same. In this case, $f(x^{\text{sol}}) = 1$ and $f(x^{\text{opt}}) = 2 + 1/2 + 1/2(2/3) = 17/6$, so $f(x^{\text{sol}})/f(x^{\text{opt}}) = 6/17$. Theorem 5.3 also states that this is the max possible efficiency, i.e. $\text{PoA}(\mathcal{H}_{\text{MC},W^M}(\mathcal{C})) = 6/17$.

Here we give an example H to serve as an upper bound on $\text{PoA}(\mathcal{H}_{U,W}(\mathcal{C}))$, and show that it is the exact expression in (5.27). Assume to begin that W^M and MC are used.

We introduce some new notation, for convenience. Let $f(i) = f(x_i^{\text{sol}})$ and let $f(i^*) = f(x_i^{\text{opt}})$. For set of agents J, let $f(J) = f\big(\bigcup_{i\in i} x_i^{\text{sol}}\big)$. For $J = \{a, a+1, ..., b\}$, let $f(a:b) = f(J)$. Let $Z_C = \bigcup_{C'<C} m_{C'}$, i.e., Z_C is the set of messages from prior cliques

that are shared with the agents in clique C. Finally, let the clique that sequentially comes before C be denoted as $C-1$.

Assume that S is a set boxes, and that f is simply the area of the boxes covered by the choices of the agents. Suppose there is a box u, where $f(u)=1$, which will be chosen by the agents in the worst case. There are also boxes $v_1, ..., v_k$, and the optimal choices will allow agents to cover all boxes.

Each clique C will be able to choose between some portion of u, called u_C, and v_C, where $f(u_C) = f(v_C)$. Additionally each u_C and v_C are divided up equally into $|C|$ parts, so that the value of each agent's choice within the clique is the same. We define $u_1 = u$, and u_C is the portion of u_{C-1} not covered by m_{C-1}. Thus

$$f(u_C) = \frac{|C-1|-1}{|C-1|} f(u_{C-1}). \tag{5.33}$$

In the last clique, the first agent can choose between u_p and v_p, and the second agent can only choose u. All other agents in the last clique are ignored. See Figure 5.7 for an example.

Agents are equally incentivized to choose their portion in u, so $f(x^{\text{sol}}) = 1$. As stated, the optimal choices yield the set $\{u, v_1, ..., v_k\}$. Each v_C is such that

$$f(v_C) = \prod_{C'<C} \frac{|C'|-1}{|C'|}, \tag{5.34}$$

where $v_1 = 1$ by convention. Therefore, if $\mathcal{C} = \{C_1, .., C_M\}$, then $f(x^{\text{opt}}) = f(u) + \sum_C f(v_C) = 2 + \sum_{i=1}^{p-1} \prod_{j=1}^{i} \frac{|C_j|-1}{|C_j|}$.

Although we intially assumed to use $\{W^M, \text{MC}\}$, we now make the claim that this canonical example serves as an upper bound on $\text{PoA}(\mathcal{H}_{\text{MC},W^M}(\mathcal{C}))$ for any W and U. From (5.23), it is clear that information sharing strategy W_C can only leverage information from

past cliques and the action sets of agents in C. *Based on this information, all agents in C are equivalent.* Any choice of m_C will yield the same efficiency guarantee (refer again to Figure 5.7). Therefore $\text{PoA}(\mathcal{H}_{U,W}(\mathcal{C}))$ is the same for any W. A similar argument is made for U. We conclude that the upper bound found by this canonical example is an upper bound on the efficiency for any strategy used.

Lower Bound

We assume that $W = W^M$ and $U = \text{MC}$. Let $\mathcal{C} = \{C_1, ..., C_p\}$, and we will use the notation that $m_{C_k} = m_k$, and likewise for Z_k. For some C_k, let j be the agent whose decision is m_k. Then:

$$\Delta(C_k|Z_k) = \sum_{i \in C_k} \Delta(i|\mathcal{N}_i \cup Z_k), \tag{5.35}$$

by application of (2.2) and since W^M uses (5.25) and j is chosen by W^M. Leveraging the definition of Δ, and the fact that $Z_{k+1} = Z_k \cup m_k$, we see that

$$f(Z_{k+1}) \geq \frac{1}{|C_k|} f(C_k, Z_k) + \left(1 - \frac{1}{|C_k|}\right) f(Z_k) \tag{5.36}$$

For simplicity, let $a_k = 1 - 1/|C_k|$. Then (5.36) becomes

$$f(Z_{k+1}) \geq (1 - a_k) f(K_k, Z_k) + a_k f(Z_k) \tag{5.37}$$

Begin with the following inequality (agent i is in clique $C_{k(i)}$):

$$f(x^{\text{opt}}) \leq 2f(x^{\text{sol}}) + \sum_{k=1}^{p-1} f(C_k, Z_k) - \sum_{k=1}^{p-1} f(Z_{k+1}), \tag{5.38}$$

which can be shown by leveraging submodularity, (2.2), our defined π, and the definition

of Δ and Z_{k+1}.

Notice that the two sums have the same number of terms, and we can apply (5.37) to each term in the second sum and get the following:

$$\begin{aligned} f(x^{\text{opt}}) \leq & (2 + a_{p-1}) f(x^{\text{sol}}) \\ & + \sum_{k=1}^{p-2} a_k f(C_k, Z_k) - \sum_{k=1}^{p-2} a_{k+1} f(Z_{k+1}) \end{aligned} \tag{5.39}$$

Again we see that both sums have the same number of terms and we apply (5.37) to get:

$$\begin{aligned} & f(x^{\text{opt}}) \leq (2 + a_{p-2} + a_{p-1}a_{p-2}) f(x^{\text{sol}}) + \\ & \sum_{k=1}^{p-3} a_k - a_{k+1}(1 - a_k) f(C_k, Z_k) - \sum_{k=1}^{p-3} a_{k+2} a_{k+1} f(Z_{k+1}) \end{aligned} \tag{5.40}$$

Notice that each application of (5.37) adds some positive term to the coefficient of $f(x^{\text{sol}})$ and drops a term from each of the sums. Let b_j be the term added to the coefficient of $f(x^{\text{sol}})$ after applying (5.37) j times. In other words, $b_1 = a_{p-1}$, $b_2 = a_{p-2} + a_{p-1}a_{p-2} - a_{p-1}$, etc. After applying (5.37) $p - 1$ times, we see that

$$f(x^{\text{opt}}) \leq \left(2 + \sum_{j=1}^{p-1} b_j\right) f(x^{\text{sol}}) \tag{5.41}$$

Thus to find the lower bound on efficiency, we need to find $\sum_j b_j$. Following the pattern, each b_j can be defined as follows:

$$b_j = \sum_{i=p-j}^{p-1} \prod_{d=p-j}^{i} a_d - \sum_{i=p-j+1}^{k} \prod_{d=p-j+1}^{i} a_d \tag{5.42}$$

Notice that the second sum for b_j is the negative of the first sum for b_{j-1}. Thus

$$\sum_{j=1}^{p-1} b_j = \sum_{i=1}^{p-1} \prod_{d=1}^{i} a_d. \tag{5.43}$$

Therefore, the lower bound meets the upper bound. ■

Chapter 6

Augmenting Action Sets

The previous chapter has considered scenarios where information sharing has been limited compared to the nominal greedy algorithm given in (5.1), i.e., we have removed edges from a fully-connected DAG and described how that affects the resulting performance guarantees. This chapter, on the other hand, explores the scenario where more information sharing is permitted via Type 2 information sharing constraints. We further investigate the idea that information sharing need not be limited to an agent's chosen action. Here we allow agents to pass additional elements of S to future agents in the sequence in an attempt to see how this type of *element passing* can be exploited to give higher performance guarantees.

As an example, consider a scenario where two flying vehicles are trying to identify the positions of a set of targets. Each vehicle captures many images, but can only send k of them to a central satellite, which uses the sent measurements to estimate the location of the targets. In this scenario, vehicle 1 can use the local communication network to share with vehicle 2 the k images which it sent to the satellite, and can additionally send p more. Vehicle 2 could then send to the satellite any combination of k images from its original set as well as these new communicated measurements. In the extreme case where

vehicle 2 was not able to capture any "valuable" images by itself, it could still send to the fusion center the p that came from vehicle 1, thus offsetting a potentially poor system performance.

Result	Measure	Lower bound	Upper bound
Theorem 6.1	$\frac{\text{AUG. GREEDY}}{\text{NOMINAL GREEDY}}$	$\frac{1}{2 - \frac{(\min(p/k,1))^{n-1}}{\sum_{i=0}^{n-1}(\min(p/k,1))^{i}}}$	$2+\min(n-1,p/k)$
Theorem 6.2	$\frac{\beta\text{– APPROX. AUG. GREEDY}}{\beta\text{– APPROX. NOM. GREEDY}}$	$\frac{1}{1+\frac{1}{\beta} - \frac{(\beta\min(p/k,1))^{n-1}}{\beta\sum_{i=0}^{n-1}(\beta\min(p/k,1))^{i}}}$	$2/\beta + p/k$, if $p \le k$ $1+\frac{1}{\beta}\left(1+\min\left(n-1,\frac{p}{k}\right)\right)$, if $p > k$

Figure 6.1: A brief summary of the results from this paper, where an agent selects up to k elements as its "action" and up to p elements to share. We explore how well an augmented greedy algorithm, which includes element passing, performs. Theorem 6.1 provides a range as to how well the augmented greedy algorithm can perform versus the nominal greedy algorithm. Theorem 6.2 describes the scenario where each agent can only approximate the solution to its local problem within a factor of β.

Theorem 6.1 directly addresses the increased system efficiency by giving bounds on how well an optimal element passing policy can perform compared to the nominal greedy algorithm: it shows that there exist problem instances for high p where element passing can outperform the greedy algorithm by a multiplicative factor of $n+1$. For smaller p, that factor reduces to $2+p/k$. Theorem 6.1 also shows that any element passing algorithm can be outperformed by the greedy algorithm for carefully chosen problem instances, but always by a factor less than 2, and as low as $2-1/n$ for high p.

Theorem 6.2 addresses the practical issue that solving the "local" problem that each agent must solve of (i) selecting the k "best" elements as its action (ii) selecting the p "best" elements to forward to other sensors, can be by themselves intractable problems. This is typically the case for the example scenario described above, when each of the flying vehicles has available a large collection of local measurements. Realistically, the optimal solution to the local problem must be approximated using some computationally-

feasible algorithm. Assuming that agents can approximate the solution to the local problem within a factor of β, Theorem 6.2 shows how these local approximations impact the results of Theorem 6.1. Interestingly, these local approximations do not affect the performance guarantees in a significant way. A summary of the theoretical results can be found in Figure 6.1.

Finally, this chapter provides a numerical example of flying vehicles to show how element passing can improve performance. We show that on average, element passing helps the most with few vehicles (e.g., $n = 3$) and many targets (e.g., 10 targets). However, even with a large number of vehicles and fewer targets (e.g., $n = 10$ and 4 targets), element passing improves performance in the vast majority of the cases (over 99%).

6.1 Chapter Model

Consider a system $H \in \mathcal{H}$ where f is submodular, monotone, and normalized. The underlying set of elements S is partitioned into sets $S_1, \ldots, S_n$, and the initial action set for agent i is $(S_i)^k := \{S' \subseteq S : |S'| \leq k\}$, i.e., each action is a subset of elements of size no larger than k. While the value of k is important to the system, we are interested in system performance as k varies, thus each problem instance of this type can be defined by the tuple $I = (f, S_1, \ldots, S_n)$. We also assume for this chapter that the information sharing constraint graph G is a fully-connected DAG.

Example 6.1 (Flying Vehicles [66]) *Consider the scenario where the agents are cameras carried on board n flying vehicles that capture images of ground targets and return their pixel coordinates. Each vehicle $i \in N$ has access to a large collection of pixel coordinate measurements taken by its own camera, which comprise the local element set S_i. However, each vehicle i needs to select a much smaller subset of these measurements (no*

more than k) to send to a satellite for data fusion. The goal of the vehicles is to select the best set of k measurements that each vehicle should send to the satellite so that an optimal estimate $\hat{\theta}$ of the targets' positions θ can be recovered by fusing the measurements received from all the vehicles.

To facilitate this goal, one can employ the use of the Fisher Information Matrix $\mathrm{FIM}(x)$ *for a set of measurements $x \subseteq S$, which is defined as follows:*

$$\mathrm{FIM}(x) := Q_0 + \sum_{s \in x} Q_s, \tag{6.1}$$

where

$$Q_0 := \frac{\partial \log p(\theta)}{\partial \theta} \cdot \frac{\partial \log p(\theta)}{\partial \theta}^T, \tag{6.2}$$

$$Q_s := \mathbb{E}_\theta \left[\frac{\partial \log p(s|\theta)}{\partial \theta} \cdot \frac{\partial \log p(s|\theta)}{\partial \theta}^T \middle| \theta \right], \tag{6.3}$$

and $p(\theta)$ is the a-priori probability density function of θ and $p(s|\theta)$ is the likelihood of measurement $s \in x$. The positive semidefinite matrices Q_0 and Q_s encode the prior information and the informative contribution of measurement s, respectively. The FIM is helpful in the current setting, given the Cramér-Rao lower bound (CRLB), which states that for an unbiased estimator,

$$\mathbb{E}\left[(\hat{\theta}(x) - \theta)(\hat{\theta}(x) - \theta)^T\right] \geq \mathrm{FIM}(x)^{-1}, \tag{6.4}$$

where we use $\geq$ in the sense that if $A \geq B$, then $A - B$ is positive semidefinite. According to (6.4)*, for any optimal estimator that achieves the CRLB, a set x of measurements that "minimizes"* $\mathrm{FIM}(x)^{-1}$ *also minimizes the error covariance. A scalar metric that is commonly used to measure the information content of a set of measurements is the*

D-optimality [67], which in our context can be defined by

$$f(x) := \log \frac{\det(\text{FIM}(x))}{\det(\text{FIM}(\emptyset))}, \tag{6.5}$$

which has been shown to be submodular [17, 68].

6.1.1 Element Passing

This chapter explores element passing as an extension of the greedy algorithm, where we relax the constraint that the k elements chosen by agent i are a subset of S_i. Specifically, we consider the case where agent i's message is of the form $m_i = \{x_i, z_i\}$, where $z_i \in (S_i)^p$ is a set of up to $p \geq 0$ of the elements in S_i which are sent to the forthcoming agents $j > i$. The subsequent agents $j > i$ can then select their choices from among their original set S_i, but also can include some of the shared elements $z_1 \cup \cdots \cup z_{j-1}$ from previous agents in the sequence. Element passing effectively generalizes the optimization in (1.1) to allow solutions which are not in the original action set $X = (S_1)^k \times \cdots \times (S_n)^k$. This new optimization can be stated as

$$\max_{\substack{z_i \in (S_i)^p \\ x_i \in (S_i \cup z_1 \cup \cdots \cup z_{i-1})^k}} f(x) \tag{6.6}$$

When $p = 0$, (6.6) is equivalent to (1.1), therefore it is also at least NP-Hard in general. [1]

In this chapter we consider decision-making algorithms of the form $\pi = (\pi_1, \ldots, \pi_n)$, where π_i is a rule employed by agent i to select $k \geq 0$ elements and "communicate"

[1] Like (1.1) for submodular f, it can be shown that (6.6) is a submodular maximization problem subject to a matroid constraint, for which there exist generic centralized algorithms that approximate the solution [54, 59]. However, this particular subclass of matroid, which is sequentially constructed, has not been explicitly studied as a submodular maximization constraint, as far as we know.

$p \geq 0$ elements. Such a policy π_i can be informed by the meta-utility function V_i, yet again it is convenient to view decisions as being made via a deterministic policy due to the sequential nature of the system. Specifically, given the decisions of previous agents $x_1^\pi, \ldots, x_{i-1}^\pi$ and elements passed by previous agents $z_1^\pi, \cdots, z_{i-1}^\pi$, π_i specifies both action and communication of appropriate dimension, i.e.,

$$\{x_i^\pi, z_i^\pi\} \in \pi_i^{k,p}(S_i, x_{1:i-1}^\pi, z_{1:i-1}^\pi), \tag{6.7}$$

subject to the constraint that $x_i^\pi \in (S_i \cup z_1^\pi \cup \cdots \cup z_{i-1}^\pi)^k$ and $z_i^\pi \in (S_i)^p$. This constraint ensures that each agent i can only select elements either from its own set S_i or elements shared by previous agents $j < i$. We denote $x^\pi(I)$ to be the resulting decision set for policy π on problem instance I, and abuse notation so that $f(x^\pi(I)) = f(x_1^\pi \cup \cdots \cup x_n^\pi)$. Here we give two example policies to illustrate.

Definition 6.1 (Extended Greedy Policy) *A policy $\pi \in \Pi^{k,p}$ is an extended greedy policy if each agent $i \in \{1, \ldots, n\}$ is associated with a selection rule π_i of the form*

$$x_i^\pi \in \operatorname*{arg\,max}_{\tilde{x} \in (S_i \cup z_1 \cup \cdots \cup z_{i-1})^k} f(x_1^\pi \cup \cdots \cup x_{i-1}^\pi \cup \tilde{x}) \tag{6.8a}$$

$$z_i^\pi \in \operatorname*{arg\,max}_{\tilde{z} \in (S_i)^p} f(x_1^\pi \cup \cdots \cup x_i^\pi \cup \tilde{z}). \tag{6.8b}$$

Here each agent greedily selects the k best elements for x_i based on what elements have previously been selected, and then greedily selects the next best p elements to share as z_i. We note that like (5.1), the rules in (6.8a)–(6.8b) are not deterministic: the arg max may be multivalued. Thus there are many extended greedy policies that satisfy (6.8a)–(6.8b) in conjunction with some tiebreaking rule.

Definition 6.2 (Augmented Greedy Policy) *A policy $\pi \in \Pi^{k,p}$ is an augmented*

greedy policy if each agent $i \in \{1, \ldots, n\}$ *is associated with a selection rule* π_i *of the form*

$$x_i^\pi \in \underset{\tilde{x} \subseteq (S_i \cup z_1^\pi \cup \cdots \cup z_{i-1}^\pi)^k}{\arg\max} f(x_1^\pi \cup \cdots \cup x_{i-1}^\pi \cup \tilde{x}) \tag{6.9a}$$

$$z_i^\pi = z_i^k \cup z_i^{p-k}, \text{ where} \tag{6.9b}$$

$$z_i^k \in \underset{\tilde{z} \in (S_i)^{\min(p,k)}}{\arg\max} f(x_1^\pi \cup \cdots \cup x_i^\pi \cup \tilde{z}) \tag{6.9c}$$

$$z_i^{p-k} \in \underset{\tilde{z} \in (S_i)^{\max(p-k,0)}}{\arg\max} f(x_1^\pi \cup \cdots \cup x_i^\pi \cup z_i^k \cup \tilde{z}) \tag{6.9d}$$

Again, since the local optimizations in (6.9a), (6.9c), and (6.9d) are multi-valued, there are many such policies. An augmented greedy policy can be seen as a modification to the extended greedy algorithm; the difference is in how each z_i is selected when $p > k$. Here, rather than simply selecting the "next best" p elements, the augmented greedy algorithm selects the "next best" k elements followed by the following "next best" $p - k$ elements. When $p \leq k$, the two policies are equivalent. See Figure 6.2 for an example problem instance where an augmented greedy policy is used.

6.2 Comparison to Greedy

In this section, we compare the nominal greedy policies to augmented greedy policies. Since nominal greedy policies are equivalent to the subset of augmented greedy policies where $p = 0$, this comparison illustrates how element passing can improve system performance.

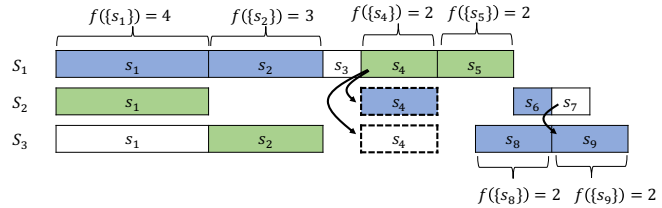

(a) An example problem, where $n = 3$, $k = 2$, $m = 1$. Each box represents an element of S, and each row represents S_i for each agent, i.e., the local elements to which the agent has access. The function f is represented by the width of each box, where the width of elements not specifically labeled in the diagram is 1. For $A \subseteq S$, $f(A)$ is the total amount of horizontal space covered by the elements in A; clearly f is submodular. For instance, $f(\{s_6, s_8\}) = 2$ and $f(\{s_5, s_6, s_8\}) = 3$. Here we assume that π is an augmented greedy policy. The arrows indicate the element passing dictated by π, for instance $z_1^\pi = \{s_4\}$. The boxes with the dashed outline indicate that s_4 is not in S_2 or S_3, but is included as part of the agents' augmented decision set, should they choose to use it. The boxes shaded in blue indicate the elements x_i^π chosen by π, and the boxes shaded in green are the optimal choices, where those differ.

Solution method	x_1	x_2	x_3	$f(x)$
Optimal	$\{s_4, s_5\}$	$\{s_1, s_6\}$	$\{s_2, s_9\}$	14
Greedy	$\{s_1, s_2\}$	$\{s_6, s_7\}$	$\{s_8, s_9\}$	11
Augmented Greedy	$\{s_1, s_2\}$	$\{s_4, s_6\}$	$\{s_8, s_9\}$	13

(b) A table representing the performance for 3 different solution methods. First, the optimal solution to (6.6) is given. Then, the nominal greedy algorithm, where agents choose according to (5.1), is shown. Finally, the last row assumes agents choose according to an augmented greedy policy.

Figure 6.2: An example problem illustrating element passing introduced in Section 6.1.1

6.2.1 Direct Comparison to Greedy

The first comparison we make between the two classes of policies is direct: we compare the ratio between the two for any given problem instance. Given k, p, we denote $\Pi^{k,p}$ to be the set of all admissible policies, i.e., all policies that satisfy (6.7). We also denote $\Pi^{k,p}_{\text{ag}}$ to be the set of augmented greedy policies and denote $\Pi^{k,p}_{\text{ng}}$ to be the set of nominal greedy policies that satisfy (5.1). Note that $\Pi^{k,0}_{\text{ag}} = \Pi^{k,0}_{\text{ng}}$.

Theorem 6.1 *Consider the element selection problem with n agents. Then for any $k \geq 1, p \geq 0$ the best-case gain in performance and worst-case loss in performance associated*

with the optimal element passing policy within $\Pi^{k,p}$ *satisfies*

$$\max_{\pi\in\Pi^{k,p}} \max_{I\in\mathcal{I},\rho\in\Pi^{k,p}_{\text{ng}}} \frac{f(x^\pi(I))}{f(x^\rho(I))} \leq 2+\min(p/k, n-1), \tag{6.10}$$

$$\max_{\pi\in\Pi^{k,p}} \min_{I\in\mathcal{I},\rho\in\Pi^{k,p}_{\text{ng}}} \frac{f(x^\pi(I))}{f(x^\rho(I))} \leq \frac{1}{2-\frac{\min((n-1)p/k,1)}{n-1+\min((n-1)p/k,1)}}, \tag{6.11}$$

where $\mathcal{I}$ *is the set of all system instances* $I=(f,S_1,\ldots,S_n)$*. When restricting attention to augmented greedy policies, the best-case gain in performance and worst-case loss in performance associated with any* $\pi\in\Pi^{k,p}_{\text{ag}}$ *satisfies*

$$\max_{I\in\mathcal{I},\rho\in\Pi^{k,p}_{\text{ng}}} \frac{f(x^\pi(I))}{f(x^\rho(I))} \geq 2+\min(p/k, n-1-1/k), \tag{6.12}$$

$$\min_{I\in\mathcal{I},\rho\in\Pi^{k,p}_{\text{ng}}} \frac{f(x^\pi(I))}{f(x^\rho(I))} \geq \frac{1}{2-\frac{(\min(p/k,1))^{n-1}}{\sum_{i=0}^{n-1}(\min(p/k,1))^i}} \tag{6.13}$$

where the bound in (6.12) *becomes an equality of the form* (6.10) *when* $p\leq nk-k-1$ *and the bound in* (6.13) *becomes an equality of the form* (6.11) *when* $p\geq k$*.*

The theorem proof is given in the next subsection. The bounds given in Theorem 6.1 represent a range of possible values for $f(x^\pi(I))/f(x^\rho(I))$ when $\rho\in\Pi^{k,p}_{\text{ng}}$, $\pi\in\Pi^{k,p}_{\text{ag}}$ and for any problem instance I. If $p=0$, i.e., π is equivalent to a nominal greedy policy, then $2\geq f(x^\pi(I))/f(x^\rho(I))\geq 1/2$, since there exist problem instances where there are (at least) two possible outcomes for the greedy algorithm: the solution to (1.1), and the other a worst-case outcome, which has 1/2 the value of the first outcome. Therefore, one would hope that element passing can increase this upper bound above 2 and lower bound above 1/2. Theorem 6.1 shows that this is the case.

Equation (6.10) gives the upper bound on $f(x^\pi(I))/f(x^\rho(I))$ for any π. As one might expect, this upper bound increases with p: the more element passing is permitted, the higher the possible performance increase. However, when $p>k(n-1)$, the upper

bound remains constant: increasing p above this value no longer increases potential improvement. Equation (6.12) shows that any augmented greedy policy is optimal in this sense when $p \leq nk - k - 1$. Furthermore, the expressions in (6.12) and (6.10) are always within an additive factor of $1/k$, so any augmented greedy policy is also at least near-optimal in this sense.

Equation (6.11) shows that, regardless of policy, one can always carefully construct problem instances where a nominal greedy policy will perform better. In fact, no policy can guarantee that this worst-case performance loss is higher than the expression in (6.11). Again, we see that increasing p increases this lower bound, although here one sees no additional increase when $p > k$. Different from the upper bound is that the expression in (6.11) decreases to $1/2$ as $n \to \infty$, regardless of the value of p. Theorem 6.1 states that any augmented greedy policy is optimal in this sense when $p \geq k$, it is also optimal when $n = 2$ or as $n \to \infty$.

Figure 6.3 gives an illustration of Theorem 6.1 for the case where $k = 8$ and for $n = 3, 4, 5, 6$. The solid colored lines indicate the "optimal" bounds given in (6.10)–(6.11), and the dashed lines indicate the bounds for any augmented greedy algorithm as given in (6.12)–(6.13). For instance, the shaded blue region represents, for $n = 3$, the possible values of $f(x^{\pi}(I))/f(x^{\rho}(I))$ for any $\pi \in \Pi_{\text{ag}}^{k,p}$, $\pi_{\text{ng}}^{k,p}$, and $I \in \mathcal{I}$. The lowest solid black line indicates what is described above: that when $p = 0$, the values of the ratio range between 2 and $1/2$. The middle solid black line represents the range of values when $p = k$; note that in this region any augmented greedy policy is optimal in that no other policy can provide a higher upper bound or a higher lower bound. Finally, the highest black like represents the range of values when $p = 4k$, where the lower bound is still optimal, but the upper bound is only guaranteed to be near-optimal.

While the performance increase that results from element passing is notable, there is some tradeoff with runtime. Note that the nominal greedy policy rule in (5.1) does not

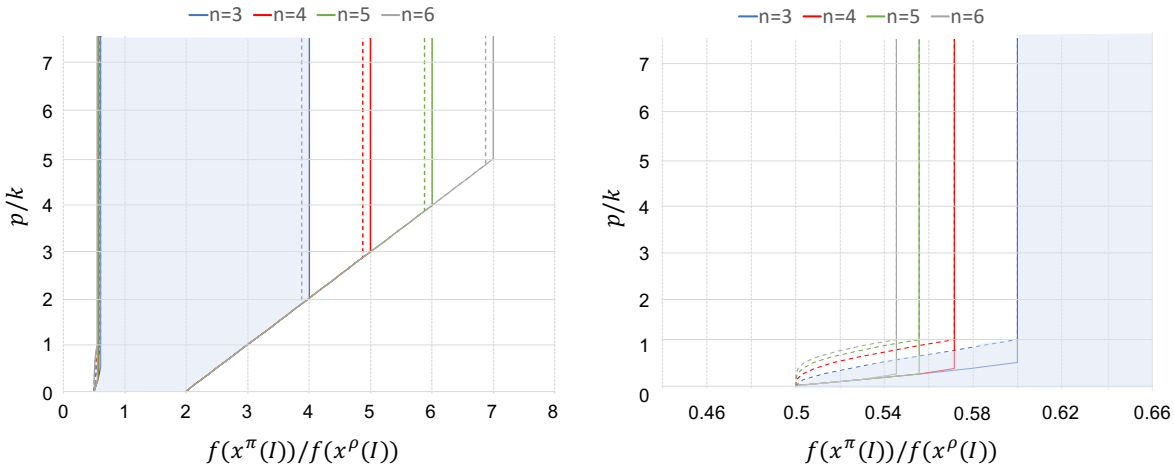

(a) The bounds from Theorem 6.1 for different values of n. For any augmented greedy policy π, the shaded blue region represents all possible values of $f(x^\pi(I))/f(x^\rho(I))$ for the corresponding value of p/k when $n = 3$. For instance, the middle of the 3 black solid lines indicates that when $p = k$, $3 \geq f(x^\pi(I))/f(x^\rho(I)) \geq 0.6$.

(b) A zoomed-in view of the plot in (a) that shows the lower bounds in (6.11) and (6.13).

Figure 6.3: The results of Theorem 6.1, illustrated for the case where $k = 8$ and for $n = 3, 4, 5, 6$. The solid lines indicate the bounds for an optimal message passing policy as given in (6.10) and (6.11) respectively, and the dashed lines indicate the proven lower bounds for any augmented greedy algorithm as given in (6.12) and (6.13). The plot in (b) is a close-up of the various values for in part (a).

prescribe how to solve the local optimization problem, which is intractable in general. A full implementation of the nominal greedy algorithm will require $O\left(n \cdot \binom{\max_i |S_i|}{k}\right)$ number of calls to f. An augmented greedy policy, by comparison, will require $O\left(n \cdot \binom{pn+\max_i |S_i|}{\max(k,p-k)}\right)$. We address this intractability in Section 6.3.

Finally, we note that in the examples which serve to prove the bounds in (6.11) and (6.12), there is some reliance on overlap among the local element sets $S_1, \ldots, S_n$. Clearly, when $S_i = S_j$ for all i, j, then all agents have access to the same information, thus element passing is futile. However, the extent to which this is the case is a topic of future work.

6.2.2 Proof for Theorem 6.1

We begin with two lemmas that, that, given two policies $\pi, \rho \in \Pi^{k,p}$, show how marginal contributions for x_i^π, z_i^π, and x_i^ρ affects $f(x^\pi)/f(x^\rho)$. Denote $x_{a:b} = \cup_{a \le i \le b} x_i$, and likewise for $z_{a:b}$.

Lemma 6.1 *Assume that policies $\rho \in \Pi^{k,p}$ is applied to instance $I \in \mathcal{I}$ and that there exists $\alpha \ge 1$ such that*

$$\alpha \Delta(x_i^\rho | x_{1:i-1}^\rho) \ge \max_{\tilde{x} \in (S_i)^k} \Delta(\tilde{x} | x_{1:i-1}^\rho), \ \forall i \tag{6.14}$$

Then for any $\pi \in \Pi^{k,p}$,

$$\frac{f(x^\pi(I))}{f(x^\rho(I))} \le \begin{cases} 2\alpha + \frac{p}{k}, & \text{if } p \le k \\ 1 + \alpha \left(1 + \min\left(\frac{p}{k}, n-1\right)\right), & \text{if } p > k. \end{cases} \tag{6.15}$$

Lemma 6.2 *Assume that policy $\pi \in \Pi^{k,p}$ is applied to instance $I \in \mathcal{I}$ and that there exist $\alpha_1, \alpha_2 \ge 1$ such that*

$$\alpha_1 \Delta(x_i^\pi | x_{1:i-1}^\pi) \ge \max_{\tilde{x} \in (S_i)^k} \Delta(\tilde{x} | x_{1:i-1}^\pi) \tag{6.16a}$$

$$\alpha_2 \cdot \max_{\tilde{z} \in (z_i^\pi)^k} \Delta(\tilde{z} | x_{1:i}^\pi) \ge \max_{\tilde{x} \in (S_i)^k} \Delta(\tilde{x} | x_{1:i-1}^\pi). \tag{6.16b}$$

Then for any $\rho \in \Pi^{k,p}$ such that $x_i^\rho \subseteq S_i$ for all i,

$$\frac{f(x^\pi(I))}{f(x^\rho(I))} \ge \frac{1}{1 + \alpha_1 - \frac{1}{\sum_{i=0}^{n-1} \alpha_2^i}} \tag{6.17}$$

The proofs for Lemma 6.1 and Lemma 6.2 are given in Appendix-A.6 and Appendix-A.7, respectively. We now prove each statement of the Theorem separately.

Equation (6.10)

Here we invoke Lemma 6.1. Let $\pi \in \Pi^{k,p}$ and $\rho \in \Pi^{k,p}_{\text{ng}}$. Using ρ, agents make choices according to (5.1), therefore for any I, let $\alpha = 1$. Then both expressions in (6.15) are equivalent: $f(x^\pi(I))/f(x^\rho(I)) \leq 2 + \min(p/k, n-1)$.

Equation (6.11)

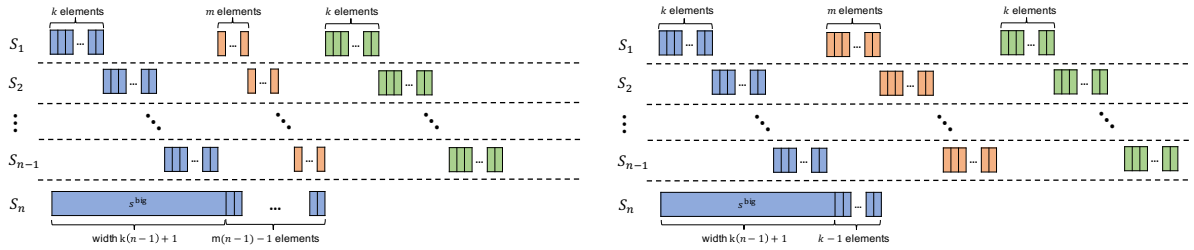

(a) An example for proving (6.11) when $p(n-1) \leq k$. The key is that $|S_n| \leq k$, and the orange elements offer no value beyond S_n. Each small box has width 1 and s^{big} has width $k(n-1)+1$. Thus $f(x^\pi(I)) = k(n-1) + p(n-1)$ and $f(x^\rho(I)) = 2k(n-1) + p(n-1)$.

(b) An example for proving (6.11) when $p(n-1) \geq k$. Unlike the example above, here S_n consists of exactly k elements, thus $f(x^\pi(I)) = k(n-1) + k$ and $f(x^\rho(I)) = 2k(n-1) + k$.

Figure 6.4: An example for proving (6.12). Here the greedy algorithm chooses the green elements which are "covered" by s^{big}, thus $f(x^\rho(I)) = k$ and $f(x^\pi(I)) = 2k + \min(p, k(n-1) - 1)$.

Fix $\pi \in \Pi^{k,p}$ and assume first that $(n-1)p \leq k$. Suppose that f and $S_1, \ldots, S_n$ are as represented in Figure 6.4a. Here the format of example is the same as in Figure 6.2. We assume that all of the small rectangles are of width 1, and that the large rectangle s^{big} is of length $k(n-1)+1$. Essentially, for agent $i \in \{1, \ldots, n-1\}$, all elements in S_i are identical according to f, since there is no horizontal overlap among them, and none of these agents is aware of the elements in S_n.

Assume that π_i selects the blue elements for x_i^π and the orange elements for z_i^π, $i = 1, \ldots, n-1$. Note that $|S_n| = p(n-1) \leq k$, so $x_n^\pi = S_n$ is feasible, and an optimal choice regardless of the previous agents' decisions. This implies that $f(x^\pi(I)) =$

$k(n-1)+p(n-1)$.

On the other hand, consider the decision set of some $\rho \in \Pi_{\text{ng}}^{k,p}$: the rectangles shaded in green for $i < n$, and the blue rectangles for $i = n$. Thus $f(x^\rho(I)) = 2k(n-1)+p(n-1)$, and for this problem instance I,

$$\begin{aligned}\frac{f(x^\pi(I))}{f(x^\rho(I))} =& \frac{k(n-1)+p(n-1)}{2k(n-1)+p(n-1)} \\ =& \frac{1}{2-\frac{(n-1)p/k}{n-1+(n-1)p/k}}.\end{aligned} \tag{6.18}$$

In the case where $p(n-1) \geq k$, consider the example in Figure 6.4b. Here S_i are the same as in Figure 6.4a, for $i = 1, \ldots, n-1$, implying again without that a possible choice for $x_1^\pi, \ldots, x_{n-1}^\pi$ are the respective blue elements. In this case, however, $|S_n| = k$, but note that $f(x_1^\pi, \ldots, x_{n-1}^\pi, x_n) = kn$ for any $x_n \in (S_n)^k$, thus $f(x^\pi(I)) = kn$. The green elements are again a possible greedy policy selection, where different from the blue, showing that $f(x^\rho(I)) = 2k(n-1)+k$. In this case we see that

$$\frac{f(x^\pi(I))}{f(x^\rho(I))} = \frac{kn}{2k(n-1)+k} = \frac{1}{2-1/n}. \tag{6.19}$$

Equations (6.18) and (6.19) establish (6.11) for all cases.

Equation (6.12)

We appeal to an example of the same style as in Figure 6.2, which is illustrated in Figure 6.4. Here S_1 is the union of a set of k green elements, a set of k blue elements, and a set of p orange elements. The sets $S_2, \ldots, S_{n-1}$ are all empty, and $S_n = \{s^{\text{big}}\}$. All elements have value 1, except s^{big}, which has value k. The element s^{big} "covers" the set of green elements, i.e., if A is the set of green elements, then $f(B, \{s^{\text{big}}\}) = k$ for any $B \subseteq A$.

Assume that using a nominal greedy policy ρ, agent 1 selects the k green elements. Then $f(x^\rho(I)) = k$. However, there exists an augmented greedy policy π such that agent 1 selects the k blue elements as x_1^π and the p orange elements as z_1^π. Since the remaining agents have no other alternatives, $\min\{p, k(n-1)-1\}$ orange elements are chosen for $x_2^\pi, \ldots, x_n^\pi$. This implies that $f(x^\pi(I)) = 2k + \min\{p, k(n-1)-1\}$, and that for this problem instance I

$$\frac{f(x^\pi(I))}{f(x^\rho(I))} = 2 + \min\{p/k, n-1-1/k\}. \tag{6.20}$$

Equation (6.13)

We invoke Lemma 6.2 by finding acceptable values of α_1, α_2 which hold for any I. For any $\pi \in \Pi_{\text{ag}}^{k,p}$, (6.9a) implies that $\alpha_1 = 1$ is a valid parameter choice. When $p \geq k$, $\alpha_2 = 1$, since $\max_{\tilde{z} \in (z_i^\pi)^k} \Delta(\tilde{z}|x_{1:i}^\pi) = z_i^k$ from (6.9c). When $p < k$, the following holds:

$$\max_{\tilde{z} \in (z_i^\pi)^k} \Delta(\tilde{z}|x_{1:i}^\pi) = \max_{\tilde{z} \in (S_i)^p} \Delta(\tilde{z}|x_{1:i}^\pi) \tag{6.21a}$$

$$\geq (p/k) \cdot \max_{\tilde{x} \in (S_i)^k} \Delta(\tilde{x}|x_{1:i}^\pi). \tag{6.21b}$$

Therefore, $\alpha_2 = 1/\min(p/k, 1)$ is an acceptable value. Since $\rho \in \Pi_{\text{ng}}^{k,p}$ satisfies that $x_i^\rho \subseteq S_i$, one can use these values for α_1, α_2 (combined with some algebraic manipulation), so that Lemma 6.2 implies (6.13).

■

6.2.3 Benchmark Comparison to Greedy

In this section we compare augmented greedy policies to nominal greedy policies by comparing them each to a the optimal solution to (1.1), to see how measurement sharing increases performance guarantees. When π is a nominal greedy policy, as has

been stated, the resulting solution is within a factor of 1/2 of this benchmark. Therefore, we are interested in how the ability to share measurements can increase the 1/2 bound:

Corollary 6.1 *Consider the measurement selection problem with n sensors. Then for any $k \geq 1$, $p \geq 0$, the worst-case performance guarantee associated with the optimal measurement passing policy within $\Pi^{k,p}$ satisfies*

$$\max_{\pi \in \Pi^{k,p}} \min_{I \in \mathcal{I}} \frac{f(x^\pi(I))}{\mathrm{OPT}(I,k)} \leq \frac{1}{2 - \frac{\min((n-1)p/k,1)}{n-1+\min((n-1)p/k,1)}}, \tag{6.22}$$

where $\mathrm{OPT}(I,k)$ is the value of the solution to (1.1). When restricting attention to augmented greedy policies, the worst-case performance guarantee associated with any $\pi \in \Pi^{k,p}$ satisfies

$$\max_{\pi \in \Pi^{k,p}_{\mathrm{ag}}} \min_{I \in \mathcal{I}} \frac{f(x^\pi(I))}{\mathrm{OPT}(I,k)} \geq \frac{1}{2 - \frac{(\min(p/k,1))^{n-1}}{\sum_{i=0}^{n-1}(\min(p/k,1))^i}}. \tag{6.23}$$

where the bound in (6.23) becomes an equality of the form (6.22) when $p \geq k$.

This result follows from Theorem 5.1, since the defining example in Figures 6.4a and 6.4b can be altered so that the green measurements are the solution to (1.1), and since $f(x^\rho(I)) \leq \mathrm{OPT}(I,k)$ for $\rho \in \Pi^{k,p}_{\mathrm{ng}}$. The main takeaway from this corollary is that when compared to the benchmark (assuming $p > 0$), an augmented greedy policy is strictly better than any nominal greedy policy. In other words, measurement passing always helps. For instance, when $n = 2$ and $p \geq k$, $f(x^\pi(I)) \geq 2/3 \cdot \mathrm{OPT}(I,k)$ for any I and $\pi \in \Pi^{k,p}_{\mathrm{ag}}$, as compared to 1/2 for any nominal greedy policy.

6.3 Suboptimal Selections

To implement an augmented greedy policy, each agent is required to solve the optimizations in (6.9a) and (6.9d), both of which are NP-Hard, in terms of k, assuming

large $|S_i|$. This means that, when k and $|S_i|$ are large, executing an augmented greedy policy may become computationally infeasible—an observation which is well-known for the nominal greedy algorithm [69]. Such scenarios are typical for the the motivating example in Section 6.1, in which the flying vehicles may need to select a large number of images from a much larger set of total images taken.

Real-world implementations of the augmented greedy algorithm must thus approximate (6.9a)–(6.9d). We devote this section to understanding how approximating the solution to such optimizations affects the message passing. The key observation from the results in that follow is that while this approximation increases the range of possible values for $f(x^\pi(I))/f(x^\rho(I))$ (as one might expect), the lower bound decreases (roughly) linearly as a factor of the approximation error. The idea of using approximate maximization for the nominal greedy algorithm has been used previously in the literature (see, for instance [69, 70]) for analyzing approximations to the nominal greedy algorithm, which model and results we extend here for augmented greedy policies.

Definition 6.3 (β-Greedy Policy) *A policy $\pi \in \Pi^{k,p}$ is a β-greedy policy for some $\beta \geq 1$ if each agent $i \in \{1, \ldots, n\}$ is associated with a selection rule π_i of the form*

$$\Delta(x_i^\pi | x_{1:i-1}^\pi) \geq \beta \cdot \max_{\tilde{x} \in (S_i)^k} \Delta(\tilde{x} | x_{1:i-1}^\pi). \tag{6.24}$$

Note that when $\beta = 1$, the nominal greedy policy defined by (5.1) *is recovered.*

An analogous approximation can be defined for the augmented greedy algorithm:

Definition 6.4 ((β_k, β_p)-Augmented Greedy Policy) *A policy $\pi \in \Pi^{k,p}$ is a (β_k, β_p)-greedy policy for some $\beta_k, \beta_p \geq 1$ if each agent $i \in \{1, \ldots, n\}$ is associated with a selection*

rule π_i *of the form*

$$\Delta(x_i^\pi | x_{1:i-1}^\pi) \geq \beta_k \cdot \max_{\tilde{x} \in (S_i \cup z_1^\pi \cup \cdots \cup z_{i-1}^\pi)^k} \Delta(\tilde{x} | x_{1:i-1}^\pi), \tag{6.25a}$$

$$z_i^\pi = z_i^k \cup z_i^{p-k}, \text{ where} \tag{6.25b}$$

$$\Delta(z_i^k | x_{1:i}^\pi) \geq \beta_p \cdot \max_{\tilde{z} \in (S_i)^{\min(m,k)}} \Delta(\tilde{z} | x_{1:i}^\pi), \tag{6.25c}$$

$$\Delta(z_i^{p-k} | x_{1:i}^\pi, z_i^k) \geq \beta_p \cdot \max_{\tilde{z} \in (S_i)^{\max(p-k,0)}} \Delta(\tilde{z} | x_{1:i}^\pi, z_i^k) \tag{6.25d}$$

In essence, a (β_k, β_p)*-augmented greedy policy is a policy which approximates an augmented greedy policy by finding a solution to* (6.9a) *and* (6.9c) *within a factor of* β_k *and* β_p*, respectively, of the optimal. When* $\beta_k = \beta_p = 1$*, the original augmented greedy policy is recovered.*

Theorem 6.2 *Consider a system* $I = (f, S_1, \ldots, S_n)$*. Then for any* (β_k, β_p)*-augmented greedy policy* π*, any* β_k*-greedy policy* ρ*, and any problem instance* I*,*

$$\frac{f(x^\pi(I))}{f(x^\rho(I))} \leq \begin{cases} \frac{2}{\beta_k} + \frac{p}{k}, & \text{if } p \leq k \\ 1 + \frac{1}{\beta_k}\left(1 + \min\left(\frac{p}{k}, n-1\right)\right), & \text{if } p > k, \end{cases} . \tag{6.26}$$

$$\frac{f(x^\pi(I))}{f(x^\rho(I))} \geq \frac{1}{1 + \frac{1}{\beta_k} - \frac{(\beta_p \min(p/k,1))^{n-1}}{\sum_{i=0}^{n-1}(\beta_p \min(p/k,1))^i}}. \tag{6.27}$$

Observe that when $\beta_k = \beta_p = 1$, the results are equivalent to (6.10) and (6.13) in Theorem 6.1. Here we forgo analogous results to (6.11) and (6.12), since the emphasis of this theorem is that while the approximations to the local optimization problems increase the range of possible values for $f(x^\pi(I))/f(x^\rho(I))$, this range is still desirable. For instance, if $\beta_k = \beta_p = 1/2$, regardless of the number of agents, Theorem 6.2 shows that $f(x^\pi(I))/f(x^\rho(I)) \geq 1/3$, as compared to the $1/2$ bound from Theorem 6.1, i.e., the potential loss in performance decreases only a moderate amount. On the other hand,

	$n=2$		$n=3$		$n=4$		$n=5$		$n=6$	
m	upper	lower	upper	lower	upper	lower	upper	lower	upper	lower
0	2.667	0.4286	2.667	0.4286	2.667	0.4286	2.667	0.4286	2.667	0.4286
2	3.667	0.5676	3.667	0.4978	3.667	0.4700	3.667	0.4556	3.667	0.4470
4	3.667	0.5580	5	0.4893	5	0.4628	5	0.4494	5	0.4419
6	3.667	0.5553	5	0.4870	6.333	0.4608	6.333	0.4479	6.333	0.4406
8	3.667	0.5540	5	0.4859	6.333	0.4600	7.667	0.4471	7.667	0.4400

Figure 6.5: Some examples that showcase the results of Theorem 6.2 for $k = 2$. It is assumed that each algorithm is implemented with the sequential greedy rule in (6.28), i.e., $\beta_k = 1 - (1 - 1/k)^k$ and $\beta_p = 1 - (1 - 1/p)^p$.

$f(x^\pi(I))/f(x^\rho(I)) \leq 2n + 1$, an increase over the $n + 1$ bound from Theorem 6.1. In words, element passing can offer a potentially larger benefit without a much higher risk.

An example of a (β_k, β_p)-augmented greedy policy π is where agent i chooses $x_i^\pi = \{s_1, \ldots, s_k\}$ by sequentially selecting element s_l with the following method:

$$s_l \in \underset{\tilde{s} \in S_i \cup z_1^\pi \cup \cdots z_{l-1}^\pi}{\arg\max} \Delta(\{\tilde{s}\}, s_{1:l-1} | x_{1:i-1}^\pi). \tag{6.28}$$

This is yet another variation of greedy algorithm, this time for choosing the k elements of x_i^π. The guarantees for this algorithm are such that $\beta_k = 1 - (1 - 1/k)^k$ [53]. Using a similar method to choose z_i^π yields $\beta_p = 1 - (1 - 1/p)^p$, so that both β_k and β_p are greater than $1 - 1/e \approx 0.63$. The $(1 - 1/e, 1 - 1/e)$-augmented greedy algorithm can now be implemented using $O(pn^2 + \sum_i |S_i|)$ calls to f, and the $(1 - 1/e)$-greedy algorithm can be implemented using $O(\sum_i |S_i|)$ calls to f.

Figure 6.5 illustrates the upper bound shown in (6.26) and the lower bound shown in (6.27) for various values of n and m when $k = 2$. Here we assume that (6.28) is used to implement both algorithms, i.e., $\beta_k = 1 - (1 - 1/k)^k$ and $\beta_p = 1 - (1 - 1/p)^p$. Since β_p decreases as m increases, the lower bound decreases when $p > k$. The upper bound, however, continues to increase in a similar manner to that of Theorem 6.1.

We now give the proof for Theorem 6.2:

Proof: We first show (6.27) by invoking Lemma 6.2: it suffices to show valid values for α_1, α_2 in order to show the lower bound in (6.27). An immediate consequence of (6.25a) is that $\alpha_1 = 1/\beta_k$ holds for all I. Likewise, one can use a similar argument to (6.21a)–(6.21b) to show that $\alpha_3 = 1/(\beta_p \min(p/k, 1))$ holds for all I. Then by Lemma 6.2 (since again ρ is such that $x_i^\rho \subseteq S_i$),

$$\begin{aligned}\frac{f(x^\pi(I))}{f(x^\rho(I))} \geq & \frac{1}{1 + \frac{1}{\beta_k} - \frac{1}{\sum_{i=0}^{n-1}(\beta_p \min(p/k,1))^i}} \\ = & \frac{1}{1 + \frac{1}{\beta_k} - \frac{(\beta_p \min(p/k,1))^{n-1}}{\sum_{i=0}^{n-1}(\beta_p \min(p/k,1))^i}}.\end{aligned}$$

Equation (6.26) can be shown similarly using Lemma 6.1, where π is the (β_k, β_p)-augmented greedy algorithm, ρ is the β_k-greedy algorithm, and, by (6.24), $\alpha = 1/\beta_k$ for all I. ■

6.4 Numerical Example

In this section, we present results for instances of the flying vehicles problem in Section 6.1, where $n = 2$ flying vehicles move on a curved path, each carrying a side-looking camera with a 90° field of view, a 50 pixel focal length, and measurement noise in the image plane with standard deviation $\sigma = 1$ pixel. There are two stationary ground targets whose 2-D location we would like to estimate using the images collected by the flying vehicles. A large number of instances were created with the two targets uniformly randomly placed in the square $[-100, 100] \times [-100, 100]$. The start position, direction, and turn rate of the each flying vehicle's path were also chosen uniformly randomly. Each flying vehicle moves at a constant forward speed and collects 100 independent measurements uniformly along its path. Details of how to construct the corresponding

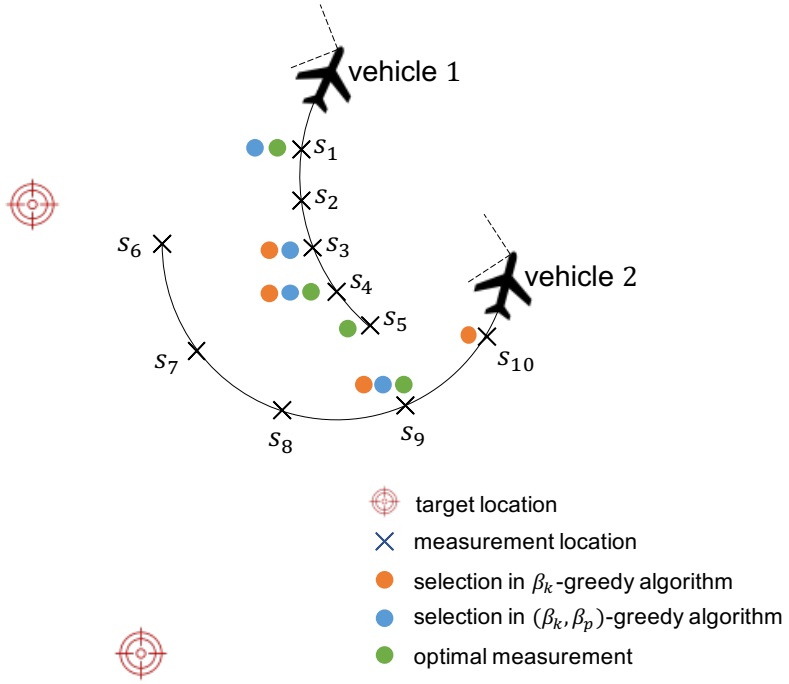

Figure 6.6: An instance of the flying vehicles problem where $n = 2$, $k = 2$, $p = 1$, and $|S_1| = |S_2| = 5$. The black solid lines indicate each vehicle's path, and the black dashed lines indicate each vehicles field of view. The black $\times$'s are where measurements of the red targets are taken; of course, both targets are not always in view of both vehicles. The β_k-greedy (orange dots) and (β_k, β_p)-augmented (blue dots) greedy policies are used in conjunction with the sequential selection rule in (6.28). The green dots indicate measurements that are optimal in the sense that they are a solution to (6.6).

matrices Q_0 and Q_s in (6.2) and (6.3) that quantify the information gain of camera measurements are found in [66]. See Figure 6.6 for an example.

Figure 6.7 summarizes the results in terms of the ratio between the performance of a $(1-(1-1/k)^k, 1-(1-1/p)^p)$-augmented greedy policy and a $1-(1-1/k)^k$-greedy policy, where $k = m = n = 2$ and $|S_1| = |S_2| = 100$. Because the number of measurements is very large, the optimizations in (6.9a), (6.9c), and (5.1) are all approximated by the sequential algorithm in (6.28), where all ties are broken by the index of the measurement.

Additional trials were run for different combinations of the number of flying vehicles and targets. Each combination was repeated 10^5 times with random position and paths. The results are summarized in Figure 6.8. The setup is the same as above, with the

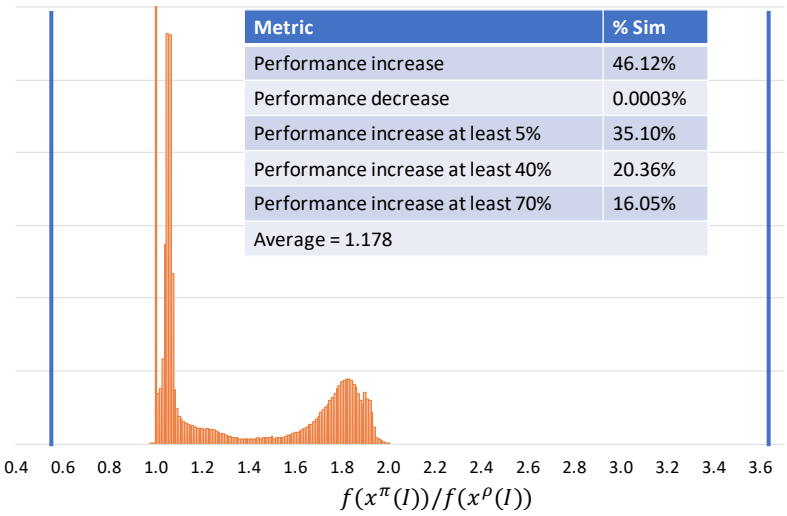

Figure 6.7: Histogram of the relative performance for 10^6 random simulations of the 2 vehicle measurement selection problem. A large number of samples fall at 1, which is interpreted as the two policies performing the same. Instances where the value is greater than one indicate the (β_k, β_p)-augmented greedy performed better than the β_k-greedy in the simulated scenario. The height of the bin corresponding to the ratio of 1 has been cropped. The solid blue bars represent the range of values for $f(x^{\pi}(I))/f(x^{\rho}(I))$ as determined by Theorem 6.2 and in Figure 6.5.

exception that, for the purposes of computation the flying vehicles gather 20 measurements over the course of the simulation rather than 100. One can see that when there are more sensors, the mean benefit of measurement passing decreases, since the measurement passing is most useful when a large percentage of flying vehicles cannot observe any targets. This is more likely to happen with fewer flying vehicles. However, one can also observe from the column labeled "% > 1" that measurement passing is more likely to give some benefit (although smaller) when there are more vehicles, simply because more measurements are being passed, increasing the likelihood that some vehicle will use another's measurement. Likewise, it can be observed from the last column that systems with many sensors are less likely to see a decrease in performance due to measurement passing.

While measurement selection is based on the submodular function (6.5), it is worth

Trials Measuring $f(x^{\pi}(I))/f(x^{\rho}(I))$

# Vehicles	# Targets	Mean	% > 1	% ≥ 1
3	10	1.05	0.625	0.97
3	3	1.045	0.667	0.985
4	4	1.044	0.835	0.972
10	4	1.024	0.996	0.997
2	10	1.041	0.33	1
5	10	1.042	0.895	0.945

Figure 6.8: Results from many different types of trials comparing nominal greedy to augmented greedy. Each row represents a different set of numbers of agents and targets, and there were 10^5 trials run for each row. One can observe that measurement passing helps on average more when there are fewer agents, however, many agents ensure that measurement passing helps more frequently.

noting that the actual improvement in target estimation, such as that described by the commonly used D-optimality estimation criterion of $\det\left(Q_0 + \sum_{s\in S} Q_s\right)$ (without the log function), can be much greater than that shown in Figure 6.7.

Chapter 7

Network Security

We now depart from the common assumptions of the previous chapters: that f is submodular and that there exists some information sharing constraint graph. Instead, we focus on a noncooperative scenario, where the system designer cannot assign the utility function of one of the players. Specifically, we address a network routing game between an agent that wants to route legitimate traffic from a source node to a destination node and another agent that wants to block traffic by flooding the network with malicious traffic. We refer to these players as the *router* and the *attacker*. Motivated by network security problems, we are interested in scenarios of asymmetric information, where the router knows only only the action space of the attacker, but must expose its action to the attacker before the attacker needs to select its action. The problem formulation considered here is motivated by the so-called Crossfire attack in which an attacker persistently degrades network connectivity by targeting a selected set of links within the network, while adjusting to changes in routing policies [71]. The defense against such attacks has been the subject of recent work [72, 73, 74, 75].

The Nash equilibrium is an attractive solution concept for noncooperative systems because it leads to very strong notions of equilibria, in that neither player regrets its

choice after the outcome of the game is revealed [45]. However, such equilibria often do not exist in problems of asymmetric information. The Stackelberg equilibrium is an alternative solution concept where one player (the leader) must select and reveal its policy before the other player (the follower) makes a decision [76]. This type of equilibrium specifically addresses the information asymmetry that we consider here and has been applied to domains closely related to the problem considered in this paper, including network routing [77], scheduling [78], and channel allocation for cognitive radios [79], but also has application in supply chain and marketing channels [80] among other fields. The Stackelberg equilibrium is a concept that is also well-suited for security of critical infrastructure systems [81] and has been applied to surveillance problems that include the ARMOR program at the Los Angeles International Airport [82], the IRIS program used by the US Federal Air Marshals [83], power grid security [84], and defending oil reserves [81]. These two types of equilibria have also been studied extensively for various types of security games [85].

This chapter includes three main contributions:

1. Theorems 7.1 and 7.2 establish that finding the best policy, or action, is an NP-Hard problem for both the router and the attacker.

2. Theorem 7.3 determines conditions on the network under which Stackelberg equilibria lead to no-regret policies (i.e., are also Nash).

3. Section 7.3 explores how uncertainty in knowledge about the capabilities of the attacker translates into performance loss for the router. Theorem 7.4 provides a closed-form expression which quantifies this for a two-link network.

We focus on a network consisting solely of p parallel links that directly connect source and destination. Even within this simple set of networks, the computation of the optimal

attack policy turns out to have higher complexity than one might expect. For any fixed routing policy, we show in Section 7.2 that the computation of the "optimal" distribution of a fixed budget of attack traffic among the parallel links is an NP-hard problem with respect to the scaling parameter p. From the attacker's perspective, "optimal" means that the attacker can prevent as much traffic as possible from reaching the destination, by flooding network links so that legitimate traffic in excess the links' capacity is dropped.

As noted above, Nash equilibria have the desirable feature that they lead to no regret by both players, a feature that is generally not shared by Stackelberg equilibria. It turns out that in the network routing games considered here, Stackelberg equilibria only lead to no-regret (i.e., are also Nash equilibria) in the extreme cases where the attacker controls a very large or a very small amount of traffic. We show this to be true for parallel networks in Section 7.3. For these two extreme cases, we actually provide explicit formulas for the optimal Stackelberg/Nash routing policies. Not surprisingly in view of the NP-hardness result, no explicit formulas are provided for intermediate levels of attack traffic.

Motivated by the nontrivial dependence of the Stackelberg policy on the total amount of traffic r^a controlled by the attacker, we also study how uncertainty in r^a affects routing performance. Previous work in this area has modeled this type of uncertainty as a distribution over the possible values of r^a, giving rise to routing policies that give an optimal expected value on the cost function [86]. However, in this work, we define a metric for the "value of information" about the power of the attacker that compares the amount of traffic that the attacker could block if the router knew precisely r^a versus the amount of traffic it could block if the router had to select a policy without precise knowledge of r^a. The latter scenario generally leads to an increase in blocked traffic. We show in Section 7.4 a closed-form expression for the value of information in two-link networks.

Finally, we explore the scenario where the router knows nothing of r^a. What route

should should be chosen then? We define a notion of robustness and show in Theorem 7.5 that a simple policy maximizes this type of robustness.

7.1 Model

This chapter focuses on a two-player network routing game where the system designer is tasked with deriving a routing policy to maximize the throughput of a given single source / single destination parallel network in the presence of an adversary. The network is comprised of a set of edges E, where each edge $e \in E$ is associated with a given capacity $c_e \geq 0$, and we denote $C(E')$ as the sum of the capacities of all edges in $E' \subseteq E$. Agent 1, which we will henceforth refer to as the *router*, is controlled by the system designer, therefore we consider them to be the same entity. The router must choose a routing profile, or action, $f = \{f_e\}_{e \in E}$ which routes $r \geq 0$ units of traffic across this network. A feasible routing profile satisfies $\sum_{e \in E} f_e = r$ and $0 \leq f_e \leq c_e$ for all edges $e \in E$. We denote the convex set of all admissible routing profiles as $\mathcal{F}(c, r)$ where $c = \{c_e\}_{e \in E}$ denotes the capacities of all edges.

This work considers the existence of an attacker whose goal is to block as much routed traffic as possible by reducing the capacities of the edges in the network through a cross-fire style attack where the attacker can send up to $r_a \geq 0$ units of non-responsive traffic on various edges in the network. An adversarial attack can be characterized by a routing profile $f^a = \{f_e^a\}_{e \in E}$ which satisfies $\sum_{e \in E} f_e^a = r^a$ and $0 \leq f_e^a \leq c_e$ for all edges $e \in E$. We denote the set of all admissible adversarial attack policies as $\mathcal{F}^a(c, r^a)$. We will often refer to r^a as the attack budget of the adversary. Given an admissible routing profile, or action, $f \in \mathcal{F}(c, r)$ and an adversarial attack $f^a \in \mathcal{F}^a(c, r^a)$, the amount of legitimate

traffic blocked on any edge $e \in E$ is defined as

$$B_e(f, f^a, c) := \max\{f_e + f_e^a - c_e, 0\}, \tag{7.1}$$

and the total blocked traffic in the system as $B(f, f^a, c) = \sum_{e \in E} B_e(f, f^a, c)$. Since the routing policy is non-responsive, the adversarial choice effectively reduces the capacity on each edge e from c_e to $c_e - f_e^a$. Lastly, we will often omit highlighting the functional dependence on the parameters c, r, and r^a for brevity, e.g., express $\mathcal{F}^a(c, r^a)$ as merely $\mathcal{F}^a$, when this dependence is clear.

One focus of this paper is to characterize different forms of equilibria in this two-player network routing game. In general, we will assume that a router is required to choose the routing strategy first and the adversary can respond accordingly. The most natural class of equilibria that captures this phenomena is that of Stackelberg equilibria (SE), which consists of any pair of routing profiles (f, f^a) such that

$$f \in \arg\inf_{\bar{f} \in \mathcal{F}} \sup_{\bar{f}^a \in \mathcal{F}^a} B(\bar{f}, \bar{f}^a, c), \tag{7.2}$$

$$f^a \in \arg\sup_{\bar{f}^a \in \mathcal{F}^a} B(f, \bar{f}^a, c). \tag{7.3}$$

If f^a satisfies (7.3), we refer to f^a as a *best response attack* to f. A second class of equilibria that we focus on is Nash equilibria (NE), which focuses on situations where both the router and adversary are required to select their strategy without knowledge of the other's choice. A NE is defined as any pair of profiles (f, f^a) such that

$$f \in \arg\inf_{\bar{f} \in \mathcal{F}} B(\bar{f}, f^a, c), \tag{7.4}$$

$$f^a \in \arg\sup_{\bar{f}^a \in \mathcal{F}^a} B(f, \bar{f}^a, c). \tag{7.5}$$

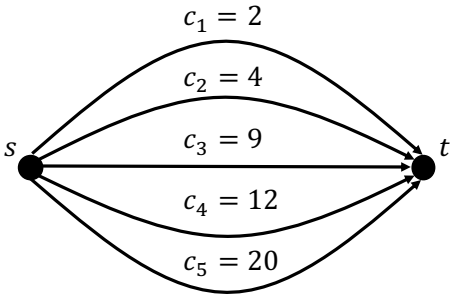

(a) An example network. Suppose that $f = \{1, 1, 5, 10, 8\}$ and $f^a = \{2, 4, 4, 4, 6\}$. Then, for instance on edge 4, since the capacity is 12, 2 units of traffic are blocked. In total we see that $B_1 = 1$, $B_2 = 1$, $B_3 = 0$, $B_4 = 2$, and $B_5 = 0$, which results in $B(f, f^a, c) = 4$.

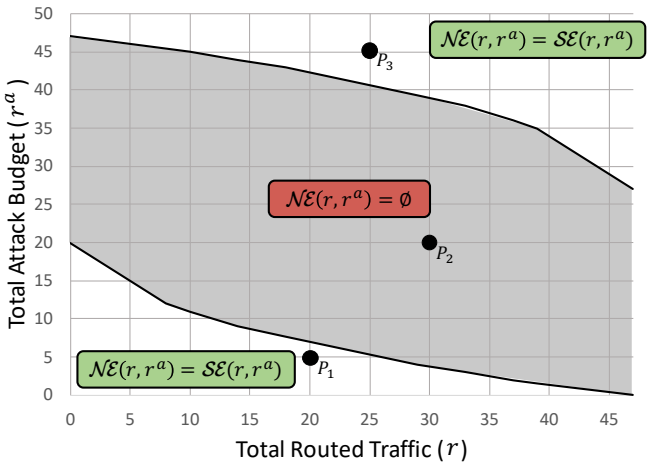

(b) This figure showcases one of the contributions of this paper: a characterization of where no NE exist (gray region) and when all SE are also NE (white regions) for the network in (a). For any values (r, r^a), one of those two properties must hold. See Example 7.1 and Theorem 7.3 for more details.

Figure 7.1: An example network showcasing the model and regions of (r, r^a) where NE exist.

We refer to $\mathcal{SE}(r, r^a)$ as the set of all SE for values r, r^a, and likewise $\mathcal{NE}(r, r^a)$ for NE. Note that given the definitions above, $\mathcal{NE}(r, r^a) \subseteq \mathcal{SE}(r, r^a)$. In the event where $\mathcal{NE}(r, r^a) = \mathcal{SE}(r, r^a)$, this implies that the router is not strategically disadvantaged by having to reveal its choice before the adversary selects its policy. However, while a SE will always exist, the same does not hold true for NE. Furthermore, this chapter will address how knowledge of the exact value of r^a impacts the existence and efficacy of such equilibria.

The set of systems described above can be mapped into our general model presented in Section 1.1, where $N = \{1, 2\}$, $X_1 = \mathcal{F}$, $X_2 = \mathcal{F}^a$, $U_1 = -B$, $U_2 = B$, $m_1 = f$, and $m_2 = r^a$. However, the notation and language presented above is more natural for this setting and its proofs and results.

Example 7.1 *We begin with the following example highlighting the complexity of com-*

puting NE and SE in such a routing game. To that end, consider the example shown in Figure 7.1a with $r = 25$ *and* $r^a = 20$ *and denote the edge set as* $E = \{1, 2, \ldots, 5\}$ *and edge capacities as* $c = \{2, 4, 9, 12, 20\}$*. Given a routing profile* $f = \{1, 1, 5, 10, 8\}$ *and an attack profile* $f^a = \{2, 4, 4, 4, 6\}$*, it follows from* (7.1) *that the traffic blocked on each edge is 1, 1, 0, 2 and 0, respectively. Note that these strategy profiles* (f, f^a) *neither capture a NE or SE as there are numerous adversarial strategies that could increase the total blocked traffic given the routing profile* f*, e.g.,* $\bar{f}^a = \{0, 0, 8, 12, 0\}$*.*

The plot in Figure 7.1b highlights the distinction between NE and SE for the considered routing problem for all pairs (r, r^a) *satisfying* $47 \geq r, r^a \geq 0$*. For instance, when* $r = 20$ *and* $r^a = 5$ *(see point* P_1 *in Figure 7.1b), any SE is also a NE. One such routing profile is* $f = \{0, 0, 0, 5, 15\}$*, as this does not allow the attacker to block any traffic. When* $r = 25$ *and* $r^a = 45$ *(point* P_3*), we see a similar phenomenon, where the attacker has much more power. In fact, observe that the routing profile* $f = \{2, 4, 6.\bar{3}, 6.\bar{3}, 6.\bar{3}\}$ *and attack profile* $f^a = \{0, 4, 9, 12, 20\}$ *constitute both a SE and NE. The router is able to design a policy such that the attacker can only block* $\sum_{e \in E} c_e - r^a$ *traffic, the best the router can achieve given* r^a*. Thus the router has no incentive to deviate, and clearly the attacker cannot. Lastly, when* $r = 30$ *and* $r^a = 20$ *(point* P_2*) we begin to notice a discrepancy between NE and SE in the sense that given any profiles* (f, f^a)*, if* (7.5) *is satisfied then* (7.4) *is not satisfied. For example, consider the profiles* $f = \{1.4, 4, 6.4, 6.4, 11.8\}$ *and* $f^a = \{0, 0, 0, 0, 20\}$ *and note that* f^a *satisfies* (7.5)*. If the attacker implements this policy, then* (f, f^a) *is not a NE, since the router would benefit unilaterally by moving some traffic from edge 5 to another unblocked edge. The forthcoming Theorem 7.3 provides the characterization shown in Figure 7.1b.*

7.2 Problem Hardness

In this section, we show that finding the best response attack policy and finding the optimal routing policy are both NP-Hard. We begin with the attacker:

Problem 7.1 *Given a parallel network with edges E, corresponding capacities c, a routing policy f, and attack power r^a, find f^a which satisfies* (7.3)*, i.e., a best response attack policy.*

Note that an instance of the problem can be defined by (E, c, f, r^a), and we show how the complexity of the problem scales with the number of edges in the parallel network.

Theorem 7.1 *Problem 7.1 is NP-Hard on the scaling variable $|E|$.*

The theorem is proved by reducing the 0-1 Knapsack Problem (KP), a known NP-Hard problem, to Problem 7.1. We do this by showing that if all f_e are "sufficiently small", then any best response attack must either block all traffic on an edge or block none of it. Thus finding the best response attack is simply finding the set of edges to fully block, corresponding to the discrete nature of the items in the 0-1 KP. This implies any method for solving these instances of Problem 7.1 will also solve the 0-1 KP.

The following lemma defines "sufficiently small" in this context:

Lemma 7.1 *Consider an instance of Problem 7.1 (E, c, f, r^a), where*

$$f_e < \min_{E' \subseteq E : r^a - C(E') > 0} r^a - C(E'), \tag{7.6}$$

for some $e \in E$. Then $B_e(f, f^a) \in \{0, f_e\}$ for any f^a which is a solution to Problem 7.1.

Proof: We prove the contrapositive statement. Let e be such that $B_e(f, f^a) \notin \{0, f_e\}$. Define $E^{\text{block}} := \{e' \in E : f_{e'} + f^a_{e'} > c_e\}$, and observe by definition that $e \in E^{\text{block}}$. Then

it must be true that $f_e > r^a - C(E^{\text{block}} \setminus \{e\}) > 0$*, otherwise the attacker could block more routed traffic by redistributing as much attack traffic as possible from* e *to the other edges in* E^{block}*. Therefore,* (7.6) *must be false.* ■

Given this, we proceed with the proof of Theorem 7.1. The 0-1 KP can be defined as follows: assume we have n items, where each item e has a cost w_e and a value v_e. Given a total cost constraint Q, find the combination of items with maximum total which does not exceed W. More formally stated, determine

$$\begin{aligned} \underset{x}{\text{maximize}} \quad & \sum_e v_e x_e \\ \text{subject to} \quad & \sum_e x_e, w_e \leq Q, \quad x_e \in \{0,1\}, \end{aligned} \tag{7.7}$$

where $x := [x_e]$. This problem is known to be NP-Hard in the number of items [87].

Mapping a 0-1 KP to Problem 7.1 can be done with the following method: let every item be mapped to an edge in a parallel network, $r^a = Q$, $c_e = w_e$, and $f_e = \varepsilon v_e$, where $\varepsilon > 0$ satisfies

$$\varepsilon v_e < \min_{E' \subseteq E : r^a - C(E') > 0} r^a - C(E'), \tag{7.8}$$

for all $e \in E$. By Lemma 7.1, we know that any solution to this subset of instances of Problem 7.1 has the property that every edge will either have all routed traffic blocked or none. Therefore, the problem can be reformulated as

$$\begin{aligned} \underset{x}{\text{maximize}} \quad & \sum_e f_e x_e \\ \text{subject to} \quad & \sum_e x_e c_e \leq r^a, \quad x_e \in \{0,1\}. \end{aligned} \tag{7.9}$$

This problem yields an equivalent solution to that in (7.7), since the constraints are the same, and each objective function is a scaled version of the other. Thus solving this

instance of Problem 7.1 will also solve 0-1 KP and shows that Problem 7.1 is NP-Hard. ■

Next we show that finding the SE route in (7.2) is also an NP-Hard problem. Although we have shown above that finding the inner supremum in (7.2) is NP-Hard for the space of all routing policies, it does not follow immediately that (7.2) is as difficult. The hardness of (7.2) is shown by reducing the partition problem to it. The problem is formally defined as follows:

Problem 7.2 *Given a parallel network with edges E, corresponding capacities c, total traffic to be routed r, attack power r^a, find f which satisfies* (7.2)*, i.e., a routing policy which minimizes the amount of traffic that can be blocked by the attacker.*

An instance of Problem 7.2 is defined by the tuple (E, c, r, r^a), and we will again show how the complexity of the problem scales with the number of edges in a parallel network.

Theorem 7.2 *Problem 7.2 is NP-Hard on the scaling variable $|E|$.*

In order to prove hardness, we will leverage the following two lemmas.

Lemma 7.2 *Consider the* even routing policy*, which is defined as $f_e = c_e r/C(E)$ for all $e \in E$. Then for any attack $f^a \in \mathcal{F}^a$*

$$\frac{r^a r}{C(E)} \geq B(f, f^a) \geq B^{\mathrm{SE}}(r, r^a), \tag{7.10}$$

where $B^{\mathrm{SE}}(r, r^a)$ is the amount of traffic blocked in the SE.

Lemma 7.3 *Consider a parallel network where (f, f^a) is a SE and denote $E(f_e^a) = \{e \in E : B_e(f, f^a) > 0\}$. Then there exists another attack policy $\tilde{f}^a$, where $E(f^a) \neq E(\tilde{f}^a)$, such that $(f, \tilde{f}^a)$ is also a SE.*

Lemma 7.2 provides an upper bound on how much traffic can be blocked when the router employs the even routing policy in terms of system parameters r, r^a, and $C(E)$. Of course, this is also an upper bound for B^{SE}. Lemma 7.3 gives insight into optimal attack policy: that there must be at least two such policies which block traffic on a different set of edges. We now proceed with the proof for Theorem 7.2.

Proof: The partition problem, a known NP-complete problem, can be reduced to Problem 7.2 in polynomial time. The partition problem is, given a multiset of positive integers $S = \{s_1, \ldots, s_n\}$, to determine whether there exists a partition S_1, S_2 on S such that the sum of the integers in S_1 equals the sum of the integers in S_2. An instance of the partition problem can thus be defined by the multiset S. Note that if such a partition exists, then the sum of each partition is $T/2$, where T is the sum of all the integers in S. If $T/2$ is non-integer, then immediately the answer to the partition problem is "no", therefore, we assume that $T/2$ is an integer.

An instance of the partition problem can be mapped to an instance of Problem 7.2 by letting $r = 1$, $r^a = T/2$, and E be a set of edges such that $|E| = |S| + 1$. Then $c_i = s_i$ for $i \leq n$ and $c_{n+1} = T/2$. Note that this mapping is completed in a polynomial number of steps with respect to the scaling variable. Also, since all capacities are integer and $r = 1$, Lemma 7.1 implies that $B_e(f, f^a) \in \{0, f_e\}$ for e, for all f, and for all f^a which satisfy (7.3).

We claim that for a solution f to the instance of Problem 7.2 described above, $f_{n+1} = 1/3$ if and only if the answer to the corresponding partition problem is "'yes". Suppose first that $f_{n+1} = 1/3$. One possible attack policy f^a is $f^a_{n+1} = T/2$ and $f^a_e = 0$ for $e \leq n$, which blocks $1/3$ of the routed traffic; thus the optimal attack must block at least $1/3$. However, since $r^a r/C(E) = 1/3$, it follows from Lemma 7.2 that 1) $B^{\text{SE}}(r, r^a) = 1/3$ (i.e., f^a is an optimal attack) and 2) the even policy is an optimal routing policy. Even though f might not be the even policy, one can use the fact that the even policy is

optimal to reason about the edge capacities: suppose that the router selected the even policy, denoted $\tilde{f}$. Lemma 7.3 implies that there must be another optimal attack $\tilde{f}^a$ that excludes edge $n+1$, but still blocks 1/3. Partition the edges in $E \setminus \{n+1\}$ into two sets S_1, S_2, where $S_1 = \{e \in E : B(\tilde{f}, \tilde{f}^a) = \tilde{f}_e\}$. Since 1/3 is blocked,

$$\frac{1}{3} = \sum_{e \in S_1} B_e(\tilde{f}, \tilde{f}^a) \tag{7.11}$$

$$= \sum_{e \in S_1} \tilde{f}_e \tag{7.12}$$

$$= \sum_{e \in S_1} \frac{c_e r}{C(E)} \tag{7.13}$$

$$\frac{C(E)}{3r} = \sum_{e \in S_1} c_e \tag{7.14}$$

$$\frac{T}{2} = \sum_{e \in S_1} c_e, \tag{7.15}$$

therefore S_1, S_2 is a partition that satisfies the conditions of the partition problem.

The other direction can be proven by contradiction: suppose that $f_{n+1} \neq 1/3$ and that the answer to the partition problem is "yes", with a partition S_1, S_2. By Lemma 7.2, $B^{\text{SE}}(r, r^a) \leq 1/3$, so it follows that $f_{n+1} < 1/3$. This implies that $\sum_{e \in S_1 \cup S_2} f_e > 2/3$, so either $\sum_{e \in S_1} f_e > 1/3$ or $\sum_{e \in S_2} f_e > 1/3$. Then there exists an attack f^a such that $B(f, f^a) > 1/3$, a contradiction.

■

7.3 Equilibria

In this section, we present results that describe precisely the relationship between SE and NE in our model. We then give a two examples: one which illustrates this relationship and one which showcases why generalizing these results to even slightly more complex

networks is nontrivial.

Theorem 7.3 *Consider a parallel network with capacities c, routing demand r, and adversarial routing power r^a. The set of Nash Equilibria $\mathcal{NE}(r, r^a)$ is nonempty and $\mathcal{NE}(r, r^a) = \mathcal{SE}(r, r^a)$ if and only if one of the following is satisfied:* [1]

$$r^a \leq \max_{E' \subseteq E} \frac{C(E') - r}{|E'|} \tag{7.16}$$

$$r^a \geq C(E) - \max_{E' \subseteq E} \frac{r - C(E \setminus E')}{|E'|}. \tag{7.17}$$

The proof for Theorem 7.3 can be found in AppendixA.11. An implication of this proof is that the following routing profiles are of some importance:

$$f_e^{\text{lo}} := \max \left\{ c_e - \max_{E' \subseteq E} \frac{C(E') - r}{|E'|}, 0 \right\} \tag{7.18}$$

$$f_e^{\text{hi}} := \min \left\{ c_e, \max_{E' \subseteq E} \frac{r - C(E \setminus E')}{|E'|} \right\}, \tag{7.19}$$

namely, that when (7.16) holds, then (7.18) is a SE (and NE) routing profile; and that when (7.17) holds, (7.19) is a SE (and NE) routing profile.

Refer again to the network in Figure 7.1. At point P_1, $r = 20$ and $r^a = 5$. Here we calculate $\max_{E' \subseteq E}(C(E') - r)/|E'| = 7$, which means that r^a satisfies (7.16). Thus the router can use the policy $f^{\text{lo}} = \{0, 0, 2, 5, 13\}$ to ensure that the attacker cannot block any traffic. By Theorem 7.3, this also implies that (f, f^a) is both a SE and a NE for any $f^a \in \mathcal{F}^a(c, r)$. At point P_3, $r = 25$ and $r^a = 45$. Here we calculate $C(E) - \max_{E' \subseteq E}(r - C(E \setminus E'))/|E'| = 40.\bar{6}$, which means that r^a satisfies (7.17). Thus

[1] While finding the maxima in (7.16) and (7.17) may appear to be computationally intractable given the number of edges in the network, it is true that the maximizing E' for both (7.16) and (7.17) is of the form $\{1, 2, \ldots, k\}$, where the edges are ordered starting with highest capacity to the lowest. Therefore, finding either maxima is equivalent to finding the best value of k, which can be completed in linear time.

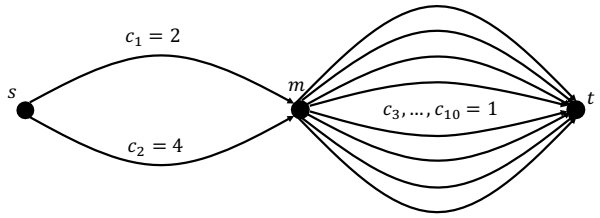

Figure 7.2: Two parallel networks in series. We use this example to illustrate the complexities for finding SE and NE in more general networks than just parallel. For instance, one cannot simply decompose the optimal attack problem into either attacking the set of edges between s and m, and attacking the edges between m and t. Even if we limited our scope to such attacks, which set of edges to attack depends on the value of r^a, not merely on f and c. See Example 7.2 for more details.

the router can use the policy $f^{\text{hi}} = \{2, 4, 6.\bar{3}, 6.\bar{3}, 6.\bar{3}\}$, and from Theorem 7.3, (f, f^a) is a NE and SE for any $f^a \in \mathcal{F}^a(c, r)$. At point P_2, $r = 30$ and $r^a = 20$. We calculate that $\max_{E' \subseteq E}(C(E') - r)/|E'| = 3.75$ and $C(E) - \max_{E' \subseteq E}(r - C(E \setminus E'))/|E'| = 59$, therefore r^a does not satisfy (7.16) or (7.17). By Theorem 7.3, we know that no NE can exist at this point.

Example 7.2 *Consider now the example in Figure 7.2, a graph where two parallel networks are connected in series. We present this as a simple example to showcase the complexities that arise when studying the SE of non-parallel networks. For more complex networks, one might think that finding a best response attack could be limited to attacking a minimal cut-set in the network. However, even in this very simple example, we show that this is not the case, and in fact, a best response attack will often incorporate edges of multiple cut-sets in the network. Thus investigating parallel networks in this paper gives a natural simplification of the problem in order to address the questions of interest.*

In Figure 7.2, denote E_{sm} as the cut-set of edges between s and m and E_{mt} as the cut-set of edges between m and t. Observe that regardless of the attacker's capability, there always exists a SE route where all edges in E_{mt} have the same amount of traffic routed on them. We assume in the following cases that the router always uses such a policy, and therefore, we need only focus on the routing strategy across E_{sm}.

Let $r = 2$ and $r^a = 5$. If the attacker restricts its attacks to a single cut-set E_{sm} or E_{mt}, then the router can choose its policy accordingly, for instance $f_e = 1$ for $e \in E_{sm}$, and $f_e = 0.25$ for $e \in E_{mt}$. Note that across each cut-set, this route satisfies (7.2). Attacking only E_{sm}, the attacker can block 1 unit of traffic, but attacking only E_{mt}, the attacker can block 1.25 units of traffic. This may seem unintuitive, since the total capacity of E_{sm} is less than that of E_{mt}. Furthermore, the best response for the attacker is to block some traffic on E_{sm} and some on E_{mt}. For instance, the attacker could block the 1 unit of traffic on edge 1, and then block all traffic on 3 of the edges in E_{mt}. Assuming that the router evenly distributes the remaining 1 unit of routed traffic that arrives at node m, this attack would block 1.375 units of traffic. Therefore, solving for a SE must include all attacks across multiple cut-sets.

Given these complexities with even very simple non-parallel networks, the characterizations of SE and NE in Theorem 7.3 only apply to parallel networks. While this class of networks is sufficiently rich to ask the questions and showcase the phenomena that are relevant to this work, future work can ask similar questions in a broader setting.

7.4 The Value of Information

In this section, we present preliminary results about the value to the router of knowing information about the attack power r^a. In order to do this, we introduce some notation. We define

$$B^*(f, r^a) := \max_{f^a \in \mathcal{F}^a(r^a)} B(f, f^a), \tag{7.20}$$

in other words, $B^*(f, r^a)$ measures how much traffic is blocked in the attacker's best response to f, given r^a. We also define

$$B^{\mathrm{SE}}(r, r^a) := B(f, f^a), \tag{7.21}$$

where $(f, f^a) \in \mathcal{F}(r) \times \mathcal{F}^a(r^a)$ is a SE. Recall that for the pair (r, r^a) the same amount of traffic will be blocked by any SE (f, f^a).

As an example of both these functions, consider the plot in Figure 7.3 for a three-link parallel network where $c = \{2, 3, 5\}$ and $r = 5$. For the fixed route $f = \{0.5, 2, 3.5\}$, the gray line represents how $B^*(f, r^a)$ changes as a function of r^a. Likewise, the orange line showcases $B^{\text{SE}}(r, r^a)$ as a function of r^a. Observe that $B^*(f, r^a) \geq B^{\text{SE}}(r, r^a)$ for all values of r^a.

7.4.1 Limited information

We limit the router's knowledge of r^a by stating that the router only knows that r^a is in some interval $\pi^a = [\underline{\pi}^a, \overline{\pi}^a]$. In light of this uncertainty, if the router chooses policy f, then we can define the risk of f on interval π^a as

$$R(f, \pi^a) := \max_{r^a \in \pi^a} \left(B^*(f, r^a) - B^{\text{SE}}(r, r^a) \right). \tag{7.22}$$

Intuitively, the value $B^*(f, r^a) - B^{\text{SE}}(r, r^a)$ represents how much more traffic the attacker is able to block because the router chose policy f instead of a SE policy for that value of r^a. Thus the risk $R(f, \pi^a)$ is the maximum such value across all $r^a \in \pi^a$. In other words, this measurement of risk shows, in the worst case, the advantage that the attacker gains by the router not knowing the true value of r^a.

As an example, consider again the plot in Figure 7.3. If we assume that the router has no knowledge of r^a (i.e., $\pi^a = [0, 10]$), then the risk associated with the route $f = \{0.5, 2, 3.5\}$ is the maximum difference between the gray and orange lines, which is achieved at $r^a = 8$. Therefore, in this case we see that $R(f, \pi^a) = 1.5$.

It turns out that the maximization in (7.22) can be restricted to a finite set of points in π^a.

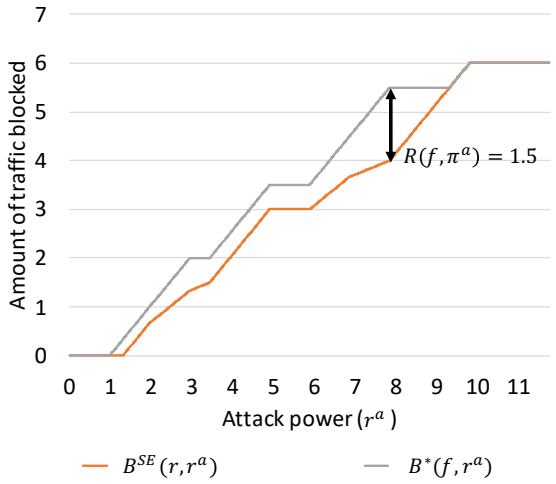

Figure 7.3: A plot showing the amount of traffic blocked by an optimal attack for a SE routing policy (orange) versus the traffic blocked when the router selects specific routing policy f (regardless of the value of r^a). Here $c = \{2, 3, 5\}$, and $r = 5$. The fixed policy represented by the gray line is $f = \{0.5, 2, 3.5\}$. We show the values of the risk $R(f, \pi^a)$ for $\pi^a = [0, 10]$.

Lemma 7.4 *For a parallel network,*

$$R(f, \pi^a) = \max_{r^a \in (\alpha \cap \pi^a) \cup \{\underline{\pi}^a, \overline{\pi}^a\}} B^*(f, r^a) - B^{\mathrm{SE}}(r, r^a), \tag{7.23}$$

where α is the finite set $\{r^a : \exists E' \subseteq E$ where $r^a = C(E')\}$, which has at most $2^{|E|}$ elements.

The full proof is given in AppendixA.10, however here we provide some intuition: consider the plot in Figure 7.3. The orange line, $B^{\mathrm{SE}}(r, r^a)$, is piecewise linear, with no line slope being greater than 1. The grey line, $B^*(f, r^a)$, is also a piecewise linear function, with lines slopes either 0 or 1. The value of the risk $R(f, \pi^a)$ is incurred at $r^a = C(\{2, 3\}) = 8$, where the attacker's best response against f is to fully block edges 2 and 3. Because the two lines are piecewise linear, the largest distance must take place at one of the points of discontinuity for the gray line inside the interval π^a.

Finally, we define the *value of information* to the router for an interval π^a as the

minimum amount of risk that can be incurred for any routing policy. More formally stated,

$$V(\pi^a) := \min_{f \in \mathcal{F}} R(f, \pi^a) \tag{7.24}$$

$$= \min_{f \in \mathcal{F}} \max_{r^a \in \pi^a} \left(B^*(f, r^a) - B^{\mathrm{SE}}(r, r^a) \right) \tag{7.25}$$

We also denote the routing policy which minimizes (7.24) by f^π. This value of information is meant to reflect how valuable (i.e., how much less traffic would be blocked) if the router knew the exact value of r^a. For instance, if $V(\pi^a) = 0$, then there exists a route which satisfies (7.2) for any value of $r^a \in \pi^a$, thus the router does not need to know the exact value. However, when $V(\pi^a)$ is high, knowing r^a would allow the router to ensure that less traffic is blocked. Figure 7.4 shows $V(\pi^a)$ and f^π for a two-link network.

7.4.2 The Value of Information in Two-Link Networks

Lemma 7.4 provides a numerical procedure to compute the risk for a routing policy f against an attack power interval π^a for general parallel networks. For two-link networks, this means there exists a closed-form solution for $R(f, \pi^a)$ and subsequently $V(\pi^a)$.

Theorem 7.4 *Consider a two-link parallel network, where $c_1 \leq c_2$. Suppose that the router only knows that $r^a \in \pi^a = [\underline{\pi}^a, \overline{\pi}^a]$. Then the value of information is*

$$V(\pi^r) = \begin{cases} 0, \text{ if } \pi^a \cap [c_1, c_2] = \emptyset \\ \frac{1}{4}(\min\{\overline{\pi}^a, c_2\} - \max\{\underline{\pi}^a, c_1\}), \text{ otherwise.} \end{cases} \tag{7.26}$$

The theorem proof is given in AppendixA.12, but here we give an example to provide some intuition. Consider the plot in Figure 7.4. In this network, $c = \{3, 6\}$, and $r = 5$. If the router knows the exact value of r^a, it can choose a SE routing policy, which will

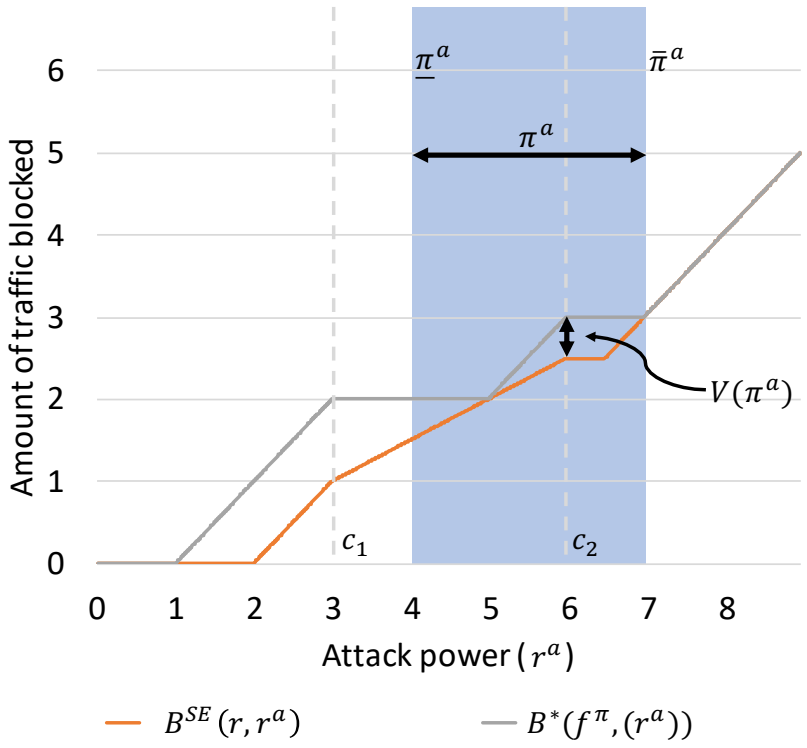

Figure 7.4: A plot with an example two-link parallel network that shows a graphical interpretation for $V(\pi^a)$. The edge capacities in the network are $\{3, 6\}$, and $r = 5$. The orange line represents how much traffic is blocked at the SE for each value of r^a, and the gray line is how much is blocked by a best response attack against $f^\pi = \{2, 3\}$ for each value of r^a. The interval $\pi^a = [4, 7]$ is the blue shaded region. The value of information $V(\pi^a)$ is the maximum difference between the two lines within the blue region.

make the difference between the lines 0 at that value of r^a. If we assume that the router only knows that $r^a \in \pi^a = [4, 7]$, then it must choose a policy to mitigate the risk associated with that loss of information. In this scenario, the router's best option is to use $f^\pi = \{2, 3\}$ (gray line), which minimizes the maximum difference between the two lines on π^a. The value of the router knowing r^a is then this minimum maximum difference, i.e., $V(\pi^a) = 0.5$.

7.5 A Robust Policy

In this section, we assume that the router has access to even less information: the attacker's budget (or presence) is unknown. Under this assumption, the router must fix its policy in the presence of high uncertainty. Of course, if the router takes an optimistic view and assumes that r^a is low, the routing policy f^{lo} is a good choice. If the router takes a pessimistic view and assumes that r^a is high, then f^{hi} is a good choice. However, if the router is wrong about its view, i.e., thinking r^a is low when it's high, this could lead to a big loss in traffic. Therefore, the router must choose a policy that is robust in some sense to any possible value of $r^a \in [0, C(E)]$. Here we measure robustness of a routing policy $T(f)$ as

$$T(f) = \inf_{r^a \in (0, C(E)]} \frac{r^a}{B^*(f, r^a)} \geq 1. \tag{7.27}$$

A high value of $T(f)$ ensures that a small percentage of the attack budget is actually blocking traffic for any value of r^a. On the other hand, if $T(f) = 1$, that means that there exists some value of r^a such that $B^*(f, r^a) = r^a$. The following result shows that the even policy is uniquely the policy that maximizes $T(f)$.

Theorem 7.5 *Suppose that f is the even policy. Then $T(f) = C(E)/r$. Furthermore, no routing policy can guarantee a higher value of $T(f)$ for any network.*

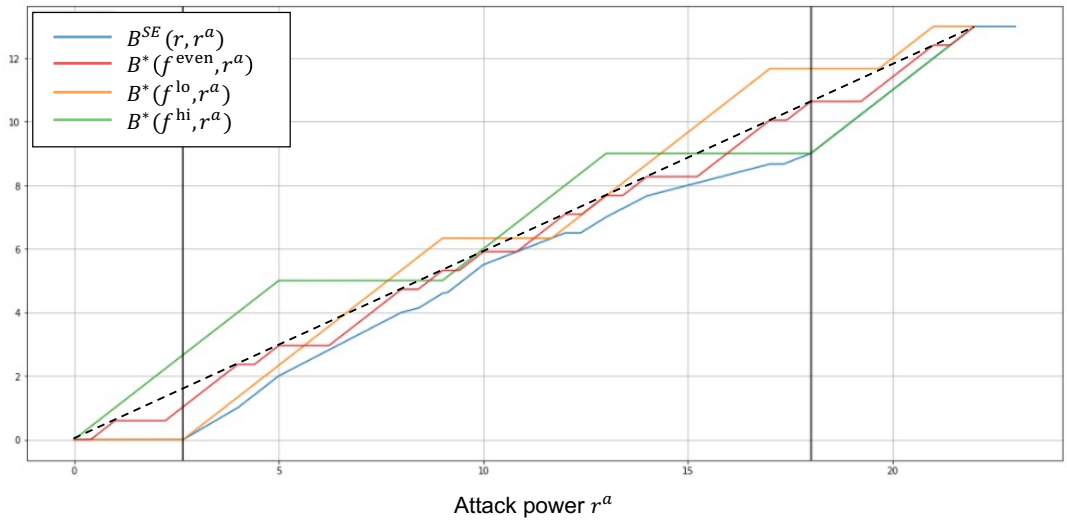

Figure 7.5: An example network showing the robustness of the even policy, which we denote as f^{even}. In this network, $c = \{4, 9, 8, 1\}$ and $r = 13$. The black dashed line represents the line $r^a r/C(E)$; notice that this is an upper bound on the blue line $B^{\text{SE}}(r, r^a)$. The red line, $B^*(f^{r^a,\text{even}}$ is always between the two, as explained in Lemma 7.2. The lines associated with f^{lo} and f^{hi} go above the dashed black line for certain values of r^a. One way to interpret Theorem 7.5 is that any plot for $f' \neq f^{\text{even}}$ will be above the black dashed line for some values of r^a.

Proof: Assume that f is the even policy. Lemma 7.2 shows that $T(f) \geq C(E)/r$. To see that this inequality is tight, consider $r^a = C(E')$ for some $E' \subseteq E$. If the attacker uses the policy $f^a_e = c_e$ for $e \in E'$ and $f^a_e = 0$ otherwise, it will block all traffic on the edges in E', which, by definition of the even policy, sums to $B(f, f^a) = C(E')r/C(E)$. Thus $r^a/B(f, f^a) = C(E)/r$.

Now assume the existence of $f' \in \mathcal{F}$ such that $T(f) > T(f)$. If $r^a = C(E')$ for some $E' \subset E$, then we know from the above discussion that $B^*(r^a, f) = r^a r/C(E)$, based on the attacker putting all its traffic on the edges in E'. However, the $B^*(r^a, f')$ must be strictly less than this, meaning that $\sum_{e \in E'} f'_e < \sum_{e \in E'} f_e$. If $r^a = C(E \setminus E')$, then again $B^*(r^a, f) = r^a r/C(E)$. However, since $\sum_{e \in E \setminus E'} f'_e > \sum_{e \in E \setminus E'}$, the attacker can attack the edges in $E \setminus E'$ and block strictly more than $r^a r/C(E)$, a contradiction. ■

Aside from robustness, using the even routing policy has two advantages. First, it

is easy to implement: the router simply portions the traffic according to the relative capacity of a link. Second, when the router uses this policy, it can be inferred from the proof of Theorem 7.2 that finding the optimal attack is NP-Hard. There is not guarantee that the attacker can block the optimal amount of traffic. This is not the case with f^{lo}, where an optimal attack policy is to attack the higher-capacity edges first, or with f^{hi}, where an optimal attack policy is to attack the lower-capacity edges first.

Chapter 8

Conclusions

In this work, we have endeavored to understand the impact of information in large-scale autonomous systems. Recall from Chapter 1 that we posed the following questions:

1. How do a set of information constraints impact the resulting emergent system-wide behavior?

We studied this question extensively in a variety of settings. In Chapter 3, we studied information sharing constraints that arose from the presence of an external attacker, who had the ability to compromise k agents, showing that the price of anarchy was (roughly) inversely proportional to k across the set of valid utilities. In Chapter 4, we considered a more general scenario where the constraints come in the form of a graph. Across the set of valid utilities, it was determined that the price of anarchy degrades quickly to $1/(n+1)$ as edges are removed from the graph. Furthermore, no choice of utility can guarantee a price of anarchy above the inverse of the independence number of G. Chapter 5 investigates the case where the graph is a DAG and the marginal contribution utility is employed. Here the price of anarchy was shown to degrade with the fractional independence number. In Chapter 6, we relax the assumption that agents share their chosen action - instead they may pass along additional information to augment other

agents' action sets. We described how this increased sharing yields better performance guarantees, as one might expect. Lastly, in Chapter 7, we studied a network security problem modeled as a Stackelberg game. We showed how knowledge of the attacker's action set affects the rational decision of the router, and how valuable it is know the attacker's exact budget.

2. How can a system designer strategically set decision-making rules for the components to offset the effects of information constraints?

This question was also addressed throughout this work. In Chapter 3, we showed that in the presence of compromised agents, the marginal contribution utility provided an increased guarantee in performance. Chapter 4 then showed that consistent utility functions mitigate the loss in performance shown to be present across the set of valid utilities. In Chapter 5 this is explored in several aspects. First, we show, if the system designer is permitted to build the graph with number of edges, what structures will yield the best price of anarchy. Second, we showed for a "clique of cliques" graph that marginal contribution is the optimal utility, and can be used to calculate the optimal information sharing policy. The augmented greedy algorithm presented in Chapter 6 is shown to be optimal given the information sharing constraints in certain settings, and near-optimal in others. Finally, Chapter 7 showcases the optimal routing policies to be used in the presence of attacker uncertainty.

8.1 Future Work and Open Questions

The study of large-scale autonomous systems is of great import - fortunately there is no shortage of interesting problems. First, we have only considered a small subset of the types of information sharing constraints that are possible. For instance, what if agents are allowed to send messages not within the set of actions? What information would be

helpful?

In this work we have only considered the effect of information sharing constraints on the resulting equilibria, however one could ask similar questions about system dynamics. For instance, how do information sharing constraints affect the transient guarantees under a common model of dynamics, such as best response?

Future work will of course focus on the elusive characterization of optimal utility design. Is there a specific utility that works well for the types of information constraints in our model? Is there a large set of utilities that could be considered optimal? For instance, in the set of valid utility games, could one choose a utility (or set of utilities) that increase the price of anarchy to $1 - 1/e$?

Appendix A

Proofs for Selected Results

A.1 Proof for Theorem 3.2

We note that the values for $\text{PoA}(\mathcal{H}_{\text{VUG}}(k))$ shown in Theorem 3.1 are lower bounds for $\text{PoA}(\mathcal{H}_{\text{MC}}(k))$. Therefore, we need only show that having one blind agent increases the lower bound for $\text{PoA}(\mathcal{H}_{\text{MC}}(k))$, then show that all these lower bounds are tight.

As with Theorem 3.1, we use submodularity and monotonicity to show that when $|B| > 0$, $\text{PoA}(\mathcal{H}_{\text{MC}}(k)) \geq \frac{1}{1+|K|}$, and then proceed to provide a canonical example to show that this lower bound on PoA is tight. As a matter of notation, for $a, b \in X$, we denote $f(a, b)$ to mean $f(c)$, where $c = \{a_i \cup b_i\}_i$. Futhermore, for some $J \subseteq N$, we use x_J to mean $\{x_i\}_{i \in J}$.

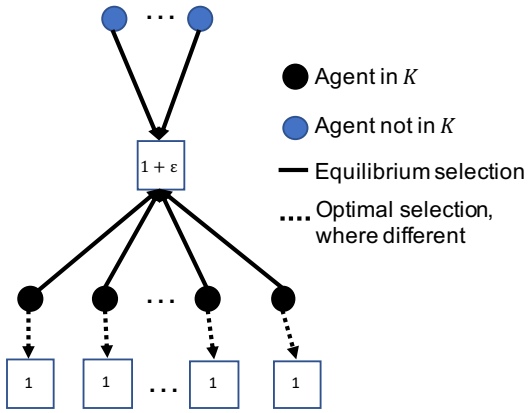

Figure A.1: An example game used in the proof for Theorem 3.2. As with the other figures, compromised agents are the black cross-hatch circles and the agents not in K are the blue circles. Each resource s has a value v_s, where $f(x) = \sum_{s \in S(x)} v_s$. A equilibrium selection yields $f(x^{\text{eq}}) = 1 + \varepsilon$, since the blue agents have no other available resources, and the agents in K act independently. The optimal selection is for the agents in K to choose their alternate resource, i.e., $f(x^{\text{opt}}) = |K| + 1 + \varepsilon$. Therefore, as $\varepsilon \to 0$, $f(x^{\text{opt}}) \to 1 + |K|$.

To see the lower bound, suppose that $B \neq \emptyset$ and consider the following:

$$
\begin{aligned}
f(x^{\text{opt}}) \leq & f\left(x^{\text{opt}}, x_B^{\text{eq}}\right), && \text{(A.1)}\\
\leq & f\left(x_{N\setminus K}^{\text{opt}}, x_B^{\text{eq}}\right) + \sum_{i \in K} f\left(x_i^{\text{opt}}\right), && \text{(A.2)}\\
= & f\left(x_{N\setminus K}^{\text{opt}}, x_B^{\text{eq}}\right) + \sum_{i \in K} f\left(x_i^{\text{eq}}\right), && \text{(A.3)}\\
\leq & f\left(x_{N\setminus K}^{\text{opt}}, x_B^{\text{eq}}\right) + f(x_B^{\text{eq}}) + (|K| - 1) f(x^{\text{eq}}) && \text{(A.4)}
\end{aligned}
$$

where (A.1) is true by monotonicity, (A.2) is true by submodularity, (A.3) is true since agents in K are either blind or isolated, and (A.4) is true by monotonicity.

Once an agent in K has chosen an action, that agent has no incentive to deviate, regardless of how the other agents behave. Therefore, we can consider a "sub system" $\bar{H} = (N \setminus K, \bar{X}, \overline{\text{MC}}, \bar{f})$ among only the non-compromised agents, assuming that the blind and isolated agents have made their choices. In this sub game, the non-compromised

agents seek to maximize the welfare function $\overline{f} : \bar{X}_i \to \mathbb{R}$, where $\bar{x} := \Pi_{i \notin K} X_i$, such that

$$\overline{f}(\bar{x}) = f(\bar{x}, x_B^{\text{eq}}) - f(x_B^{\text{eq}}), \tag{A.5}$$

for $\bar{x} \in \bar{X}_i$. Note that $\overline{f}$ is also submodular monotone, with $\overline{f}(\emptyset) = 0$. The agents are endowed with the following utility function

$$\overline{\text{MC}}_i(\bar{x}_i, \bar{x}_{-i}) = \overline{f}(\bar{x}_i, \bar{x}_{-i}) - \overline{f}(\bar{x}_{-i}). \tag{A.6}$$

It can be easily shown that is a VUG. Therefore, we know from [16] that $2\overline{f}(\bar{x}^{\text{eq}}) \geq \overline{f}(\bar{x}^{\text{opt}})$, where $\bar{x}^{\text{opt}} \in \arg\max_{\bar{x}} \overline{f}(a)$. It is also important to note that by design, $\bar{x}^{\text{eq}}$ is also a equilibrium profile of actions for agents not in K for the *original* game G, assuming that agents in B choose x_B^{eq}.

Returning to (A.4), we see that

$$\begin{aligned}
f(x^{\text{opt}}) \leq & f\left(x_{N\setminus K}^{\text{opt}}, x_B^{\text{eq}}\right) - f(x_B^{\text{eq}}) + 2f(x_B^{\text{eq}}) + (|K|-1)f(x^{\text{eq}}) && \text{(A.7)}\\
= & \overline{f}\left(x_{N\setminus K}^{\text{opt}}\right) + 2f(x_B^{\text{eq}}) + (|K|-1)f(x^{\text{eq}}) && \text{(A.8)}\\
\leq & \overline{f}\left(\bar{x}^{\text{opt}}\right) + 2f(x_B^{\text{eq}}) + (|K|-1)f(x^{\text{eq}}) && \text{(A.9)}\\
\leq & 2\overline{f}\left(\bar{x}^{\text{eq}}\right) + 2f(x_B^{\text{eq}}) + (|K|-1)f(x^{\text{eq}}) && \text{(A.10)}\\
= & 2f\left(x_{N\setminus K}^{\text{eq}}, x_B^{\text{eq}}\right) + (|K|-1)f(x_K^{\text{eq}}) && \text{(A.11)}\\
\leq & (1+|K|)f(x^{\text{eq}}), && \text{(A.12)}
\end{aligned}$$

where (A.7) is trivially true, (A.8) is true by defintion of $\overline{f}$, (A.9) is true by definition of $\bar{x}^{\text{opt}}$, (A.10) is true since $\bar{G}$ is a VUG, (A.11) is true by definition of $\overline{f}$, and (A.12) is true by monotonicity. Thus for any G that meets the requirements of the theorem statement,

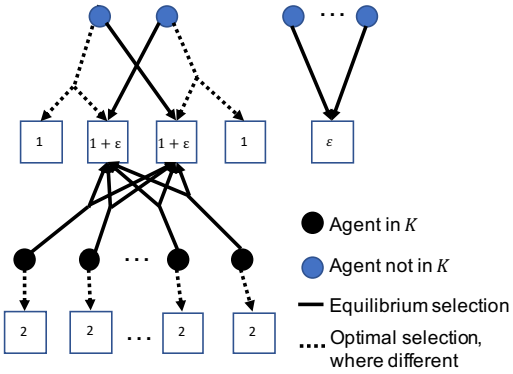

Figure A.2: An example game used in the proof for Theorem 3.2, for the case where K is only isolated agents. We use the same f and notation as Figure A.1, except with different resource values and action sets. As with previous examples, agents in K act independently, choosing the two $1+\varepsilon$ resources. The agents not in K are not informed of these decisions, so a worst-case equilibrium is where the two agents which also have access to these two resources also select them. The remaining agents not in K have only access to one ε resource. Therefore, in this case, $f(x^{\text{eq}}) = 2 + 3\varepsilon$. The optimal allocation, on the other hand, is for all resources to be selected. Thus as $\varepsilon \to 0$, $f(x^{\text{eq}})/f(x^{\text{opt}}) \to 1/(2 + |K|)$.

and for any x^{eq}, it follows that

$$\frac{f(x^{\text{eq}})}{f(x^{\text{opt}})} \geq \frac{1}{1 + |K|}, \tag{A.13}$$

implying that this is also a lower bound for $\text{PoA}(\mathcal{H}_{\text{MC}}(k))$.

We now show that the lower bounds established above are tight. First consider the problem instance in Figure A.1. Again, compromised agents are the black cross-hatch circles and the agents not in K are the blue circles. As with the other examples in this paper, each resource r has a value v_s, where $f(a) = \sum_{s \in S(x)} v_s$. It should be clear that a equilibrium selection yields $f(x^{\text{eq}}) = 1 + \varepsilon$, since the blue agents have no other available resources, and the agents in K act independently. The optimal selection is for the agents in K to choose their alternate resource, i.e., $f(x^{\text{opt}}) = |K| + 1 + \varepsilon$. Therefore, as $\varepsilon \to 0$,

we see that

$$\frac{f(x^{\text{eq}})}{f(x^{\text{opt}})} \to \frac{1}{1+|K|}. \tag{A.14}$$

Note that this holds for any combination of isolated and blind agents in K, and as long as $|K| < n$.

In the case where $K = N$, then in the example in Figure A.1 one agent will still choose the $1+\varepsilon$ resource. Here we see that as $\varepsilon \to 0$, $f(x^{\text{eq}})/f(x^{\text{opt}}) \to 1/n$.

In the case where $|K| < n-1$ and there are no blind agents in K, we invoke the example in Figure A.2, which uses the same f and notation as Figure A.1, except with different resource values and action sets. As with previous examples, agents in K act independently, choosing the two $1+\varepsilon$ resources. The agents not in K are not informed of these decisions, so a worst-case equilibrium is where the two agents which also have access to these two resources also select them. The remaining agents not in K have only access to one ε resource. Therefore, in this case, $f(x^{\text{eq}}) = 2+3\varepsilon$. The optimal allocation, on the other hand, is for all resources to be selected. Thus as $\varepsilon \to 0$,

$$\frac{f(x^{\text{eq}})}{f(x^{\text{opt}})} \to \frac{1}{2+|K|}. \tag{A.15}$$

A.2 Proof for Theorem 4.1

We first show that

$$\frac{f(x^{\text{eq}})}{f(x^{\text{opt}})} \geq \frac{1}{1+\tau(G)}, \tag{A.16}$$

and then show that for any G, there exists $f, \{X_i\}_i, \{U_i\}_i$ which make the expression tight. We denote $\mathcal{N}_T$ to mean the set of incoming neighbors common to information

group T. Begin with

$$
\begin{aligned}
f(x^{\text{opt}}) \leq & f(x^{\text{eq}}) + \Delta(x^{\text{opt}}|x^{\text{eq}}), && \text{(A.17)}\\
= & f(x^{\text{eq}}) + \sum_i \Delta(x_i^{\text{opt}}|x_{1:i-1}^{\text{opt}}, x^{\text{eq}}), && \text{(A.18)}\\
\leq & f(x^{\text{eq}}) + \sum_i \Delta(x_i^{\text{opt}}|x^{\text{eq}}), && \text{(A.19)}\\
= & f(x^{\text{opt}}) + \sum_{T\in\mathcal{T}(G)} \sum_{i\in T} \Delta(x_i^{\text{opt}}|x^{\text{eq}}), && \text{(A.20)}\\
\leq & f(x^{\text{eq}}) + \sum_{T\in\mathcal{T}(G)} \sum_{i\in T} \Delta(x_i^{\text{opt}}|x_{j\in\mathcal{N}_T}^{\text{eq}}), && \text{(A.21)}\\
\leq & f(x^{\text{eq}}) + \sum_{T\in\mathcal{T}(G)} \sum_{i\in T} U_i(x_i^{\text{opt}}, x_{j\in\mathcal{N}_T}^{\text{eq}}), && \text{(A.22)}\\
\leq & f(x^{\text{eq}}) + \sum_{T\in\mathcal{T}(G)} \sum_{i\in T} U_i(x_i^{\text{eq}}, x_{j\in\mathcal{N}_T}^{\text{eq}}), && \text{(A.23)}\\
\leq & f(x^{\text{eq}}) + \sum_{T\in\mathcal{T}(G)} f(x_T^{\text{eq}}), && \text{(A.24)}\\
\leq & f(x^{\text{eq}})(1+\tau(G)), && \text{(A.25)}
\end{aligned}
$$

where (A.17) is true by monotonicity of f, (A.18) is true by definition of $\Delta(\cdot)$, (A.19) is true by submodularity of f, (A.20) is a reorganization of the sum, (A.21) is true by submodularity, (A.22) is true by statement 2) in Definition 3.1, (A.23) is true by the defintion of equilibrium, (A.24) is true by statement 3) in the Definition 3.1, and (A.25) is true by monotonicity of f.

Next we construct an example worst-case f, $\{X_i\}_i$, $\{U_i\}_i$ such that

$$
\frac{f(x^{\text{eq}})}{f(x^{\text{opt}})} = \frac{1}{1+\tau(G)} \tag{A.26}
$$

for any G. Let $S = \{s_\varepsilon, s_{\text{sm}}, s_{\text{big}}, s_2, \ldots, s_\tau\}$ be a set of (possibly overlapping) 2-D boxes,

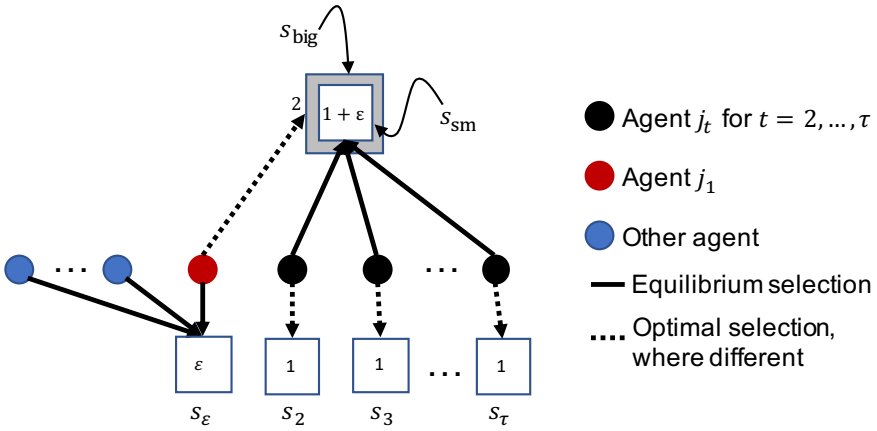

Figure A.3: Example used in the proof for Theorem 4.1. While agent j_1's choice in equilibrium may seem unintuitive, it is based on a carefully crafted utility function, such that Defintion 3.1 is still satisfied for the system.

as shown in Figure 4.3. Let $f(x)$ be the total area covered by the boxes in $S(x)$: this function is normalized, submodular, and monotone. For some small $\varepsilon > 0$, let $f(\{s_\varepsilon\}) = \varepsilon$, $f(\{s_{\text{sm}}\}) = 1 + \varepsilon$, $f(\{s_{\text{big}}\}) = 2$, and $f(\{s_2\}) = \cdots = f(\{s_\tau\}) = 1$. The box s_{big} "covers" the box s_{sm}, i.e., $f(\{s_{\text{sm}}, s_{\text{big}}\}) = 2$. The remaining pairs of boxes are disjoint.

From each information group $T_1, \ldots, T_\tau$ choose a representative agent $j_1, \ldots, j_\tau$. Since the label order is arbitrary, we assume without loss of generality that there exists incoming edges from agent j_2 to the agents in T_1, i.e., $j_2 \in \mathcal{N}(T_1)$.[1] The action sets are allocated as

$$X_i = \begin{cases} \{\{s_\varepsilon\}, \{s_{\text{big}}\}\} & \text{if } i = j_1, \\ \{\{s_{\text{sm}}\}, \{s_t\}\} & \text{if } i = j_t \text{ and } t > 1, \\ \{\{s_\varepsilon\}\} & \text{otherwise.} \end{cases} \tag{A.27}$$

Again, both f and $\{X_i\}_i$ are represented in Figure A.3.

[1] In the case where no such edges exists, i.e., G is a set of disconnected cliques, tightness can be shown from Corollary 5.1.

In order to define the utilities, we first define the action profile x^{eq}:

$$x^{\text{eq}} = \begin{cases} \{s_{\text{sm}}\} & \text{if } i \in \{j_2, \ldots, j_\tau\}, \\ \{s_\varepsilon\} & \text{if } i \notin \{j_2, \ldots, j_\tau\}. \end{cases} \tag{A.28}$$

As the notation implies, we will design the utilities so that this action profile is an equilibrium. For every $T_t \in \mathcal{T}(G)$ define x^t, where

$$x_i^t = \begin{cases} x_i^{\text{eq}} & \text{if } i \in T_t \cup \mathcal{N}(T_t), \\ \emptyset & \text{if } i \notin T_t \cup \mathcal{N}(T_t). \end{cases} \tag{A.29}$$

In other words, x^t is the set of actions in x^{eq}, with the exception that all agents not in T_t choose the empty set. It is important to note that due to the graph constraints, the utility of agent $i \in T_t$ for action profile x^{eq} is $U_i(x^t)$. It can also be observed that $f(x^t) = 1 + \varepsilon + p_t\varepsilon$, where $p_t \in \{0, 1\}$ is an indicator: $p_t = 1$ if $j_1 \in \mathcal{N}(T_t)$ or $|T_t| > 1$ (i.e., s_ε is chosen by some agent in $T_t \cup \mathcal{N}(T_t)$), and $p_t = 0$ otherwise.

The utility functions are as follows:

$$U_i(x_i, x_{\mathcal{N}_i}) = \begin{cases} 1 + \varepsilon & \text{if } x = x^t \text{ and } i = j_t \text{ for some } T_t \in \mathcal{T}(G) \\ f(x_i, x_{-i}) - f(x_{-i}) & \text{otherwise.} \end{cases} \tag{A.30}$$

We claim that $f, \{X_i\}, \{U_i\}_i$ is a VUG. Since $U_i = \text{MC}_i$ for all action profiles except when $x = x^t$ and $i = j_t$, we need only prove that the statements 2) and 3) in the VUG definition are satisfied for these exceptions. Statement 2) holds since $f(x^t) - f(x^t_{i \neq j_t}) \leq$

$1 + \varepsilon = U_{j_t}(x^t)$. Statement 3) holds, since:

$$\sum_i U_i(x^t) = U_{j_t}(x^t) + \sum_{i \notin \{j_t, j_1\}} U_i(x^t) \tag{A.31}$$

$$\leq 1 + \varepsilon + p_t\varepsilon \tag{A.32}$$

$$= f(x^t). \tag{A.33}$$

Recall that agent j_1's action set is $\{\{s_{\text{big}}\}, \{s_\varepsilon\}\}$, implying that

$$U_{j_1}(x^1) = 1 + \varepsilon \tag{A.34}$$

$$> 2 + p_t\varepsilon - (1 + \varepsilon + p_t\varepsilon) \tag{A.35}$$

$$= f(\{s_{\text{big}}\}, x^1_{-j_1}) - f(x^1_{-j_1}) \tag{A.36}$$

$$= U_{j_1}(\{s_{\text{big}}\}, x^1_{-j_1}) \tag{A.37}$$

For agent $j_t, t > 1$, the action set is $\{\{s_{\text{sm}}\}, \{s_t\}\}$, implying that

$$U_{j_t}(x^t) = 1 + \varepsilon \tag{A.38}$$

$$> 1 + p_t\varepsilon - p_t\varepsilon \tag{A.39}$$

$$= f(\{s_t\}, x^t_{-j_t}) - f(x^t_{-j_t}) \tag{A.40}$$

$$= U_{j_t}(\{s_t\}, x^t_{-j_t}) \tag{A.41}$$

Since all other agents have only a single action in their action sets, we conclude that x^{eq} is an equilibrium action profile. The optimal profile x^{opt} is where j_1 chooses $\{s_{\text{big}}\}$, j_t chooses $\{s_t\}$ for $t > 1$, and all other agents choose $\{s_\varepsilon\}$, implying that $f(x^{\text{opt}}) =$

$2 + \tau(G) - 1 + \varepsilon$. Therefore,

$$\frac{f(x^{\text{eq}})}{f(x^{\text{opt}})} = \frac{1 + 2\varepsilon}{1 + \tau(G) + \varepsilon}. \tag{A.42}$$

As $\varepsilon \to 0$, we see that (A.26) holds.

A.2.1 Proof for Theorem 5.1

The upper bound is a consequence of Proposition 4.1. For the lower bound, let $x_i = x_i^{\text{sol}}$ and $x = x^{\text{sol}}$ for ease of notation. Then consider the following:

$$f(x^{\text{opt}}) \leq f(x, x^{\text{opt}}) = f(x) + \sum_i \Delta\left(x_i^{\text{opt}} \middle| \bigcup_{j<i} x_j^{\text{opt}} \cup x\right), \tag{A.43}$$

$$\leq f(x) + \sum_i \Delta(x_i^{\text{opt}} | x_{\mathcal{N}_i}), \tag{A.44}$$

$$\leq f(x) + \sum_i \Delta(x_i | x_{\mathcal{N}_i}), \tag{A.45}$$

where (A.44) follows from submodularity and (A.45) follows from the decision-making rule in (5.1).

Consider a set of scalars $\{y_k\}_{k \in K(G)}$ such that $y_k \geq 0$ for all k and $\sum_{k:i \in k} y_k \geq 1$ for

all i. Then

$$\begin{aligned}
f(x^{\text{opt}}) \leq & f(x) + \sum_i \Delta(x_i|x_{\mathcal{N}_i}) \left(\sum_{k:i\in k} y_k \right) && \text{(A.46)}\\
= & f(x) + \sum_i \sum_{k:i\in k} y_k \Delta(x_i|x_{\mathcal{N}_i}) && \text{(A.47)}\\
= & f(x) + \sum_{k\in K(G)} \sum_{i\in k} y_k \Delta(x_i|x_{\mathcal{N}_i}) && \text{(A.48)}\\
= & f(x) + \sum_{k\in K(G)} y_k \sum_{i\in k} \Delta(x_i|x_{\mathcal{N}_i}) && \text{(A.49)}\\
\leq & f(x) + \sum_{k\in K(G)} y_k \sum_{i\in k} \Delta(x_i|x_{j:j<i,j\in k}) && \text{(A.50)}\\
= & f(x) + \sum_{k\in K(G)} y_k f(x_k) && \text{(A.51)}\\
\leq & f(x) + f(x) \sum_{k\in K(G)} y_k, && \text{(A.52)}
\end{aligned}$$

where (A.50) holds by submodularity and (A.52) holds by monotonicity. Then

$$\frac{f(x)}{f(x^{\text{opt}})} \geq \frac{1}{1 + \sum_{k\in K(G)} y_k}. \tag{A.53}$$

To make the bound in (A.53) as tight as possible, one can solve the following optimization:

$$\begin{aligned}
\min_y \quad & \textstyle\sum_{k\in K(G)} y_k \\
\text{subject to} \quad & \textstyle\sum_{k\in K(G):i\in k} y_k \geq 1, \ \forall i \\
& y_k \geq 0, \ \forall k
\end{aligned} \tag{A.54}$$

This is the same optimization problem as that in (4.4), the value of which is defined as $k^*(G)$, and by duality, $k^*(G) = \alpha^*(G)$. ■

A.3 Proof for Lower Bound in Example 5.3

We begin with the following inequality, where $x_{1:2} = x_1 \cup x_2$:

$$
\begin{aligned}
f(x^{\text{opt}}) \leq & f(x^{\text{opt}}, x_{1:2}) && \text{(A.55)}\\
= & f(x_{1:2}) + \Delta(x_1^{\text{opt}}|x_{1:2}) + \Delta(x_2^{\text{opt}}|x_1^{\text{opt}}, x_{1:2}) \\
& + \Delta(x_3^{\text{opt}}|x_{1:2}^{\text{opt}}, x_{1:2}) + \Delta(x_4^{\text{opt}}|x_{1:3}^{\text{opt}}, x_{1:2}) && \text{(A.56)}\\
\leq & f(x_{1:2}) + f(x_1^{\text{opt}}) + \Delta(x_2^{\text{opt}}|x_1) + \Delta(x_3^{\text{opt}}|x_{1:2}) \\
& + \Delta(x_4^{\text{opt}}|x_{1:2}) && \text{(A.57)}\\
\leq & f(x_{1:2}) + f(x_1) + \Delta(x_2|x_1) + \Delta(x_3|x_{1:2}) \\
& + \Delta(x_4|x_{1:2}) && \text{(A.58)}\\
= & f(x_{1:3}) + f(x_{1:2}, x_4) \leq 2f(x_{1:4}), && \text{(A.59)}
\end{aligned}
$$

where (A.55) is true by submodularity, (A.56) is true by definition of Δ, (A.57) is true by submodularity, (A.58) is true since agents choose according to (5.1), (A.59) is true by definition of Δ and submodularity. Therefore we see $\gamma(G) = 1/2$. ■

A.4 Proof for Lemma 5.3

In this proof, we construct a G such that (5.11) is at equality, then reason that no other graph with n nodes, independence number r, and without the Sibling Property can have fewer edges. The proof also leverages the properties for a graph without the Sibling Property, found in Lemma 5.1. Let J be the unique maximum independent set (Property 1) in G, and let G' be the induced subgraph of G created by removing the nodes in J. Then we know that $\alpha(G') < \alpha(G)$ (Property 3). From (5.5) and Lemma 5.2, the minimum number of edges that such a G' can have is $M(n-r, r-1)$. Finally,

every node in G' must have outgoing edges to at least two nodes in J, therefore, G must have an additional $2(n-r)$ edges. Thus the minimum number of edges to construct G is given in (5.11). ■

A.5 Proof for Lemma 5.2

This can be shown by construction. Recall that $\overline{T(n,r)}$ is a set of disconnected cliques, of as close to equal size as possible. Making purely equal-sized cliques would mean that each clique is of size $\lfloor n/r \rfloor$, with $n \bmod r$ nodes left over. If each of these remaining nodes is added to a different clique, then G consists of $n \bmod r$ cliques of size $\lceil n/r \rceil$ and the rest of size $\lfloor n/r \rfloor$. Since a clique of size p contains $\frac{1}{2}p(p-1)$ edges, we can see that the first line in (5.10) is the number of edges in all the larger cliques, and the second line is the number of edges in all the smaller cliques. ■

A.6 Proof for Lemma 6.1

Begin with the following:

$$
\begin{aligned}
f(x^\pi) \leq & f(x^\rho) + \Delta(x^\pi | x^\rho) && \text{(A.60a)}\\
\leq & f(x^\rho) + \Delta(z^\pi_{1:n-1} \cap x^\pi | x^\rho) + \Delta(x^\pi \setminus z^\pi_{1:n-1} | x^\rho) && \text{(A.60b)}\\
= & f(x^\rho) + \Delta(z^\pi_{1:n-1} \cap x^\pi | x^\rho) \\
& + \sum_i \Delta(x^\pi_i \setminus z^\pi_{1:n-1} | x^\rho, x^\pi_{1:i-1} \setminus z^\pi_{1:n-1}) && \text{(A.60c)}\\
\leq & f(x^\rho) + \Delta(z^\pi_{1:n-1} \cap x^\pi | x^\rho) \\
& + \sum_i \Delta(x^\pi_i \setminus z^\pi_{1:n-1} | x^\rho_{1:i-1}) && \text{(A.60d)}\\
\leq & f(x^\rho) + \Delta(z^\pi_{1:n-1} \cap x^\pi | x^\rho) + \sum_i \alpha \Delta(x^\rho_i | x^\rho_{1:i-1}) && \text{(A.60e)}\\
= & (1+\alpha) f(x^\rho) + \Delta(z^\pi_{1:n-1} \cap x^\pi | x^\rho) && \text{(A.60f)}
\end{aligned}
$$

where (A.60a), (A.60b), (A.60d) are true by submodularity of f, and (A.60e) is true by (6.14).

We denote $\tilde{z}^\pi_i$ to mean $z^\pi_i \cap x^\pi$, and suppose that there exists β such that $\Delta(\tilde{z}^\pi_i | x^\rho_{1:i}) \leq \beta \Delta(x^\rho_i | x^\rho_{1:i-1})$. Then the second term in (A.60f) can be upper bounded by the following:

$$
\begin{aligned}
\Delta(\tilde{z}_{1:n-1} | x^\rho) \leq & \sum_{i=1}^{n-1} \Delta(\tilde{z}^\pi_i | \tilde{z}^\pi_{1:i-1}, x^\rho) && \text{(A.61a)}\\
\leq & \sum_{i=1}^{n-1} \Delta(\tilde{z}^\pi_i | x^\rho_{1:i}) && \text{(A.61b)}\\
\leq & \sum_{i=1}^{n-1} \beta \Delta(x^\rho_i | x^\rho_{1:i}) && \text{(A.61c)}\\
= & \beta f(x^\rho_{1:n-1}) \leq \beta f(x^\rho), && \text{(A.61d)}
\end{aligned}
$$

where (A.61b) and (A.61d) are true by the submodularity of f. Substituting this upper bound back into (A.60f), we see that

$$f(x^\pi) \leq (1 + \alpha + \beta) f(x^\rho).$$

We now show that such a β exists and define it for two cases: when $p \leq k$ and when $p \geq k$. First suppose that $p \leq k$. Denote $x_i^{k-p} \in \arg\max_{\tilde{x} \in (x_i^\rho)^{k-p}} \Delta(\tilde{x} | x_{1:i-1}^\rho)$. Then

$$\begin{aligned}
\Delta(x_i^\rho | x_{1:i-1}^\rho) &\geq (1/\alpha) \Delta(\tilde{z}_i^\pi \cup x_i^{k-p} | x_{1:i-1}^\rho) && \text{(A.62a)} \\
&= (1/\alpha) \Delta(\tilde{z}_i^\pi | x_{1:i-1}^\rho, x_i^{k-p}) + (1/\alpha) \Delta(x_i^{k-p} | x_{1:i-1}^\rho) && \text{(A.62b)} \\
&\geq (1/\alpha) \Delta(\tilde{z}_i^\pi | x_{1:i}^\rho) + (1/\alpha) \Delta(x_i^{k-p} | x_{1:i-1}^\rho) \implies && \text{(A.62c)}
\end{aligned}$$

$$\begin{aligned}
\Delta(\tilde{z}_i^\pi | x_{1:i}^\rho) &\leq \alpha \Delta(x_i^\rho | x_{1:i-1}^\rho) - \Delta(x_i^{k-p} | x_{1:i-1}^\rho) && \text{(A.62d)} \\
&\leq \alpha \Delta(x_i^\rho | x_{1:i-1}^\rho) - \frac{k-p}{k} \Delta(x_i^\rho | x_{1:i-1}^\rho) && \text{(A.62e)} \\
&= (\alpha - 1 + p/k) \Delta(x_i^\rho | x_{1:i-1}^\rho), && \text{(A.62f)}
\end{aligned}$$

where (A.62a) is true by (6.14), (A.62c) is true by submodularity of f. We conclude that when $p \leq k$, $\beta = (\alpha - 1 + p/k)$, implying that for this case

$$f(x^\pi) \leq (2\alpha + p/k) f(x^\rho).$$

Next suppose that $p \geq k$. Observe that $|\tilde{z}_i^\pi| \leq k(n-1)$, since using the approximated augmented greedy policy, no more than $k(n-1)$ elements of z_i^π can be chosen by other

agents. This implies the following:

$$\Delta(\tilde{z}_i^\pi | x_{1:i}^\rho) \leq (1/\min(k/|\tilde{z}_i^\pi|, 1)) \cdot \max_{z \in (\tilde{z}_i^\pi)^k} \Delta(z | x_{1:i}^\rho) \tag{A.63a}$$

$$\leq \max(|\tilde{z}_i^\pi|/k, 1) \cdot \max_{z \in (\tilde{z}_i^\pi)^k} \Delta(z | x_{1:i}^\rho) \tag{A.63b}$$

$$\leq \max(\min(k(n-1), p)/k, 1) \cdot \max_{z \in (\tilde{z}_i^\pi)^k} \Delta(z | x_{1:i}^\rho) \tag{A.63c}$$

$$= \min(n-1, p/k) \cdot \max_{z \in (\tilde{z}_i^\pi)^k} \Delta(z | x_{1:i}^\rho) \tag{A.63d}$$

$$\leq \alpha \cdot \min(n-1, p/k) \Delta(x_i^\rho | x_{1:i-1}^\rho) \tag{A.63e}$$

We conclude that when $p \geq k$, $\beta = \alpha \cdot \min(n-1, p/k)$, implying that for this case:

$$f(x^\pi) \leq (1 + \alpha(1 + \min(n-1, p/k))) f(x^\rho).$$

■

A.7 Proof for Lemma 6.2

We begin with

$$f(x^\rho) \leq f(x_{1:n-1}^\pi) + \Delta(x^\rho | x_{1:n-1}^\pi) \tag{A.64a}$$

$$= f(x_{1:n-1}^\pi) + \Delta(x_n^\rho | x_{1:n-1}^\rho, x_{1:n-1}^\pi) + \sum_{i=1}^{n-1} \Delta(x_i^\rho | x_{1:i-1}^\rho, x_{1:n-1}^\pi) \tag{A.64b}$$

$$\leq f(x_{1:n-1}^\pi) + \Delta(x_n^\rho | x_{1:n-1}^\pi) + \sum_{i=1}^{n-1} \Delta(x_i^\rho | x_{1:i}^\pi) \tag{A.64c}$$

$$\leq f(x_{1:n-1}^\pi) + \alpha_1 \Delta(x_n^\pi | x_{1:n-1}^\pi) + \sum_{i=1}^{n-1} \Delta(x_i^\rho | x_{1:i}^\pi) \tag{A.64d}$$

where (A.64a) and (A.64c) follow from submodularity of f, (A.64b) follows from the definition of $\Delta(\cdot)$, and (A.64d) follows from (6.16a). Focusing on the sum in (A.64d), for any $0 \leq \varepsilon_1, \dots \varepsilon_{n-1} \leq 1$ (and defining $\varepsilon_0 = 0$), we see that

$$\sum_{i=1}^{n-1} \Delta(x_i^\rho | x_{1:i}^\pi) = \sum_{i=1}^{n-1} (1-\varepsilon_i) \Delta(x_i^\rho | x_{1:i}^\pi) + \sum_{i=1}^{n-1} \varepsilon_i \Delta(x_i^\rho | x_{1:i}^\pi) \tag{A.65a}$$

$$\leq \sum_{i=1}^{n-1} (1-\varepsilon_i) \Delta(x_i^\rho | x_{1:i-1}^\pi) + \sum_{i=1}^{n-1} \alpha_2 \varepsilon_i \Delta^k(z_i^\pi | x_{1:i}^\pi) \tag{A.65b}$$

$$\leq \sum_{i=1}^{n-1} \alpha_1 (1-\varepsilon_i) \Delta(x_i^\pi | x_{1:i-1}^\pi) + \sum_{i=1}^{n-1} \alpha_1 \alpha_2 \varepsilon_i \Delta(x_{i+1}^\pi | x_{1:i}^\pi) \tag{A.65c}$$

$$= \alpha_1 \alpha_2 \varepsilon_{n-1} \Delta(x_n^\pi | x_{1:n-1}^\pi) + \sum_{i=1}^{n-1} \alpha_1 (1 - \varepsilon_i + \alpha_2 \varepsilon_{i-1}) \Delta(x_i^\pi | x_{1:i-1}^\pi), \tag{A.65d}$$

where (A.65b) is true by submodularity of f (1st term) and (6.16b) (2nd term), (A.65c) is true by (6.16a), and (A.65d) is just a rearrangement of the terms. Applying this to (A.64d) yields

$$f(x^\rho) \leq f(x_{1:n-1}^\pi) + (\alpha_1 + \alpha_1 \alpha_2 \varepsilon_{n-1}) \Delta(x_n^\pi | x_{1:n-1}^\pi) + \sum_{i=1}^{n-1} (\alpha_1 - \alpha_1 \varepsilon_i + \alpha_1 \alpha_2 \varepsilon_{i-1}) \Delta(x_i^\pi | x_{1:i-1}^\pi) \tag{A.66a}$$

$$= (\alpha_1 + \alpha_1 \alpha_2 \varepsilon_{n-1}) \Delta(x_n^\pi | x_{1:n-1}^\pi) + \sum_{i=1}^{n-1} (1 + \alpha_1 - \alpha_1 \varepsilon_i + \alpha_1 \alpha_2 \varepsilon_{i-1}) \Delta(x_i^\pi | x_{1:i-1}^\pi). \tag{A.66b}$$

Suppose that for a particular choice of ε_i, we let

$$\varepsilon_i = \frac{(1/\alpha_1) \sum_{j=0}^{i-1} \alpha_2^j}{\sum_{j=0}^{n-1} \alpha_2^j}. \tag{A.67}$$

Since $\alpha_1, \alpha_2 \geq 1$, this satisfies the requirement that $0 \leq \varepsilon_i \leq 1$ for $i \in \{1, \ldots, n-1\}$. Then

$$-\alpha_1\varepsilon_i + \alpha_1\alpha_2\varepsilon_{i-1} = -\frac{\sum_{j=0}^{i-1}\alpha_2^j}{\sum_{j=0}^{n-1}\alpha_2^j} + \frac{\sum_{j=1}^{i-1}\alpha_2^j}{\sum_{j=0}^{n-1}\alpha_2^j} \tag{A.68}$$

$$= -\frac{1}{\sum_{j=0}^{n-1}\alpha_2^j} \tag{A.69}$$

Likewise

$$\alpha_1\alpha_2\varepsilon_{n-1} = \frac{\alpha_2\sum_{j=0}^{n-2}\alpha_2^j}{\sum_{j=0}^{n-1}\alpha_2^j} = \frac{\sum_{j=1}^{n-1}\alpha_2^j}{\sum_{j=0}^{n-1}\alpha_2^j} \tag{A.70}$$

$$= 1 - \frac{1}{\sum_{j=0}^{n-1}\alpha_2^j} \tag{A.71}$$

Applying (A.69) and (A.71) to (A.66b) yields

$$f(x^\rho) \leq \left(\alpha_1 + 1 - \frac{1}{\sum_{j=0}^{n-1}\alpha_2^j}\right)\Delta(x_n^\pi|x_{1:n-1}^\pi) + \sum_{i=1}^{n-1}\left(1 + \alpha_1 - \frac{1}{\sum_{j=0}^{n-1}\alpha_2^j}\right)\Delta(x_i^\pi|x_{1:i-1}^\pi) \tag{A.72}$$

$$= \left(1 + \alpha_1 - \frac{1}{\sum_{j=0}^{n-1}\alpha_2^j}\right) f(x^\pi) \tag{A.73}$$

■

A.8 Proof for Lemma 7.2

The right inequality is true by definition of SE, so we focus on the left. Denote $A = \{e \in E : B_e(f, f^a) > 0\}$

$$\begin{aligned} B(f, f^a) =& \sum_{e \in A} B_e(f, f^a) && \text{(A.74)} \\ =& \sum_{e \in A} f_e + f_e^a - c_e && \text{(A.75)} \\ =& \sum_{e \in A} c_e r / C(E) + f_e^a - c_e && \text{(A.76)} \\ =& \sum_{e \in A} c_e (r / C(E) - 1) + f_e^a && \text{(A.77)} \\ \leq& \sum_{e \in A} f_e^a (r / C(E) - 1) + f_e^a && \text{(A.78)} \\ =& \sum_{e \in A} r f_e^a / C(E) && \text{(A.79)} \\ \leq& r r^a / C(E) && \text{(A.80)} \end{aligned}$$

■

A.9 Proof for Lemma 7.3

This can be proven by contradiction: suppose that f^a is the unique optimal attack against an optimal routing policy f which is a solution to (7.2). Let $\tilde{f}^a$ be a "next best" attack, i.e.,

$$\tilde{f}^a \in \arg\max_{x \in \mathcal{F}^a \setminus f^a} B(f, x). \quad \text{(A.81)}$$

For some very small $\varepsilon > 0$, $d \in E(\tilde{f}^a) \setminus E(f^a)$, and $d' \in E(\tilde{f}^a) \setminus E(f^a)$, construct the following routing policy

$$\tilde{f}_e = \begin{cases} f_e + \varepsilon & \text{if } e = d, \\ f_e - \varepsilon & \text{if } e = d', \\ f_e, & \text{if } e \in E \setminus \{d, d'\}. \end{cases} \tag{A.82}$$

Since ε is small, f^a remains the optimal attack: $B(\tilde{f}, f^a) > B(\tilde{f}, \tilde{f}^a)$. However, $B(f, f^a) = B(\tilde{f}, f^a) + \varepsilon$, which implies that f is not a solution to (7.2), a contradiction.

■

A.10 Proof for Lemma 7.4

Fix r and $f \in \mathcal{F}(r)$. Since all parameters except r^a are fixed, we use the notation $B^*(r^a)$ and $B^{\text{SE}}(r^a)$ to emphasize that we are considering how much traffic is blocked as r^a varies.

To prove this lemma, we claim the following to be true:

1. $B^*(r^a)$ is a continuous function.

2. Suppose r^a is such that there exists a best response f^a and $e \in E$ where $c_e - f_e \leq f_e^a < c_e$. Then there exists $\varepsilon > 0$ such that

$$\frac{B^*(r^a + \delta) - B^*(r^a)}{\delta} = 1 \text{ for all } 0 < \delta < \varepsilon. \tag{A.83}$$

Otherwise, if no such f^a, e exist, then there is $\varepsilon > 0$ such that

$$\frac{B^*(r^a + \delta) - B^*(r^a)}{\delta} = 0 \text{ for all } 0 < \delta < \varepsilon. \tag{A.84}$$

In words, r^a is the lower boundary of a neighborhood where the derivative of $B^*(r^a)$ is either 1 for all points in the neighborhood or 0 for all points in the neighborhood.

3. If there exists $\varepsilon > 0$ such that

$$\frac{B^*(r^a + \delta) - B^*(r^a)}{\delta} = 0, \text{ and} \tag{A.85}$$

$$\frac{B^*(r^a) - B^*(r^a - \delta)}{\delta} = 1, \tag{A.86}$$

for all $0 < \delta \leq \varepsilon$, then $r^a \in \alpha$.

4. On a plot of $B^{\text{SE}}(r^a)$ vs r^a, the slope of the line between any two points is in the interval $[0, 1]$.

Assuming the claims are true, claims 1 and 2 imply that $B^*(r^a)$ is a continuous piecewise linear function, where the slope of each line is either 1 or 0. By claim 4, $B^*(r^a) - B^{\text{SE}}(r^a)$ is increasing when the slope of $B^*(r^a)$ is 1, and decreasing when the slope of $B^*(r^a)$ is 0. Therefore, the max of $B^*(r^a) - B^{\text{SE}}(r^a)$ must occur at some value of r^a where the slope of $B^*(r^a)$ changes from 1 to 0. By claim 3, all such values of r^a are contained in α. In the case where $\alpha \setminus \pi^a$ is nonempty, we include the boundary points $\underline{\pi}^a$ and $\overline{\pi}^a$ as possible values where the max on the interval π^a can occur.

Now we prove each of the claims. First we show that $B^*(r^a)$ is continuous. Observe that when r^a increases (decreases) by $\varepsilon > 0$, $B^*(r^a)$ can increase (decrease) by no more than ε. More formally,

$$|r^a - \hat{r}^a| < \varepsilon \implies |B^*(r^a) - B^*(\hat{r}^a)| < \varepsilon, \tag{A.87}$$

and thus the function is continuous.

To show claim 2, suppose that r^a is such that there exists a best response attack f^a

where $c_e - f_e \leq f_e^a < c_e$ for some $e \in E$. Increasing r^a (and f_e^a) by δ allows the attacker to increase $B(f, f^a)$ by δ. Therefore, $B^*(r^a + \delta) = B^*(r^a) + \delta$, which implies (A.83).

Now suppose that r^a is such that no such f^a, e exist, i.e., that for any best response attack policy f^a and for all $e \in E$, either

$$f_e^a = c_e \text{ or} \tag{A.88}$$

$$f_e^a < c_e - f_e, \tag{A.89}$$

If δ is small enough so that (A.89) can be replaced with $f_e^a < c_e - f_e - \delta$, then increasing r^a by δ cannot increase $B^*(r^a)$. This implies (A.84).

To prove claim 3, we state an implication of claim 2: if (A.85) is satisfied, then no best response f^a, e exist where $c_e - f_e \leq f_e^a < c_e$. However, (A.86) implies that r^a is also an upper boundary of a neighborhood where such an f^a and e exist. The only both statements can be true is if $f_e^a \in \{0, c_e\}$ for all e and for all optimal f^a. This implies that $r^a = C(E')$ for some $E' \subseteq E$, i.e., that $r^a \in \pi^a$.

We now prove claim 4. The function $B^{\text{SE}}(r^a)$ must be nondecreasing, since any attack policy that can be implemented with low r^a can also be carried out with high r^a. Equation (A.87) also applies to B^{SE}, so the slope of the line between any two points on $B^{\text{SE}}(r^a)$ is ≤ 1. We thus conclude that the claim holds. ■

A.11 Proof for Theorem 7.3

Note that since $B(f, f^a) = \sum_e \max\{f_e + f_e^a - c_e, 0\}$, then a lower bound on $B(f, f^a)$ is

$$B(f, f^a) \geq r + r^a - C(E). \tag{A.90}$$

We now begin with a few observations about router best responses:

1. For a policy pair (f, f^a), if $f_e + f_e^a \leq c_e$ for all e, then $B(f, f^a) = 0$, and the router has no incentive to deviate. If f^a satisfies (7.5), then (f, f^a) is both a SE and a NE.

2. For a policy pair (f, f^a), if $f_e + f_e^a \geq c_e$ for all e, then $B(f, f^a) = r + r^a - C(E)$, the lower bound in (A.90). Thus the router has no incentive to deviate. If f^a also satisfies (7.5), then (f, f^a) is both a SE and a NE.

3. For a policy pair (f, f^a), if there exist $e, e' \in E$ such that $f_e + f_e^a > c_e$ and $f_{e'} + f_{e'}^a < c_{e'}$, then (7.4) is not satisfied. Therefore, (f, f^a) is not a NE.

We now proceed with proving Theorem 7.3. To that end, consider the routing policy f^{lo}, where

$$f_e^{\text{lo}} := \max\left\{c_e - \max_{E' \subseteq E} \frac{C(E') - r}{|E'|}, 0\right\}. \tag{A.91}$$

We first show that f^{lo} is feasible. Let $E^* \in \arg\max_{E' \subseteq E}(C(E') - r)/|E'|$ be the highest-cardinality set in that family. Then $e \in E^*$ if and only if

$$\frac{C(E^*) - r}{|E^*|} \geq \frac{C(E^* \setminus \{e\}) - r}{|E^* \setminus \{e\}|}. \tag{A.92}$$

$$= \frac{C(E^*) - c_e - r}{|E^*| - 1} \implies \tag{A.93}$$

$$c_e \geq \frac{C(E^*) - r}{|E^*|}, \tag{A.94}$$

where we define $r/0 = \infty$, thus the capacity constraint is always respected. Since this is true, it follows that

$$\sum_{e\in E} f_e^{\text{lo}} = \sum_{e\in E^*} f_e^{\text{lo}} \tag{A.95}$$

$$= \sum_{e\in E^*} c_e - \frac{C(E^*) - r}{|E^*|} \tag{A.96}$$

$$= C(E^*) - |E^*|\frac{C(E^*) - r}{|E^*|} = r, \tag{A.97}$$

and thus f^{lo} is feasible.

Note that if r^a satisfies (7.16), then for any allowable attack f^a, $f_e^{\text{lo}} + f_e^a \leq c_e$ for all e. Hence by observation 1, (f, f^a) is a SE and a NE. Since $B(f, f^a) = 0$ must hold for any SE, we conclude that $\mathcal{SE}(r, r^a) = \mathcal{NE}(r, r^a)$.

We now turn our attention the case in (7.17). To this end, consider the routing policy f^{hi}, where

$$f_e^{\text{hi}} := \min\left\{c_e, \max_{E'\subseteq E} \frac{r - C(E \setminus E')}{|E'|}\right\}. \tag{A.98}$$

This policy is feasible, which can be shown using a similar argument as that given above for the feasibility of f^{lo}. If r^a satisfies (7.17), then for any allowable attack f^a, $f_e^{\text{hi}} + f_e^a \geq c_e$ for all e. By observation 2, (f^{hi}, f^a) is a SE and a NE. Since $B(f, f^a) = r + r^a - C(E)$ for any SE, we conclude that $\mathcal{NE}(r, r^a) = \mathcal{SE}(r, r^a)$.

Suppose that r^a does not satisfy (7.16). Let $f \in \mathcal{F}(c, r)$ and denote $E^{\text{flow}} = \{e : f_e > 0\}$. Then

$$r^a > \max_{E'\subseteq E} \frac{C(E') - r}{|E'|} \geq \frac{C(E^{\text{flow}}) - r}{|E^{\text{flow}}|} \geq \min_{e\in E^{\text{flow}}} c_e - f_e. \tag{A.99}$$

If e' minimizes the expression in the righthand side of (A.99), then there exists an f^a such that $f_{e'} + f_{e'}^a > c_{e'}$, implying that $B(f, f^a) > 0$. It must then be true that for any

SE (f, f^a), there exists an edge e where $f_e + f_e^a > c_e$.

Suppose that r^a does not satisfy (7.17). Let $f \in \mathcal{F}(c, r)$ and denote $E^{\text{part}} = \{e : f_e < c_e\}$. Then

$$r^a < C(E) - \max_{E' \subseteq E} \frac{r - C(E \subseteq E')}{|E'|} \leq C(E) - \min_{e \in E^{\text{part}}} f_e, \tag{A.100}$$

If e' minimizes the rightmost expression in (A.100), then there must exist an attack policy f^a where $f_{e'} + f_{e'}^a < c_{e'}$. Since $B(f, f^a) > r + r^a - c$, it must be true that for any SE (f, f^a), there must be an edge e where $f_e + f_e^a < c_e$. Therefore, by observation 3 we conclude that when r^a satisfies neither (7.16) nor (7.17), no NE can exist. ■

A.12 Proof for Theorem 7.4

To prove Theorem 7.4, we first show that we need only consider two attacks as best response.

Lemma A.1 *Consider a two-link network. For any f,*

$$B^*(f, r^a) = \max_{f^a \in \{f^{a1}(r^a), f^{a2}(r^a)\}} B(f, f^a), \tag{A.101}$$

where

$$f^{a1}(r^a) := \left(\min\{r^a, c_1\}, \max\{r^a - c_1, 0\}\right), \tag{A.102}$$

$$f^{a2}(r^a) := \left(\max\{r^a - c_2, 0\}, \min\{r^a, c_2\}\right). \tag{A.103}$$

In other words, there always exists a best response attack policy where either (1) the attacker puts as much attack traffic as possible on edge 1 and the reminder on edge 2 (i.e., $f^{a1}(r^a)$); or (2) vice versa (i.e., $f^{a2}(r^a)$).

Proof: *Let f^a be a best response attack policy to f. If $B(f, f^a) = 0$, then the lemma is trivially true. Therefore, let e be an edge where $B_e(f, f^a) > 0$, then one can create a new attack policy $\hat{f}^a$ by redistributing as much attack traffic as possible from the other edge e' to e. Let this amount be δ, so $\hat{f}^a_e = f^a_e + \delta$. Then $B_e(f, \hat{f}^a) = B_e(f, f^a) + \delta$ and $B_{e'}(f, \hat{f}^a) \geq B_{e'}(f, f^a) - \delta$. This implies $B(f, \hat{f}^a) \geq B(f, f^a)$, which is at equality since f^a is a best response. Since $\hat{f}^a \in \{f^{a1}, f^{a2}\}$, we conclude the proof.* ■

Lemma A.1 allows us to only consider two attack policies when solving for the best response, but it also gives us a simple way to solve for a SE. In the two-link case, f is a SE routing policy if

$$f \in \arg\min_{f \in \mathcal{F}} \max \left(B(f, f^{a1}(r^a)), B(f, f^{a2}(r^a)) \right). \tag{A.104}$$

Observe that if $B(f, f^{a1}(r^a)) = B(f, f^{a2}(r^a))$ then f satisfies (A.104), since moving traffic between the edges can only increase $B(f, f^{a1}(r^a))$ or $B(f, f^{a2}(r^a))$. We will leverage this observation to find $B^{\text{SE}}(r, r^a)$ in the following proof.

Now we prove Theorem 7.4, beginning with the case when $\pi^a \cap [c_1, c_2] = \emptyset$. First let $r^a < c_1$, and denote g as the value of the maximization in (7.16). When $r^a \leq g$, we know from the proof of Theorem 7.3 that $B^*(f^{\text{lo}}, r^a) = B^{SE}(r^a) = 0$. When $g < r^a < c_1$, then $B(f^{\text{lo}}, f^{a1}(r^a)) = B(f^{\text{lo}}, f^{a2}(r^a)) = r^a - g$, therefore by the observation above, f^{lo} is a SE routing policy, and $B^*(f^{\text{lo}}, r^a) = B^{SE}(r^a)$.

We now let $r^a > c_2$ - the other possible scenario when $\pi^a \cap [c_1, c_2] = \emptyset$. Here we denote h as the value of the maximization in (7.17). When $r^a \geq C(E) - h$, we know from the proof of Theorem 7.3 that $B^*(f^{\text{hi}}, r^a) = B^{SE}(r^a) = r + r^a - C(E)$. When $c_2 < r^a < h$, then Theorem 7.3 also informs that there must always be an edge e where $B_e = 0$, in the two-link case, one edge is fully blocked and the other has no routed traffic blocked. It follows then that $B(f^{\text{hi}}, f^{a1}(r^a)) = B(f^{\text{hi}}, f^{a2}(r^a)) = h$, and f^{hi} is a SE routing policy.

We conclude that when $\pi^a \cap [c_1, c_2] = \emptyset$, then $V(\pi^a) = 0$.

For the remainder of the proof, we consider the case where $\pi^a \cap [c_1, c_2]$ is nonempty. We leverage the following lemma which simplifies the expression for $B^*(f, r^a) - B^{\mathrm{SE}}(r^a)$.

Lemma A.2 *For a two-link network, if $r^a \in [c_1, c_2]$, then for any f,*

$$B^*(f, r^a) - B^{\mathrm{SE}}(r^a) = |f_1 - (r + r^a - c_2)/2| \tag{A.105}$$

Proof: When $r^a \in [c_1, c_2]$, then we know from Lemma A.1 that for any f,

$$\begin{aligned} B^*(f, r^a) &= \max\{B(f, f^{a1}(r^a), B(f, f^{a2}(r^a))\} && \text{(A.106)} \\ &= \max\{f_1 + \max\{f_2 + \tilde{\gamma}_i - c_1 - c_2, 0\}, \\ &\qquad \max\{f_2 + \tilde{\gamma}_i - c_2, 0\}\}, && \text{(A.107)} \\ &= \max\{f_1, r + \tilde{\gamma}_i - c_1 - c_2, f_2 + \tilde{\gamma}_i - c_2\} && \text{(A.108)} \\ &= \max\{f_1, r - f_1 + r^a - c_2\} && \text{(A.109)} \end{aligned}$$

From the observation made above, a SE routing policy is therefore one where $f_1 = r - f_1 + r^a - c_2$, i.e., f such that

$$f_1 = (r + r^a - c_2)/2, \;\; f_2 = (r - r^a + c_2)/2 \tag{A.110}$$

satisfies (A.104)*. It follows then for any f that*

$$\begin{aligned} B^*(f, r^a) - B^{SE}(r^a) &= \max\{f_1, r - f_1 + r^a - c_2\} \\ &\qquad - (r + r^a - c_2)/2, && \text{(A.111)} \\ &= |f_1 - (r + r^a - c_2)/2|. && \text{(A.112)} \end{aligned}$$

■

As argued in the proof of Lemma 7.4, $\underline{\pi}^a$ need not be included in the maximization in (7.23) if $\underline{\pi}^a \leq c_1$ and $\overline{\pi}^a$ need not be included if $\overline{\pi}^a \geq c_2$. Therefore, our calculation of $V(\pi^a)$ can be further simplified:

$$\begin{aligned} V(\pi^a) &= \min_{f \in \mathcal{F}} \max\{B^*(f, \underline{r}^a) - B^{SE}(\underline{r}^a), \\ &\qquad\qquad B^*(f, \overline{r}^a) - B^{SE}(\overline{r}^a)\}, \qquad \text{(A.113)} \\ &= \min_{f \in \mathcal{F}} \max\{|f_1 - (r + \underline{r}^a - c_2)/2|, \\ &\qquad\qquad |f_1 - (r + \overline{r}^a - c_2)/2|\}, \qquad \text{(A.114)} \end{aligned}$$

where $\underline{r}^a := \max\{c_1, \underline{\pi}^a\}$ and $\overline{r}^a := \min\{c_2, \overline{\pi}^a\}$. This implies that the minimizing value of f_1 in (A.114) is halfway between $(r + \underline{r}^a - c_2)/2$ and $(r + \overline{r}^a - c_2)/2$, i.e.,

$$f_1 = (2r + \underline{r}^a + \overline{r}^a - 2c_2)/4, \qquad \text{(A.115)}$$

which implies that $V(\pi^a) = (\overline{r}^a - \underline{r}^a)/4$. ■

Bibliography

[1] R. V. B. Vinod, P. Vadakkepat, M. Sundaram, K. Sujatha, and J. J. Brislin, *Advancements in Automation, Robotics and Sensing*. Springer, 2016.

[2] S. L. Young and F. J. Pierce, *Automation: The future of weed control in cropping systems*. Springer, 2013.

[3] R. T. Yarlagadda, *Ai automation and it's future in the unitedstates*, *International Journal of Creative Research Thoughts (IJCRT), ISSN* (2017) 2320–2882.

[4] M. D. Lytras, V. Raghavan, and E. Damiani, *Big data and data analytics research: From metaphors to value space for collective wisdom in human decision making and smart machines*, *International Journal on Semantic Web and Information Systems (IJSWIS)* **13** (2017), no. 1 1–10.

[5] Y. Karasawa, H. Nakayama, and S. Dohi, *Trade-off analysis for optimal design of automated warehouses*, *International Journal of Systems Science* **11** (1980), no. 5 567–576.

[6] A. E. Turgut, H. Çelikkanat, F. Gökçe, and E. Şahin, *Self-organized flocking in mobile robot swarms*, *Swarm Intelligence* **2** (2008), no. 2 97–120.

[7] V. R. Lesser, *Multiagent systems: An emerging subdiscipline of ai*, *ACM Computing Surveys (CSUR)* **27** (1995), no. 3 340–342.

[8] D. Grimsman, V. Chetty, N. Woodbury, E. Vaziripour, S. Roy, D. Zappala, and S. Warnick, *A case study of a systematic attack design method for critical infrastructure cyber-physical systems*, in *2016 American Control Conference (ACC)*, pp. 296–301, IEEE, 2016.

[9] D. Ganesan, B. Krishnamachari, A. Woo, D. Culler, D. Estrin, and S. Wicker, *Complex behavior at scale: An experimental study of low-power wireless sensor networks*, tech. rep., Technical Report UCLA/CSD-TR 02, 2002.

[10] A. Meissner, T. Luckenbach, T. Risse, T. Kirste, and H. Kirchner, *Design challenges for an integrated disaster management communication and information system*, in *The First IEEE Workshop on Disaster Recovery Networks (DIREN 2002)*, vol. 24, pp. 1–7, 2002.

[11] D. Albani, T. Manoni, D. Nardi, and V. Trianni, *Dynamic uav swarm deployment for non-uniform coverage*, in *Proceedings of the 17th international conference on autonomous agents and multiagent systems*, pp. 523–531, 2018.

[12] R. Shokri and V. Shmatikov, *Privacy-preserving deep learning*, in *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, pp. 1310–1321, 2015.

[13] D. Grimsman, J. H. Seaton, J. R. Marden, and P. N. Brown, *The cost of denied observation in multiagent submodular optimization*, in *2020 59th IEEE Conference on Decision and Control (CDC)*, pp. 1666–1671, IEEE, 2020.

[14] J. R. Marden and J. S. Shamma, *Game theory and distributed control*, *Handbook of Game Theory with Economic Applications* **4** (2015), no. 1 861–899.

[15] M. Gairing, *Covering games: Approximation through non-cooperation*, in *Lecture Notes in Computer Science*, vol. 5929, pp. 184–195, 2009.

[16] A. Vetta, *Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions*, in *Annual Symposium on Foundations of Computer Science*, pp. 416–425, IEEE, 2002.

[17] A. Krause, A. Singh, and C. Guestrin, *Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies*, *Journal of Machine Learning Research* **9** (2008) 235–284.

[18] D. Kempe, J. Kleinberg, and É. Tardos, *Maximizing the spread of influence through a social network*, in *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 137–146, 2003. arXiv:0806.2034.

[19] M. Gomez-Rodriguez, J. Leskovec, and A. Krause, *Inferring networks of diffusion and influence*, *ACM Transactions on Knowledge Discovery from Data* **5** (2012), no. 4.

[20] P. Kohli, M. Pawan Kumar, and P. H. Torr, *P3 & beyond: Move making algorithms for solving higher order functions*, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **31** (2009), no. 9 1645–1656.

[21] O. Barinova, V. Lempitsky, and P. Kholi, *On detection of multiple object instances using hough transforms*, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **34** (2012), no. 9 1773–1784.

[22] H. Lin and J. Bilmes, *A class of submodular functions for document summarization*, in *Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, vol. 1, pp. 510–520, 2011.

[23] A. Singh, W. Kaiser, M. Batalin, A. Krause, and C. Guestrin, *Efficient planning of informative paths for multiple robots*, in *International Joint Conference on Artificial Intelligence*, pp. 2204–2211, 2007.

[24] A. Krause, R. Rajagopal, A. Gupta, and C. Guestrin, *Simultaneous placement and scheduling of sensors*, in *International Conference on Information Processing in Sensor Networks*, pp. 181–192, 2009.

[25] J. R. Marden, *The role of information in distributed resource allocation*, *IEEE Transactions on Control of Network Systems* **4** (2017), no. 3 654–664.

[26] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. Vanbriesen, and N. Glance, *Cost-effective outbreak detection in networks*, in *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 420–429, 2007.

[27] B. Mirzasoleiman, A. Karbasi, R. Sarkar, and A. Krause, *Distributed submodular maximization*, *Journal of Machine Learning Research* **17** (2016), no. 1 8330–8373.

[28] G. Qu, D. Brown, and N. Li, *Distributed greedy algorithm for multi-agent task assignment problem with submodular utility functions*, *Automatica* **105** (2019), no. February 206–215.

[29] A. Clark and R. Poovendran, *A submodular optimization framework for leader selection in linear multi-agent systems*, in *IEEE Conference on Decision and Control*, pp. 3614–3621, IEEE, 2011.

[30] M. Grötschel, L. Lovász, and A. Schrijver, *The ellipsoid method and its consequences in combinatorial optimization*, *Combinatorica* **1** (1981), no. 2 169–197.

[31] S. Iwata, L. Fleischer, and S. Fujishige, *A combinatorial, strongly polynomial-time algorithm for minimizing submodular functions*, *Journal of the ACM* **48** (2001), no. 4 761–777.

[32] A. Schrijver, *A combinatorial algorithm minimizing submodular functions in strongly polynomial time*, *Journal of Combinatorial Theory. Series B* **80** (2000), no. 2 346–355.

[33] L. Lovász, *Submodular functions and convexity*, in *Mathematical Programming The State of the Art*, pp. 235–257. 1983.

[34] C. H. Papadimitriou, *Algorithms, games, and the internet*, in *Annual ACM symposium on Theory of Computing*, pp. 749–753, 2001.

[35] E. Koutsoupias and C. Papadimitriou, *Worst-case equilibria*, in *Annual Symposium on Theoretical Aspects of Computer Science*, pp. 404–413, Springer, 1999.

[36] T. Roughgarden, *Selfish routing and the price of anarchy*. MIT press, 2005.

[37] N. Andelman, M. Feldman, and Y. Mansour, *Strong price of anarchy*, *Games and Economic Behavior* **65** (2009), no. 2 289–317.

[38] P. N. Brown and J. R. Marden, *Studies on robust social influence mechanisms: Incentives for efficient network routing in uncertain settings*, *IEEE Control Systems Magazine* **37** (2017), no. 1 98–115.

[39] R. P. Leme and E. Tardos, *Pure and bayes-nash price of anarchy for generalized second price auction*, in *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pp. 735–744, IEEE, 2010.

[40] P. N. Brown, H. P. Borowski, and J. R. Marden, *Projecting network games onto sparse graphs*, in *Asilomar Conference on Signals, Systems and Computers*, pp. 307–309, 2018.

[41] H. Jaleel, W. Abbas, and J. S. Shamma, *Robustness of Stochastic Learning Dynamics to Player Heterogeneity in Games*, in *IEEE Conference on Decision and Control*, pp. 5002–5007, 2019.

[42] B. Gharesifard and S. L. Smith, *Distributed submodular maximization with limited information*, *IEEE Transactions on Control of Network Systems* **5** (2017), no. 4 1635–1645, [arXiv:1706.0408].

[43] P. N. Brown and J. R. Marden, *On the feasibility of local utility redesign for multiagent optimization*, in *18th European Control Conference*, pp. 3396–3401, 2019.

[44] R. Chandan, D. Paccagnan, and J. R. Marden, *Optimal price of anarchy in cost-sharing games*, in *American Control Conference*, pp. 2277–2282, 2019.

[45] J. Nash, *Non-Cooperative Games*, *Annals of Mathematics* **54** (1951), no. 2 286–295.

[46] R. Gopalakrishnan, J. R. Marden, and A. Wierman, *An architectural view of game theoretic control*, *ACM SIGMETRICS Performance Evaluation Review* **38** (2010), no. 3 31–36.

[47] J. P. Hespanha, *Noncooperative game theory: an introduction for engineers and computer scientists*. 2016.

[48] C. Alós-Ferrer and N. Netzer, *The logit-response dynamics*, *Games and Economic Behavior* **68** (2010), no. 2 413–427.

[49] J. R. Marden and J. S. Shamma, *Revisiting log-linear learning: asynchrony, completeness and payoff-based implementation*, *Games and Economic Behavior* **75** (2012), no. 2 788–808.

[50] C. Godsil and G. Royle, *Algebraic graph theory*. Springer, 2001.

[51] S. Arumugam and K. Reji Kumar, *Fractional independence and fractional domination chain in graphs.*, *AKCE International Journal of Graphs and Combinatorics* **4** (2007), no. 2 161–169.

[52] J. Matousek and B. Gärtner, *Understanding and using linear programming*. Springer Science & Business Media, 2007.

[53] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, *An analysis of approximations for maximizing submodular set functions-I*, *Mathematical Programming* **14** (1978), no. 1 265–294.

[54] M. L. Fisher, G. L. Nemhauser, and L. A. Wolsey, *An analysis of approximations for maximizing submodular set functions-II*, *Polyhedral Combinatorics* (1978) 73–87.

[55] M. Minoux, *Accelerated greedy algorithms for maximizing submodular set functions*, in *Optimization Techniques*, pp. 234–243. Springer Berlin Heidelberg, 1978.

[56] N. Buchbinder, M. Feldman, J. Naor, and R. Schwartz, *A tight linear time (1/2)-approximation for unconstrained submodular maximization*, *SIAM Journal on Computing* **44** (2015), no. 5 1384–1402.

[57] J. Vondrák, *Optimal approximation for the submodular welfare problem in the value oracle model*, in *ACM Symposium on Theory of Computing*, pp. 67–74, 2008.

[58] M. Sviridenko, *A note on maximizing a submodular set function subject to a knapsack constraint*, *Operations Research Letters* **32** (2004), no. 1 41–43.

[59] G. Calinescu, C. Chekuri, M. Pál, and J. Vondrák, *Maximizing a monotone submodular function subject to a matroid constraint*, *SIAM Journal on Computing* **40** (2011), no. 6 1740–1766, [9780201398298].

[60] Y. Filmus and J. Ward, *The power of local search: maximum coverage over a matroid*, in *Symposium on Theoretical Aspects of Computer Science*, pp. 601–612, LIPIcs, 2012.

[61] U. Feige, *A threshold of ln n for approximating set cover*, *Journal of the ACM* **45** (1998), no. 4 634–652.

[62] B. Mirzasoleiman, A. Karbasi, R. Sarkar, and A. Krause, *Distributed submodular maximization: Identifying representative elements in massive data*, in *Advances in Neural Information Processing Systems*, pp. 2049–2057, 2013.

[63] G. Arslan, J. R. Marden, and J. S. Shamma, *Autonomous vehicle-target assignment: a game-theoretical formulation*, *Journal of Dynamic Systems, Measurement and Control, Transactions of the ASME* **129** (2007), no. 5 584–596.

[64] P. Turán, *On an extremal problem in graph theory*, *Mat. Fiz. Lapok* **48** (1941), no. 436-452 137.

[65] L. Volkmann, *On perfect and unique maximum independent sets in graphs.*, *Mathematica Bohemica* **129** (2004), no. 3 273–282.

[66] M. R. Kirchner, J. P. Hespanha, and D. Garagić, *Heterogeneous measurement selection for vehicle tracking using submodular optimization*, in *IEEE Aerospace Conference*, pp. 1–10, 2020. arXiv:1910.0914.

[67] A. Pázman, *Foundations of optimum experimental design*, vol. 14. Springer, 1986.

[68] T. H. Summers, F. L. Cortesi, and J. Lygeros, *On Submodularity and Controllability in Complex Dynamical Networks*, *IEEE Transactions on Control of Network Systems* **3** (2016), no. 1 91–101, [arXiv:1404.7665].

[69] P. Goundan and A. Schulz, *Revisiting the greedy approach to submodular set function maximization*, *Optimization online* (2007), no. 1984 1–25.

[70] B. Lehmann, D. Lehmann, and N. Nisan, *Combinatorial auctions with decreasing marginal utilities*, *Games and Economic Behavior* **55** (2006), no. 2 270–296.

[71] M. S. Kang, S. B. Lee, and V. D. Gligor, *The crossfire attack*, *Proceedings - IEEE Symposium on Security and Privacy* (2013) 127–141.

[72] D. Gkounis, V. Kotronis, and X. Dimitropoulos, *Towards Defeating the Crossfire Attack using SDN*, arXiv:1412.2013.

[73] A. Aydeger, N. Saputro, K. Akkaya, and M. Rahman, *Mitigating Crossfire Attacks Using SDN-Based Moving Target Defense*, in *Conference on Local Computer Networks*, pp. 627–630, 2016.

[74] D. Gkounis, V. Kotronis, C. Liaskos, and X. Dimitropoulos, *On the interplay of link-flooding attacks and traffic engineering*, *Computer Communication Review* **46** (2016), no. 2 5–11, [arXiv:1611.0248].

[75] A. R. Narayanadoss, T. Truong-Huu, P. M. Mohan, and M. Gurusamy, *Crossfire attack detection using deep learning in software defined its networks*, *IEEE Vehicular Technology Conference* **2019-April** (2019) [1812.0363].

[76] H. Von Stackelberg, *Marktform und gleichgewicht*. J. springer, 1934.

[77] Y. A. Korilis, A. A. Lazar, and A. Orda, *Stackelberg Routing Strategies*, *IEEE/ACM Transactions on Networking* **5** (1997), no. 1 161–173.

[78] T. Roughgarden, *Stackelberg scheduling strategies*, *SIAM Journal on Computing* **33** (2004), no. 2 332–350.

[79] M. Bloem, T. Alpcan, and T. Basar, *A Stackelberg Game for Power Control and Channel Allocation in Cognitive Radio Networks*, in *VALUETOOLS*, pp. 1–9, 2007.

[80] X. He, A. Prasad, S. P. Sethi, and G. J. Gutierrez, *A survey of Stackelberg differential game models in supply and marketing channels*, *Journal of Systems Science and Systems Engineering* **16** (2007), no. 4 385–413.

[81] G. Brown, M. Carlyle, J. Salmerón, and K. Wood, *Defending critical infrastructure*, *Interfaces* **36** (2006), no. 6 530–544.

[82] J. Pita, M. Jain, J. Marecki, F. Ordóñez, C. Portway, M. Tambe, C. Western, P. Paruchuri, and S. Kraus, *Deployed ARMOR protection: The application of a game-theoretic model for security at the Los Angeles international airport*, in *Conf. on Autonomous Agents and Multiagent Systems*, 2008.

[83] M. Jain, J. Tsai, J. Pita, C. Kiekintveld, S. Rathi, M. Tambe, and F. Ordóñez, *Software assistants for randomized patrol planning for the lax airport police and the Federal Air Marshal Service*, *Interfaces* **40** (2010), no. 4 267–290.

[84] G. G. Brown, W. M. Carlyle, J. Salmerón, and K. Wood, *Analyzing the Vulnerability of Critical Infrastructure to Attack and Planning Defenses*, *Emerging Theory, Methods, and Applications* (2005) 102–123.

[85] D. Korzhyk, Z. Yin, C. Kiekintveld, V. Conitzer, and M. Tambe, *Stackelberg vs. nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness*, *Journal of Artificial Intelligence Research* **41** (2011) 297–327.

[86] P. Paruchuri, J. P. Pearce, J. Marecki, M. Tambe, F. Ordonez, and S. Kraus, *Playing games for security: An efficient exact algorithm for solving Bayesian Stackelberg games*, in *International Joint Conference on Autonomous Agents and Multiagent Systems*, vol. 2, pp. 877–884, 2008.

[87] H. Kellerer, U. Pferschy, and D. Pisinger, *Knapsack problems*. Springer, Berlin, 2004.