Network Security and Min-Cost Max-Flow Problem

by

Mathieu Dahan

B.Eng., École Centrale Paris (2013)
B.S., Université Paris-Sud XI (2013)

Submitted to the Center for Computational Engineering
in partial fulfillment of the requirements for the degree of

Master of Science in Computation for Design and Optimization

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2016

Author . . . .
Center for Computational Engineering
January 20, 2016

Certified by . . . .
Saurabh Amin
Assistant Professor of Civil and Environmental Engineering
Thesis Supervisor

Accepted by . . . .
Nicolas Hadjiconstantinou
Professor of Mechanical Engineering
Co-Director, Computation for Design and Optimization

Network Security and Min-Cost Max-Flow Problem

by

Mathieu Dahan

Submitted to the Center for Computational Engineering
on January 20, 2016, in partial fulfillment of the
requirements for the degree of
Master of Science in Computation for Design and Optimization

Abstract

Network optimization has widely been studied in the literature for a variety of design and operational problems. This has resulted in the development of computational algorithms for the study of classical operations research problems such as the maximum flow problem, the shortest path problem, and the network interdiction problem. However, in environments where network components are subject to adversarial failures, the network operator needs to strategically allocate at least some of her resources (e.g., link capacities, network flows, etc.) while accounting for the presence of a strategic adversary. This motivates the study of network security games. This thesis considers a class of network security games on flow networks, and focuses on utilizing well-known results in network optimization toward the characterization of Nash equilibria of this class of games.

Specifically, we consider a 2-player strategic game for network routing under link disruptions. Player 1 (defender) routes flow through a network to maximize her value of effective flow while facing transportation costs. Player 2 (attacker) simultaneously disrupts one or more links to maximize her value of lost flow but also faces cost of disrupting links. Linear programming duality and the Max-Flow Min-Cut Theorem are applied to obtain properties that are satisfied in any Nash equilibrium. Using graph theoretic arguments, we give a characterization of the support of the equilibrium strategies. Finally, we study the conditions under which these results extend to a revised version of the game where both players face budget constraints. Thus, our contribution can be viewed as a generalization of the classical minimum cost maximum flow problem and the minimum cut problem to adversarial environments.

Thesis Supervisor: Saurabh Amin
Title: Assistant Professor of Civil and Environmental Engineering

Acknowledgments

I sincerely thank my advisor, Professor Saurabh Amin, for giving me the opportunity to study at MIT and to work with him. I am really thankful for all the time he devoted to the completion of this work and for his help during my Master's studies. Working under his supervision has been an amazing experience.

I would also like to thank Kate Nelson at the CDO office for her kindness and her help regarding the administrative formalities.

Most importantly, I am deeply grateful to my parents, to whom I owe everything. I would not be here if it were not for them, and I cannot thank them enough for everything they have done for me.

Finally, I would like to acknowledge the financial support of FORCES (Foundations Of Resilient CybEr-Physical Systems), which receives support from the National Science Foundation (NSF award numbers CNS-1238959, CNS-1238962, CNS-1239054, CNS-1239166), NSF CAREER award CNS-1453126, and the AFRL LABLET - Science of Secure and Resilient Cyber-Physical Systems (Contract ID: FA8750-14-2-0180, SUB 2784-018400).

Contents

1 Introduction **13**

1.1 Overview of the Problem . . . 13

1.2 Related Work . . . 14

1.3 Main Contributions . . . 15

2 Game Theoretic Model **19**

2.1 Preliminaries . . . 19

2.2 The Game . . . 21

2.3 Standard Definitions and Main Assumption . . . 24

3 Characterization of Nash Equilibria **29**

3.1 Preliminary Results . . . 29

3.2 Main Theorem . . . 37

3.3 Necessary Conditions . . . 49

4 Budget-Constrained Game **61**

4.1 Revised Model . . . 62

4.2 **P1**'s Budget . . . 63

4.3 **P2**'s Budget . . . 65

4.3.1 Partition-Based Equilibrium Strategies . . . 66

4.3.2 Optimization Problem . . . 75

5 Some Extensions **83**

5.1 Relaxing Assumption 1 . . . 83

5.2 Transportation Cost 85
5.3 Cost of Attack 86

6 Conclusion **89**
6.1 Summary of the Results 89
6.2 Future Work 90

List of Figures

2-1 Example graph. 23

2-2 One shot of game Γ played on a given graph. 23

2-3 Min-cost max-flow (drawn in blue) in a graph satisfying (A1). The labels of each edge correspond to the flow it carries (blue), its capacity (red) and its transportation cost (green). 26

3-1 Support of equilibrium strategies in Regions **I-III**. 32

3-2 Example network. Edge capacities and transportation costs are labeled in red and green colors respectively. 37

3-3 NE described in Props. 1, 2 and 3. The min-cost max-flow (resp. min-cut set attack) is in bold blue (resp. dotted red). 38

3-4 Example network containing edges outside a min-cut set that are disrupted in equilibrium. 51

3-5 Min-cost max-flow (bold blue) and min-cut set attack (dotted red). The labels in the boxes represent the edges that are saturated by the min-cost max-flow. 54

4-1 Equilibrium strategy described in Prop. 8 when $p_1 > 3$, $p_2 = 4$ and $b_1 = 4.5$. $x^\dagger$ is drawn in blue. 65

4-2 Support of partition-based equilibrium strategies in Regions **I-III**. . . 67

4-3 Probability distribution of **P2**'s partition based equilibrium strategies. 67

4-4 NE described in Prop. 9 based on the partition $\{\mu^1, \mu^2\}$. 75

4-5 Enumeration of the equilibrium strategies in $\mathcal{S}_{p_1}$. 81

5-1 Initial graph. 83

5-2 NE in the case $p_1 = p_2 = 6$. 84

5-3 Removal of paths that are too costly for **P1**. 85

5-4 Min-cost max-flow (bold blue) and min-cut set (dotted red) of a graph satisfying (A2). 86

5-5 Graph with general cost of attack. Edge capacities, transportation cost and cost of attack are denoted in red, green and orange colored labels respectively. 87

5-6 NE in the case $p_1 > 2$ and $p_2 > 3$ for a graph with general cost of attack. 87

List of Tables

3.1 Players' payoffs of a zero-sum game. 47

3.2 Players' payoffs of a game strategically equivalent to the zero-sum game presented in Table 3.1. 47

Chapter 1

Introduction

1.1 Overview of the Problem

In this thesis, we study network flow routing in the wake of disruptions caused by strategic adversaries. Our work is motivated by the need of improving the operational resilience of transportation networks and other flow networks under strategic disruptions. Our goal in this thesis is to give a realistic model of an attacker-defender game played on a network and investigate the physical properties of the game in equilibrium. More particularly, we relate the structure of the players' equilibrium strategies to the solutions of classical routing problems.

We study a two-player non-cooperative game over a directed network in which Player 1 (defender or operator) chooses a flow to be routed from a source node to a destination node, and Player 2 (attacker or interdictor) chooses to disrupt one or more edges. In our model, we account the value of the effective flow and the transportation cost for the defender, and the value of the lost flow and the cost of attack for the attacker. Specifically, Player 1's payoff linearly increases in the amount of effective flow that reaches the destination node, but decreases with the cost of transporting the initial flow chosen by her. Player 2's payoff linearly increases in the amount of lost flow as a result of an attack and decreases with the attacking cost. This two-player game is not a zero-sum game and the payoff structures are motivated by the previous formulations in both network interdiction problems (see Wood and Kevin

[25], Cormican et al. [10], Bertsimas et al. [9], Avenhaus and Canty [4]) and network security games (see Gueye et al. [17], Szeto [24], Baykal-Gürsoy et al. [6], Dahan and Amin [11]).

1.2 Related Work

Network interdiction problems have already been widely studied, but our focus is to extend these formulations to simultaneous game settings. Related to our approach is the article by Washburn and Wood [3]. The authors model a sequential game, where the defender (leader) chooses one $s-t$ path and then the interdictor (follower) inspects one edge. The objective of this sequential game model is to maximize the probability with which the operator is detected by the interdictor. Our game differs from the model by Washburn and Wood [3] in that we model a simultaneous game, which captures each player's strategic uncertainty about her opponent. Secondly, we allow both players to have a much larger set of actions (feasible flow that may contain many $s-t$ paths and loops, and attacks that can simultaneously disrupt several edges). Finally, we account for the attacker's cost of attack as well as the defender's cost of transporting flow through the network.

The work by Avenhaus and Canty [4] presents several models of inspection games. One of them considers a two-player game between a passenger of a subway system and the local transit authority. Simultaneously, the passenger chooses whether she pays the ticket or not, and the transit authority decides whether to inspect the passenger or not. This model was motivated by the fact that inspecting all the time is either too costly or not worthwhile. Actually, this model is analogous to our game played on a single link (where the transit authority is an interdictor). By adjusting the parameters, we find that these models are strategically equivalent and our results applied to a graph composed of a single link coincide. Thus, our model can be viewed as a generalization of their inspection game to a network setting where the interdictor chooses the probability of inspecting specific links on the network.

Another related line of work in network interdiction games is by Bertsimas et al.

[8]. In this sequential game, the operator first chooses a feasible flow, and then the interdictor disrupts a fixed number of edges. The goal is to minimize the largest amount of flow that reaches the destination node. The authors consider two different models for the disruption: an arc-based formulation where the flow can be rerouted when there is an attack, and a path-based formulation where the flow carried by a disrupted edge is lost. Our formulation is related to the path-based formulation. Since we model a simultaneous game, it is reasonable to assume that the flow through disrupted edges is lost and cannot be re-routed. Although, in [8], the interdictor can disrupt several edges at the same time, she must always disrupt the same number of edges for every action, which is still a restriction of the set of actions we considered for our interdictor.

Our model is also related to the work of Hong and Wooders [18] and Gueye et al. [17] (also see [20]). In these papers the authors model simultaneous attacker-defender games, where the defender (operator) chooses a feasible flow and the attacker (interdictor) disrupts edges of the network, preventing the flow from reaching the destination node. The major differences with Gueye et al. [17] is that their attacker can only disrupt one edge, and they consider an uncapacitated graph with given supplies and demands, while we consider a capacitated graph with no constraint on the supplies and demands.

We note that Goyal and Vigier [16], and Acemoglu et al. [1] studied network security as well, but from another perspective. Indeed, in their models, the attacker targets the nodes of a network previously chosen by the defender. The latter chooses the network and the allocation of defense resources in order to minimize the cascading effect due to the attack. In our model, we decided that the flow that was supposed to take an edge that is disrupted is simply lost, so there is no contagion.

1.3 Main Contributions

For the sake of simplicity, we restrict the class of graphs we study in this thesis (even though our study applies to a much larger class of graphs). This enables us to develop a rather complete characterization of the equilibria of our game. Our first result is that, given the different characteristics of the graph (attacking and transportation cost), we give a tractable formulation of one Nash equilibrium that is based on minimum cost maximum flows for Player 1 and on minimum cut sets for Player 2, extending the results of Avenhaus and Canty [4] and of Hong and Wooders [18].

We also present theoretical properties satisfied by all the Nash equilibria of our game. The most interesting property is that each player has a unique payoff value in *all* equilibria, and we were able to analytically compute the values of effective and lost flow and the costs of transportation and attack in terms of the parameters of the game. The interest comes from the fact that these characteristics, that are common for all equilibria, need not full enumeration of the equilibria; they can be derived using a combination of game theoretic and network optimization ideas, one of them being the Max-Flow Min-Cut Theorem.

We show that the support of the strategies that can be potentially Nash equilibria can be restricted using graph theoretic properties of the network. We know that computing a Nash equilibrium is hard. While Daskalakis et al. [12] showed that the computation of a Nash equilibrium is PPAD-complete for a general two-player game, Von Neumann [21] showed that, for a zero-sum two-player game, finding Nash equilibria is equivalent to solving a linear programming problem. We show that our game is equivalent to a zero-sum game, thus we can use linear programming techniques to compute a Nash equilibrium efficiently. By restricting the support of the strategies that can potentially be Nash equilibria, we decrease the number of variables and constraints of the linear problems, thus speeding the computation of Nash equilibria. We also give an alternative to computing a Nash equilibrium. Indeed, we computed in closed form an equilibrium based on a minimum cost maximum flow for the defender and on a minimum cut set for the attacker. Thus, in order to find one Nash equilibrium, one only need to compute a minimum cost maximum flow and a minimum cut set of a network, which can be done efficiently by viewing the minimum cost maximum flow problem as a minimum cost circulation problem (see Goldberg and Tarjan [15]). This shows how we can use and extend classical routing problems to adversarial environments.

Lastly, we study a generalization of the game where both players face budget constraints. Specifically, we view the transportation cost of a flow (resp. cost of an attack) as a resource that needs to be available to the defender (resp. attacker) in order to send the flow (resp. lead the attack). We compute the minimum budget that players must have to ensure that the equilibrium properties derived for the previous game still hold for the new game. Using the infiniteness of the defender's set of actions, we give a tight lower bound for the defender's budget for transporting flows. However, the attacker's set of action is discrete and we cannot derive an analogous lower bound for the budget needed by the attacker to conduct attacks. We restrict our attention to a subset of the attacker's equilibrium strategies: we compute in closed form equilibrium strategies for the attacker that are constructed from the partitions of the minimum cut sets. Then we find the equilibrium strategies in this subset that require the lowest budget. It turns out that we can formulate this problem as an integer programming problem whose optimal solution gives a bound (maybe not the best one) on the attacker's budget for which our analysis holds.

The rest of the thesis is organized as follows: in Chapter 2 we discuss the main assumptions and present our game model. Chapter 3 presents our main results on the characterization of Nash equilibria of the game and their relations with classical routing problems. Then in Chapter 4 we extend our results to a budget-constrained game. Lastly, the implications of relaxing some of the modeling assumptions are discussed in Chapter 5, before briefly discussing how to extend the game settings in Chapter 6.

Chapter 2

Game Theoretic Model

In this chapter, we recall some standard definitions and results in game theory and network optimization such as the Minimax Theorem for zero-sum games and the Max-Flow Min-Cut Theorem for directed graphs. Then we set up the game theoretic model that we solve in this thesis, involving a defender and an attacker playing on a flow network. Finally, we present an assumption on the class of graphs we focus on for this thesis.

2.1 Preliminaries

Consider a capacitated directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where $\mathcal{V}$ (resp. $\mathcal{E}$) represents the set of nodes (resp. the set of edges) of $\mathcal{G}$. For each edge $(i,j) \in \mathcal{E}$, let $c_{ij} \in \mathbb{R}^+$ denote its capacity. Let $s \in \mathcal{V}$ denote a source node and $t \in \mathcal{V}$ a destination node. A flow, defined by the function $x : \mathcal{E} \to \mathbb{R}^+$, can only enter the network from s and leave from t. There is no demand or supply at other nodes. A flow x is said to be feasible if it satisfies flow conservation at each node and if the flow through each edge does not exceed its capacity:

$$\forall i \in \mathcal{V} \backslash \{s,t\}, \sum_{(j,i) \in \mathcal{E}} x(j,i) = \sum_{(i,j) \in \mathcal{E}} x(i,j)$$

$$\forall (i,j) \in \mathcal{E},\ 0 \leq x(i,j) \leq c_{ij}.$$

Let $\mathcal{F}$ denote the set of feasible flows, and Λ the set containing all the loops and source-destination $s-t$ paths of the network. Let $x_{ij} := x(i,j)$ denote the flow through edge (i,j), and x_λ the quantity of flow of x sent through $\lambda \in \Lambda$. The edge flows x_{ij} and loop/path flows x_λ satisfy:

$$\forall (i,j) \in \mathcal{E},\ x_{ij} = \sum_{\{\lambda \in \Lambda \,|\, (i,j) \in \lambda\}} x_\lambda. \tag{2.1}$$

An $s-t$ *cut* is a partition $\{S,T\}$ of $\mathcal{V}$, such that $s \in S$ and $t \in T$. The *cut-set* of $\{S,T\}$ and its *capacity* are defined as $E(\{S,T\}) = \{(i,j) \in \mathcal{E} \,|\, i \in S,\ j \in T\}$ and $C(\{S,T\}) = \sum_{(i,j) \in E(\{S,T\})} c_{ij}$. Let $\mathrm{F}(x) = \sum_{\{i \in \mathcal{V} \,|\, (i,t) \in \mathcal{E}\}} x_{it}$ denote the amount of flow passing from the source s to the sink t. We recall the *max-flow problem*:

$$\begin{aligned} (\mathcal{P}_1) \quad & \text{maximize} && \mathrm{F}(x) \\ & \text{subject to} && x \in \mathcal{F}. \end{aligned}$$

The well-known *Max-Flow Min-Cut Theorem* by Ford and Fulkerson [14] states that the optimal value of the maximum flow problem is equal to the minimum capacity over all $s-t$ cuts. We call *min-cut set*, the cut-set of a minimum capacity $s-t$ cut. We also state the *minimum cost maximum flow problem* by Edmonds and Karp [13]:

$$\begin{aligned} (\mathcal{P}_2) \quad & \text{minimize} && \sum_{(i,j) \in \mathcal{E}} b_{ij} x_{ij} \\ & \text{subject to} && x \in \mathcal{F} \\ & && \mathrm{F}(x) \geq \mathrm{F}(x'), \quad \forall x' \in \mathcal{F}, \end{aligned}$$

where for every edge $(i,j) \in \mathcal{E}$, $b_{ij} \in \mathbb{R}^+$ denotes the cost of transporting a unit flow through (i,j).

We use Θ_1 (resp. Ω_1) to denote the optimal value (resp. optimal solution set) of the max-flow problem $(\mathcal{P}_1)$. Similarly, we denote the optimal value (resp. the set of optimal solutions) of problem $(\mathcal{P}_2)$ by Θ_2 (resp. Ω_2).

2.2 The Game

We focus on a simultaneous, semi-infinite, two-player strategic game denoted $\Gamma := \langle\{1,2\},(\mathcal{F},\mathcal{A}),(u_1,u_2)\rangle$ defined as follows: player 1 (**P1**) is the defender (operator) who chooses to route a flow $x \in \mathcal{F}$ through the network, and player 2 (**P2**) is the attacker (interdictor) who chooses an attack μ to disrupt a subset of edges of the graph $\mathcal{G}$. The action set for **P1** (resp. **P2**) is given by $\mathcal{F}$ (resp. $\mathcal{A} := \{0,1\}^{\mathcal{E}}$).

An attack μ is a function from $\mathcal{E}$ to $\{0,1\}$ defined as follows:

$$\mu_{ij} := \mu(i,j) = \begin{cases} 1 & \text{if } (i,j) \text{ is disrupted,} \\ 0 & \text{otherwise.} \end{cases} \tag{2.2}$$

Note that **P2** can disrupt multiple edges of the network by choosing a single attack μ.

We use the following notation to describe certain specific player actions: x^0 the action of not sending flow in the network, x^* an optimal solution of $(\mathcal{P}_2)$ i.e., a min-cost max-flow, μ^0 the action of not attacking any edge of the network, and μ^{min} the action that disrupts all the edges of a min-cut set of the network.

Since Γ is a simultaneous game, it is reasonable to assume that after an edge is disrupted, the flow that was supposed to cross this edge (if there were no attack) is lost and it is not re-routed. For the sake of simplicity, we do not consider attacks that can only result in partially disrupted edges and might still permit some flow to pass through the attacked edges. Thus, the *effective flow*, denoted x^μ, when a flow x is chosen by **P1** and an attack μ is chosen by **P2** can be expressed as follows:

$$\forall (i,j) \in \mathcal{E},\ x^\mu_{ij} = \sum_{\lambda \in \Lambda^\mu_{ij}} x_\lambda,$$

where $\Lambda^\mu_{ij} := \{\lambda \in \Lambda \mid (i,j) \in \lambda \text{ and } \forall (i',j') \in \lambda, \mu_{i'j'} = 0\}$. That is, the effective flow through an edge (i,j) is the sum of all the initial path flows through edge (i,j) that do not contain any attacked edge. The effective flow x^μ can be viewed as a feasible flow in $\mathcal{F}$ that succesfully carries the amount of flow from x that is not lost due to the attack μ.

In this model, the payoff of **P1** is defined as the value of effective flow assessed by **P1** net the cost of transporting the initial flow:

$$u_1(x, \mu) = \underbrace{p_1 \mathrm{F}(x^\mu)}_{\text{value of effective flow}} - \underbrace{\mathrm{C}_1(x)}_{\text{transportation cost}} \tag{2.3}$$

where $p_1 \in \mathbb{R}^+$ is the marginal value of the flow for **P1**, and $\mathrm{C}_1(x) := \sum_{(i,j)\in\mathcal{E}} b_{ij} x_{ij}$ is the cost of transporting the initial flow x. Thus, when one additional unit of flow reaches t, **P1**'s payoff increases by p_1 and at the same time decreases by its transportation cost.

Similarly, the payoff of **P2** is defined as the value of lost flow assessed by **P2** net the cost of executing the attack:

$$u_2(x, \mu) = \underbrace{p_2 \mathrm{F}(x - x^\mu)}_{\text{value of lost flow}} - \underbrace{\mathrm{C}_2(\mu)}_{\text{cost of attack}} \tag{2.4}$$

where $p_2 \in \mathbb{R}^+$ is the marginal value of the lost flow for **P2** (in general, $p_1 \neq p_2$), and $\mathrm{C}_2(\mu) := \sum_{(i,j)\in\mathcal{E}} c_{ij}\mu_{ij}$ is the cost of the attack μ. Thus, if the disruption of an edge induces the loss of one unit of flow, the payoff of **P2** increases by p_2, and at the same time decreases by the cost of attack. In this model, we suppose that the cost of attacking an edge is proportional to its capacity. After rescaling **P2**'s payoff, we assume without loss of generality that the cost of attacking an edge is equal to its capacity.

Notice that in this model, F and C_1 are "linear forms" on $\mathcal{F}$ (i.e., $\mathcal{F}$ is not a vector space). Similarly, C_2 is a "linear form" on $\mathcal{A}$ (i.e., $\mathcal{A}$ is not a vector space).

We allow both players to randomize over their set of pure actions. Let $\Delta(\mathcal{F})$ and $\Delta(\mathcal{A})$ denote the mixed extensions of **P1**'s and **P2**'s pure strategies, respectively, i.e.:

$$\Delta(\mathcal{F}) = \left\{\sigma^1 \in [0,1]^{\mathcal{F}} \;\middle|\; \sum_{x\in\mathcal{F}} \sigma^1(x) = 1\right\}, \quad \Delta(\mathcal{A}) = \left\{\sigma^2 \in [0,1]^{\mathcal{A}} \;\middle|\; \sum_{\mu\in\mathcal{A}} \sigma^2(\mu) = 1\right\}.$$

For notational simplicity, we define $\sigma^1_x := \sigma^1(x)$ and $\sigma^2_\mu := \sigma^2(\mu)$. Given any function

$\varphi : \mathcal{F} \times \mathcal{A} \longrightarrow \mathbb{R}$ and a strategy profile $\sigma = (\sigma^1, \sigma^2) \in \Delta(\mathcal{F}) \times \Delta(\mathcal{A})$, we denote $\mathbb{E}_\sigma[\varphi(x,\mu)] := \sum_{x\in\mathcal{F}} \sigma^1_x \sum_{\mu\in\mathcal{A}} \sigma^2_\mu \varphi(x,\mu)$ the expectation of φ with respect to σ.

Given a strategy profile $\sigma = (\sigma^1, \sigma^2) \in \Delta(\mathcal{F}) \times \Delta(\mathcal{A})$, the respective player expected payoffs can be expressed as:

$$U_1(\sigma^1, \sigma^2) = p_1 \mathbb{E}_\sigma[\mathrm{F}(x^\mu)] - \mathbb{E}_\sigma[\mathrm{C}_1(x)] \tag{2.5}$$

$$U_2(\sigma^1, \sigma^2) = p_2 (\mathbb{E}_\sigma[\mathrm{F}(x)] - \mathbb{E}_\sigma[\mathrm{F}(x^\mu)]) - \mathbb{E}_\sigma[\mathrm{C}_2(\mu)]. \tag{2.6}$$

We will use the notation $U_i(x, \sigma^2) = U_i(\mathbb{1}_{\{x\}}, \sigma^2)$ and $U_i(\sigma^1, \mu) = U_i(\sigma^1, \mathbb{1}_{\{\mu\}})$ for $i \in \{1, 2\}$. The mixed extension of game Γ is given by $\langle \{1,2\}, (\Delta(\mathcal{F}), \Delta(\mathcal{A})), (U_1, U_2) \rangle$.

Let us illustrate this model through an example.

Example 1. Consider the network shown in Fig. 2-1. The edge labels give the capacities and transportation costs.

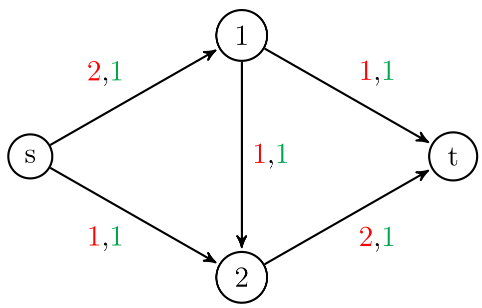

Figure 2-1: Example graph.

Both players play one shot of the game according to Fig. 2-2a.

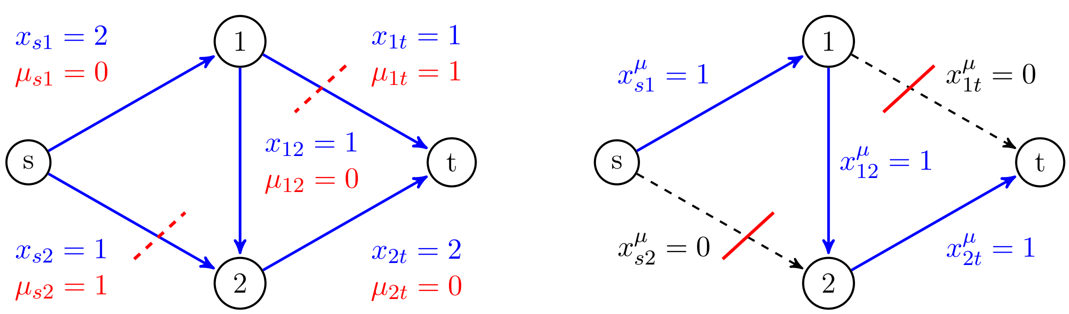

(a) Initial flow and attack

(b) Resulting effective flow

Figure 2-2: One shot of game Γ played on a given graph.

In this example, **P1** sends one unit of flow through each of the $s-t$ paths $\{s,1,t\},\{s,1,2,t\}$ and $\{s,2,t\}$, and **P2** disrupts edges $(1,t)$ and $(s,2)$. Therefore, the flows through paths $\{s,1,t\}$ and $\{s,2,t\}$ are lost and the effective flow, shown in Fig 2-2b, consists of the unit flow through the path $\{s,1,2,t\}$, i.e., $\mathrm{F}(x^{\mu})=1$. Since each edge (i,j) has a transportation cost $b_{ij}=1$, the cost of transporting the initial flow x is $\mathrm{C}_1(x)=7$. Thus, **P1**'s payoff is $u_1(x,\mu)=p_1-7$.

The amount of lost flow $\mathrm{F}(x-x^{\mu})$ that results from the attack μ is equal to 2. Since **P2** disrupted 2 edges of capacity 1 each, the cost of attack $\mathrm{C}_2(\mu)=2$. Thus, **P2**'s payoff is $u_2(x,\mu)=2p_2-2$.

2.3 Standard Definitions and Main Assumption

Let us recall the following standard definitions:

A mixed strategy profile $(\sigma^{1^*},\sigma^{2^*})\in\Delta(\mathcal{F})\times\Delta(\mathcal{A})$ is a Nash Equilibrium (NE) if and only if:

$$\forall\sigma^1\in\Delta(\mathcal{F}),\ U_1(\sigma^{1^*},\sigma^{2^*})\geq U_1(\sigma^1,\sigma^{2^*}), \tag{2.7}$$

$$\forall\sigma^2\in\Delta(\mathcal{A}),\ U_2(\sigma^{1^*},\sigma^{2^*})\geq U_2(\sigma^{1^*},\sigma^2). \tag{2.8}$$

We denote $\mathcal{S}_\Gamma$ the set of NE of the game Γ. Equivalently, at a NE $(\sigma^{1^*},\sigma^{2^*})$, σ^{1^*} (resp. σ^{2^*}) is a Best Response (BR) to σ^{2^*} (resp. σ^{1^*}). The *support* of σ^1 (resp. σ^2) is $\mathrm{supp}(\sigma^1)=\{x\in\mathcal{F}\mid\sigma^1_x>0\}$ (resp. $\mathrm{supp}(\sigma^2)=\{\mu\in\mathcal{A}\mid\sigma^2_\mu>0\}$).

A two-player game is a *strictly competitive game* (SCG) if, when both players change their mixed strategies, either the expected payoffs remain the same or one of the expected payoffs strictly increases and the other strictly decreases. In particular, a *zero-sum* game (i.e., $u_1=-u_2$) is an SCG. Adler et al. [2] define SCG using the notion of affine variance: u_1 is an affine variant of $-u_2$ if and only if $\exists\,(\lambda,\beta)\in\mathbb{R}_+^*\times\mathbb{R}\mid\forall(x,\mu)\in\mathcal{F}\times\mathcal{A},\ u_1(x,\mu)=-\lambda u_2(x,\mu)+\beta$. The game Γ is an SCG if and only if u_1 is an affine variant of $-u_2$. Besides, recall that the Minimax Theorem by

Von Neumann [21] for a zero-sum game $\widetilde{\Gamma} = \langle \{1,2\}, (\mathcal{F},\mathcal{A}), (\widetilde{u}_1, -\widetilde{u}_1)\rangle$ states that:

$$\max_{\sigma^1\in\Delta(\mathcal{F})} \min_{\sigma^2\in\Delta(\mathcal{A})} \widetilde{U}_1(\sigma^1,\sigma^2) = \min_{\sigma^2\in\Delta(\mathcal{A})} \max_{\sigma^1\in\Delta(\mathcal{F})} \widetilde{U}_1(\sigma^1,\sigma^2)$$

and:

$$(\sigma^{1^*},\sigma^{2^*}) \in \mathcal{S}_{\widetilde{\Gamma}} \iff \begin{cases} \sigma^{1^*} \in \arg\max_{\sigma^1\in\Delta(\mathcal{F})} \min_{\sigma^2\in\Delta(\mathcal{A})} \widetilde{U}_1(\sigma^1,\sigma^2), \\ \sigma^{2^*} \in \arg\min_{\sigma^2\in\Delta(\mathcal{A})} \max_{\sigma^1\in\Delta(\mathcal{F})} \widetilde{U}_1(\sigma^1,\sigma^2). \end{cases}$$

Consider two games $\Gamma = \langle \{1,2\}, (\mathcal{F},\mathcal{A}), (u_1,u_2)\rangle$, $\widetilde{\Gamma} = \langle \{1,2\}, (\mathcal{F},\mathcal{A}), (\widetilde{u}_1,\widetilde{u}_2)\rangle$ with:

$$\widetilde{u}_1(x,\mu) = a_1 u_1(x,\mu) + g(\mu)$$
$$\widetilde{u}_2(x,\mu) = a_2 u_2(x,\mu) + h(x)$$

where $(a_1,a_2) \in (\mathbb{R}_+^*)^2$, $g : \mathcal{A} \to \mathbb{R}$ and $h : \mathcal{F} \to \mathbb{R}$. Then, Γ and $\widetilde{\Gamma}$ are strategically equivalent (i.e., they have the same set of equilibria).

In this thesis, we restrict the class of graphs that we study. This enables us to develop a rather complete characterization of the equilibria of game Γ and helps us relate them to the solutions of classical routing problems. Specifically, we consider the graphs that satisfy the following assumption:

Assumption 1. *Let $\alpha := \min_{\lambda\in\Lambda_{path}} \sum_{(i,j)\in\lambda} b_{ij}$. There exists an optimal solution of $(\mathcal{P}_2)$, $x^* \in \Omega_2$, that takes $s-t$ paths with marginal transportation cost equal to α, i.e.,*

$$\exists\, x^* \in \Omega_2 \text{ s.t. } \forall \lambda \in \Lambda_{path} :\ x_\lambda > 0 \implies \sum_{(i,j)\in\lambda} b_{ij} = \alpha,$$

where Λ_{path} is the set containing all the $s-t$ paths of the network.

This assumption, noted (A1), implies that if $x^* \in \Omega_2$ denotes a min-cost max-flow, the cost of transporting a unit flow through each $s-t$ path taken by x^* is identically equal to α. By definition of α, every other path in the network cannot have a smaller marginal transportation cost. Notice that if such an x^* exists, then (A1) will be satisfied for any optimal solution of $(\mathcal{P}_2)$. The case when every $s-t$ has an identical marginal transportation cost is a special case of this assumption. In Section 5.1, we will discuss the implications of relaxing (A1). We illustrate (A1) with the following example:

Example 2. Consider the network flow problem in Fig. 2-3. There is a unique min-cost max-flow x^*, which carries 1 unit of flow through paths $\{s, 2, 4, t\}$, $\{s, 2, 3, t\}$ and $\{s, 1, t\}$. Thus, the total amount of flow is equal to 3 units. In this network, $\alpha = 3$, and each path taken by x^* has a marginal transportation cost equal to 3. Thus, the cost of transporting x^* is equal to 9. The remaining paths that are not taken by x^* are $\{s, 4, t\}$ with a transportation cost 4, and $\{s, 1, 3, t\}$ with a transportation cost 3. Thus, (A1) is satisfied.

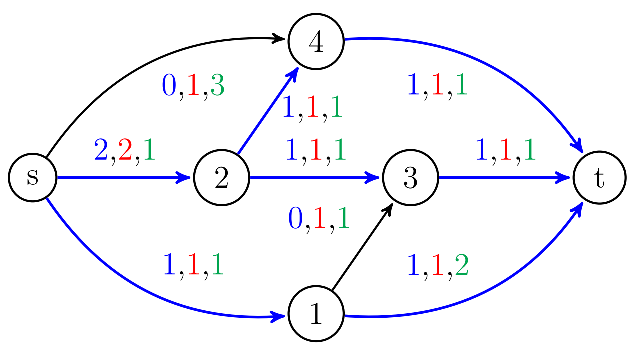

Figure 2-3: Min-cost max-flow (drawn in blue) in a graph satisfying (A1). The labels of each edge correspond to the flow it carries (blue), its capacity (red) and its transportation cost (green).

Remark 1. (A1) implies that for all $x \in \mathcal{F}$, $\mathrm{C}_1(x) \geq \alpha \, \mathrm{F}(x)$ and $\Theta_2 = \alpha\Theta_1$. To show this, we note b_λ the cost of transporting one unit of flow through path $\lambda \in \Lambda_{path}$, then we obtain:

$$\forall x \in \mathcal{F},\ \mathrm{C}_1(x) = \sum_{\lambda\in\Lambda} b_\lambda x_\lambda \geq \sum_{\lambda\in\Lambda_{path}} b_\lambda x_\lambda \overset{(\mathrm{A1})}{\geq} \alpha \sum_{\lambda\in\Lambda_{path}} x_\lambda = \alpha \, \mathrm{F}(x), \tag{2.9}$$

and

$$\Theta_2 = \mathrm{C}_1\left(x^*\right) = \sum_{\lambda\in\Lambda_{path}} b_\lambda x^*_\lambda \overset{(\mathrm{A1})}{=} \alpha \sum_{\lambda\in\Lambda_{path}} x^*_\lambda = \alpha\,\mathrm{F}\left(x^*\right) = \alpha\Theta_1, \tag{2.10}$$

where we used the fact that any min-cost max-flow does not send flow in any loop.

Note that our setup can be easily extended to networks with multiple sources and multiple destination nodes, but still satisfying (A1). For such a network, one needs to add an extra source (resp. destination) node and connect it to every existing source (resp. destination) node with an uncapacitated edge of cost of transportation equal to 0. This modification gives a new network with single source and single destination. The NE of the game defined for the original network remain the same as that of the game defined for the new network.

Now that the game is set up, the objective of this thesis is to solve Γ in closed form and characterize its NE. We want to give structural insights on the set of NE and relate these equilibria to classical network routing problems. More particularly, our objective is to utilize the equilibrium properties to relate the support of equilibrium strategies of **P1** (resp. **P2**) with the solutions of $(\mathcal{P}_2)$ (resp. min-cut sets).

Chapter 3

Characterization of Nash Equilibria

In this chapter, we present theoretical properties satisfied by the NE of the game Γ. Given p_1 and p_2, we give one NE of Γ that is based on min-cost max-flows for **P1** and on min-cut sets for **P2**. Then we focus on the main region, $p_1 > \alpha$, $p_2 > 1$, and we compute in closed form certain physical quantities of interest at any NE using a combination of game theoretic arguments and network optimization results. Specifically, we prove that each player has a unique payoff value in all NE and we characterize the value of effective (resp. lost) flow and the cost of transportation (resp. cost of attack) in terms of the parameters of Γ: p_1, p_2, the maximum amount of flow in the network Θ_1 and the smallest transportation cost of the max-flows Θ_2. Finally, we relate the mixed strategy NE of Γ to the solutions of the minimum cost maximum flow problem ($\mathcal{P}_2$) and to the minimum cut sets, and we show how we can restrict the support of the strategies that can be NE using graph theoretic properties of the network.

3.1 Preliminary Results

The following lemma states that even though Γ is not a zero-sum game, one can specify a zero-sum game, $\widetilde{\Gamma}$, that is strategically equivalent to Γ.

Lemma 1. Γ *is not a zero-sum game, but is strategically equivalent to the zero-sum* *game* $\widetilde{\Gamma} := \langle \{1,2\}, (\mathcal{F}, \mathcal{A}), (\widetilde{u}_1, -\widetilde{u}_1) \rangle$ *where:*

$$\forall (x, \mu) \in \mathcal{F} \times \mathcal{A},\ \widetilde{u}_1(x, \mu) = \mathrm{F}\left(x^{\mu}\right) - \frac{1}{p_1}\mathrm{C}_1\left(x\right) + \frac{1}{p_2}\mathrm{C}_2\left(\mu\right) \tag{3.1}$$

Therefore, the NE of Γ *can be obtained by solving the following two linear programming problems:*

$$\boxed{\begin{array}{lll} (LP_1) & \textit{maximize} & z \\ & \textit{subject to} & \widetilde{U}_1(\sigma^1, \mu) \geq z, \quad \forall \mu \in \mathcal{A} \\ & & \sigma^1 \in \Delta(\mathcal{F}) \end{array}}$$

$$\boxed{\begin{array}{lll} (LP_2) & \textit{maximize} & z' \\ & \textit{subject to} & \widetilde{U}_2(x, \sigma^2) \geq z', \quad \forall x \in \mathcal{F} \\ & & \sigma^2 \in \Delta(\mathcal{A}) \end{array}}$$

If $(\sigma^{1^*}, \sigma^{2^*}) \in \mathcal{S}_\Gamma$ *and* $(\sigma^{1^\dagger}, \sigma^{2^\dagger}) \in \mathcal{S}_\Gamma$*, then* $(\sigma^{1^*}, \sigma^{2^\dagger}) \in \mathcal{S}_\Gamma$ *and* $(\sigma^{1^\dagger}, \sigma^{2^*}) \in \mathcal{S}_\Gamma$ *(interchangeability). Furthermore,* $\mathcal{S}_\Gamma$ *is a convex set.*

Proof of Lemma 1. u_1 is not an affine variant of $-u_2$. Indeed, let us suppose the contrary:

$$\exists\, (\lambda, \beta) \in \mathbb{R}_+^* \times \mathbb{R} \mid \forall (x, \mu) \in \mathcal{F} \times \mathcal{A},\ u_1(x, \mu) = -\lambda u_2(x, \mu) + \beta$$

Then we have the following contradiction:

$$0 = u_1(x^0, \mu^0) + \lambda u_2(x^0, \mu^0) = \beta = u_1(x^*, \mu^0) + \lambda u_2(x^*, \mu^0) \neq 0.$$

Therefore, u_1 is not an affine variant of $-u_2$ and Γ is not an SCG (and a fortiori not a zero-sum game either).

However, the following transformations preserve the set of NE:

$$\frac{1}{p_1}u_1(x,\mu)+\frac{1}{p_2}\,\mathrm{C}_2\,(\mu)=\mathrm{F}\,(x^{\mu})-\frac{1}{p_1}\,\mathrm{C}_1\,(x)+\frac{1}{p_2}\,\mathrm{C}_2\,(\mu)=\widetilde{u}_1 \tag{3.2}$$

$$\frac{1}{p_2}u_2(x,\mu)-\mathrm{F}\,(x)+\frac{1}{p_1}\,\mathrm{C}_1\,(x)=-\,\mathrm{F}\,(x^{\mu})+\frac{1}{p_1}\,\mathrm{C}_1\,(x)-\frac{1}{p_2}\,\mathrm{C}_2\,(\mu)=-\widetilde{u}_1 \tag{3.3}$$

So Γ is strategically equivalent to a zero-sum game $\widetilde{\Gamma}$ and $\mathcal{S}_{\Gamma}=\mathcal{S}_{\widetilde{\Gamma}}$.

□

The following lemma states that **P1**'s all strategies containing loops are strictly dominated.

Lemma 2. *Any flow containing loops is not a BR for* ***P1****.*

Proof of Lemma 2. Suppose that **P1** chooses a flow x containing a loop l, i.e., the flow $x_l > 0$ stays in the loop and never reaches t. Therefore x_l induces a disutility because of its transportation cost, and for any attack μ, $u_1(x,\mu) < u_1(x - x_l, \mu)$, where $x - x_l$ is the flow resulting from removing the part of x that goes through l. Thus, any flow containing loops is strictly dominated. □

The intuition behind this result is that if **P1** sends flow in a loop, then she will pay an extra cost without increasing the amount of flow that can reach the terminal node. Therefore **P1** has no incentive to send flow in any loop. Then, Λ in (2.1) can be restricted to the set of $s-t$ paths and $\mathcal{F}$ can be restricted to the set of feasible flows that do not take any loop.

Props. 1–3 below provide that, for given p_1, p_2 and α, the game Γ admits qualitatively different equilibria in regions $0 < p_1 < \alpha$ and $p_2 > 0$ (Region **I**), $p_1 > \alpha$ and $0 < p_2 < 1$ (Region **II**), and $p_1 > \alpha$ and $p_2 > 1$ (Region **III**). These regions are illustrated in Fig. 3-1.

The following result states that no flow and no attack is the unique NE of Γ in Region **I**:

Proposition 1 (Region **I**)**.** *If* $p_1 < \alpha$*, then* $\mathcal{S}_{\Gamma} = \{(x^0,\mu^0)\}$*, with* $u_1(x^0,\mu^0)=0$ *and* $u_2(x^0,\mu^0)=0$.

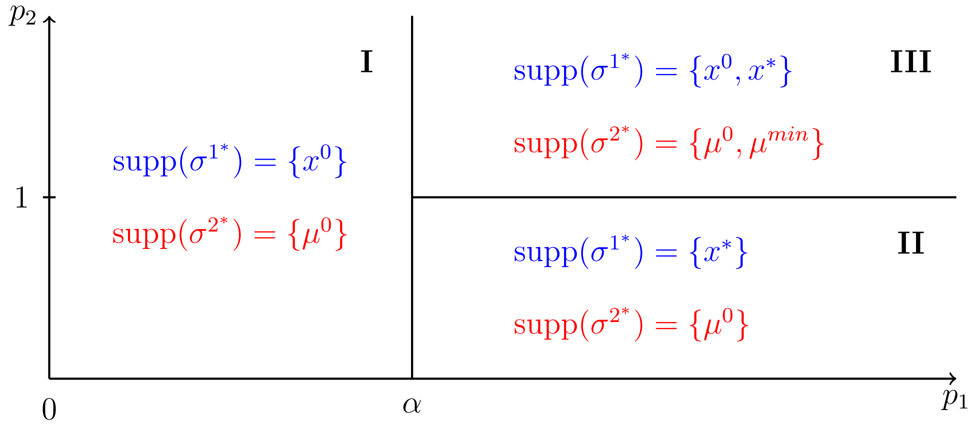

Figure 3-1: Support of equilibrium strategies in Regions **I-III**.

Proof of Proposition 1. First, note that $u_1(x^0, \mu^0) = 0$. Second,

$$\forall x \in \mathcal{F},\ u_1(x, \mu^0) = p_1 \,\mathrm{F}\left(x^{\mu^0}\right) - \mathrm{C}_1(x) = p_1 \,\mathrm{F}(x) - \mathrm{C}_1(x) \overset{(2.9)}{\leq} (p_1 - \alpha)\,\mathrm{F}(x) \leq 0.$$

So x^0 is a BR for **P1**.

Similarly, $u_2(x^0, \mu^0) = p_2 \,\mathrm{F}\left(x^0 - (x^0)^{\mu^0}\right) - \mathrm{C}_2(\mu^0) = 0$, and $\forall \mu \in \mathcal{A},\ u_2(x^0, \mu) = -\mathrm{C}_2(\mu) \leq 0$. Therefore μ^0 is a BR for **P2**. Thus, (x^0, μ^0) is a NE.

Lastly, let us argue that this NE is unique using the iterated elimination of strictly dominated strategies.

$$\begin{aligned}\forall (x, \mu) \in \mathcal{F} \times \mathcal{A},\ u_1(x, \mu) = p_1 \,\mathrm{F}(x^\mu) - \mathrm{C}_1(x) &\overset{(2.9)}{\leq} p_1 \,\mathrm{F}(x^\mu) - \alpha \,\mathrm{F}(x) \\ &\leq (p_1 - \alpha)\,\mathrm{F}(x) \leq 0.\end{aligned}$$

If $x \neq x^0$, then $\forall \mu \in \mathcal{A},\ u_1(x, \mu) < 0 = u_1(x^0, \mu)$. Therefore $x \neq x^0$ is strictly dominated and cannot be in the support of any NE.

Since $\forall \mu \in \mathcal{A},\ u_2(x^0, \mu) = -\mathrm{C}_2(\mu)$, then, $\forall \mu \in \mathcal{A} \backslash \{\mu^0\},\ u_2(x^0, \mu) = -\mathrm{C}_2(\mu) < 0 = u_2(x^0, \mu^0)$. Hence, $\mu \neq \mu^0$ is now strictly dominated and cannot be in the support of any NE. Thus, (x^0, μ^0) is the unique NE when $p_1 < \alpha$. $\square$

Intuitively, when $0 < p_1 < \alpha$, the marginal value of effective flow that reaches the destination node t is less than the marginal transportation cost for every s-t path.

Therefore, **P1** will face negative utility if she sends flow through the network. Thus, in this case, her BR is not to route any flow. Since no flow is sent by **P1**, **P2**'s BR is not to attack, otherwise she faces the cost of attack without gaining any value from lost flow.

Next, for Region **II**, we obtain that min-cost max-flow and no attack is a pure NE.

Proposition 2 (Region **II**). *If $p_1 > \alpha$ and $p_2 < 1$, then $\forall x^* \in \Omega_2,\ \{x^*, \mu^0\} \in \mathcal{S}_\Gamma$. The equilibrium payoffs are $u_1(x^*, \mu^0) = (p_1 - \alpha)\Theta_1$ and $u_2(x^*, \mu^0) = 0$.*

Proof of Proposition 2. First, $u_1(x^*, \mu^0) = p_1 \,\mathrm{F}\left((x^*)^{\mu^0}\right) - \mathrm{C}_1(x^*) \overset{(2.10)}{=} (p_1 - \alpha)\Theta_1$.

Second, $\forall x \in \mathcal{F},\ u_1(x, \mu^0) = p_1 \,\mathrm{F}\left(x^{\mu^0}\right) - \mathrm{C}_1(x) \overset{(2.9)}{\leq} (p_1 - \alpha)\,\mathrm{F}(x) \leq (p_1 - \alpha)\Theta_1$. So x^* is a BR for **P1**.

Similarly, $u_2(x^*, \mu^0) = p_2 \,\mathrm{F}\left(x^* - (x^*)^{\mu^0}\right) - \mathrm{C}_2(\mu^0) = 0$ because $(x^*)^{\mu^0} = x^*$. Besides,

$$\forall \mu \in \mathcal{A},\ u_2(x^*, \mu) = p_2 \,\mathrm{F}\left(x^* - (x^*)^{\mu}\right) - \mathrm{C}_2(\mu) \leq \mathrm{F}\left(x^* - (x^*)^{\mu}\right) - \mathrm{C}_2(\mu)$$

since $p_2 \leq 1$. $\mathrm{F}\left(x^* - (x^*)^{\mu}\right)$ is the loss induced by the attack μ when x^* is in the network, and $\mathrm{C}_2(\mu)$ is the cost of the attack μ which can also be viewed as the maximum amount of flow that can be lost because of the attack. Thus: $\forall \mu \in \mathcal{A},\ u_2(x^*, \mu) \leq 0$ and μ^0 is a BR for **P2**.

Therefore $\forall x^* \in \Omega_2,\ (x^*, \mu^0)$ is a NE. □

This result can be explained as follows: on one hand, since **P2**'s valuation of lost flow is small ($p_2 < 1$), for any attack, the utility gained from the lost flow is always lower than the cost of attack. Therefore, **P2**'s BR is not to attack any edge. On the other hand, **P1**'s valuation of effective flow reaching t is higher than the disutility it faces in transportation costs ($p_1 > \alpha$). Since **P2** does not disrupt any edge, every flow sent through the network reaches t; thus, **P1**'s BR is to send a maximum flow. Among the different maximum flows, a min-cost max-flow maximizes **P1**'s equilibrium payoff. Note that if $p_1 = \alpha$ and $p_2 < 1$, then both (x^0, μ^0) and (x^*, μ^0) are NE. The equilibrium payoffs are still $(0, 0)$.

The following Proposition 3 focuses on Region **III**. It shows that given $p_1 > \alpha$, $p_2 > 1$, and a graph $\mathcal{G}$ satisfying (A1), Γ admits a NE whose support is based on a min-cost max-flow for **P1**, and on a min-cut set for **P2**.

Proposition 3 (Region **III**). *If $p_1 > \alpha$ and $p_2 > 1$, then Γ has no pure NE. Furthermore, $\exists\, \tilde{\sigma} = (\tilde{\sigma}^1, \tilde{\sigma}^2) \in \mathcal{S}_\Gamma$ such that $U_1(\tilde{\sigma}^1, \tilde{\sigma}^2) = U_2(\tilde{\sigma}^1, \tilde{\sigma}^2) = 0$, and* $\mathrm{supp}(\tilde{\sigma}^1) = \{x^0, x^*\}$ *and* $\mathrm{supp}(\tilde{\sigma}^2) = \{\mu^0, \mu^{min}\}$. *The corresponding probabilities are given by:*

$$\tilde{\sigma}^1_{x^0} = 1 - \frac{1}{p_2}, \quad \tilde{\sigma}^1_{x^*} = \frac{1}{p_2}, \tag{3.4}$$

$$\tilde{\sigma}^2_{\mu^0} = \frac{\alpha}{p_1}, \quad \tilde{\sigma}^2_{\mu^{min}} = 1 - \frac{\alpha}{p_1}. \tag{3.5}$$

Proof of Proposition 3.

– First, let us show that any pure strategy is not a NE.

Let us suppose that **P2** chooses an attack μ that disrupts m edges $e_1, \ldots, e_m$ of $\mathcal{G}$ (m can be equal to 0). Now, let us suppose that **P1** chooses a flow x that crosses one of the attacked edges, for instance e_1. If we note x_p the part of x that goes through e_1, then x_p will be lost because of the attack so the value of effective flow obtained by routing x is the same as if $x - x_p$ had been routed. However, the transportation cost of x is strictly greater than the transportation cost of $x - x_p$. Therefore $u_1(x, \mu) < u_1(x - x_p, \mu)$.

Thus, a BR for **P1** does not take paths containing at least one attacked edge. Now, let us suppose that **P1** chooses only such flows, the utility becomes $u_1(x, \mu) = p_1 \,\mathrm{F}\,(x) - \mathrm{C}_1\,(x)$. Let us note $\mathcal{G}^\mu = (\mathcal{V}, \mathcal{E}\backslash\{e_1, \ldots, e_m\})$. Due to the latest comment, **P1**'s BR is a feasible flow in $\mathcal{G}^\mu$. Now there are two cases:

Case 1: there is no path in $\mathcal{G}^\mu$ with marginal transportation cost less than p_1. Then, **P1**'s BR is x^0 (no flow). However, if **P1** chooses x^0, then it's easy to see that **P2**'s BR is μ^0 (no attack), which means that the initial μ is not a BR for

P2 in this case. Remark that if the initial μ is μ^0, then Case 1 is not satisfied (there are paths in $\mathcal{G}$ of marginal transportation cost less than p_1, because of the definition of α).

Case 2: there exists at least one path in $\mathcal{G}^\mu$ with marginal transportation cost less than p_1. Then, **P1**'s BR is to send as much flow as she can along the paths with maginal transportation cost less than p_1 (if there are different such flows with the same value, then **P1**'s BR is the one with least transportation cost). However, if **P1** chooses this BR, then the initial attack μ does not induce any loss, and **P2** will have an incentive to disrupt some edges of $\mathcal{G}^\mu$ instead (for instance we can prove that at least one edge is saturated by **P1**'s BR so **P2** will gain utility by attacking that edge since $p_2 > 1$). Thus, **P2**'s BR is different from μ.

Therefore, every pure strategy is not a NE.

– Now let us prove that $\tilde{\sigma}$ is a NE.

$$\begin{aligned}\forall \sigma^1 \in \Delta(\mathcal{F}),\ U_1(\sigma^1, \tilde{\sigma}^2) &\overset{(2.5)}{=} p_1 \frac{\alpha}{p_1} \mathbb{E}_\sigma \left[\mathrm{F}\left(x\right)\right] - \mathbb{E}_\sigma \left[\mathrm{C}_1\left(x\right)\right] \\ &\overset{(2.9)}{\leq} \alpha \mathbb{E}_\sigma \left[\mathrm{F}\left(x\right)\right] - \alpha \mathbb{E}_\sigma \left[\mathrm{F}\left(x\right)\right] = 0\end{aligned}$$

Besides: $U_1(\tilde{\sigma}^1, \tilde{\sigma}^2) = \alpha \frac{1}{p_2} \mathrm{F}\left(x^*\right) - \frac{1}{p_2} \mathrm{C}_1\left(x^*\right) \overset{(2.10)}{=} \frac{\alpha}{p_2} \mathrm{F}\left(x^*\right) - \frac{\alpha}{p_2} \mathrm{F}\left(x^*\right) = 0$

Similarly:

$$\begin{aligned}\forall \sigma^2 \in \Delta(\mathcal{A}),\ U_2(\tilde{\sigma}^1, \sigma^2) &\overset{(2.6)}{=} \mathrm{F}\left(x^*\right) - \mathbb{E}_\sigma \left[\mathrm{F}\left((x^*)^\mu\right)\right] - \mathbb{E}_\sigma \left[\mathrm{C}_2\left(\mu\right)\right] \\ &= \mathbb{E}_\sigma \left[\mathrm{F}\left(x^* - (x^*)^\mu\right) - \mathrm{C}_2\left(\mu\right)\right] \leq 0\end{aligned}$$

where the inequality follows from the fact that for any attack μ, $\mathrm{F}\left(x^* - (x^*)^\mu\right)$ is the loss induced by μ when x^* is in the network, and $\mathrm{C}_2\left(\mu\right)$ is the cost of μ which can also be viewed as the maximum amount of flow that can be lost because of the attack.

Besides: $U_2(\tilde{\sigma}^1, \tilde{\sigma}^2) = \mathrm{F}(x^*) - \frac{\alpha}{p_1}\mathrm{F}(x^*) - \left(1 - \frac{\alpha}{p_1}\right)\mathrm{C}_2(\mu^{min}) = 0$ thanks to the Max-Flow Min-Cut Theorem. Thus, $(\tilde{\sigma}^1, \tilde{\sigma}^2)$ is a NE.

□

We can make a few useful observations from this result. First, in contrast to Props. 1 and 2, in Region **III**, both players randomize their actions in any equilibrium. Indeed, if **P1**'s pure strategy is to route some flow x in the network, then **P2**'s BR is to disrupt some edges taken by x in order to induce the maximum loss. Then **P1** has an incentive to change her strategy and route another flow that takes other paths not disrupted by **P2**. Therefore, for every pure action profile, at least one of the players has an incentive to deviate, preventing them from reaching an equilibrium.

Second, the mixed equilibrium $(\tilde{\sigma}^1, \tilde{\sigma}^2)$, as defined by (3.4) and (3.5), can be obtained from a solution of problem $(\mathcal{P}_2)$ and a min-cut set of the graph $\mathcal{G}$. Note that $(\tilde{\sigma}^1, \tilde{\sigma}^2)$ has a particularly simple structure, i.e., **P1** either sends a whole min-cost max-flow, or does not send any flow in the network. Similarly, **P2** either disrupts all the edges of a min-cut set, or does not attack any edge of the network.

Finally, Prop. 3 provides a game-theoretic intuition: **P1**'s equilibrium strategy $\tilde{\sigma}^1$ is characterized by p_2, and similarly, **P2**'s equilibrium strategy $\tilde{\sigma}^2$ is characterized by p_1 and α (given and fixed under (A1)). This can be explained as follows: as p_2 increases, $\tilde{\sigma}^1_{x^*}$ decreases while $\tilde{\sigma}^1_{x^0}$ increases. When **P2**'s valuation of lost flow is large, she has more incentive to attack, so any flow sent by **P1** will be more likely to be lost. Thus, **P1** chooses not to send any flow with higher probability than sending x^*. Likewise, as p_1 increases, $\tilde{\sigma}^2_{\mu^{min}}$ increases while $\tilde{\sigma}^2_{\mu^0}$ decreases. Again, when the marginal valuation of effective flow is large, **P1** will prefer to send as much flow as she can. Thus, **P2** will be more likely to attack a min-cut set.

The following example applies the results of Props. 1-3:

Example 3. Consider the graph in Fig. 3-2. We can see that $\alpha = 3$, and that the min-cost max-flow sends 1 unit of flow through $\{s, 1, 3, t\}$, $\{s, 2, 3, t\}$ and $\{s, 2, 4, t\}$, and only takes paths with transportation cost equal to 3. Thus, (A1) is satisfied.

The min-cut set is given by $\{(1,3),(2,3),(2,4)\}$. The NE described in Props. 1-3 are illustrated in Fig. 3-3.

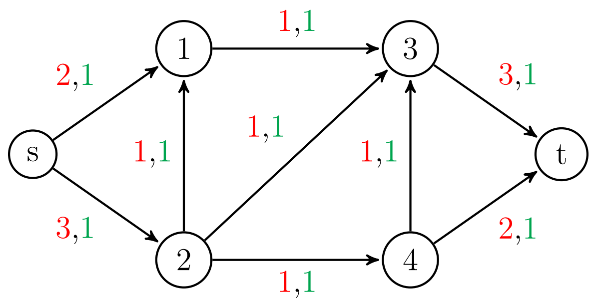

Figure 3-2: Example network. Edge capacities and transportation costs are labeled in red and green colors respectively.

Thanks to Props. 1-3, **P1** (or **P2**) has a strategy which achieves a stable outcome for any p_1 and p_2, and any graph satisfying (A1). However, **P1** may be interested in an equilibrium strategy that maximizes the expected amount of effective flow that successfully crosses the network after attack. Similarly, the interdictor may be interested in an equilibrium strategy that maximizes the expected amount of flow that is lost. To solve this problem, we need to further analyze the set of NE. Although Regions **I** and **II** do not require further study, Region **III** hosts many nice properties that will be used in our subsequent analysis of equilibria in budget-constrained environments. We now present the main results focusing on NE in Region **III**.

3.2 Main Theorem

We now present our main result focusing on NE in Region **III**.

Theorem 1. *If $p_1 > \alpha$, $p_2 > 1$, and under (A1), then for any $\sigma^* \in \mathcal{S}_\Gamma$:*

(i) Both players' equilibrium payoffs are equal to 0, i.e.:

$$U_1(\sigma^{1^*}, \sigma^{2^*}) \equiv 0 \tag{3.6}$$

$$U_2(\sigma^{1^*}, \sigma^{2^*}) \equiv 0 \tag{3.7}$$

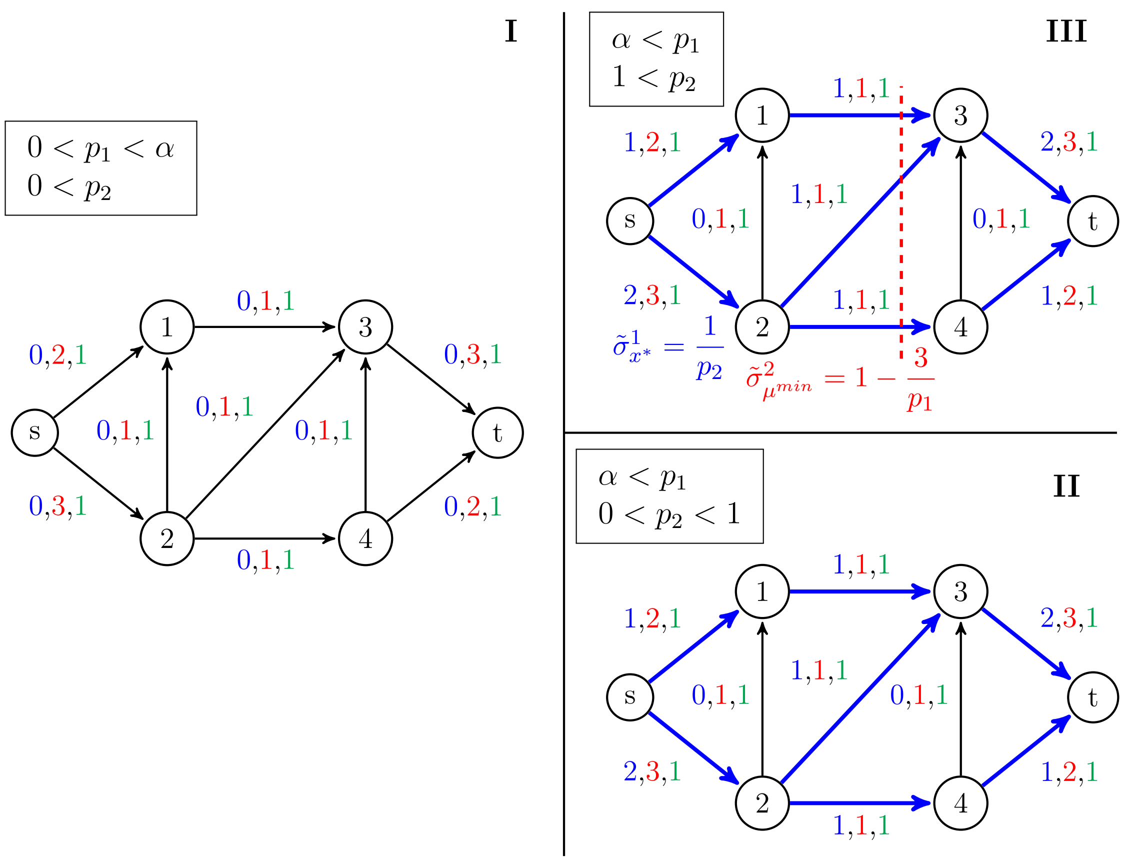

Figure 3-3: NE described in Props. 1, 2 and 3. The min-cost max-flow (resp. min-cut set attack) is in bold blue (resp. dotted red).

(ii) The expected amount of flow sent in the network is given by:

$$\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] \equiv \frac{1}{p_2}\Theta_1 \tag{3.8}$$

and the expected transportation cost is given by:

$$\mathbb{E}_{\sigma^*}\left[\mathrm{C}_1\left(x\right)\right] \equiv \frac{1}{p_2}\Theta_2 \tag{3.9}$$

(iii) The expected cost of attack is given by:

$$\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] \equiv \Theta_1 - \frac{1}{p_1}\Theta_2 \tag{3.10}$$

(iv) The expected amount of effective flow (that reaches t) is given by:

$$\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x^{\mu}\right)\right] \equiv \frac{1}{p_1 p_2}\Theta_2. \tag{3.11}$$

We derive two proofs of Thm. 1, by combining the Max-Flow Min-Cut Theorem with best response inequalities (in the first proof), and with linear programming duality (in the second proof). In order to prove Thm. 1, we first need the following lemma:

Lemma 3.

$$\forall(\sigma^{1*}, \sigma^{2*}) \in \mathcal{S}_{\Gamma},\ \mathbb{E}_{\sigma^*}\left[\mathrm{F}\left((x^*)^{\mu}\right)\right] = \mathrm{F}\left(x^*\right) - \mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right]. \tag{3.12}$$

Proof of Lemma 3. We can find a link between both players' expected payoffs that we can write in two different ways:

$$U_1(\sigma^1, \sigma^2) = p_1\mathbb{E}_{\sigma}\left[\mathrm{F}\left(x\right)\right] - \mathbb{E}_{\sigma}\left[\mathrm{C}_1\left(x\right)\right] - \frac{p_1}{p_2}\mathbb{E}_{\sigma}\left[\mathrm{C}_2\left(\mu\right)\right] - \frac{p_1}{p_2}U_2(\sigma^1, \sigma^2) \tag{3.13}$$

$$U_2(\sigma^1, \sigma^2) = -\mathbb{E}_{\sigma}\left[\mathrm{C}_2\left(\mu\right)\right] + p_2\mathbb{E}_{\sigma}\left[\mathrm{F}\left(x\right)\right] - \frac{p_2}{p_1}\mathbb{E}_{\sigma}\left[\mathrm{C}_1\left(x\right)\right] - \frac{p_2}{p_1}U_1(\sigma^1, \sigma^2) \tag{3.14}$$

Let $\sigma^* = (\sigma^{1*}, \sigma^{2*}) \in \mathcal{S}_\Gamma$. Since $(\tilde{\sigma}^1, \tilde{\sigma}^2) \in \mathcal{S}_\Gamma$ (Prop. 3), we have:

$$0 = U_2(\tilde{\sigma}^1, \tilde{\sigma}^2) \overset{(2.8)}{\geq} U_2(\tilde{\sigma}^1, \sigma^{2*}) \overset{(2.6)}{=} \mathrm{F}\left(x^*\right) - \mathbb{E}_{\sigma^*}\left[\mathrm{F}\left((x^*)^\mu\right)\right] - \mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right]$$

So we get the first inequality:

$$\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left((x^*)^\mu\right)\right] \geq \mathrm{F}\left(x^*\right) - \mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] \tag{3.15}$$

Now, since $(\sigma^{1*}, \sigma^{2*}) \in \mathcal{S}_\Gamma$:

$$\begin{aligned} U_1(\sigma^{1*}, \sigma^{2*}) \overset{(2.7)}{\geq} U_1(\tilde{\sigma}^1, \sigma^{2*}) &\overset{(2.5)}{=} \frac{p_1}{p_2}\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left((x^*)^\mu\right)\right] - \frac{1}{p_2}\mathrm{C}_1\left(x^*\right) \\ &\overset{(2.10)}{=} \frac{p_1}{p_2}\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left((x^*)^\mu\right)\right] - \frac{\alpha}{p_2}\mathrm{F}\left(x^*\right) \end{aligned} \tag{3.16}$$

By combining (3.13) and (2.8), using $\tilde{\sigma}^2$, we get:

$$\begin{aligned} U_1(\sigma^{1*}, \sigma^{2*}) &\overset{(3.13)}{=} p_1\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] - \mathbb{E}_{\sigma^*}\left[\mathrm{C}_1\left(x\right)\right] - \frac{p_1}{p_2}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] - \frac{p_1}{p_2}U_2(\sigma^{1*}, \sigma^{2*}) \\ &\overset{(2.9)}{\leq} (p_1 - \alpha)\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] - \frac{p_1}{p_2}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] - \frac{p_1}{p_2}U_2(\sigma^{1*}, \sigma^{2*}) \\ &\overset{(2.8)}{\leq} (p_1 - \alpha)\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] - \frac{p_1}{p_2}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] - \frac{p_1}{p_2}U_2(\sigma^{1*}, \tilde{\sigma}^2) \\ &\overset{(2.6)}{=} (p_1 - \alpha)\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] - \frac{p_1}{p_2}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] - p_1\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] \\ &\quad + \alpha\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] + \left(1 - \frac{\alpha}{p_1}\right)\frac{p_1}{p_2}\mathrm{C}_2\left(\mu^{min}\right) \end{aligned}$$

Therefore:

$$U_1(\sigma^{1*}, \sigma^{2*}) \leq \frac{p_1}{p_2}\left(\mathrm{C}_2\left(\mu^{min}\right) - \mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right]\right) - \frac{\alpha}{p_2}\mathrm{C}_2\left(\mu^{min}\right) \tag{3.17}$$

By combining (3.16) and (3.17), and using the Max-Flow Min-Cut Theorem, we obtain:

$$\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left((x^*)^\mu\right)\right] \leq \mathrm{F}\left(x^*\right) - \mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] \tag{3.18}$$

Equations (3.15) and (3.18) lead to:

$$\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left((x^*)^\mu\right)\right] = \mathrm{F}\left(x^*\right) - \mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right]$$

thus proving Lemma 3. □

Now we can prove Thm. 1:

First Proof of Theorem 1. Let $\sigma^* = (\sigma^{1^*}, \sigma^{2^*}) \in \mathcal{S}_\Gamma$. Let us start by showing (iii), but we will need a few intermediate equations before.

First, let us prove that $U_1(\sigma^{1^*}, \sigma^{2^*}) \geq 0$:

$$U_1(\sigma^{1^*}, \sigma^{2^*}) \overset{(2.7)}{\geq} U_1(x^0, \sigma^{2^*}) = 0 \tag{3.19}$$

By combining (3.19) and (3.17), we obtain:

$$\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] \leq \left(1 - \frac{\alpha}{p_1}\right) \mathrm{C}_2\left(\mu^{min}\right) \tag{3.20}$$

In order to get the reverse inequality, let us consider the strategy σ^1_ϵ defined by $\sigma^1_{x^*} = \dfrac{1+\epsilon}{p_2}$ and $\sigma^1_{x^0} = 1 - \dfrac{1+\epsilon}{p_2}$ for an ϵ small enough (we can find such an ϵ and still have a probability distribution since $p_2 > 1$):

$$\begin{aligned} U_1(\sigma^{1^*}, \sigma^{2^*}) \overset{(2.7)}{\geq} U_1(\sigma^1_\epsilon, \sigma^{2^*}) &\overset{(2.5)}{=} \frac{p_1(1+\epsilon)}{p_2}\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left((x^*)^\mu\right)\right] - \frac{1+\epsilon}{p_2}\mathrm{C}_1\left(x^*\right) \\ &\overset{(2.10)}{=} \frac{p_1(1+\epsilon)}{p_2}\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left((x^*)^\mu\right)\right] - \frac{\alpha(1+\epsilon)}{p_2}\mathrm{F}\left(x^*\right) \end{aligned}$$

Equation (3.12) gives us:

$$U_1(\sigma^{1^*}, \sigma^{2^*}) \geq \frac{p_1(1+\epsilon)}{p_2}\left(\mathrm{F}\left(x^*\right) - \mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right]\right) - \frac{\alpha(1+\epsilon)}{p_2}\mathrm{F}\left(x^*\right) \tag{3.21}$$

We just have to combine (3.17) and (3.21) in order to get:

$$\frac{p_1\epsilon}{p_2}\mathrm{F}\left(x^*\right) - \frac{p_1\epsilon}{p_2}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] - \frac{\alpha\epsilon}{p_2}\mathrm{F}\left(x^*\right) \leq 0$$

which is equivalent to:

$$\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] \geq \left(1-\frac{\alpha}{p_1}\right)\mathrm{F}\left(x^*\right) \tag{3.22}$$

Equations (3.20), (3.22) and the Max-Flow Min-Cut Theorem give us (iii):

$$\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] = \left(1-\frac{\alpha}{p_1}\right)\mathrm{C}_2\left(\mu^{min}\right) = \left(1-\frac{\alpha}{p_1}\right)\Theta_1 \overset{(2.10)}{=} \Theta_1 - \frac{1}{p_1}\Theta_2 \tag{3.23}$$

We can now use this equation in order to prove that **P1**'s payoff is equal to 0 at equilibrium: by combining (3.23) and (3.17), we obtain:

$$U_1(\sigma^{1^*}, \sigma^{2^*}) \leq 0 \tag{3.24}$$

Therefore, (3.24) and (3.19) give:

$$U_1(\sigma^{1^*}, \sigma^{2^*}) = 0 \tag{3.25}$$

Let us now prove (3.8) from (ii). Similarly, let us first prove that $U_2(\sigma^{1^*}, \sigma^{2^*}) \geq 0$:

$$U_2(\sigma^{1^*}, \sigma^{2^*}) \overset{(2.8)}{\geq} U_2(\sigma^{1^*}, \mu^0) = 0 \tag{3.26}$$

We can use previous results:

$$\begin{aligned} U_2(\sigma^{1^*}, \sigma^{2^*}) &\overset{(3.14)}{=} -\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] + p_2\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] - \frac{p_2}{p_1}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_1\left(x\right)\right] - \frac{p_2}{p_1}U_1(\sigma^{1^*}, \sigma^{2^*}) \\ &\overset{(3.25)}{=} -\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] + p_2\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] - \frac{p_2}{p_1}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_1\left(x\right)\right] \\ &\overset{(2.9)}{\leq} -\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] + p_2\left(1-\frac{\alpha}{p_1}\right)\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] \\ &\overset{(3.23)}{=} \left(1-\frac{\alpha}{p_1}\right)\left(p_2\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] - \mathrm{C}_2\left(\mu^{min}\right)\right) \end{aligned} \tag{3.27}$$

By combining (3.27) and (3.26), we obtain:

$$\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] \geq \frac{1}{p_2}\,\mathrm{C}_2\left(\mu^{min}\right) \tag{3.28}$$

In order to get the reverse inequality, let us consider the strategy σ^2_ϵ defined by $\sigma^2_{\mu^0} = \dfrac{\alpha-\epsilon}{p_1}$ and $\sigma^2_{\mu^{min}} = 1 - \dfrac{\alpha-\epsilon}{p_1}$, for an ϵ small enough (we can find such an ϵ and still have a probability distribution since $p_1 > \alpha$):

$$\begin{aligned}
U_1(\sigma^{1^*},\sigma^{2^*}) &\overset{(3.13)}{=} p_1\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] - \mathbb{E}_{\sigma^*}\left[\mathrm{C}_1\left(x\right)\right] - \frac{p_1}{p_2}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] - \frac{p_1}{p_2}U_2(\sigma^{1^*},\sigma^{2^*})\\
&\overset{(2.9)}{\leq} (p_1-\alpha)\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] - \frac{p_1}{p_2}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] - \frac{p_1}{p_2}U_2(\sigma^{1^*},\sigma^{2^*})\\
&\overset{(2.8)}{\leq} (p_1-\alpha)\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] - \frac{p_1}{p_2}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] - \frac{p_1}{p_2}U_2(\sigma^{1^*},\sigma^2_\epsilon)\\
&\overset{(2.6)}{=} (p_1-\alpha)\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] - \frac{p_1}{p_2}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] - p_1\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right]\\
&\qquad + (\alpha-\epsilon)\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] + \left(1-\frac{\alpha-\epsilon}{p_1}\right)\frac{p_1}{p_2}\,\mathrm{C}_2\left(\mu^{min}\right)\\
&= \frac{p_1}{p_2}\left(\mathrm{C}_2\left(\mu^{min}\right) - \mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right]\right) - \frac{\alpha-\epsilon}{p_2}\,\mathrm{C}_2\left(\mu^{min}\right) - \epsilon\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right]\\
&\overset{(3.23)}{=} \left(\frac{p_1}{p_2} - \frac{p_1}{p_2}\left(1-\frac{\alpha}{p_1}\right) - \frac{\alpha-\epsilon}{p_2}\right)\mathrm{C}_2\left(\mu^{min}\right) - \epsilon\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right]\\
&= \frac{\epsilon}{p_2}\,\mathrm{C}_2\left(\mu^{min}\right) - \epsilon\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right]
\end{aligned} \tag{3.29}$$

Equations (3.29) and (3.19) give us:

$$0 \leq \frac{\epsilon}{p_2}\,\mathrm{C}_2\left(\mu^{min}\right) - \epsilon\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right]$$

Thus:

$$\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] \leq \frac{1}{p_2}\,\mathrm{C}_2\left(\mu^{min}\right) \tag{3.30}$$

Equations (3.28), (3.30) and the Max-Flow Min-Cut Theorem give us:

$$\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] = \frac{1}{p_2}\,\mathrm{F}\left(x^*\right) = \frac{1}{p_2}\Theta_1 \tag{3.31}$$

Likewise, by combining (3.31) and (3.27), we get:

$$U_2(\sigma^{1^*}, \sigma^{2^*}) \leq 0 \tag{3.32}$$

Equations (3.32) and (3.26) give us:

$$U_2(\sigma^{1^*}, \sigma^{1^*}) = 0 \tag{3.33}$$

thus proving (i).

Now, by combining (2.6), (3.31), (3.23) and (3.33), we can prove (iv):

$$\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x^{\mu}\right)\right] = \frac{1}{p_2}\Theta_1 - \frac{1}{p_2}\left(1 - \frac{\alpha}{p_1}\right)\Theta_1 = \frac{\alpha}{p_1 p_2}\Theta_1 \overset{(2.10)}{=} \frac{1}{p_1 p_2}\Theta_2 \tag{3.34}$$

Lastly, by combining (2.5), (3.34) and (3.25), we can finish proving (ii):

$$\mathbb{E}_{\sigma^*}\left[\mathrm{C}_1\left(x\right)\right] = p_1 \frac{\alpha}{p_1 p_2}\Theta_1 = \frac{\alpha}{p_2}\Theta_1 \overset{(2.10)}{=} \frac{1}{p_2}\Theta_2 \tag{3.35}$$

thus ending the first proof of Thm. 1. □

Second Proof of Theorem 1. Let us now derive a second proof using the strategically equivalent zero-sum game $\widetilde{\Gamma}$ (recall that $\mathcal{S}_{\Gamma} = \mathcal{S}_{\widetilde{\Gamma}}$).

We know that in a zero-sum game, each player's payoff is constant for any NE. Prop. 3 tells us that $\tilde{\sigma} \in \mathcal{S}_{\widetilde{\Gamma}}$. Therefore:

$$\begin{aligned}
\forall(\sigma^{1^*}, \sigma^{2^*}) \in \mathcal{S}_{\widetilde{\Gamma}}, \widetilde{U}_1(\sigma^{1^*}, \sigma^{2^*}) &= \widetilde{U}_1(\tilde{\sigma}^1, \tilde{\sigma}^2) \\
&\overset{(3.1)}{=} \frac{1}{p_2}\frac{\alpha}{p_1}\,\mathrm{F}\left(x^*\right) - \frac{1}{p_1}\frac{1}{p_2}\,\mathrm{C}_1\left(x^{max}\right) + \frac{1}{p_2}\left(1 - \frac{\alpha}{p_1}\right)\mathrm{C}_2\left(\mu^{min}\right) \\
&= \frac{1}{p_2}\left(1 - \frac{\alpha}{p_1}\right)\Theta_1
\end{aligned} \tag{3.36}$$

and this quantity is the optimal value of (LP_1).

Let $\sigma^* = (\sigma^{1^*}, \sigma^{2^*}) \in \mathcal{S}_{\Gamma}$. First, let us prove again (3.12) from Lemma 3:

By interchangeability, since $\tilde{\sigma} \in \mathcal{S}_{\widetilde{\Gamma}}$, then:

$$\begin{aligned}\frac{1}{p_2}\left(1-\frac{\alpha}{p_1}\right)\Theta_1 &\overset{(3.36)}{=} \widetilde{U}_1(\tilde{\sigma}^1, \tilde{\sigma}^2) = \widetilde{U}_1(\tilde{\sigma}^1, \sigma^{2*}) \\ &= \frac{1}{p_2}\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left((x^*)^{\mu}\right)\right] - \frac{\alpha}{p_1 p_2}\mathrm{F}\left(x^*\right) + \frac{1}{p_2}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right]\end{aligned}$$

which directly gives the result.

Now, we prove the equalities thanks to complementary slackness: recall that (LP$_2$) is the dual of (LP$_1$) (and vice-versa), which means that $\forall \mu \in \mathcal{A},\ \sigma^2_\mu$ is the dual variable associated with the constraint $\widetilde{U}_1(\sigma^1, \mu) \geq z$. Similarly, $\forall x \in \mathcal{F},\ \sigma^1_x$ is the dual variable associated with the constraint $\widetilde{U}_2(x, \sigma^2) \geq z'$. We know that the optimal value of (LP$_1$) (resp. (LP$_2$)) is $\frac{1}{p_2}\left(1-\frac{\alpha}{p_1}\right)\Theta_1$ (resp. $-\frac{1}{p_2}\left(1-\frac{\alpha}{p_1}\right)\Theta_1$). Therefore, since NE are the optimal solutions of these LPs, complementary slackness gives us, $\forall (\sigma^{1*}, \sigma^{2*}) \in \mathcal{S}_{\widetilde{\Gamma}}$:

$$\forall x \in \mathcal{F},\ \sigma^{1*}_x\left(\widetilde{U}_2(x, \sigma^{2*}) + \frac{1}{p_2}\left(1-\frac{\alpha}{p_1}\right)\Theta_1\right) = 0 \tag{3.37}$$

$$\forall \mu \in \mathcal{A},\ \sigma^{2*}_\mu\left(\widetilde{U}_1(\sigma^{1*}, \mu) - \frac{1}{p_2}\left(1-\frac{\alpha}{p_1}\right)\Theta_1\right) = 0 \tag{3.38}$$

Prop. 3 tells us that $\mu^{min} \in \text{supp}(\tilde{\sigma}^2)$ (or equivalently $\tilde{\sigma}^2_{\mu^{min}} > 0$), therefore, complementary slackness gives us:

$$\begin{aligned}\frac{1}{p_2}\left(1-\frac{\alpha}{p_1}\right)\Theta_1 &\overset{(3.38)}{=} \widetilde{U}_1(\sigma^{1*}, \mu^{min}) \overset{(3.1)}{=} -\frac{1}{p_1}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_1\left(x\right)\right] + \frac{1}{p_2}\Theta_1 \\ \Longleftrightarrow \mathbb{E}_{\sigma^*}\left[\mathrm{C}_1\left(x\right)\right] &= \frac{\alpha}{p_2}\Theta_1 \overset{(2.10)}{=} \frac{1}{p_2}\Theta_2\end{aligned} \tag{3.39}$$

Likewise, $\mu^0 \in \text{supp}(\tilde{\sigma}^2)$ therefore we can prove (ii):

$$\begin{aligned}\frac{1}{p_2}\left(1-\frac{\alpha}{p_1}\right)\Theta_1 &\overset{(3.38)}{=} \widetilde{U}_1(\sigma^{1*}, \mu^0) \overset{(3.1)}{=} \mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] - \frac{1}{p_1}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_1\left(x\right)\right] \\ \Longleftrightarrow \mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] &\overset{(3.39)}{=} \frac{1}{p_2}\left(1-\frac{\alpha}{p_1}\right)\Theta_1 + \frac{\alpha}{p_1 p_2}\Theta_1 = \frac{1}{p_2}\Theta_1\end{aligned} \tag{3.40}$$

Similarly, $x^* \in \text{supp}(\tilde{\sigma}^1)$ (or equivalently $\tilde{\sigma}^1_{x^*} > 0$), therefore, complementary slackness gives us:

$$\begin{aligned}
-\frac{1}{p_2}\left(1-\frac{\alpha}{p_1}\right)\Theta_1 &\overset{(3.37)}{=} \widetilde{U}_2(x^*,\sigma^{2*}) = -\widetilde{U}_1(x^*,\sigma^{2*}) \\
&\overset{(3.1)}{=} -\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left((x^*)^{\mu}\right)\right] + \frac{1}{p_1}\,\mathrm{C}_1\left(x^*\right) - \frac{1}{p_2}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] \\
&\overset{(L3)}{=} -\left(\Theta_1 - \mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right]\right) + \frac{\alpha}{p_1}\Theta_1 - \frac{1}{p_2}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right]
\end{aligned}$$

which is equivalent to:

$$\left(1-\frac{1}{p_2}\right)\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] = \left(1-\frac{\alpha}{p_1}\right)\Theta_1 - \frac{1}{p_2}\left(1-\frac{\alpha}{p_1}\right)\Theta_1 = \left(1-\frac{1}{p_2}\right)\left(1-\frac{\alpha}{p_1}\right)\Theta_1$$

Therefore, we obtain (iii) again:

$$\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] = \left(1-\frac{\alpha}{p_1}\right)\Theta_1 \overset{(2.10)}{=} \Theta_1 - \frac{1}{p_1}\Theta_2 \tag{3.41}$$

We can now deduce (iv):

$$\begin{aligned}
\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x^{\mu}\right)\right] &\overset{(2.10)}{=} \widetilde{U}_1(\sigma^{1*},\sigma^{2*}) + \frac{1}{p_1}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_1\left(x\right)\right] - \frac{1}{p_2}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] \\
&\overset{(3.36),(3.39),(3.41)}{=} \frac{1}{p_2}\left(1-\frac{\alpha}{p_1}\right)\Theta_1 + \frac{1}{p_1}\frac{\alpha}{p_2}\Theta_1 - \frac{1}{p_2}\left(1-\frac{\alpha}{p_1}\right)\Theta_1 \\
&= \frac{\alpha}{p_1 p_2}\Theta_1 \overset{(2.10)}{=} \frac{1}{p_1 p_2}\Theta_2
\end{aligned}$$

Lastly, let us derive (i):

$$\begin{aligned}
U_1(\sigma^{1*},\sigma^{2*}) &\overset{(3.2)}{=} p_1\widetilde{U}_1(\sigma^{1*},\sigma^{2*}) - \frac{p_1}{p_2}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right] \\
&\overset{(3.41),(3.36)}{=} \frac{p_1}{p_2}\left(1-\frac{\alpha}{p_1}\right)\Theta_1 - \frac{p_1}{p_2}\left(1-\frac{\alpha}{p_1}\right)\Theta_1 = 0
\end{aligned}$$

$$\begin{aligned}
U_2(\sigma^{1*},\sigma^{2*}) &\overset{(3.3)}{=} p_2\widetilde{U}_2(\sigma^{1*},\sigma^{2*}) + p_2\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] - \frac{p_2}{p_1}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_1\left(x\right)\right] \\
&\overset{(3.40),(3.39),(3.36)}{=} -\left(1-\frac{\alpha}{p_1}\right)\Theta_1 + \Theta_1 - \frac{\alpha}{p_1}\Theta_1 = 0
\end{aligned}$$

□

From Thm. 1 we observe that, in any equilibrium for Region **III**, the expected amount of initial and effective flow, and expected transportation cost to **P1** and the expected cost of attack to **P2**, can be computed in closed form using the parameters p_1, p_2, the optimal values Θ_1 (the maximum amount of flow in the network), and Θ_2 (the smallest transportation cost among the max-flows). It is easy to check that Thm. 1 is satisfied by $(\tilde{\sigma}^1, \tilde{\sigma}^2)$ defined by (3.4) and (3.5).

Interestingly, the payoffs of both players are zero for any NE (ref. (3.6) and (3.7)). Note that in general, a game that is strategically equivalent to a zero-sum game has different NE that lead to different payoffs. To get oneself convinced, consider the following example:

Example 4. Consider the following zero-sum game:

P1 \ **P2**	L	R
U	$1, -1$	$2, -2$
D	$1, -1$	$3, -3$

Table 3.1: Players' payoffs of a zero-sum game.

In the game presented in Table 3.1, **P1**'s (resp. **P2**'s) pure actions are U and D (resp. L and R). In each cell, the first number corresponds to **P1**'s payoff and the second number corresponds to **P2**'s payoff. In this game, it is easy to see that R is strictly dominated by L so **P2**'s BR is L (no matter what **P1** chooses). One can see that (U, L) and (D, L) are NE, and they give a payoff of 1 (resp. -1) to **P1** (resp. **P2**). Now, consider the following game.

P1 \ **P2**	L	R
U	$1, 0$	$2, -1$
D	$1, -1$	$3, -3$

Table 3.2: Players' payoffs of a game strategically equivalent to the zero-sum game presented in Table 3.1.

The game presented in Table 3.2 is strategically equivalent to the zero-sum game presented in Table 3.1 (we added $\mathbb{1}_{\{U\}}$ to **P2**'s payoff, which is a function that only depends on **P1**). Therefore, (U, L) and (D, L) are NE of this new game (one can easily check it from Table 3.2). However, **P2**'s payoff in equilibrium (U, L) is 0, and her payoff in equilibrium (D, L) is -1. So this game has different NE that lead to different payoffs.

Thus, (i) cannot be entirely derived from the equivalent zero-sum game $\widetilde{\Gamma}$ and requires other results such as the Max-Flow Min-Cut Theorem (see the proof of Thm. 1). Below we further explain (3.8)-(3.11).

Following (3.8) (resp. (3.9)), the expected amount of initial flow (resp. the expected cost of transportation) at equilibrium is equal to some fraction of the value (resp. transporting cost) of the min-cost max-flows, and these expectations decrease with p_2.

Following (3.10), the expected cost of attack at any NE is a constant. In fact, under (A1), we know that $\Theta_2 = \alpha\Theta_1$ (see (2.10)). Therefore, the expected cost of attack at equilibrium becomes $(1-\frac{\alpha}{p_1})\Theta_1$. Applying the Max-Flow Min-Cut Theorem, we obtain that the expected cost of attack is equal to some fraction of the cost of attacking a min-cut set, and this fraction increases with p_1.

Following (3.11), the expected amount of effective flow is again a constant at equilibrium. Under (A1), this quantity becomes $\frac{\alpha}{p_1 p_2}\Theta_1$. Although the amount of effective flow depends on both players' strategies, in any NE, its expectation is always equal to some fraction of the amount of max-flow. Since $\frac{\alpha}{p_1 p_2} < \frac{1}{p_2}$, this flow is always smaller than the expected amount of initial flow. Interestingly, the expected amount of effective flow decreases when p_1 and/or p_2 increase. This result can be explained by noting that when p_1 increases, the disruption caused by **P2** increases, so there is more lost flow and the expected effective flow decreases.

Thm. 1 enables estimation of the expected amount of lost flow and the *yield* of **P1** in any NE. We define yield as the ratio of the expected amount of effective flow and the expected amount of initial flow in the network. We have the following corollary:

Corollary 1. *The expected amount of lost flow is given by:*

$$\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x - x^{\mu}\right)\right] \equiv \frac{1}{p_2}\left(\Theta_1 - \frac{1}{p_1}\Theta_2\right) \overset{(2.10)}{=} \frac{1}{p_2}\left(1 - \frac{1}{p_1}\right)\Theta_1, \tag{3.42}$$

and the yield is given by:

$$\frac{\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x^{\mu}\right)\right]}{\mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right]} \equiv \frac{\Theta_2}{p_1\Theta_1} \overset{(2.10)}{=} \frac{\alpha}{p_1}. \tag{3.43}$$

From (3.42), in any equilibrium, the expected amount of lost flow is equal to some fraction of the amount of max-flow. The corresponding coefficient increases with p_1, because when p_1 is large, **P1** sends more flow and **P2** disrupts more edges. However, the coefficient decreases when p_2 increases, because when p_2 is large, **P2** causes more disruption and **P1** sends less flow in the network. Finally, from (3.43), in any equilibrium, the yield decreases in the ratio $\frac{p_1}{\alpha}$, but it does not depend on p_2 or the maximum amount of flow Θ_1. When $\frac{p_1}{\alpha}$ is large, **P1** has more incentive to send flow in the network and **P2** will attack more frequently, resulting in a lower yield of the network.

Thus, Thm. 1 provides many properties that are satisfied by any NE in Region **III**. Next, we study the support of NE and relate it to optimal solutions of $(\mathcal{P}_2)$ and the min-cut sets.

3.3 Necessary Conditions

Recall the NE in Prop. 3 which has a support based on a min-cost max-flow (for **P1**) and on a min-cut set (for **P2**). We now investigate the generality of this result to other NE. This leads to additional properties satisfied by all NE, and also eases the computation of NE.

Let us present a result regarding the paths taken by the flows in the support of **P1**'s strategy at equilibrium.

Lemma 4. *Every flow in the support of a NE only takes paths whose marginal trans-*

portation cost is α.

$$\forall(\sigma^{1^*},\sigma^{2^*})\in\mathcal{S}_\Gamma,\forall x\in\operatorname{supp}(\sigma^{1^*}),\ \forall\lambda\in\Lambda,\ x_\lambda>0\Longrightarrow\sum_{(i,j)\in\lambda}b_{ij}=\alpha. \tag{3.44}$$

Proof of Lemma 4. Let us consider $\sigma^*=(\sigma^{1^*},\sigma^{2^*})$ a NE . Since $(\tilde{\sigma}^1,\tilde{\sigma}^2)$ is also a NE (Prop. 3), then, by interchangeability, $(\sigma^{1^*},\tilde{\sigma}^2)$ is a NE as well. So, thanks to Thm. 1, we have $\forall x\in\operatorname{supp}(\sigma^{1^*}),\ 0=U_1(x,\tilde{\sigma}^2)=\alpha\,\mathrm{F}\,(x)-\mathrm{C}_1\,(x)$ where the first equality follows from (3.6). Therefore, $\forall x\in\operatorname{supp}(\sigma^{1^*}),\ \mathrm{C}_1\,(x)=\alpha\,\mathrm{F}\,(x)$. Since every path has a marginal transportation cost at least equal to α, then the last equality entails that any flow in the support of a NE takes paths that induce a marginal transportation cost equal to α. □

In other words, the paths that induce a transportation cost strictly greater than α are not chosen in equilibrium. Notice that this lemma implies that any max-flow that is not a min-cost max-flow is not in the support of any NE. This lemma is useful for constructing **P1**'s equilibrium strategies when the network has only a small number of paths of marginal transportation cost α. In the case when the only paths of marginal transportation cost equal to α are the ones taken by the min-cost max-flow x^*, we can deduce that all the equilibrium strategies of **P1** can be constructed from x^* (or sub-flows of x^*). This lemma is of limited use if most of the paths of the network have the same smallest transportation cost.

The following result characterizes the support of **P2**'s equilibrium strategies.

Proposition 4. *Every attack in the support of a NE has a cost at most equal to the cost of attacking a min-cut set, and disrupts edges that are saturated by every min-cost max-flow:*

$$\begin{aligned}\forall(\sigma^{1^*},\sigma^{2^*})\in\mathcal{S}_\Gamma,\ \forall\mu\in\operatorname{supp}(\sigma^{2^*}):\ &\mathrm{C}_2\,(\mu)\le\mathrm{C}_2\left(\mu^{min}\right)=\Theta_1 \\ &\forall(i,j)\in\mathcal{E},\ \mu_{ij}=1\Longrightarrow\forall x^*\in\Omega_2,\ x^*_{ij}=c_{ij}.\end{aligned} \tag{3.45}$$

Proof of Proposition 4. Let us consider $(\sigma^{1^*},\sigma^{2^*})\in\mathcal{S}_\Gamma$. We know that $(\tilde{\sigma}^1,\sigma^{2^*})\in\mathcal{S}_\Gamma$ too (where $\tilde{\sigma}^1$ is defined by (3.4)). Then we can deduce that $\forall\mu\in\operatorname{supp}(\sigma^{2^*}),\ 0=$

$U_2(\tilde{\sigma}^1, \mu) = \mathrm{F}\,(x^* - (x^*)^\mu) - \mathrm{C}_2\,(\mu)$, where the first equality is a consequence of (3.7). Therefore, the Max-Flow Min-Cut Theorem gives us: $\forall \mu \in \mathrm{supp}(\sigma^{2^*}),\ \mathrm{C}_2\,(\mu) = \mathrm{F}\,(x^* - (x^*)^\mu) \leq \mathrm{F}\,(x^*) = \mathrm{C}_2\,(\mu^{min})$.

Besides, since $\forall \mu \in \mathrm{supp}(\sigma^{2^*}),\ \mathrm{C}_2\,(\mu) = \mathrm{F}\,(x^* - (x^*)^\mu)$, then this means that the cost of conducting an attack that is in the support of a NE is equal to the loss it induces to any min-cost max-flow. Since the loss induced by an attack is never greater than the cost of the latter, it means that each edge disrupted by an attack in the support of a NE is saturated by every min-cost max-flow. $\square$

This result tells us that the attacks that require a cost that is strictly greater than the cost of disrupting a min-cut set are not chosen in any equilibrium. Further, if an edge is not saturated by at least one min-cost max-flow, then it is not disrupted in equilibrium. Recall that **P2**'s set of action is isomorphic to the power set of $\mathcal{E}$, that has $2^{|\mathcal{E}|}$ elements which can be huge. Therefore, Prop. 4 enables us to drastically restrict **P2**'s set of actions that can be potentially chosen in equilibrium.

However, since the edges that are part of a min-cut set are saturated by every (min-cost) max-flow, we cannot restrict the set of edges that can be potentially disrupted in equilibrium beyond the min-cut sets. Nevertheless, one can find NE where edges that are not part of any min-cut set are disrupted with positive probability. To get oneself convinced, one can consider the following example:

Example 5. Consider the graph in Fig. 3-4:

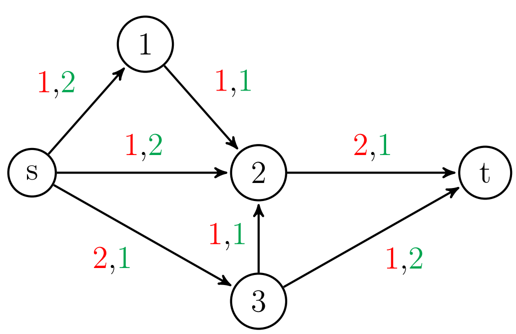

Figure 3-4: Example network containing edges outside a min-cut set that are disrupted in equilibrium.

We can see that $\alpha = 3$, and the min-cost max-flow sends one unit of flow through $\{s,2,t\}$, $\{s,3,2,t\}$ and $\{s,3,t\}$, and only takes paths with transportation cost equal to 3. Thus, (A1) is satisfied. In this graph, there is a unique min-cut set given by $\{(2,t),(3,t)\}$. Let $\mu' = \mathbb{1}_{\{(s,2),(s,3)\}}$ the attack that disrupts edges $(s,2)$ and $(s,3)$. In the case when $3 < p_1 < 4$ and $p_2 > 1$, one can see that there exists an equilibrium strategy σ^{2^*} for **P2** defined by $\sigma^{2^*}_{\mu^0} = \frac{\alpha}{p_1}$ and $\sigma^{2^*}_{\mu'} = 1 - \frac{\alpha}{p_1}$. However, $(s,2)$ and $(s,3)$ are not part of the min-cut set.

Finally, since there are equilibrium strategies that disrupt at least one min-cut set (see e.g. Prop. 3), it is useful to estimate the amount of lost flow for each edge that is attacked in a min-cut set. Similarly, the probability with which each edge of a min-cut set will be disrupted is also of interest because it can be interpreted as the probability with which the flow routed by **P1** is lost when **P2**'s equilibrium strategy only involves edges belonging to a min-cut set. The following proposition answers these questions.

Proposition 5. *Consider a min-cut set $E(\{S,T\})$, then:*

$$\forall (\sigma^{1^*}, \sigma^{2^*}) \in \mathcal{S}_\Gamma,\ \forall (i,j) \in E(\{S,T\}),\ \mathbb{E}_{\sigma^*}[x_{ij}] = \frac{c_{ij}}{p_2}. \tag{3.46}$$

Furthermore, for any NE whose support only contains attacks that disrupt edges of $E(\{S,T\})$ we have:

$$\forall (i,j) \in E(\{S,T\}),\ \mathbb{P}\left((i,j) \text{ is disrupted}\right) = 1 - \frac{\alpha}{p_1}. \tag{3.47}$$

Proof of Proposition 5. Consider a min-cut set $E(\{S,T\})$. Given $p_1 > \alpha$ and $p_2 > 1$, one can find a NE $(\sigma^{1^\dagger}, \sigma^{2^\dagger})$ such that $\forall (i,j) \in E(\{S,T\}),\ \mathbb{1}_{\{(i,j)\}} \in \operatorname{supp}(\sigma^{2^\dagger})$ (for a full characterization of such strategy, refer to Chapter 4).

Consider a NE $(\sigma^{1^*}, \sigma^{2^*})$, we know by interchangeability that $(\sigma^{1^*}, \sigma^{2^\dagger})$ is also a

NE. Therefore:

$$\begin{aligned}
\forall (i,j) \in E(\{S,T\}),\, U_2(\sigma^{1^*}, \mathbb{1}_{\{(i,j)\}}) &\overset{(3.7)}{=} 0 \\
&\iff p_2 \mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x - x^{\mathbb{1}_{\{(i,j)\}}}\right)\right] - \mathrm{C}_2\left(\mathbb{1}_{\{(i,j)\}}\right) = 0 \\
&\iff \mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x - x^{\mathbb{1}_{\{(i,j)\}}}\right)\right] = \frac{c_{ij}}{p_2} \\
&\iff \mathbb{E}_{\sigma^*}\left[x_{ij}\right] = \frac{c_{ij}}{p_2}
\end{aligned}$$

Now, suppose that σ^{2^*} does not contain attacks that disrupt edges outside of $E(\{S,T\})$. First, notice that an edge e is disrupted if and only if it is attacked by at least one attack. Therefore:

$$\mathbb{P}\left(\{e \text{ is disrupted}\}\right) = \sum_{\mu \in \mathcal{A}} \sigma_\mu^{2^*} \mathbb{1}_{\{\mu_e = 1\}} = \sum_{\{\mu \in \mathcal{A} \,|\, \mu_e = 1\}} \sigma_\mu^{2^*}$$

Similarly, one can find a NE $(\sigma^{1'}, \sigma^{2'})$ such that $\forall e \in E(\{S,T\})$, there is a flow $x^e \in \text{supp}(\sigma^{1^*})$ that crosses $E(\{S,T\})$ only at edge e and that takes paths of marginal transportation cost equal to α. We know by interchangeability that $(\sigma^{1'}, \sigma^{2^*})$ is also a NE. Therefore:

$$\begin{aligned}
\forall e \in E(\{S,T\}),\, U_1(x^e, \sigma^{2^*}) &\overset{(3.6)}{=} 0 \\
&\iff p_1 \mathbb{E}_{\sigma^*}\left[\mathrm{F}\left((x^e)^\mu\right)\right] - \alpha\, \mathrm{F}\left(x^e\right) = 0 \\
&\iff \sum_{\mu \in \mathcal{A}} \sigma_\mu^{2^*} \mathrm{F}\left((x^e)^\mu\right) = \frac{\alpha}{p_1} \mathrm{F}\left(x^e\right) \\
&\iff \sum_{\{\mu \in \mathcal{A} \,|\, \mu_e = 1\}} \sigma_\mu^{2^*} \underbrace{\mathrm{F}\left((x^e)^\mu\right)}_{=0} + \sum_{\{\mu \in \mathcal{A} \,|\, \mu_e = 0\}} \sigma_\mu^{2^*} \underbrace{\mathrm{F}\left((x^e)^\mu\right)}_{=\mathrm{F}(x^e)} = \frac{\alpha}{p_1} \mathrm{F}\left(x^e\right) \\
&\iff \mathrm{F}\left(x^e\right) \sum_{\{\mu \in \mathcal{A} \,|\, \mu_e = 0\}} \sigma_\mu^{2^*} = \frac{\alpha}{p_1} \mathrm{F}\left(x^e\right) \\
&\iff \frac{\alpha}{p_1} = \sum_{\{\mu \in \mathcal{A} \,|\, \mu_e = 0\}} \sigma_\mu^{2^*} = 1 - \sum_{\{\mu \in \mathcal{A} \,|\, \mu_e = 1\}} \sigma_\mu^{2^*} \\
&\iff 1 - \frac{\alpha}{p_1} = \sum_{\{\mu \in \mathcal{A} \,|\, \mu_e = 1\}} \sigma_\mu^{2^*} = \mathbb{P}(\{e \text{ is disrupted}\})
\end{aligned}$$

$\square$

From (3.46), at any NE, the expected amount of flow that goes through any edge of a min-cut set is always equal to a constant fraction of its capacity. And from (3.47), if **P2**'s equilibrium strategy only disrupts edges of one min-cut set, then the probability with which an edge is disrupted is constant for all the edges of that min-cut set, irrespective of the capacities of these edges. We can deduce the following corollary that directly follows from Prop. 5:

Corollary 2.

$$\forall (\sigma^{1^*}, \sigma^{2^*}) \in \mathcal{S}_\Gamma,\ \forall \textit{ min-cut set } E(\{S,T\}),\ \forall (i,j) \in E(\{S,T\}) :$$
$$\exists\, x \in \operatorname{supp}(\sigma^{1^*}) \textit{ s.t. } x_{ij} > 0.$$

That is, for any NE and for any edge of a min-cut set, there exists a flow chosen with non-zero probability that passes through that edge.

We apply the previous results to the following example.

Example 6. Consider the network in Fig. 3-2. The only min-cost max-flow is the flow x^* that sends 1 unit of flow through $\{s, 1, 3, t\}$, $\{s, 2, 3, t\}$ and $\{s, 2, 4, t\}$, and the only min-cut set is $\{(1, 2), (2, 3), (2, 4)\}$; see Fig. 3-5.

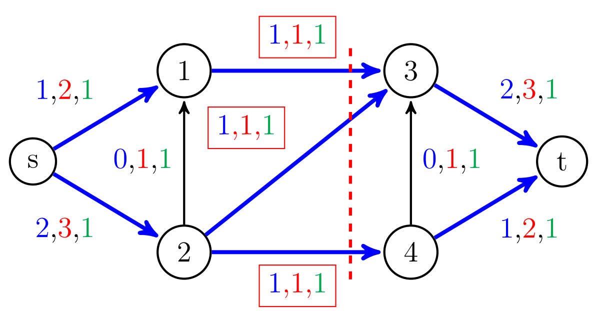

Figure 3-5: Min-cost max-flow (bold blue) and min-cut set attack (dotted red). The labels in the boxes represent the edges that are saturated by the min-cost max-flow.

In this example, the $s - t$ paths that induce the smallest transportation cost are the ones taken by x^*; thus (A1) is satisfied. Lemma 4 tells us that the flows sent with positive probability in equilibrium only take paths taken by x^*. By combining this fact with Prop. 5, we conclude that in any equilibrium, the expected amount of flow in each of the paths $\{s, 1, 3, t\}$, $\{s, 2, 3, t\}$ and $\{s, 2, 4, t\}$ is $\frac{1}{p_2}$.

Besides, the only edges that are saturated by x^* are edges $(1, 3)$, $(2, 3)$ and $(2, 4)$. From Prop. 4, we obtain that only these three edges can be disrupted with nonzero probability in equilibrium. Hence, any **P2**'s equilibrium strategy σ^{2^*} is supported over at most $2^3 = 8$ pure actions instead of $2^9 = 512$ initial pure actions. These edges are exactly the min-cut set of the graph in Fig. 3-5. Therefore, from Prop. 5, each of these edges is disrupted with probability $1 - \frac{3}{p_1}$.

In this example, we showed that the necessary conditions derived above help us restrict the pure actions that support equilibrium strategies to a significant extent. These properties also enable us to derive the following bounds on the probabilities with which certain actions are chosen at equilibrium.

Proposition 6. *Consider $(\sigma^{1^*}, \sigma^{2^*}) \in \mathcal{S}_\Gamma$. Then we have the following bounds:*

(i) If $x^0 \in \operatorname{supp}(\sigma^{1^*})$, then $\sigma^{1^*}_{x^0} \leq 1 - \dfrac{1}{p_2}$

(ii) If $x^* \in \operatorname{supp}(\sigma^{1^*})$, then $\sigma^{1^*}_{x^*} \leq \dfrac{1}{p_2}$

(iii) If $\mu^{min} \in \operatorname{supp}(\sigma^{2^*})$, then $\sigma^{2^*}_{\mu^{min}} \leq 1 - \dfrac{\alpha}{p_1}$

(iv) If $\mu^0 \in \operatorname{supp}(\sigma^{2^*})$, then $\sigma^{2^*}_{\mu^0} \leq \dfrac{\alpha}{p_1}$

Proof of Proposition 6. First, let us derive the bound for x^0:

$$\frac{1}{p_2}\mathrm{F}\left(x^*\right) \overset{(3.8)}{=} \mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] = \sum_{x \in \mathcal{F}\setminus\{x^0\}} \sigma^{1^*}_x \mathrm{F}\left(x\right) \leq \mathrm{F}\left(x^*\right) \sum_{x \in \mathcal{F}\setminus\{x^0\}} \sigma^{1^*}_x = \left(1 - \sigma^{1^*}_{x^0}\right)\mathrm{F}\left(x^*\right)$$

Therefore, $\sigma^{1^*}_{x^0} \leq 1 - \dfrac{1}{p_2}$.

Then, let us derive the bound for any min-cost max-flow x^*:

$$\frac{1}{p_2}\mathrm{F}\left(x^*\right) \overset{(3.8)}{=} \mathbb{E}_{\sigma^*}\left[\mathrm{F}\left(x\right)\right] = \sigma^{1^*}_{x^*}\mathrm{F}\left(x^*\right) + \sum_{x \in \mathcal{F}\setminus\{x^*\}} \sigma^{1^*}_x \mathrm{F}\left(x\right) \geq \sigma^{1^*}_{x^*}\mathrm{F}\left(x^*\right)$$

Therefore, $\sigma^{1^*}_{x^*} \leq \dfrac{1}{p_2}$.

Similarly, for any min-cut set attack μ^{min}, we have:

$$\left(1-\frac{\alpha}{p_1}\right)\mathrm{C}_2\left(\mu^{min}\right)\overset{(3.10)}{=}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right]=\sigma^{2^*}_{\mu^{min}}\,\mathrm{C}_2\left(\mu^{min}\right)+\underbrace{\sum_{\mu\in\mathcal{A}\setminus\{\mu^{min}\}}\sigma^{2^*}_{\mu}\,\mathrm{C}_2\left(\mu\right)}_{\geq 0}$$

Therefore, $\sigma^{2^*}_{\mu^{min}}\leq 1-\dfrac{\alpha}{p_1}$.

Finally, we can derive the same type of bound for μ^0:

$$\begin{aligned}\left(1-\frac{\alpha}{p_1}\right)\mathrm{C}_2\left(\mu^{min}\right)&\overset{(3.10)}{=}\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2\left(\mu\right)\right]=\sum_{\mu\in\mathcal{A}\setminus\{\mu^0\}}\sigma^{2^*}_{\mu}\,\mathrm{C}_2\left(\mu\right)\\&\overset{(3.45)}{\leq}\mathrm{C}_2\left(\mu^{min}\right)\sum_{\mu\in\mathcal{A}\setminus\{\mu^0\}}\sigma^{2^*}_{\mu}=\left(1-\sigma^{2^*}_{\mu^0}\right)\mathrm{C}_2\left(\mu^{min}\right)\end{aligned}$$

Therefore, $\sigma^{2^*}_{\mu^0}\leq\dfrac{\alpha}{p_1}$.

$\square$

This proposition gives upper bounds on the probability with which $x^0, x^*, \mu^0, \mu^{min}$ can be chosen in equilibrium. The NE $(\tilde{\sigma}^1, \tilde{\sigma}^2)$ derived in Prop. 3 attains these bounds. From these upper bounds, we see that when p_2 is close to 1, the probability with which x^0 can be chosen is very small. In contrast, when p_2 is large, x^* can be chosen only with small probability. Similarly, when p_1 is close to α, μ^{min} can be chosen only with small probability, and when p_1 is large, μ^0 can be chosen only with a small probability.

Lastly, we present a result analogous to the Minimax Theorem by Von Neumann [21] for zero-sum games. Recall that the Minimax Theorem is generally not true for games that are strategically equivalent to a zero-sum game; however, Γ satisfies some features of the Minimax Theorem.

Proposition 7. *Each player's payoffs for both maximinimizing and minimaximizing strategies are equal to the payoff at NE, i.e.,*

$$\max_{\sigma^1\in\Delta(\mathcal{F})}\min_{\sigma^2\in\Delta(\mathcal{A})}U_1(\sigma^1,\sigma^2)=0=\min_{\sigma^2\in\Delta(\mathcal{A})}\max_{\sigma^1\in\Delta(\mathcal{F})}U_1(\sigma^1,\sigma^2)$$

$$\max_{\sigma^2\in\Delta(\mathcal{A})}\min_{\sigma^1\in\Delta(\mathcal{F})}U_2(\sigma^1,\sigma^2)=0=\min_{\sigma^1\in\Delta(\mathcal{F})}\max_{\sigma^2\in\Delta(\mathcal{A})}U_2(\sigma^1,\sigma^2)$$

Furthermore, the set of minimaximizers is a superset of $\mathcal{S}_\Gamma$, i.e., any NE is a minimaximizer.

Proof of Proposition 7. First, let us prove that $\max_{\sigma^1} \min_{\sigma^2} U_1(\sigma^1, \sigma^2) = 0$ by directly exhibiting a maximinimizer.

– First, let us prove that $\forall \sigma^1 \in \Delta(\mathcal{F}),\ \min_{\sigma^2 \in \Delta(\mathcal{A})} U_1(\sigma^1, \sigma^2) = -\mathbb{E}_\sigma\left[\mathrm{C}_1\left(x\right)\right]$.

$$\begin{aligned}\forall (\sigma^1, \sigma^2) \in \Delta(\mathcal{F}) \times \Delta(\mathcal{A}),\ U_1(\sigma^1, \sigma^2) &\overset{(2.5)}{=} p_1 \underbrace{\mathbb{E}_\sigma\left[\mathrm{F}\left(x^\mu\right)\right]}_{\geq 0} - \mathbb{E}_\sigma\left[\mathrm{C}_1\left(x\right)\right] \\ &\geq -\mathbb{E}_\sigma\left[\mathrm{C}_1\left(x\right)\right]\end{aligned}$$

which is independent of σ^2. Therefore: $\forall \sigma^1 \in \Delta(\mathcal{F}),\ \min_{\sigma^2 \in \Delta(\mathcal{A})} U_1(\sigma^1, \sigma^2) \geq -\mathbb{E}_\sigma\left[\mathrm{C}_1\left(x\right)\right]$.

Besides $\forall \sigma^1 \in \Delta(\mathcal{F}),\ U_1(\sigma^1, \mu^{min}) = -\mathbb{E}_\sigma\left[\mathrm{C}_1\left(x\right)\right]$. Therefore:

$$\forall \sigma^1 \in \Delta(\mathcal{F}),\ \min_{\sigma^2 \in \Delta(\mathcal{A})} U_1(\sigma^1, \sigma^2) = -\mathbb{E}_\sigma\left[\mathrm{C}_1\left(x\right)\right] \leq 0.$$

This inequality tells us that $\max_{\sigma^1 \in \Delta(\mathcal{F})} \min_{\sigma^2 \in \Delta(\mathcal{A})} U_1(\sigma^1, \sigma^2) \leq 0$. Now it is easy to see that x^0 is a maximinimizer of U_1:

$$\max_{\sigma^1 \in \Delta(\mathcal{F})} \min_{\sigma^2 \in \Delta(\mathcal{A})} U_1(\sigma^1, \sigma^2) = \min_{\sigma^2 \in \Delta(\mathcal{A})} U_1(x^0, \sigma^2) = 0 \tag{3.48}$$

– Now, let us prove that $\min_{\sigma^2} \max_{\sigma^1} U_1(\sigma^1, \sigma^2) \leq 0$ using the definition of a NE: let $(\sigma^{1*}, \sigma^{2*}) \in \mathcal{S}_\Gamma$, then: $\forall \sigma^1 \in \Delta(\mathcal{F}),\ 0 = U_1(\sigma^{1*}, \sigma^{2*}) \geq U_1(\sigma^1, \sigma^{2*})$ where the equality follows from (3.6) and the inequality follows from (2.7). Therefore, $\max_{\sigma^1 \in \Delta(\mathcal{F})} U_1(\sigma^1, \sigma^{2*}) = 0$. Then:

$$\min_{\sigma^2 \in \Delta(\mathcal{A})} \max_{\sigma^1 \in \Delta(\mathcal{F})} U_1(\sigma^1, \sigma^2) \leq \max_{\sigma^1 \in \Delta(\mathcal{F})} U_1(\sigma^1, \sigma^{2*}) = 0 \tag{3.49}$$

– Now, we can get the reverse inequality thanks to the following inequality: $\max_{\sigma^1} \min_{\sigma^2} U_1(\sigma^1, \sigma^2) \leq \min_{\sigma^2} \max_{\sigma^1} U_1(\sigma^1, \sigma^2)$ (inequality that is true for any function of two variables):

$$0 \overset{(3.48)}{=} \max_{\sigma^1 \in \Delta(\mathcal{F})} \min_{\sigma^2 \in \Delta(\mathcal{A})} U_1(\sigma^1, \sigma^2) \leq \min_{\sigma^2 \in \Delta(\mathcal{A})} \max_{\sigma^1 \in \Delta(\mathcal{F})} U_1(\sigma^1, \sigma^2) \tag{3.50}$$

Therefore, (3.49) and (3.50) lead to:

$$\min_{\sigma^2 \in \Delta(\mathcal{A})} \max_{\sigma^1 \in \Delta(\mathcal{F})} U_1(\sigma^1, \sigma^2) = 0 \tag{3.51}$$

– The proof for U_2 is similar, but in this case μ^0 is the maximinimizer of U_2.

– Let us prove that any NE of Γ is a minimaximizer: let $(\sigma^{1*}, \sigma^{2*}) \in \mathcal{S}_\Gamma$, then:

$$0 \overset{(3.51)}{=} \min_{\sigma^2 \in \Delta(\mathcal{A})} \max_{\sigma^1 \in \Delta(\mathcal{F})} U_1(\sigma^1, \sigma^2) \leq \max_{\sigma^1 \in \Delta(\mathcal{F})} U_1(\sigma^1, \sigma^{2*}) \overset{(2.7)}{=} U_1(\sigma^{1*}, \sigma^{2*}) \overset{(3.6)}{=} 0$$

Therefore, $\min_{\sigma^2} \max_{\sigma^1} U_1(\sigma^1, \sigma^2) = \max_{\sigma^1} U_1(\sigma^1, \sigma^{2*})$, i.e., σ^{2*} is a *minimaximizer* of U_1.

A similar proof tells us that σ^{1*} is a *minimaximizer* of U_2.

□

As in the Minimax Theorem, maximinimizing or minimaximizing each player's payoff gives the value of the game at equilibrium. In addition, Prop. 7 tells us that NE are minimaximizers. Interestingly, however, one can find minimaximizers that are not NE, which differs from the Minimax Theorem. Finally the proof of this proposition also implies that NE are not maximinimizers.

In this chapter, we gave a rather full characterization of the NE of game Γ and we saw how they are related to classical network optimization problems such as the minimum cost maximum flow problem and the minimum cut problem. However, in our model, we supposed that **P1** could send any feasible flow in the network and **P2** could disrupt any subset of edges of the network. What if they cannot? In the next chapter, we focus on a more general model where **P1** (resp. **P2**) cannot send every feasible flow (resp. attack any subset of edges) due to budget constraints, and we study the cases where we can still apply the results derived in this chapter to the budget-constrained game.

Chapter 4

Budget-Constrained Game

From Thm. 1, we obtain that the expected cost of transportation (for **P1**) and the expected cost of attack (for **P2**) are constant in any NE. However, NE might differ from each other in the maximum cost of the actions chosen with positive probability. In this chapter, we view these costs as "budget expenditures" of the respective players. Recall the NE $(\tilde{\sigma}^1, \tilde{\sigma}^2)$ in Prop. 3, in which **P1** randomizes between x^0 and x^*, and in which **P2** randomizes between μ^0 and μ^{min} in Region **III**. To play the strategy $\tilde{\sigma}^1$ (resp. $\tilde{\sigma}^2$) in (3.4) (resp. (3.5)), **P1** (resp. **P2**) needs a budget of Θ_2 (resp. Θ_1) for sending a min-cost max-flow (resp. attacking a min-cut set). In this chapter, we study the implications of the players not having a budget high enough to perform $(\tilde{\sigma}^1, \tilde{\sigma}^2)$, and more particularly we characterize the budgets for which the results derived in Sections 3.2 and 3.3 still hold (with minor changes).

For **P1**, we use the infiniteness of her set of actions to find the lowest budget for which we can apply the results in Sections 3.2 and 3.3. However, **P2**'s set of action is discrete and we cannot derive the same bound. We restrict our attention to a subset of **P2**'s equilibrium strategies; in particular, we consider NE we can construct from partitioning the min-cut sets. Prop. 9 below provides our explicit construction of such NE. Next, for this subset of NE, we formulate a problem for computing minimum budget equilibrium strategies as an integer programming problem.

4.1 Revised Model

From now on, let us assume that both players face budget constraints, noted b_1 and b_2 respectively: **P1** (resp. **P2**) can only send flows with transportation cost less than or equal to b_1 (resp. choose an attack with cost less than or equal to b_2). We revise the action sets in game Γ to include the budget constraints b_1 and b_2 as follows:

$$\mathcal{F}_{b_1} = \{x \in \mathcal{F} \mid \mathrm{C}_1(x) \leq b_1\}, \qquad \mathcal{A}_{b_2} = \{\mu \in \mathcal{A} \mid \mathrm{C}_2(\mu) \leq b_2\}.$$

Thus, we consider a more general game $\Gamma_{b_1,b_2} = \{\{1,2\}, (\mathcal{F}_{b_1}, \mathcal{A}_{b_2}), (u_1, u_2)\}$ where u_1 (resp. u_2) is given by (2.3) (resp. (2.4)). As previously, we denote $\Delta(\mathcal{F}_{b_1})$ and $\Delta(\mathcal{A}_{b_2})$ the set of probability distributions over $\mathcal{F}_{b_1}$ and $\mathcal{A}_{b_2}$.

The purpose of this chapter is to find the minimum budgets b_1^* and b_2^* for which the results derived in Sections 3.2 and 3.3 for Γ can still be applicable to $\Gamma_{b_1^*,b_2^*}$ under (A1).

One way to tackle this problem is, given b_1 and b_2, to find a NE of game Γ that satisfies the budget constraints. Indeed, note that $\mathcal{F}_{b_1} \subseteq \mathcal{F}$ and $\mathcal{A}_{b_2} \subseteq \mathcal{A}$. Therefore, if $\sigma^* \in \mathcal{S}_\Gamma \cap (\Delta(\mathcal{F}_{b_1}) \times \Delta(\mathcal{A}_{b_2}))$, then $\sigma^* \in \mathcal{S}_{\Gamma_{b_1,b_2}}$ (follows from (2.7) and (2.8)). In other words, a NE of game Γ that satisfies the budget constraints b_1 and b_2 is a NE of game Γ_{b_1,b_2} and all the results derived in Sections 3.2 and 3.3 are applicable to Γ_{b_1,b_2}.

Note that for the case when b_1 and b_2 are large enough so both players can send any flow in the network and attack any subset of edges (the budget constraints are not binding), $\mathcal{F}_{b_1} = \mathcal{F}$ and $\mathcal{A}_{b_2} = \mathcal{A}$, so $\Gamma_{b_1,b_2} = \Gamma$ and all the results derived so far are applicable.

We now focus on the more interesting case where the budget constraints are binding ($\mathcal{F}_{b_1} \subsetneq \mathcal{F}$ and $\mathcal{A}_{b_2} \subsetneq \mathcal{A}$). Because of the interchangeability of the NE, we can investigate each player's case independently while assuming that the other player's budget constraint is not binding.

4.2 P1's Budget

In this section, we are looking for the lowest budget for **P1**, b_1^*, such that the results presented in Sections 3.2 and 3.3 hold for the budget-constrained game. Because of the interchangeability of the NE, we can investigate **P1**'s case while assuming that $b_2 \geq \Theta_1$. This ensures that $\tilde{\sigma}^2$ from (3.5) is an equilibrium strategy for **P2** in $\Gamma_{b_1^*, b_2}$.

First of all, note that $b_1^* \geq \frac{1}{p_2}\Theta_2$. Indeed, if there existed a NE $(\sigma^{1^*}, \sigma^{2^*}) \in \mathcal{S}_{\Gamma_{b_1,b_2}}$ with $b_1 < \frac{1}{p_2}\Theta_2$, then we would have:

$$\mathbb{E}_{\sigma^*}\left[\mathrm{C}_1(x)\right] = \sum_{x \in \mathcal{F}_{b_1}} \sigma_x^{1^*}\, \mathrm{C}_1(x) < \frac{1}{p_2}\Theta_2 \sum_{x \in \mathcal{F}_{b_1}} \sigma_x^{1^*} = \frac{1}{p_2}\Theta_2,$$

which contradicts (3.9) in Thm. 1. Therefore $b_1^* \geq \frac{1}{p_2}\Theta_2$.

Now, given any $b_1 \geq \frac{1}{p_2}\Theta_2$, we find an equilibrium strategy of Γ for **P1** that assigns positive probability on flows with transportation cost no greater than b_1. This will ensure that $b_1^* \leq \frac{1}{p_2}\Theta_2$.

When $b_1 \geq \Theta_2$, $x^* \in \mathcal{F}_{b_1}$ so $\tilde{\sigma}$ from Prop. 3 is a NE of Γ_{b_1,b_2}. When $\frac{1}{p_2}\Theta_2 \leq b_1 \leq \Theta_2$, then $x^\dagger := \frac{b_1}{\Theta_2}x^* \in \mathcal{F}_{b_1}$ and the following proposition gives a NE of Γ_{b_1,b_2}.

Proposition 8. *If $p_1 > \alpha$, $p_2 > 1$, $\frac{\Theta_2}{p_2} \leq b_1 \leq \Theta_2$, $b_2 \geq \Theta_1$, and under (A1), then $\exists\, \sigma^* = (\sigma^{1^*}, \sigma^{2^*}) \in \mathcal{S}_{\Gamma_{b_1,b_2}}$ such that $U_1(\sigma^{1^*}, \sigma^{2^*}) = U_2(\sigma^{1^*}, \sigma^{2^*}) = 0$, and $\mathrm{supp}(\sigma^{1^*}) = \{x^0, x^\dagger\}$ and $\mathrm{supp}(\sigma^{2^*}) = \{\mu^0, \mu^{min}\}$. The corresponding probabilities are given by:*

- $\sigma_{x^0}^{1^*} = 1 - \dfrac{\Theta_2}{p_2 b_1}, \quad \sigma_{x^\dagger}^{1^*} = \dfrac{\Theta_2}{p_2 b_1}$
- $\sigma_{\mu^0}^{2^*} = \dfrac{\alpha}{p_1}, \quad \sigma_{\mu^{min}}^{2^*} = 1 - \dfrac{\alpha}{p_1}.$

Proof of Proposition 8. Suppose that $p_1 > \alpha$, $p_2 > 1$, $\dfrac{\Theta_2}{p_2} \leq b_1 \leq \Theta_2$, and $b_2 \geq \Theta_1$. Let us show that σ^* is a NE.

$$\begin{aligned} \forall \sigma^1 \in \Delta(\mathcal{F}_{b_1}),\ U_1(\sigma^1, \sigma^{2^*}) &\overset{(2.5)}{=} p_1 \frac{\alpha}{p_1} \mathbb{E}_\sigma\left[\mathrm{F}(x)\right] - \mathbb{E}_\sigma\left[\mathrm{C}_1(x)\right] \\ &\overset{(2.9)}{\leq} \alpha \mathbb{E}_\sigma\left[\mathrm{F}(x)\right] - \alpha \mathbb{E}_\sigma\left[\mathrm{F}(x)\right] = 0 \end{aligned}$$

Besides: $U_1(\sigma^{1^*}, \sigma^{2^*}) = \alpha \frac{\Theta_2}{p_2 b_1} \mathrm{F}\left(\frac{b_1}{\Theta_2} x^*\right) - \frac{\Theta_2}{p_2 b_1} \mathrm{C}_1\left(\frac{b_1}{\Theta_2} x^*\right) \overset{(2.10)}{=} \frac{\alpha}{p_2} \mathrm{F}(x^*) - \frac{\alpha}{p_2} \mathrm{F}(x^*) = 0$

Similarly:

$$\begin{aligned}\forall \sigma^2 \in \Delta(\mathcal{A}_{b_2}), U_2(\sigma^{1^*}, \sigma^2) &\overset{(2.6)}{=} \frac{\Theta_2}{b_1} \mathrm{F}\left(\frac{b_1}{\Theta_2} x^*\right) - \frac{\Theta_2}{b_1} \mathbb{E}_\sigma\left[\mathrm{F}\left(\left(\frac{b_1}{\Theta_2} x^*\right)^\mu\right)\right] - \mathbb{E}_\sigma\left[\mathrm{C}_2(\mu)\right] \\ &= \mathbb{E}_\sigma\left[\mathrm{F}\left(x^* - (x^*)^\mu\right) - \mathrm{C}_2(\mu)\right] \leq 0\end{aligned}$$

where the last inequality follows from the fact that for any attack μ, $\mathrm{F}(x^* - (x^*)^\mu)$ is the loss induced by μ when x^* is in the network, and $\mathrm{C}_2(\mu)$ is the cost of μ which can also be viewed as the maximum amount of flow that can be lost because of the attack.

Besides: $U_2(\sigma^{1^*}, \sigma^{2^*}) = \mathrm{F}(x^*) - \frac{\alpha}{p_1} \mathrm{F}(x^*) - \left(1 - \frac{\alpha}{p_1}\right) \mathrm{C}_2(\mu^{min}) = 0$ thanks to the Max-Flow Min-Cut Theorem. Thus, $(\sigma^{1^*}, \sigma^{2^*})$ is a NE.

□

The difference between Prop. 8 and Prop. 3 is that since x^* is too costly to send, **P1** sends only a fraction of x^* but more frequently (so (3.8) is still satisfied). This is possible because **P1** has a continuous set of actions.

Thus, if one follows the proofs of the results given in Sections 3.2 and 3.3, it can be seen that Prop. 8 ensures that they still hold when **P1**'s budget is greater than or equal to $\frac{\Theta_2}{p_2}$. Therefore $b_1^* = \frac{\Theta_2}{p_2}$.

Let us illustrate Prop. 8 with an example.

Example 7. Consider the graph given in Fig. 3-2 and suppose that $p_1 > 3 = \alpha$, $p_2 = 4$ and $b_1 = 4.5$. Since $\mathrm{C}_1(x^*) = 9 > b_1$, then $x^* \notin \mathcal{F}_{b_1}$ and $\tilde{\sigma}^1$ defined in (3.4) is not an equilibrium strategy anymore. However, Prop. 8 gives us a new equilibrium strategy that satisfies the budget constraint. It is illustrated in Fig. 4-1.

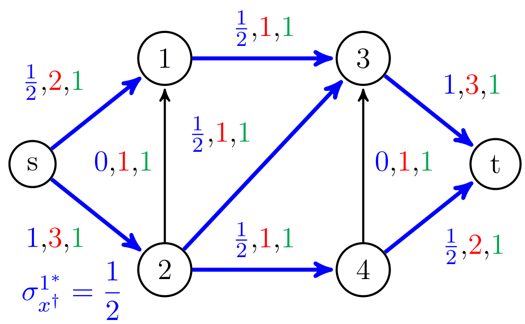

Figure 4-1: Equilibrium strategy described in Prop. 8 when $p_1 > 3$, $p_2 = 4$ and $b_1 = 4.5$. $x^\dagger$ is drawn in blue.

4.3 P2's Budget

Similarly, we are looking for the lowest attack budget for **P2**, b_2^*, such that the structural results in Sections 3.2 and 3.3 hold for the budget-constrained game. Analogously to Section 4.2, we investigate **P2**'s case while assuming that $b_1 \geq \Theta_2$. This ensures that $\tilde{\sigma}^1$ from (3.4) is an equilibrium strategy for **P1** in Γ_{b_1,b_2^*}.

Note that $b_2^* \geq \Theta_1 - \frac{1}{p_1}\Theta_2$. Indeed, if there existed a NE $(\sigma^{1^*}, \sigma^{2^*}) \in \mathcal{S}_{\Gamma_{b_1,b_2}}$ with $b_2 < \Theta_1 - \frac{1}{p_1}\Theta_2$, then we would have:

$$\mathbb{E}_{\sigma^*}\left[\mathrm{C}_2(\mu)\right] = \sum_{\mu \in \mathcal{A}_{b_2}} \sigma^{2^*}_{\mu}\, \mathrm{C}_2(\mu) < \left(\Theta_1 - \frac{1}{p_1}\Theta_2\right) \sum_{\mu \in \mathcal{A}_{b_2}} \sigma^{2^*}_{\mu} = \Theta_1 - \frac{1}{p_1}\Theta_2,$$

which contradicts (3.10). Therefore $b_2^* \geq \Theta_1 - \frac{1}{p_1}\Theta_2 = (1 - \frac{\alpha}{p_1})\Theta_1$.

Unfortunately, this bound is seldom tight due to the finiteness of **P2**'s set of actions. For example, consider the graph given in Fig. 3-2. We can notice that attacking any edge incurs a cost of at least 1. Thus, if $b_2 = (1 - \frac{\alpha}{p_1})\Theta_1 < 1$, then $\mathcal{A}_{b_2} = \{\mu^0\}$ (**P2** cannot attack) and we cannot apply any of the structural results in Sections 3.2 and 3.3.

So far, we know that $b_2^* \leq \Theta_1$ thanks to Prop. 3 (for any budget $b_2 \geq \Theta_1$, $\tilde{\sigma}^2$ is an equilibrium strategy for **P2**). We focus on computing a better upper bound on b_2^* by first considering a large subset of **P2**'s equilibrium strategies based on the partitions of the min-cut sets. Then we find, in this subset of equilibrium strategies, the ones with the lowest attack budget.

4.3.1 Partition-Based Equilibrium Strategies

Throughout this subsection, we consider a min-cut set $E(\{S,T\})$. Let $\{e_1,\dots,e_N\}$ denote the edges that constitute the min-cut set, where N is the number of edges in $E(\{S,T\})$. Recall that μ^{min} is the attack that disrupts all the edges of $E(\{S,T\})$.

A partition of $\{e_1,\dots,e_N\}$ of size n is a set $\{T_1,\dots,T_n\}$ such that:

$$\forall (i,j) \in [\![1,n]\!]^2 \mid i \neq j: \; T_i \cap T_j = \emptyset, \text{ and } \bigcup_{k\in[\![1,n]\!]} T_k = \{e_1,\dots,e_N\}.$$

Definition 1. We say that $\{\mu^1,\dots,\mu^n\}$ is a *partition of* μ^{min} if there exists a partition $\{T_1,\dots,T_n\}$ of the min-cut set $\{e_1,\dots,e_N\}$ of size n such that $\forall k \in [\![1,n]\!],\ \mu^k = \mathbb{1}_{T_k}$ is the attack that disrupts the edges of T_k, i.e.:

$$\forall k \in [\![1,n]\!],\ \mu^k_{ij} = \begin{cases} 1 & \text{if } (i,j) \in T_k, \\ 0 & \text{otherwise.} \end{cases}$$

Note that $\sum_{k=1}^n \mu^k = \mu^{min}$. The following proposition computes NE based on the partitions of μ^{min}.

Proposition 9. *If $p_1 > \alpha$, $p_2 > 1$, and under (A1), then for any partition $\{\mu^1,\dots,\mu^n\}$ of μ^{min} of size n, there exists a NE with support illustrated in Fig. 4-2, and with corresponding probabilities defined according to the following two regions:*

*Case (i) Region **III.a**: if $\alpha < p_1 < \dfrac{n\alpha}{n-1}$ and $p_2 > 1$, then:*

- $\sigma^{1^*}_{x^0} = 1 - \dfrac{1}{p_2}, \quad \sigma^{1^*}_{x^*} = \dfrac{1}{p_2}$
- $\forall k \in [\![1,n]\!],\ \sigma^{2^*}_{\mu^k} = 1 - \dfrac{\alpha}{p_1}, \quad \sigma^{2^*}_{\mu^0} = 1 - n\left(1 - \dfrac{\alpha}{p_1}\right)$

*Case (ii) Region **III.b**: if $p_1 > \dfrac{n\alpha}{n-1}$ and $p_2 > 1$, then:*

- $\sigma^{1^*}_{x^0} = 1 - \dfrac{1}{p_2}, \quad \sigma^{1^*}_{x^*} = \dfrac{1}{p_2}$
- $\forall k \in [\![1,n]\!],\ \sigma^{2^*}_{\mu^k} = \dfrac{\alpha}{p_1(n-1)}, \quad \sigma^{2^*}_{\mu^{min}} = 1 - \dfrac{n\alpha}{p_1(n-1)}$

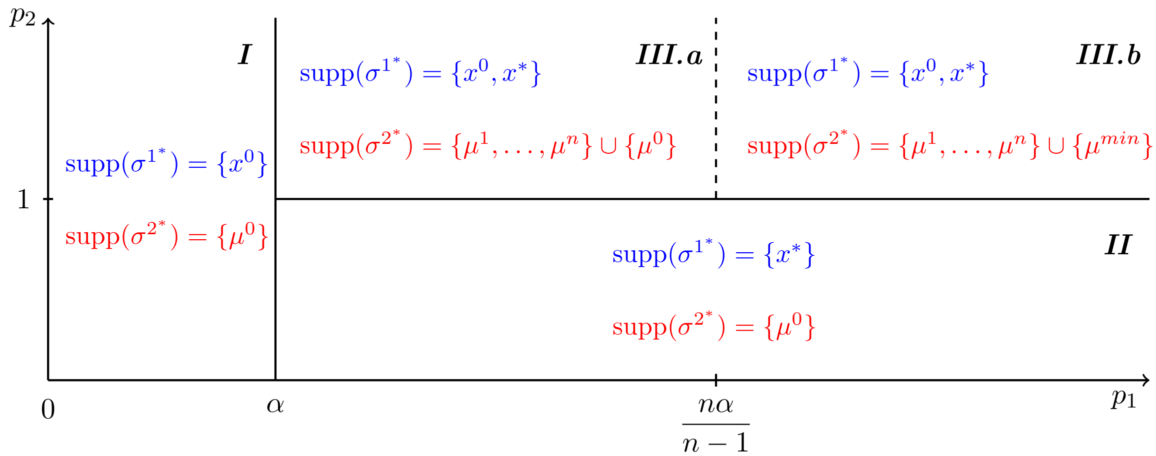

Figure 4-2: Support of partition-based equilibrium strategies in Regions **I-III**.

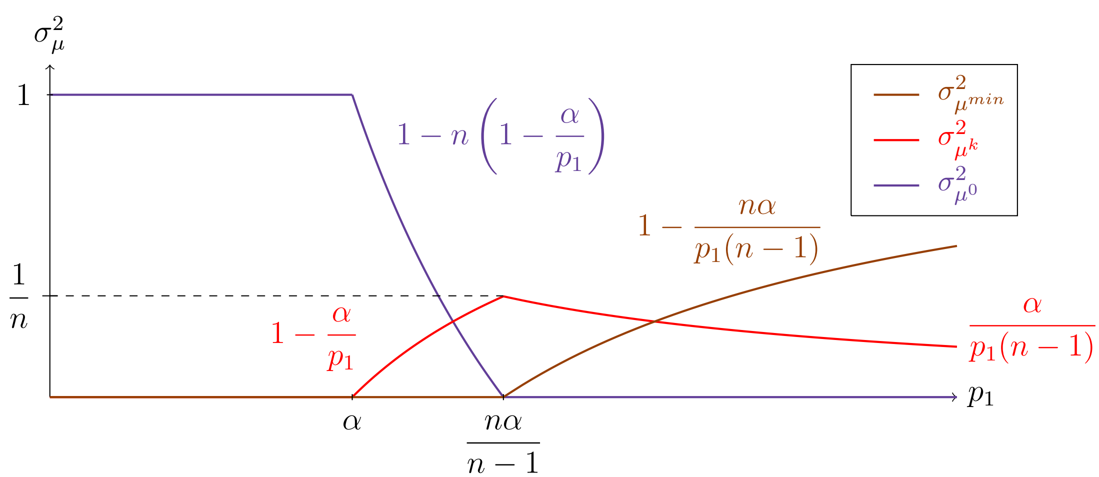

Figure 4-3: Probability distribution of **P2**'s partition based equilibrium strategies.

The probability distributions evolve as in Fig. 4-3.

Proof of Proposition 9.

Case (i): If $\alpha < p_1 < \dfrac{n\alpha}{n-1}$ and $p_2 > 1$ (region **III.a**), then let us prove that $(\{x^0, x^*\}, \{\mu^0, \mu^1, \ldots, \mu^n\})$ is the support of a NE $(\sigma^{1^*}, \sigma^{2^*})$ where:

- $\sigma^{1^*}_{x^0} = 1 - \dfrac{1}{p_2}, \quad \sigma^{1^*}_{x^*} = \dfrac{1}{p_2}$

- $\forall k \in [\![1, n]\!],\ \sigma^{2^*}_{\mu^k} = 1 - \dfrac{\alpha}{p_1}, \quad \sigma^{2^*}_{\mu^0} = 1 - n\left(1 - \dfrac{\alpha}{p_1}\right)$

Let us first prove (i) that σ^{1^*} and σ^{2^*} are probability distributions. Then let us prove (ii) $U_1(x^0, \sigma^{2^*}) = U_1(x^*, \sigma^{2^*}) = 0$ and (iii) $\forall x \in \mathcal{F},\ U_1(x, \sigma^{2^*}) \leq 0$. Similarly we prove (iv) $U_2(\sigma^{1^*}, \mu^0) = U_2(\sigma^{1^*}, \mu^1) = \cdots = U_2(\sigma^{1^*}, \mu^n) = 0$ and (v) $\forall \mu \in \mathcal{A},\ U_2(\sigma^{1^*}, \mu) \leq 0$.

(i) Let us prove that σ^{1^*} and σ^{2^*} are probability distributions:

σ^{1^*} is clearly a probability distribution since $p_2 > 1$.

- $p_1 > \alpha$ so $\sigma^{2^*}_{\mu^k} = 1 - \dfrac{\alpha}{p_1} \geq 0$
- $\sigma^{2^*}_{\mu^0} = 1 - n\left(1 - \dfrac{\alpha}{p_1}\right) = \dfrac{p_1 - np_1 + n\alpha}{p_1} = \dfrac{(n-1)(\frac{n\alpha}{n-1} - p_1)}{p_1} \geq 0$ because $p_1 \leq \dfrac{n\alpha}{n-1}$
- $\displaystyle\sum_{\mu \in \mathcal{A}} \sigma^{2^*}_{\mu} = 1$

So σ^{2^*} is a probability distribution.

(ii) Let us show that $U_1(x^0, \sigma^{2^*}) = 0$.

Since $\forall \mu \in \mathcal{A},\ u_1(x^0, \mu) = 0$, then $U_1(x^0, \sigma^{2^*}) = 0$.

Now, let us prove that $U_1(x^*, \sigma^{2^*}) = 0$.

$$
\begin{aligned}
U_1(x^*, \sigma^{2^*}) =& \sum_{k=1}^{n} \left(1 - \frac{\alpha}{p_1}\right) \left(p_1 \operatorname{F}\left((x^*)^{\mu^k}\right) - \operatorname{C}_1(x^*)\right) \\
&+ \left(1 - n\left(1 - \frac{\alpha}{p_1}\right)\right) \left(p_1 \operatorname{F}\left((x^*)^{\mu^0}\right) - \operatorname{C}_1(x^*)\right) \\
\overset{(2.10)}{=}& (p_1 - \alpha) \sum_{k=1}^{n} \operatorname{F}\left((x^*)^{\mu^k}\right) + (p_1 - n(p_1 - \alpha)) \operatorname{F}(x^*) - \alpha \operatorname{F}(x^*)
\end{aligned}
$$

We can decompose x^* into $\{x^1, \ldots, x^N\}$ where each x^l is the part of x^* that goes through e_l of the min-cut set $E(\{S, T\})$:

$$
\forall l \in [\![1, N]\!], \forall (i,j) \in \mathcal{E},\ x_{ij}^l = \sum_{\lambda \in \Lambda_{ij}^{e_l}} x_\lambda^*
$$

where $\Lambda_{ij}^{e_l} = \{\lambda \in \Lambda \,|\, (i,j) \in \lambda \text{ and } e_l \in \lambda\}$ is the set of paths that go through (i,j) and e_l.

Therefore:

$$
\sum_{k=1}^{n} \operatorname{F}\left((x^*)^{\mu^k}\right) = \sum_{k=1}^{n} \sum_{l=1}^{N} \operatorname{F}\left(\left(x^l\right)^{\mu^k}\right)
$$

Moreover, only one of the μ^k affects x^l. Therefore we get:

$$
\sum_{k=1}^{n} \operatorname{F}\left(\left(x^l\right)^{\mu^k}\right) = (n-1) \operatorname{F}\left(x^l\right)
$$

If we sum over l, we get:

$$
\sum_{k=1}^{n} \operatorname{F}\left((x^*)^{\mu^k}\right) = \sum_{l=1}^{N} \sum_{k=1}^{n} \operatorname{F}\left(\left(x^l\right)^{\mu^k}\right) = (n-1) \sum_{l=1}^{N} \operatorname{F}\left(x^l\right) = (n-1) \operatorname{F}(x^*)
$$

which leads to:

$$
U_1(x^*, \sigma^{2^*}) = (p_1 - \alpha)(n-1) \operatorname{F}(x^*) + (p_1 - n(p_1 - \alpha)) \operatorname{F}(x^*) - \alpha \operatorname{F}(x^*) = 0
$$

(iii) Now, let us show that $\forall x \in \mathcal{F},\ U_1(x, \sigma^{2^*}) \leq 0$.

$$\begin{aligned}
\forall x \in \mathcal{F},\ U_1(x, \sigma^{2^*}) &= \sum_{k=1}^{n} \left(1 - \frac{\alpha}{p_1}\right) \left(p_1 \operatorname{F}\left(x^{\mu^k}\right) - \operatorname{C}_1(x)\right) \\
&+ \left(1 - n\left(1 - \frac{\alpha}{p_1}\right)\right) \left(p_1 \operatorname{F}\left(x^{\mu^0}\right) - \operatorname{C}_1(x)\right) \\
&= (p_1 - \alpha) \sum_{k=1}^{n} \operatorname{F}\left(x^{\mu^k}\right) + (p_1 - n(p_1 - \alpha)) \operatorname{F}(x) - \operatorname{C}_1(x) \\
&\overset{(2.9)}{\leq} (p_1 - \alpha) \sum_{k=1}^{n} \operatorname{F}\left(x^{\mu^k}\right) + (p_1 - n(p_1 - \alpha)) \operatorname{F}(x) - \alpha \operatorname{F}(x)
\end{aligned}$$

Likewise: $\sum_{k=1}^{n} \operatorname{F}\left(x^{\mu^k}\right) = (n-1) \operatorname{F}(x)$. Therefore:

$$U_1(x, \sigma^{2^*}) \leq ((p_1 - \alpha)(n-1) + (p_1 - n(p_1 - \alpha)) - \alpha) \operatorname{F}(x) = 0$$

So σ^{1*} is a BR for **P1**.

(iv) Similarly, let us show that $U_2(\sigma^{1^*}, \mu^0) = 0$.

Since $\forall x \in \mathcal{F},\ u_2(x, \mu^0) = p_2 \operatorname{F}\left(x - x^{\mu^0}\right) - \operatorname{C}_2(\mu^0) = 0$ then $U_2(\sigma^{1^*}, \mu^0) = 0$.

Now, let us prove that $\forall k \in [\![1, n]\!],\ U_2(\sigma^{1^*}, \mu^k) = 0$.

$$\begin{aligned}
\forall k \in [\![1, n]\!],\ U_2(\sigma^{1^*}, \mu^k) &= \left(1 - \frac{1}{p_2}\right) \left(p_2 \times 0 - \operatorname{C}_2\left(\mu^k\right)\right) \\
&+ \frac{1}{p_2} \left(p_2 \operatorname{F}\left(x^* - (x^*)^{\mu^k}\right) - \operatorname{C}_2\left(\mu^k\right)\right) \\
&= \operatorname{F}\left(x^* - (x^*)^{\mu^k}\right) - \operatorname{C}_2\left(\mu^k\right)
\end{aligned}$$

The loss induced to x^* by attacking edges of the min-cut set is exactly equal to the capacity of the attack: $\operatorname{F}\left((x^*)^{\mu^k}\right) = \operatorname{F}(x^*) - \operatorname{C}_2\left(\mu^k\right)$. Therefore: $\forall k \in [\![1, n]\!],\ U_2(\sigma^{1^*}, \mu^k) = 0$.

(v) Lastly, let us prove that $\forall \mu \in \mathcal{A},\ U_2(\sigma^{1^*}, \mu) \leq 0$.

$$\forall \mu \in \mathcal{A},\ U_2(\sigma^{1^*}, \mu) = \mathrm{F}\left(x^* - (x^*)^{\mu}\right) - \mathrm{C}_2\left(\mu\right)$$

The first term is the loss induced by the attack μ when x^* is in the network, and the second term is the cost of the attack which can also be viewed as the maximum amount of flow that can be lost because of the attack. Thus: $\forall \mu \in \mathcal{A},\ U_2(\sigma^{1^*}, \mu) \leq 0$ and σ^{2*} is a BR for **P2**.

Therefore, $(\sigma^{1^*}, \sigma^{2^*}) \in \mathcal{S}_{\Gamma}$.

Case (ii): Now, we consider the case when $p_1 > \dfrac{n\alpha}{n-1}$ and $p_2 > 1$ (Region **III.b**). Then we prove that $(\{x^0, x^*\}, \{\mu^1, \ldots, \mu^n, \mu^{min}\})$ is the support of a NE $(\sigma^{1^*}, \sigma^{2^*})$ where:

- $\sigma^{1^*}_{x^0} = 1 - \dfrac{1}{p_2}, \quad \sigma^{1^*}_{x^*} = \dfrac{1}{p_2}$
- $\sigma^{2^*}_{\mu^k} = \dfrac{\alpha}{p_1(n-1)}\ \forall k \in [\![1, n]\!], \quad \sigma^{2^*}_{\mu^{min}} = 1 - \dfrac{n\alpha}{p_1(n-1)}$

Let us first prove (i) that σ^{1^*} and σ^{2^*} are probability distributions. Then let us prove (ii) $U_1(x^0, \sigma^{2^*}) = U_1(x^*, \sigma^{2^*}) = 0$ and (iii) $\forall x \in \mathcal{F},\ U_1(x, \sigma^{2^*}) \leq 0$. Similarly we prove (iv) $U_2(\sigma^{1^*}, \mu^1) = \cdots = U_2(\sigma^{1^*}, \mu^n) = U_2(\sigma^{1^*}, \mu^{min}) = 0$ and (v) $\forall \mu \in \mathcal{A},\ U_2(\sigma^{1^*}, \mu) \leq 0$.

(i) First, let us prove that σ^{1^*} and σ^{2^*} are probability distributions:

σ^{1^*} is clearly a probability distribution since $p_2 > 1$.

- $\sigma^{2^*}_{\mu^k} = \dfrac{\alpha}{p_1(n-1)} \geq 0$
- $\sigma^{2^*}_{\mu^{min}} = 1 - \dfrac{n\alpha}{p_1(n-1)} = \dfrac{1}{p_1}\left(p_1 - \dfrac{n\alpha}{n-1}\right) \geq 0$ because $p_1 \geq \dfrac{n\alpha}{n-1}$
- $\displaystyle\sum_{\mu \in \mathcal{A}} \sigma^{2^*}_{\mu} = 1$

So σ^{2^*} is a probability distribution.

(ii) Let us prove that $U_1(x^0, \sigma^{2^*}) = 0$.

Since $\forall \mu \in \mathcal{A},\ u_1(x^0, \mu) = 0$, then $U_1(x^0, \sigma^{2^*}) = 0$.

Now let us prove that $U_1(x^*, \sigma^{2^*}) = 0$.

$$\begin{aligned}
U_1(x^*, \sigma^{2^*}) &= \sum_{k=1}^{n} \frac{\alpha}{p_1(n-1)} \left(p_1 \, \mathrm{F}\left((x^*)^{\mu^k} \right) - \mathrm{C}_1\,(x^*) \right) \\
&\quad + \left(1 - \frac{\alpha n}{p_1(n-1)}\right) \left(p_1 \underbrace{\mathrm{F}\left((x^*)^{\mu^{min}} \right)}_{=0} - \mathrm{C}_1\,(x^*) \right) \\
&\overset{(2.10)}{=} \frac{\alpha}{n-1} \sum_{k=1}^{n} \mathrm{F}\left((x^*)^{\mu^k} \right) - \ \alpha \, \mathrm{F}\,(x^*)
\end{aligned}$$

As previously, one can show that $\sum_{k=1}^{n} \mathrm{F}\left((x^*)^{\mu^k} \right) = (n-1)\,\mathrm{F}\,(x^*)$, which leads to:

$$U_1(x^*, \sigma^{2^*}) = \frac{\alpha}{n-1}(n-1)\,\mathrm{F}\,(x^*) - \alpha\,\mathrm{F}\,(x^*) = 0$$

(iii) Now let us show that $\forall x \in \mathcal{F},\ U_1(x, \sigma^{2^*}) \leq 0$.

$$\begin{aligned}
U_1(x, \sigma^{2^*}) &= \sum_{k=1}^{n} \frac{\alpha}{p_1(n-1)} \left(p_1 \, \mathrm{F}\left(x^{\mu^k} \right) - \mathrm{C}_1\,(x) \right) \\
&\quad + \left(1 - \frac{\alpha n}{p_1(n-1)}\right) \left(p_1 \underbrace{\mathrm{F}\left(x^{\mu^{min}} \right)}_{=0} - \mathrm{C}_1\,(x) \right) \\
&= \frac{\alpha}{n-1} \sum_{k=1}^{n} \mathrm{F}\left(x^{\mu^k} \right) - \mathrm{C}_1\,(x)
\end{aligned}$$

Likewise: $\displaystyle\sum_{k=1}^{n} \mathrm{F}\left(x^{\mu^k} \right) = (n-1)\,\mathrm{F}\,(x)$.

Therefore:

$$U_1(x, \sigma^{2^*}) \overset{(2.9)}{\leq} \frac{\alpha}{n-1}(n-1)\,\mathrm{F}\,(x) - \alpha\,\mathrm{F}\,(x) = 0$$

So σ^{1*} is a BR for **P1**.

(iv) Similarly, let us show that $U_2(\sigma^{1*}, \mu^{min}) = 0$.

$$\begin{aligned} U_2(\sigma^{1^*}, \mu^{min}) = & \left(1 - \frac{1}{p_2}\right)\left(p_2 \times 0 - \mathrm{C}_2\left(\mu^{min}\right)\right) \\ & + \frac{1}{p_2}\left(p_2\,\mathrm{F}\left(x^* - (x^*)^{\mu^{min}}\right) - \mathrm{C}_2\left(\mu^{min}\right)\right) \\ = & \,\mathrm{F}\,(x^*) - \mathrm{C}_2\left(\mu^{min}\right) = 0 \end{aligned}$$

where the last equality follows from the Max-Flow Min-Cut Theorem.

Now, let us prove that $\forall k \in [\![1, n]\!],\ U_2(\sigma^{1*}, \mu^k) = 0$

$$\begin{aligned} \forall k \in [\![1, n]\!],\ U_2(\sigma^{1^*}, \mu^k) = & \left(1 - \frac{1}{p_2}\right)\left(p_2 \times 0 - \mathrm{C}_2\left(\mu^k\right)\right) \\ & + \frac{1}{p_2}\left(p_2\,\mathrm{F}\left(x^* - (x^*)^{\mu^k}\right) - \mathrm{C}_2\left(\mu^k\right)\right) \\ = & \,\mathrm{F}\left(x^* - (x^*)^{\mu^k}\right) - \mathrm{C}_2\left(\mu^k\right) \end{aligned}$$

The loss induced to x^* by attacking edges of the min-cut set is exactly equal to the capacity of the attack: $\mathrm{F}\left((x^*)^{\mu^k}\right) = \mathrm{F}\,(x^*) - \mathrm{C}_2\left(\mu^k\right)$. Therefore: $\forall k \in [\![1, n]\!],\ U_2(\sigma^{1^*}, \mu^k) = 0$

(v) We already proved in Region **III.a** that: $\forall \mu \in \mathcal{A},\ U_2(\sigma^{1^*}, \mu) \leq 0$. Thus σ^{2*} is a BR for **P2**.

Therefore, $(\sigma^{1*}, \sigma^{2*}) \in \mathcal{S}_\Gamma$.

□

Given $p_1 > \alpha$ and $n \in [\![1, N]\!]$, let us note $\mathcal{S}^n_{p_1}$ the set of **P2**'s equilibrium strategies described by Prop. 9 whose support has a size equal to $n+1$ (the support is based on a partition of μ^{min} of size n and, depending on p_1, whether includes μ^0 or μ^{min} according to Prop. 9). We denote the set of **P2**'s equilibrium strategies described by Prop. 9 (for a fixed p_1) as:

$$\mathcal{S}_{p_1} := \bigcup_{n \in [\![1,N]\!]} \mathcal{S}_{p_1}^n \tag{4.1}$$

Prop. 9 enables us to have an analytical expression of a large number of **P2**'s equilibrium strategies. Indeed, given any partition of a min-cut set, we can find a corresponding partition-based equilibrium strategy for **P2** thanks to Prop. 9. Since there are $2^N - 1$ such partitions, then $\mathcal{S}_{p_1}$ contains $2^N - 1$ equilibrium strategies for **P2**.

In Fig. 4-2, we find again Regions **I** and **II** outlined in Props. 1 and 2. For $\mathcal{S}_{p_1}$, Prop. 9 splits Region **III** into two subregions, where each region considers the partitions of μ^{min} in a specific manner. For Case 1 ($\frac{n\alpha}{n-1} > p_1$), an equilibrium strategy for **P2** randomizes over the partition $\{\mu^1, \ldots, \mu^n\}$ and μ^0. However, for Case 2 ($\frac{n\alpha}{n-1} < p_1$), an equilibrium strategy for **P2** randomizes over the same partition and μ^{min}. Intuitively, if **P2** partitions μ^{min} in too many components (i.e., $\frac{n\alpha}{n-1}$ decreases), then she assigns positive probability to the min-cut set attack. However, if she partitions μ^{min} in fewer components, then she chooses no attack action with a nonzero probability.

Remark 2. The case $n = 1$ corresponds to attacking the whole min-cut set (it's a partition of size 1). When n tends to 1 from above, $\frac{n\alpha}{n-1} \longrightarrow +\infty$. Therefore, if we draw Fig. 4-2 in the case $n = 1$, we find again Fig. 3-1 (Region **III.a** expands); thus, Prop. 3 is a particular case of Prop. 9.

Example 8. Let us illustrate Prop. 9 with the example in Fig. 3-2. Recall that the only min-cut set is $\{(1,3), (2,3), (2,4)\}$ and the only min-cost max-flow sends one unit of flow through each of the paths $\{s,1,3,t\}$, $\{s,2,3,t\}$ and $\{s,2,4,t\}$.

Let us consider one partition $\{\{(1,3),(2,3)\}\{(2,4)\}\}$ of the min-cut set. From this partition, we construct the corresponding attacks μ^1 that disrupts edges $(1,3)$ and $(2,3)$, and μ^2 that disrupts edge $(2,4)$. Thus, $\{\mu^1, \mu^2\}$ is a partition of μ^{min}. The results obtained by applying Prop. 9 to this example are presented in Fig. 4-4.

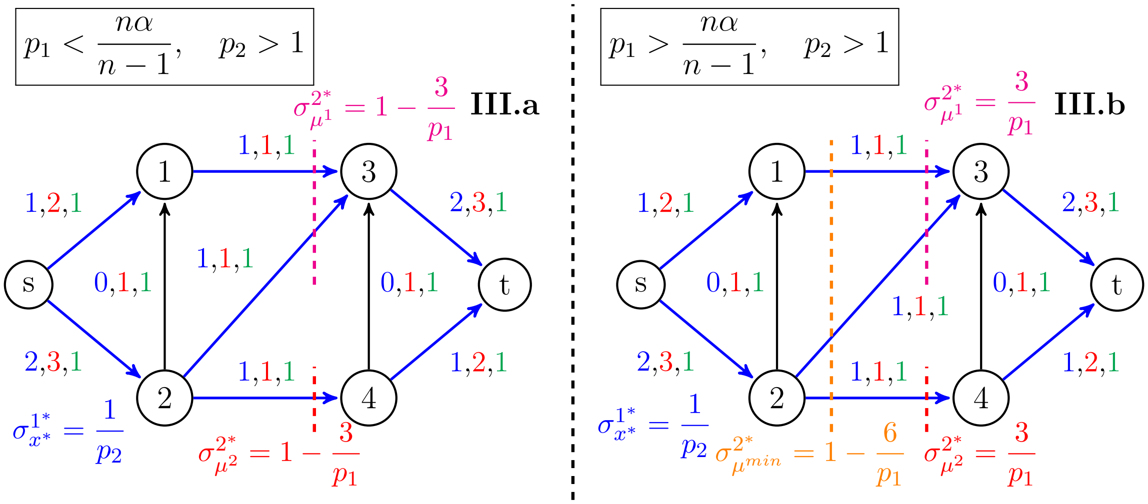

Figure 4-4: NE described in Prop. 9 based on the partition $\{\mu^1, \mu^2\}$.

Now that we can analytically compute many more NE, we can try to find the equilibrium strategies among $\mathcal{S}_{p_1}$ that require the lowest budget.

4.3.2 Optimization Problem

In the previous subsection, we saw that any partition of μ^{min} (along with μ^0, μ^{min}) can be used to explicitly construct a subset of equilibrium strategies for **P2**. Specifically, an equilibrium based on such a partition can be mapped into one of the two regions (**III.a-b**) illustrated in Fig. 4-2.

Without loss of generality, let us consider a unique min-cut set $(e_1, \ldots, e_N)$ consisting of N edges. With a slight abuse of notation, let us denote c_k the capacity of edge e_k (for all $k \in [\![1, N]\!]$).

First, note the following:

$$\sigma^2 \in \Delta(\mathcal{A}_{b_2}) \iff \forall \mu \in \text{supp}(\sigma^2),\ \mathrm{C}_2(\mu) \leq b_2 \iff \max_{\mu \in \text{supp}(\sigma^2)} \mathrm{C}_2(\mu) \leq b_2.$$

That is, a strategy satisfies the budget constraint if and only if the maximum cost of conducting an attack chosen with positive probability is no greater than the budget constraint.

Therefore, in order to get a better upper bound on the minimum budget b_2^* for which the results in Sections 3.2 and 3.3 still hold, we want to find a strategy in $\mathcal{S}_{p_1}$ (defined by (4.1)) that minimizes the maximum cost of conducting an attack chosen with a positive probability, i.e.:

$$\arg \min_{\sigma^{2*} \in \mathcal{S}_{p_1}} \max_{\mu \in \mathrm{supp}(\sigma^{2*})} \mathrm{C}_2(\mu)$$

The following proposition gives the answer.

Proposition 10. *Among the NE listed in Prop. 9, a strategy that minimizes the budget needed is based on a partition of μ^{min} of size $n^* := \min\left\{\left\lfloor \frac{p_1}{p_1-\alpha} \right\rfloor, N\right\}$, and is obtained by solving the following integer programming problem:*

$$\boxed{\begin{array}{lll} (IP) & \textit{minimize} & z \\ & \textit{subject to} & z \geq \displaystyle\sum_{l=1}^{N} c_l\, y_{lk}, \quad \forall k \in [\![1, n^*]\!] \\ & & \displaystyle\sum_{k=1}^{n^*} y_{lk} = 1, \quad \forall l \in [\![1, N]\!] \\ & & y_{lk} \in \{0,1\}, \quad \forall (l,k) \in [\![1, N]\!] \times [\![1, n^*]\!]. \end{array}}$$

Proof of Proposition 10. We want to solve the following problem:

$$\min_{\sigma^{2*} \in \mathcal{S}_{p_1}} \max_{\mu \in \mathrm{supp}(\sigma^{2*})} \mathrm{C}_2(\mu) = \min_{n \in [\![1,N]\!]} \min_{\sigma^{2*} \in \mathcal{S}_{p_1}^n} \max_{\mu \in \mathrm{supp}(\sigma^{2*})} \mathrm{C}_2(\mu). \tag{4.2}$$

First, let us start by finding $n^* \in \arg\min_{n \in [\![1,N]\!]} \min_{\sigma^{2*} \in \mathcal{S}_{p_1}^n} \max_{\mu \in \mathrm{supp}(\sigma^{2*})} \mathrm{C}_2(\mu)$. There are two cases to consider:

– Case 1: If $\dfrac{n\alpha}{n-1} < p_1$, then $\forall \sigma^{2*} \in \mathcal{S}_{p_1}^n,\ \mathrm{supp}(\sigma^{2*}) = \{\mu^1, \ldots, \mu^n\} \cup \{\mu^{min}\}$ so:

$$\forall \sigma^{2*} \in \mathcal{S}_{p_1}^n, \quad \max_{\mu \in \mathrm{supp}(\sigma^{2*})} \mathrm{C}_2(\mu) = \mathrm{C}_2\left(\mu^{min}\right) = \Theta_1.$$

Therefore, if $\dfrac{n\alpha}{n-1} < p_1$, then $\displaystyle\min_{\sigma^{2*} \in \mathcal{S}_{p_1}^n} \max_{\mu \in \mathrm{supp}(\sigma^{2*})} \mathrm{C}_2(\mu) = \Theta_1$.

– Case 2: If $\dfrac{n\alpha}{n-1} \geq p_1$, then $\forall \sigma^{2*} \in \mathcal{S}_{p_1}^n,\ \text{supp}(\sigma^{2*}) = \{\mu^1, \ldots, \mu^n\} \cup \{\mu^0\}$ so:

$$\forall \sigma^{2*} \in \mathcal{S}_{p_1}^n, \max_{\mu \in \text{supp}(\sigma^{2*})} \mathrm{C}_2(\mu) = \max_{k \in [\![1,n]\!]} \mathrm{C}_2\left(\mu^k\right).$$

For any partition $\{\mu^1, \ldots, \mu^n\}$ of μ^{min}, we note $\{T_1, \ldots, T_n\}$ the corresponding partition of the min-cut set $\{e_1, \ldots, e_N\}$ of capacities $c_1, \ldots, c_N$. Then the cost of each μ^k is equal to the sum of the capacities of the edges it disrupts, i.e.:

$$\forall k \in [\![1,n]\!],\ \mathrm{C}_2\left(\mu^k\right) = \sum_{e_l \in T_k} c_l$$

Therefore:

$$\forall \sigma^{2*} \in \mathcal{S}_{p_1}^n, \max_{\mu \in \text{supp}(\sigma^{2*})} \mathrm{C}_2(\mu) = \max_{k \in [\![1,n]\!]} \sum_{e_l \in T_k} c_l.$$

One can see that the problem is equivalent to finding a partition $\{T_1, \ldots, T_n\}$ of the min-cut set such that the maximum sum of the capacities of the edges constituting each element of the partition is minimized. n still being fixed, we want to solve the following bilevel optimization problem:

$$\psi(n) := \min_{\{T_1,\ldots,T_n\}} \max_{k \in [\![1,n]\!]} \sum_{e_l \in T_k} c_l$$

Now, let us argue that the optimal value of the previous bilevel problem, $\psi(n)$, does not increase when partitioning the min-cut set into more pieces, i.e., if $n' > n$, then $\psi(n') \leq \psi(n)$.

Indeed, let us consider $n \leq N-1$, and let us note $\{T_1^*, \ldots, T_n^*\}$ an optimal partitioning of the min-cut set of size n: $\psi(n) = \max_{k \in [\![1,n]\!]} \sum_{e_l \in T_k^*} c_l$.

Since $n \leq N-1$, then at least one of the T_k^* contains at least two edges (Dirichlet's principle). Without loss of generality, let us assume that T_n^* contains at least two edges. Let us denote e_{k_0} one of the edges of T_n^*.

Now, let us consider $\{T_1, \dots, T_{n+1}\}$ a partition of the min-cut set of size $n+1$ such that $\forall k \in [\![1, n-1]\!],\ T_k = T_k^*$, $T_n = T_n^* \backslash \{e_{k_0}\}$ and $T_{n+1} = \{e_{k_0}\}$.

Notice that $c_{k_0} \leq \sum_{e_l \in T_n^*} c_l$ and $\sum_{e_l \in T_n} c_l \leq \sum_{e_l \in T_n^*} c_l$.

Thus, we constructed a partition of the min-cut set of size $n+1$, such that:

$$\psi(n+1) \leq \max_{k \in [\![1, n+1]\!]} \sum_{e_l \in T_k} c_l \leq \max_{k \in [\![1, n]\!]} \sum_{e_l \in T_k^*} c_l = \psi(n)$$

Therefore, ψ is a non-increasing function, which means that in (4.2), we need to increase n as much as possible. However, let us not forget that we are in the case where $\dfrac{n\alpha}{n-1} \geq p_1$.

Thus, the optimal partitioning size in this case is the largest integer n^* such that $\dfrac{n\alpha}{n-1} \geq p_1$, i.e., $n^* = \max\left\{n \in [\![1, N]\!] \mid \dfrac{n\alpha}{n-1} \geq p_1\right\} = \min\left\{\left\lfloor \dfrac{p_1}{p_1 - \alpha} \right\rfloor, N\right\}$.

Finally, notice that for any partition $\{T_1, \dots, T_n\}$, we have

$$\forall k \in [\![1, n]\!],\ \sum_{e_l \in T_k} c_l \leq \sum_{l=1}^{N} c_l = \Theta_1.$$

This implies that $\psi(n^*) \leq \Theta_1$, which is the optimal value of the original problem in Case 1. Therefore n^* is the optimal partitioning size.

Now that we know n^*, we can derive an integer programming problem that provides us with an optimal partition of μ^{min}:

$$\begin{array}{lll}
\text{minimize} & z & \\
\text{subject to} & z \geq \displaystyle\sum_{l=1}^{N} c_l\, y_{lk}, & \forall k \in [\![1, n^*]\!] \\
& \displaystyle\sum_{k=1}^{n^*} y_{lk} = 1, & \forall l \in [\![1, N]\!] \\
& y_{lk} \in \{0, 1\}, & \forall (l, k) \in [\![1, N]\!] \times [\![1, n^*]\!].
\end{array}$$

One can see that this integer programming problem gives the optimal value $\psi(n^*)$ and an optimal way of partitioning the min-cut set thanks to the y_{lk}. Indeed, for every edge e_l, and any set T_k, $y_{lk} = 1$ means that edge e_l goes in set T_k. Thanks to the constraint $\sum_{k=1}^{n^*} y_{lk} = 1,\ \forall l \in [\![1, N]\!]$, each edge e_l goes in exactly one set T_k, thus creating a partition of the min-cut set.

Then, from this partition of the min-cut set, one can partition $\mu^{\min}$ accordingly, and use Prop. 9 in order to derive the corresponding probabilities. □

Given a partition $\{\mu^1, \ldots, \mu^n\}$ of μ^{min}, the support of the corresponding strategy in $\mathcal{S}_{p_1}$ is this partition along with μ^{min} or μ^0 depending on p_1. Let us note $\{T_1, \ldots, T_n\}$ the corresponding partition of the min-cut set. For every $k \in [\![1, n]\!]$, one may notice the following:

$$\mathrm{C}_2\left(\mu^k\right) = \sum_{(i,j)\in T_k} c_{ij} = \sum_{l=1}^{N} c_l \mu^k_{e_l} \tag{4.3}$$

The cost of conducting attack μ^k is equal to the sum of the capacities of the edges of the min-cut set that μ^k disrupts. Therefore, the more **P2** partitions μ^{min}, the less number of edges of the min-cut set each μ^k disrupts. In Prop. 9, we saw that when $\frac{n\alpha}{n-1} > p_1$, **P2** randomizes over the partition and no attack, so the maximum attacking cost is induced by one of the elements of the partition. Thus, **P2** needs to increase n. However, when n increases, $\frac{n\alpha}{n-1}$ decreases and we saw that when $\frac{n\alpha}{n-1} < p_1$, then μ^{min} enters the support and the budget that is needed is Θ_1 (the capacity of the min-cut). Therefore, **P2** needs to increase n until $n^* = \max\left\{n \in [\![1, N]\!] \mid \frac{n\alpha}{n-1} \geq p_1\right\} = \min\left\{\left\lfloor \frac{p_1}{p_1-\alpha} \right\rfloor, N\right\}$.

Knowing the optimal size of the partition of μ^{min}, we can find a partition of size n^* that minimizes the maximum attacking cost. Thanks to (4.3), one can see that this is equivalent to assigning N objects (the edges of the min-cut set) of value c_l each, into n^* bags such that the maximum value of the bags is minimized. This is the purpose of (IP).

The optimal value of (IP), z^*, gives a new upper bound on b_2^* (we are certain that the results in Sections 3.2 and 3.3 hold when $b_2 \geq z^* \geq b_2^*$). Depending on the min-cut set, z^* may be much smaller than Θ_1, which was the previous upper bound on b_2^* deduced from Prop. 3. In addition, the optimal solution of (IP), y_{lk}^*, gives us the corresponding way of partitioning μ^{min} ($y_{lk}^* = 1$ if and only if edge e_l is disrupted by μ^k), and Prop. 9 derives the corresponding probabilities to construct an equilibrium strategy.

Let us illustrate Prop. 10 with an example.

Example 9. Once again, consider the graph given in Fig. 3-2 and assume that $p_1 = 5$. First, let us enumerate all the equilibrium strategies of $\mathcal{S}_{p_1}$ in Fig. 4-5.

Fig. 4-5a shows that the NE of $\mathcal{S}_{p_1}^1$ contains an attack that induces a cost of 3, while Figs. 4-5b, 4-5c, and 4-5d show that the NE of $\mathcal{S}_{p_1}^2$ contain attacks that induce at most a cost of 2. However, Fig. 4-5e shows that the NE of $\mathcal{S}_{p_1}^3$ contains an attack that also has a cost of 3. Therefore, **P2**'s equilibrium strategies of $\mathcal{S}_{p_1}$ that require the lowest budget are based on a partition of size 2, which corresponds to min $\left\{ \left\lfloor \frac{5}{5-3} \right\rfloor, 3 \right\}$. Thus, we are sure for this example that the results in Sections 3.2 and 3.3 hold when $b_2 \geq 2 = z^*$.

Thus, combining Props. 9 and 10 computes a new upper bound on the lowest attack budget for **P2** to which we can apply the structural results presented in Sections 3.2 and 3.3.

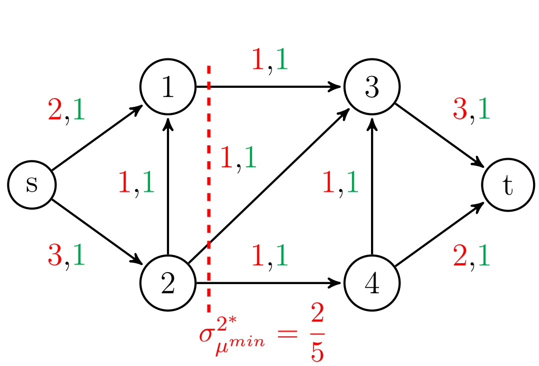

(a) Partition of size 1

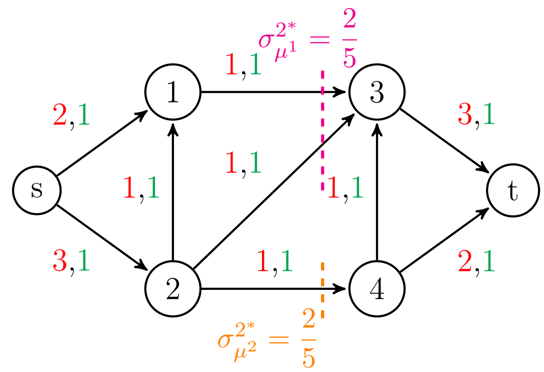

(b) Partition of size 2

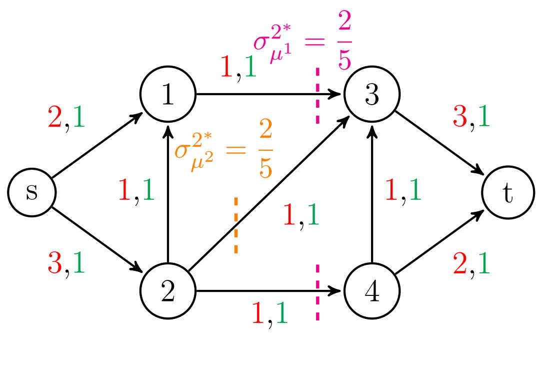

(c) Partition of size 2

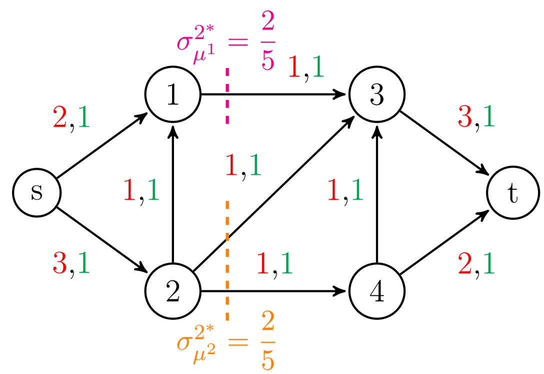

(d) Partition of size 2

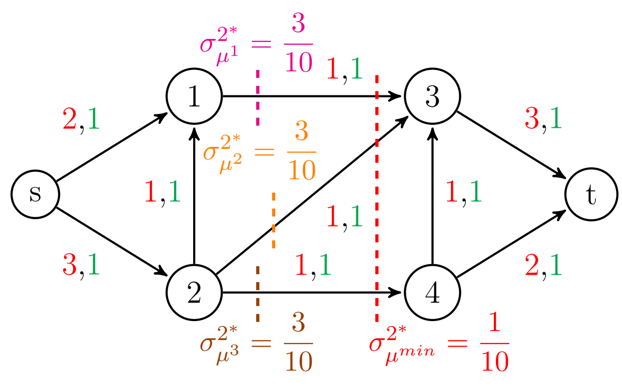

(e) Partition of size 3

Figure 4-5: Enumeration of the equilibrium strategies in $\mathcal{S}_{p_1}$.

Chapter 5

Some Extensions

Finally, we discuss the implications of relaxing some assumptions of our model. First, we consider a network that does not satisfy (A1). We use this example to see under which circumstances our results can still be applied. Then, we present a weaker assumption than (A1) for which the results derived in Chapter 3 hold as well. Following, we further discuss our assumption regarding the cost of attack: in the model studied in this thesis, we supposed that the cost of attacking an edge was proportional to its capacity. We now present a network with a general cost of attack and we see how some of the results derived in this thesis apply to such networks.

5.1 Relaxing Assumption 1

We now study the implications of relaxing (A1) by way of an example. Consider the graph given in Fig. 5-1 and consider the game Γ.

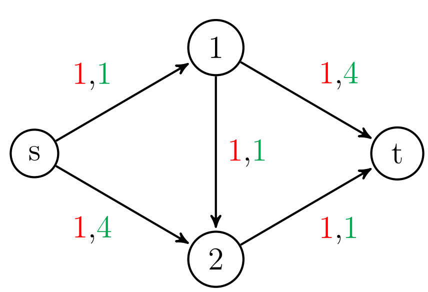

Figure 5-1: Initial graph.

The unique min-cost max-flow sends one unit of flow through each of the paths $\{s, 1, t\}$ and $\{s, 2, t\}$ whose marginal transportation cost is equal to 5. However, $\{s, 1, 2, t\}$ has a marginal transportation cost equal to 3 so (A1) does not hold. Let us assume that $p_1 = p_2 = 6$. Then, we can show that $\tilde{\sigma}$ defined by (3.4) and (3.5) is not a NE anymore. Let us note x' the flow that sends 1 unit through path $\{s, 1, 2, t\}$, $\mu^1 = \mathbb{1}_{(s,1)}$, $\mu^2 = \mathbb{1}_{(1,2)}$ and $\mu^3 = \mathbb{1}_{(2,t)}$. Then, one can show that there exists an equilibrium where **P1**'s strategy σ^{1^*} is defined by $\sigma^{1^*}_{x'} = \frac{1}{6}$ and $\sigma^{1^*}_{x^0} = \frac{5}{6}$, and **P2**'s strategy σ^{2^*} is defined by $\sigma^{2^*}_{\mu^1} = \sigma^{2^*}_{\mu^2} = \sigma^{2^*}_{\mu^3} = \frac{1}{6}$ and $\sigma^{2^*}_{\mu^0} = \frac{1}{2}$. We can see that this strategy does not rely on the min-cost max-flow and the min-cut set anymore. This NE is illustrated in Fig. 5-2.

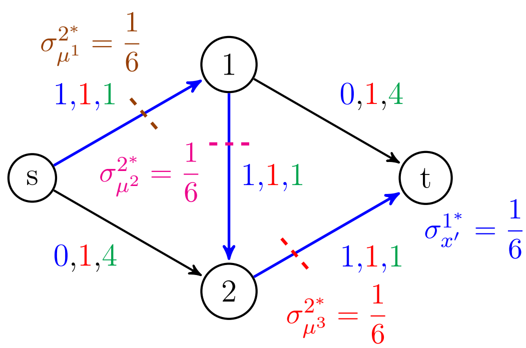

Figure 5-2: NE in the case $p_1 = p_2 = 6$.

However, if we suppose instead that $3 < p_1 < 5$ and $p_2 > 1$, then we can prove that $(\sigma^{1^*}, \sigma^{2^*})$ defined by $\sigma^{1^*}_{x^0} = 1 - \frac{1}{p_2}$, $\sigma^{1^*}_{x'} = \frac{1}{p_2}$, and $\sigma^{2^*}_{\mu^0} = \frac{3}{p_1}$, $\sigma^{2^*}_{\mu^2} = 1 - \frac{3}{p_1}$ is a NE. This result looks similar to the one we derived in Prop. 3. Actually, they are related: when $3 < p_1 < 5$, the marginal transportation costs of paths $\{s, 1, t\}$ and $\{s, 2, t\}$ are higher than the marginal value of effective flow, so **P1** has no incentive to send any flow along these paths. If we remove these paths from the graph (as in the elimination of strictly dominated strategies), we obtain the subgraph in Fig. 5-3a.

It turns out this subgraph satisfies (A1) (it's only a path). Therefore, we can apply all our results to this subgraph, and the elimination of strictly dominated strategies tells us that they will hold for the original graph (the equilibrium we found is exactly $\tilde{\sigma}$ from Prop. 3 applied to this subgraph). The NE is illustrated in Fig. 5-3b.

Thanks to this example, we see that we can relax (A1) and extend all the results

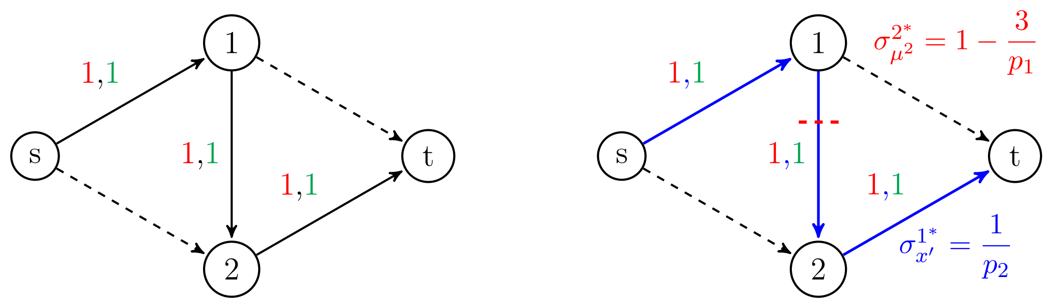

(a) Subgraph (b) NE in the case $3 < p_1 < 5$ and $p_2 > 1$

Figure 5-3: Removal of paths that are too costly for **P1**.

we derived in this thesis to the graphs whose subgraph, obtained by removing the paths that are too costly, satisfies (A1).

5.2 Transportation Cost

In this thesis, we solved game Γ for a restricted class of graphs (the ones satisfying (A1)). It turns out the characterization of the NE we gave in Chapter 3 is also valid for a larger class of graphs satisfying the following weaker assumption:

Assumption 2. *There exists an optimal solution of* $(\mathcal{P}_2)$ *denoted* $x^* \in \Omega_2$, *and there exists a min-cut set* $(e_1, \ldots, e_N)$ *with* $\alpha_k := \min_{\{\lambda\in\Lambda \,|\, e_k\in\lambda\}} \sum_{(i,j)\in\lambda} b_{ij}$, $\forall k \in [\![1, N]\!]$, *then for every* $k \in [\![1, N]\!]$, *all* $s-t$ *paths taken by* x^* *that go through* e_k *have identical marginal transportation cost* α_k, *i.e.,*

$$\exists\, x^* \in \Omega_2,\, \exists\, \textit{min-cut set}\ (e_1, \ldots, e_N) \mid \forall k \in [\![1, N]\!],\ \forall \lambda \in \Lambda \mid e_k \in \lambda :$$
$$x^*_\lambda > 0 \implies \sum_{(i,j)\in\lambda} b_{ij} = \alpha_k.$$

In contrast to (A1), this assumption considers a class of graphs whose min-cost max-flow takes paths that have different marginal transportation cost. Note that in (A2), when $\alpha_1 = \cdots = \alpha_N \equiv \alpha$, we find again (A1). Below is an instance of a graph that does not satisfy (A1) but satisfies (A2).

Example 10. Consider the network flow problem in Fig. 5-4. There is a unique min-cost max-flow x^*, which carries 1 unit of flow through paths $\{s, 2, 4, t\}$, $\{s, 2, 3, t\}$ and $\{s, 1, t\}$. Thus, the total amount of flow is equal to 3 units. However, these paths induce different transportation costs, therefore (A1) is not satisfied.

Let us note $e_1 = (s, 1)$, $e_2 = (2, 3)$ and $e_3 = (4, t)$, then (e_1, e_2, e_3) is a min-cut set. One can check that $\alpha_1 = 2$, $\alpha_2 = 3$ and $\alpha_3 = 4$. For $k \in \{1, 2, 3\}$, the paths taken by x^* that go through e_k induce a transportation cost equal to α_k. Thus, (A2) is satisfied and one can check that the results presented in Chapter 3 also apply to the game Γ defined on this network.

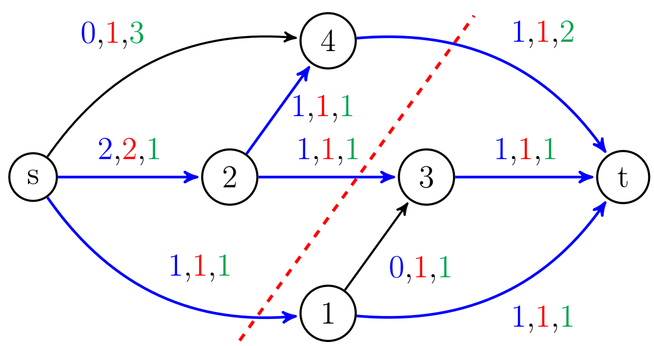

Figure 5-4: Min-cost max-flow (bold blue) and min-cut set (dotted red) of a graph satisfying (A2).

5.3 Cost of Attack

One of the main assumptions of the model is that the cost of attacking an edge is proportional to its capacity (ref. (2.4)). In this section, we want to investigate through an example the implications of considering the case of a more general cost of attack. Let us consider the graph given by Fig. 5-5.

Once again, one can show that $\tilde{\sigma}$ from (3.4) and (3.5) is not a NE anymore.

Let x^1 be the flow that sends 1 unit through path $\{s, 1, t\}$, x^2 be the flow that sends 1 unit through path $\{s, 2, t\}$, and $\mu' = \mathbb{1}_{\{(1,t),(2,t)\}}$ be the attack that disrupts edges $(1, t)$ and $(2, t)$. Then, one can show that when $p_1 > 2$ and $p_2 > 3$, $(\sigma^{1^*}, \sigma^{2^*})$ defined by $\sigma^{1^*}_{x^0} = 1 - \frac{3}{p_2}$, $\sigma^{1^*}_{x^1} = \frac{1}{p_2}$, $\sigma^{1^*}_{x^2} = \frac{2}{p_2}$ and $\sigma^{2^*}_{\mu^0} = \frac{2}{p_1}$, $\sigma^{2^*}_{\mu'} = 1 - \frac{2}{p_1}$ is a NE. This NE is illustrated in Fig. 5-6.

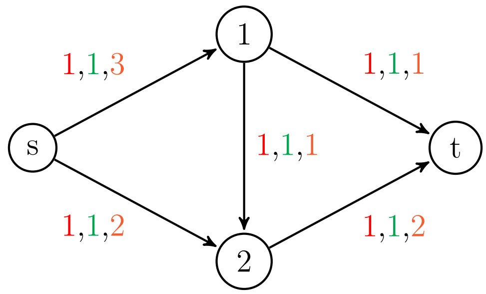

Figure 5-5: Graph with general cost of attack. Edge capacities, transportation cost and cost of attack are denoted in red, green and orange colored labels respectively.

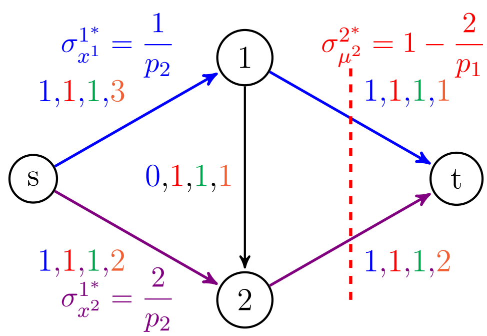

Figure 5-6: NE in the case $p_1 > 2$ and $p_2 > 3$ for a graph with general cost of attack.

In this case, one may notice that $\{(1,t),(2,t)\}$ is the cut-set of the graph that induces the smallest cost of attack. Hence, according to this NE, we conjecture that in the general cost of attack case, the notion that generalizes the min-cut set is the cut-set that induces the smallest attacking cost. Indeed, in our model, since the attacking cost was proportional to the edge capacity, then the min-cut sets were the cut-sets that induced the smallest cost of attack.

Regarding **P1**'s strategy, we can see that in this case, she routes x^2 twice as frequently as x^1, most likely because the cost of attacking $(2,t)$ is twice as the cost of attacking $(1,t)$. This result differs from the NE we found in our model: each path taken by the min-cost max-flow was taken with the same probability, $\frac{1}{p_2}$, at equilibrium.

Chapter 6

Conclusion

6.1 Summary of the Results

In this thesis, we considered a simultaneous network flow game between a defender and an attacker. Using linear programming duality for zero-sum games and network optimization ideas such as the Max-Flow Min-Cut Theorem, we gave structural insights on the set of Nash equilibria and we computed in closed form certain physical quantities of interest in equilibrium. Specifically, we showed that each player has a unique payoff value in all Nash equilibria and we analytically computed the value of effective (resp. lost) flow and the cost of transportation (resp. cost of attack) in terms of the parameters of the game. Then, we used graph theoretic properties of the network to provide a characterization of the support of the equilibria and relate it to the solutions of two classical network optimization problems that are the minimum cost maximum flow problem and the minimum cut problem.

Lastly, we studied a generalization of the game where both players face budget constraints and we looked for the minimum budget that players must have so the structural results we presented in the unconstrained case extend to the budget-constrained game. Using the infiniteness of the defender's set of actions, we computed a tight lower bound for the defender's budget for transporting flows. Unfortunately, an analogous lower bound could not be derived for the attacker's budget due to the finiteness of her set of actions. Therefore, we computed in closed form a large number of equilibrium strategies for the attacker in the unconstrained game (based on the partitions of the minimum cut sets). Then, using an integer programming problem, we found the partition-based equilibrium strategies that require the lowest budget, thus providing a bound on the attacker's budget for which the analysis in the unconstrained case holds for the budget-constrained game.

6.2 Future Work

In this thesis, we assumed that the value of effective flow was linear in the amount of effective flow and the edge transportation cost was linear in the amount of edge flow (see (2.3)). However, many network flow problems have more complex losses [5, 7, 22, 23]. Thus, it would be interesting to extend our model to the case where the edge transportation cost is $b_{ij}x_{ij}^{\gamma}$ with $\gamma > 0$.

Another extension of our work is to take into account the reliability failures of the edges. Using a stochastic modeling of the edge capacities [19], we could model an attacker-defender game on a network subject to reliability failures. It would also be interesting to analyze how the equilibria depend on the distribution of reliability failures.

Finally, we could use the model derived in this thesis and study it in a repeated setting with imperfect information, i.e., one player has more information of the state than the other. While playing the game repeatedly, the uninformed player could study the strategy of the informed player in order to learn the hidden information.

Bibliography

[1] Daron Acemoglu, Azarakhsh Malekian, and Asuman Ozdaglar. Network security and contagion. Working Paper 19174, National Bureau of Economic Research, June 2013.

[2] Ilan Adler, Constantinos Daskalakis, and Christos H. Papadimitriou. A note on strictly competitive games. In Stefano Leonardi, editor, *Internet and Network Economics, 5th International Workshop, WINE 2009, Rome, Italy, December 14-18, 2009. Proceedings*, volume 5929 of *Lecture Notes in Computer Science*, pages 471–474. Springer, 2009.

[3] Kevin Wood Alan Washburn. Two-person zero-sum games for network interdiction. *Operations Research*, 43(2):243 – 251, 1995.

[4] Rudolf Avenhaus and Morton John Canty. Inspection games. In Robert A. Meyers, editor, *Encyclopedia of Complexity and Systems Science*, pages 4855–4868. Springer, 2009.

[5] Jayanth R. Banavar, Francesca Colaiori, Alessandro Flammini, Amos Maritan, and Andrea Rinaldo. Topology of the Fittest Transportation Network. *Physical Review Letters*, 84(20):4745–4748, May 2000.

[6] Melike Baykal-Gürsoy, Zhe Duan, H. Vincent Poor, and Andrey Garnaev. Infrastructure security games. *European Journal of Operational Research*, 239(2):469–478, 2014.

[7] Dimitri P. Bertsekas, Lazaros C. Polymenakos, and Paul Tseng. An ϵ-relaxation method for separable convex cost network flow problems. *SIAM J. on Optimization*, 7(3):853–870, March 1997.

[8] Dimitris Bertsimas, Ebrahim Nasrabadi, and James B. Orlin. On the power of randomization in network interdiction. *CoRR*, abs/1312.3478, 2013.

[9] Dimitris Bertsimas, Ebrahim Nasrabadi, and Sebastian Stiller. Robust and adaptive network flows. *Operations Research*, 61(5):1218–1242, 2013.

[10] Kelly J. Cormican, David P. Morton, and R. Kevin Wood. Stochastic Network Interdiction. *Operations Research*, 46(2), 1998.

[11] Mathieu Dahan and Saurabh Amin. Network flow routing under strategic link disruptions. In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing, Allerton Park & Retreat Center, Monticello, IL, USA*, September 2015.

[12] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

[13] Jack Edmonds and Richard M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *J. ACM*, 19(2):248–264, April 1972.

[14] L. R. Ford and D. R. Fulkerson. Maximal flow through a network. *Canadian Journal of Mathematics*, 8:399–404, 1956.

[15] Andrew V Goldberg and Robert E Tarjan. Finding minimum-cost circulations by canceling negative cycles. *Journal of the ACM (JACM)*, 36(4):873–886, 1989.

[16] Sanjeev Goyal and Adrien Vigier. Attack, defence, and contagion in networks. *Review of Economic Studies*, 81(4):1518–1542, 2014.

[17] Assane Gueye, Jean C. Walrand, and Venkat Anantharam. Design of network topology in an adversarial environment. In Tansu Alpcan, Levente Buttyán, and John S. Baras, editors, *Decision and Game Theory for Security - First International Conference, GameSec 2010, Berlin, Germany, November 22-23, 2010. Proceedings*, volume 6442 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2010.

[18] Sunghoon Hong and Myrna Wooders. Strategic Network Interdiction. Vanderbilt University Department of Economics Working Papers 1010, Vanderbilt University Department of Economics, June 2010.

[19] Fernando A. Kuipers, Song Yang, Stojan Trajanovski, and Ariel Orda. Constrained maximum flow in stochastic networks. In *Proceedings of the 2014 IEEE 22nd International Conference on Network Protocols*, ICNP '14, pages 397–408, Washington, DC, USA, 2014. IEEE Computer Society.

[20] Aron Laszka and Assane Gueye. Quantifying network topology robustness under budget constraints: General model and computational complexity. In Sajal K. Das, Cristina Nita-Rotaru, and Murat Kantarcioglu, editors, *Decision and Game Theory for Security - 4th International Conference, GameSec 2013, Fort Worth, TX, USA, November 11-12, 2013. Proceedings*, volume 8252 of *Lecture Notes in Computer Science*, pages 154–174. Springer, 2013.

[21] John Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100(1):295–320, 1928.

[22] Ignacio Rodríguez-Iturbe and Andrea Rinaldo. *Fractal river basins: chance and self-organization.* Cambridge University Press, 2001.

[23] Joseph P. Straley and Stephen W. Kenkel. Percolation theory for nonlinear conductors. *Phys. Rev. B*, 29:6299–6305, Jun 1984.

[24] W.Y. Szeto. Routing and scheduling hazardous material shipments: Nash game approach. *Transportmetrica B: Transport Dynamics*, 1(3):237–260, 2013.

[25] R. Kevin Wood. Deterministic network interdiction. *Mathematical and Computer Modelling*, 17(2):1 – 18, 1993.