The concavity of atomic splittable congestion games with non-linear utility functions

DARRELL HOY, Northwestern University

Classical work in network congestion games assumes networks are deterministic and agents are risk-neutral. In many settings, this is unrealistic and players have more complicated preferences. When driving to work in the morning, a commuter may prefer a safer route, rather than the faster but riskier route. A website sending out streaming video packets may not care about packets once they are late or derive much benefit from packets arriving much earlier, but would rather prefer a more consistent delivery model.

We consider the atomic-splittable setting and model these preferences in two ways: when players have non-linear preferences over (i) the delay on every path, and (ii) on the total delay they experience over all paths. We ask - when are these games concave?

In the risk-neutral setting, the concavity of the setting underlies many results, including the existence of pure Nash equilibria. In setting (ii), when players have preferences over the total cost seen, the game is concave and pure Nash equilibria will always exist. In setting (i) however, we show that the game is no longer concave, and as a result we no longer know if pure Nash equilibria always exist. In both of these settings, we show that we can reduce questions about them in the stochastic setting to questions about them in the deterministic setting.

1. INTRODUCTION

Congestion games come up often in practice - they represent our drives to work in the morning, the performance of factories under load, and the routing of packets through computer networks.

In addition, all of these instances involve uncertainty. In the automobile traffic model, this uncertainty can come from the random assortment of drivers on the road at one time, but also the random interactions of wildlife on the road; the chances of accidents; the weather. In computer networks, links can become overloaded and see quality of service degrade, both wired and wireless.

Agents can respond to this uncertainty in different ways. When trying to get to work in time for an important meeting, a commuter might adjust her route to take a perhaps slower, but safer route. If a company is relying on having components in their supply chain ready for a product launch, they may well try and split the load across multiple factories on the hope that at least one will be able to get things done in time (a diversification effect).

We consider the *atomic-splittable* setting, where each player can divide up their flow among all possible paths. These models are useful for modeling computer networks, among others, where the individual units of flow are nearly infinitesimal, but players can control perhaps large chunks of the flow.

Our players then can either be risk-averse in their total cost, across all flow, or risk averse in the cost that any particular unit or piece of flow experiences. A coalition of commuters (as considered in [Cominetti et al. 2009]) may be more concerned with their individual risk aversion. Packets from a streaming video may be useless after a certain time, while a site may be less concerned with the exact arrival of bits for a large file download. A baker delivering loaves of fresh bread might be concerned with the time that each loaf takes to get to market. On the other hand, a supply chain with non-biodegradable parts may be more concerned with the total cost across all goods produced (or, equally, the average cost).

We model the non-linear utility over arrival times with increasing cost functions. We do this in two ways: with *per-path* cost functions, and with *total delay* cost functions. In the former, a player defines the cost of receiving flow at a time in the future and then tries to minimize that total cost. A loaf of fresh bread may be worth $4, while a loaf of

Table I: Differences between linear and non-linear cost function models.

	Concave game in parallel-link graphs	Concave game in general graphs	Stoch. Diversification Benefits (A.2.1)	Path/Edge Decomposition Difference (A.2.2)
Linear costs	Y	Y	N	N
Per-path convex costs	Y	N	N	Y
Total convex costs	Y	Y	Y	N

day-old bread may only be worth $1. Per-path costs can represent this drop in value. Total delay cost functions on the other hand better represent the total cost across all goods if there is no decrease in value of a certain good if it takes a little longer to arrive. In addition, we assume that each player has the same cost function for all of their flow.

We model a stochastic congestion game as a distribution over deterministic games. This gives us the flexibility to draw direct links between the deterministic setting and stochastic setting. In particular, when the deterministic version of a game is concave, the overall game is concave, and results relying on that will carry over.

1.1. Our Contributions

We present two models of per-user aggregate cost in network congestion games: (i) when players have per-path cost functions, and (ii) when the agents have cost functions over their total delay.

We characterize when games in both of these settings will be concave and thus have guaranteed existence of pure Nash equilibria. In particular, in setting (i), we show that when agents have per-path cost functions these games are not always concave, even when the cost function is convex (or the utility function is concave).

1.2. Prior Work

Rosenthal [1973] introduced the notion of a *congestion game*, a generalization of network flow games which abstracts the graph. He showed directly that these games were concave, and thus always featured pure Nash equilibria.

Nikolova & Stier-Moses [2011] look at nonatomic and atomic unsplittable games with added-delay independent stochasticity over the edges. In their setting, the agents have preferences over combinations of the mean and variance of their expected transit time, rather than the expected utility of our setting. They show that pure equilibria exists in the nonatomic setting, and in the atomic unsplittable setting only if the delay is *exogenous*, and the distribution is just shifted with the amount of flow on the link.

Ordonez and Stier-Moses [2010] consider the nonatomic setting, with added, random noise on edges, with players seeking to either a) minimize the ith percentile of expected traffic time, b) pad their transit time with a multiplicative factor of expected travel time or c) try to limit the number of edges which will be far away from their average. They show that pure equilibria exist in each of these instances. Nie [2011] explores ways of finding the equilibrium in the nonatomic setting, with heterogeneous users trying to minimize the ith percentile of expected travel time.

A number of papers have looked at atomic splittable routing games in the deterministic setting. The question of the welfare loss due to selfish behavior (Price of Anarchy) has been explored in [Cominetti et al. 2009; Roughgarden 2009; Roughgarden and Schoppmann 2011]. Bhaskar et al. [2009] addressed the question of uniqueness of equilibria in such settings, showing the multiple equilibria can exist in games with 3 or more players, or in graphs with two types of players that do not fall in type topological classes.

A number of papers have looked at *weighted* congestion games [Milchtaich 1996, 2009; Harks and Klimm 2010], in which players affect the total delay in different ways, and often have their own per-resource, separable cost functions. This differs principally from what we consider in that the cost in such games is still per-resource, as opposed to aggregate costs across congestion.

1.3. Structure of this Paper

In Section 2, we introduce our models of per-path and total-cost function games. In Section 3, we consider exactly when these games will be concave. In Section 4, we consider when these games will have pure equilibria. In Section 5, we conclude.

2. MODELS

2.1. Atomic-splittable preliminaries

An *atomic-splittable network congestion game* is a k-player game defined over a graph $G = (V, E)$, where each player controls flow of v_i and is routing between nodes s_i and t_i in the graph. Each player can then split their flow across all valid $s_i - t_i$ paths. Associated with every edge is a latency function $l_e(x)$, denoting the delay of the edge at any level of congestion. We'll use $\mathscr{L}$ to denote the set of allowable such latency functions. Usually, these will be increasing and *semi-convex*, where semi-convex means $xl(x)$ is convex.

In a flow f, we'll let f_p^i denote the flow that player i sends along path p. Let $f_e^i = \sum_{p \ni e} f_p^i$, the total flow from player i that travels through edge e. Let f_p, f_e and f respectively represent the sum across all players of flow on a path, edge and the total graph. Let f_p^{-i}, f_e^{-i} and f^{-i} likewise represent the flow from all players except i on the respective resource. We'll allow players to have potentially two other types of cost functions - either a *per-path* or a *total* cost function $C_i()$. This denotes the actual cost to the player of experiencing the given delay through a path or the entire graph. In this paper, we'll only refer to having one or the other, hence we'll overload $\mathscr{C}$ to denote the set of allowable cost functions in whichever setting we are discussing.

2.2. Deterministic Models

Definition 2.1. An *atomic-splittable, per-path cost functions* network congestion game is an atomic-splittable network congestion game in which players have individual cost functions $C_i() \in \mathscr{C}$ over the cost of paths. Each player's total cost is given by $\sum_{p \in \mathcal{P}} f_p^i \cdot C_i \left(\sum_{e \in p} l_e(f_e^i + f_e^{-i}) \right)$, thus giving each player the following optimization problem:

$$\begin{aligned}\min_{f^i} \sum_{p\in\mathcal{P}} \quad & f_p^i \cdot C_i\left(\sum_{e\in p} l_e(f_e^i + f_e^{-i})\right)\\ s.t. \quad & \sum_{p\ni e} f_p^i = f_e^i \quad \forall i\\ & \sum_{p\in\mathcal{P}} f_p^i = v_i, \quad f_p^i \geq 0\\ & l_e \in \mathscr{L}\ \forall e, \quad C_i \in \mathscr{C}\ \forall i\end{aligned}$$

Definition 2.2. An *atomic-splittable, total cost functions* network congestion game is an atomic-splittable network congestion game in which players have cost functions $C_i() \in \mathscr{C}$ over the total flow they experience. Each player's total cost is given by $C_i\left(\sum_{p\in\mathcal{P}} f_p^i \cdot \sum_{e\in p} l_e(f_e^i + f_e^{-i})\right)$, thus giving each player the following optimization problem:

$$\begin{aligned}\min_{f^i} \quad & C_i\left(\sum_{p\in\mathcal{P}} f_p^i \sum_{e\in p} l_e(f_e^i + f_e^{-i})\right)\\ s.t. \quad & \sum_{p\ni e} f_p^i = f_e^i \quad \forall i\\ & \sum_{p\in\mathcal{P}} f_p^i = v_i, \quad f_p^i \geq 0\\ & l_e \in \mathscr{L}\ \forall e, \quad C_i \in \mathscr{C}\ \forall i\end{aligned}$$

2.3. Stochastic, Atomic-splittable Models

When we move to the stochastic setting, we'll be thinking of this as drawing a deterministic setting from some distribution. We'll specifically implement this by drawing a shared, arbitrary-dimensional random variable from some distribution that all delay functions see. If the edges all see their randomness drawn independently, this would be the $|E|$-dimensional product distribution across all of the individual distributions. Then, we enforce for any fixed draw that the game falls into whichever groups we are concerned with.

Definition 2.3. A *stochastic atomic-splittable, per-path cost functions congestion game* is a k-player game defined over a ground set of resources E, a set of acceptable sets of resources, $\mathcal{P}$, volumes of flow v_i controlled by each player, delay functions $l_e(x, r) \in \mathscr{L}$ on each resource, player specific cost functions $C_i() \in \mathscr{C}$ over the cost of paths, and an uncertainty distribution $\mathcal{R}$. Each players action space consists of all possible splits of v_i across paths in P. Each players total cost is given by $\sum_{p\in\mathcal{P}} f_p^i \cdot C_i\left(\sum_{e\in p} l_e(f_e^i + f_e^{-i}, r)\right)$, thus giving each player the following optimization problem:

$$
\begin{aligned}
\min_{f^i} \;\; E_{r\sim\mathcal{R}} \quad & \left[\sum_{p\in\mathcal{P}} f_p^i \cdot C_i\left(\sum_{e\in p} l_e(f_e^i + f_e^{-i}, r)\right)\right] \\
s.t. \quad & \sum_{p\ni e} f_p^i = f_e^i \quad \forall i \\
& \sum_{p\in\mathcal{P}} f_p^i = v_i, \quad f_p^i \geq 0 \\
& l_e(x,r) \in \mathscr{L} \; \forall e, r, \quad C_i \in \mathscr{C} \;\; \forall i
\end{aligned}
$$

Definition 2.4. A *stochastic atomic-splittable, total cost functions* game is similar to above, except each player has a cost function over the total delay that they experience in the graph.

$$
\begin{aligned}
\min_{f^i} \quad & E_{r\sim\mathcal{R}} \left[C_i\left(\sum_{p\in\mathcal{P}} f_p^i \sum_{e\in p} l_e(f_e^i + f_e^{-i}, r)\right)\right] \\
s.t. \quad & \sum_{p\ni e} f_p^i = f_e^i \quad \forall i \\
& \sum_{p\in\mathcal{P}} f_p^i = v_i, \quad f_p^i \geq 0 \\
& l_e(x,r) \in \mathscr{L} \; \forall e, r, \quad C_i \in \mathscr{C} \;\; \forall i
\end{aligned}
$$

See Section A.2 for an illustration of diversification effects and edge-flow decomposition in these models.

3. CONCAVITY OF THE GAMES

In this section, we'll discuss exactly when atomic-splittable network congestion games with user specific aggregate cost functions are concave. By concave, we mean the following:

Definition 3.1. An n-person game is called a *concave game* if each person has utility $\psi_i(x_i, x_{-i})$, where $x_i \in S_i$, S_i is a convex, closed, bounded set, and $\psi_i(x_i, x_{-i})$ is concave in x_i for fixed x_{-i}.

In a concave game, each player is effectively solving a convex optimization problem. Moreover, we know from Rosen [1965] that this is sufficient to prove existence of a pure equilibria.

The primary result of this section is that even in the deterministic setting, with convex cost functions, the per-path cost function model *will not* guarantee a concave game, while the total-cost function model is always a concave game with convex cost functions. These results are summarized in Table II.

In the traditional model, atomic-splittable flow games are concave so long as the delay functions are semi-convex. We have a concave game so long as the delay functions are semi-convex. This comes as a result of the fact that the total cost a player experiences reduces to minimizing $\sum_{e\in E} f_e^i \cdot l_e(f_e^i + f_e^{-i})$. So, when l_e is always semi-convex, this becomes a sum of convex functions, and thus the total summation is convex, giving a concave game. We'll now address each setting individually.

Table II: Guaranteed concavity of games with per-path vs total cost functions, by setting.

	Linear $\mathscr{C}$, semi-convex $\mathscr{L}$	Convex $\mathscr{C}$, semi-convex $\mathscr{L}$, parallel-link graphs	Convex $\mathscr{C}$, semi-convex $\mathscr{L}$, general graphs	Semi-convex $\mathscr{C}$, convex $\mathscr{L}$
Per-path costs	Y	Y	N	N
Total costs	Y	Y	Y	N

3.1. Total cost setting

In the total cost setting, the game is almost trivially concave when the total cost functions are convex. When edge latencies are semi-convex, the total delay experienced by a player is the same as in the standard, linear cost model: $\sum_p f_p^i \cdot \sum_{e\in p} l_e(f_e^i + f_e^{-i}) = \sum_e f_e^i l_e(f_e^i + f_e^{-i})$. This we know is convex in f^i. Since the composition of two convex functions is convex, the total cost experienced by the player, $C_i(\sum_p f_p^i \cdot \sum_{e\in p} l_e(f_e^i + f_e^{-i}))$ is also convex, and hence it is a concave game for any graph. When C is not convex, then using linear delay functions results in a concave function, hence the game would not be concave. Thus we have the following:

LEMMA 3.2. *Atomic-splittable flow games with player-specific convex total cost functions are concave games.*

3.2. Per-path costs, parallel-link graphs

We now move to the per-path cost setting, but restrict ourselves to parallel-link graphs. In a parallel-link graph, we have only two vertices and a number of edges between them. Thus, the problem a player faces can be simplified to:

$$\min_{f^i} \sum_{e\in\mathcal{P}} \quad f_e^i \cdot C_i\left(l_e(f_e^i + f_e^{-i})\right)$$
$$s.t. \quad \sum_{e\in\mathcal{P}} f_e^i = v_i, \quad f_e^i \geq 0$$

CLAIM 1. *For any convex C_i and semi-convex l_e, $C_i(l_e(x))$ is semi-convex.*

We'll argue directly that the second-derivative is always non-negative:

$$\begin{aligned}\partial^2(xC_i(l_e(x)))/\partial x^2 &= 2C_i'(l_e(x))l_e'(x) + xC_i''(l_e(x))l_e'(x)^2 + xC_i'(l_e(x))l_e''(x) \\ &= xC_i''(l_e(x))l_e'(x)^2 + C_i'(l_e(x))\left(2l_e'(x) + xl_e''(x)\right)\end{aligned}$$

As l is semi-convex, we have $\partial^2(xl_e(x))/\partial x^2 = 2l_e'(x) + xl_e''(x) \geq 0$. The other quantities are non-negative due to the convexity and monotonicity of C_i hence the above quantity is positive, and $xC_i(l_e)$ is convex, hence $C_i(l_e)$ is semi-convex.

LEMMA 3.3. *Atomic splittable flow games in parallel-link graphs with player specific per-path cost functions are concave games.*

$$v = 10$$
$$C(x) = \begin{cases} 0 & \text{if } x \le 10, \\ x - 10 & \text{if } x > 10. \end{cases}$$

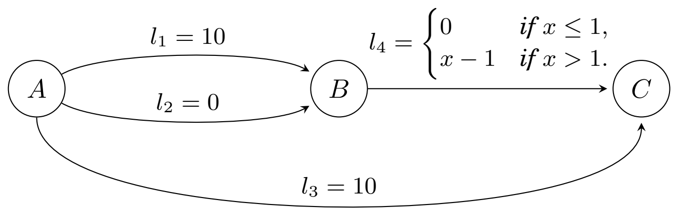

Fig. 1: An example of a non-concave game.

PROOF. This now follows directly from our above claim. For each edge, $f_e^i \cdot C_i(l_e(f_e^i + f_e^{-i})$ is convex in player i's complete flow vector f^i, hence her total cost is convex in f^i. Thus the total cost is convex, and the game is concave. □

3.3. Per-path costs, general graphs

We now move to the general graph setting. We'll show here that these games need not be concave, even for a single player. Specifically, this will show up when moving flow away from a path p, but moving even more flow to a path q, which shares edges with p. Thus, even though we are removing flow from p, the cost of p increases.

Example 3.4. Consider the setting of one player routing flow of volume 10, from A to B in the network in Figure 1. Imagine she is indifferent between arrivals until time 10, and sees linear cost thereafter. There are three possible paths for her - call them H, M, L from top to bottom. If she sends a little flow through A and a lot through C, all flow will arrive at time 10. If she sends all flow through the middle path, all that flow will arrive before time 10. Since she is indifferent until time 10, she is fine with both of these and sees a cost of 0. However, if she splits between those, there will more than one unit of flow on edge 3, hence the top path will take more than 10 and she experiences cost.

Consider specifically the two flows $f = (1, 0, 9)$, $g = (0, 10, 0)$. In both, every path sees delay of 10 or less, since at most one unit of flow is routed over edge 4. Hence $C(f) = C(g) = 0$. However, consider the flow $\frac{f+g}{2} = (0.5, 5, 4.5)$. Now edge 4 sees 5.5 units of flow, hence has a delay of 4.5, giving the flow across the top path a total delay of 14.5, hence $C(\frac{f+g}{2}) = 0.5 \cdot C(14.5) = 2.75$. Thus, the cost function cannot be convex and games of this sort are not always concave games.

In a sense though, this example relies on the player deciding not to just take the free edge, which is strictly better than the edge of cost 10. We now illustrate a slightly more complicated example in which this is not the case.

Example 3.5. Consider a user of volume 1, routing through the graph shown in Figure 2. Let her be indifferent between arriving until time 1, at which point her cost increases linearly. Now, let the delay on edges 1 and 5 be 1; let edge 3 be free and let both edges 2 and 4 be free up to traffic of $1/2$, and $x - 1/2$ thereafter. Now, either flow of $(0.5, 0, 0.5)$ and $(0, 1, 0)$ both give a total cost of 0, since every path used sees a latency of 1 and hence a cost of 0. However, moving a little bit of flow from any of those paths will increase the latency on one of the paths above 1, hence there are no local changes she can do from either position to improve her situation.

3.4. Stochastic Games

Now, we'll show that stochastic versions of the concave games already discussed are also concave. Our first result is a simple one - that if a stochastic game can be inter-

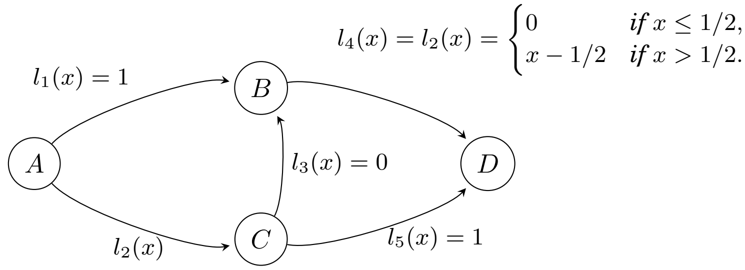

$$
\begin{aligned}
v_1 &= 1 \\
C_1(x) &= \begin{cases} 0 & \text{if } x \leq 1, \\ x - 1 & \text{if } x > 1. \end{cases}
\end{aligned}
$$

Fig. 2: Player can either put flow through middle path, or split flow along the top and bottom paths, but any flow in between yields path latencies above 1, hence a non-zero cost. Thus this is not a concave game.

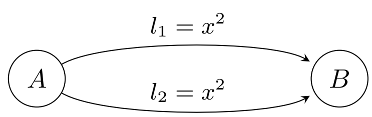

$$
\begin{aligned}
v_1 &= 1.4 \\
v_2 &= 0.2 \\
C_1(x) &= 10 - \frac{1}{x + 0.1} \\
C_2(x) &= x
\end{aligned}
$$

Fig. 3: Nonexistance of equilibria with non-convex, per-path cost functions

preted as a distribution over concave games with the same strategy space, the stochastic game is also concave.

LEMMA 3.6. *Let $\mathscr{G}$ be a set of concave, deterministic, n player games, where each game G shares the same closed, convex and bounded strategy space S_i for each player i, and each game G gives each player i utility $\psi^G(x^i, x^{-i})$. Let F be a distribution over games in $\mathscr{G}$. Let Q be a stochastic game with per player utility $\psi(x^i, x^{-i}) = E_{G \sim F}\left[\psi^G(x^i, x^{-i})\right] \ \forall x^i, x^{-i}$. Then Q is a concave game.*

This follows because any weighted combination of concave functions will be concave. So, if every time we fix the coin flips for the randomness in our universe the game is concave, then the total game is concave. Combining this with the concavity of the deterministic total-cost function game, shown in Theorem 3.2, we get the following:

LEMMA 3.7. *Stochastic atomic splittable games with convex total-cost functions are concave games.*

4. PURE EQUILIBRIA EXISTENCE

We now consider when and where pure equilibria exist. First, all classes of games that we have just shown to be concave exhibit pure equilibria by [Rosen 1965]. Now, we argue that this is roughly tight. In the per-path cost function case, we show that even in parallel-link games, when players exhibit non-convex cost functions, pure equilibria need not exist.

Example 4.1. Consider the parallel two edge graph shown in Figure 3 with one large player and lots of small players. Specifically, the large player will control flow of volume 1, and the non-atomic players will control flow of volume 0.3. The large player will have a cost function $C_1(x) = 10 - \frac{1}{x+0.1}$ - the non-atomic players will have standard, linear preferences.

First, we claim that at equilibrium, either (a) both edges will have the same traffic, or (b) all the non-atomic users will be on the cheaper edge. Otherwise, some of the non-atomic users on the pricer path would just move to the cheaper one.

Consider the first setting, where the delay and hence the total flow on each edge is the same. Thus, each edge has flow of volume 0.8 on it, and player 1 sees cost $1.4 \cdot C_1(0.8^2) \approx 12.108$. WLOG, let edge 1 be the edge on which player 1 sends at least half her flow. Now, we'll show that she always prefers to moving 0.2 units of flow from edge 1 to edge 2. Let g be his flow after the shift, hence $g_1^1 = f_1^1 - 0.2 \geq 0.5$, $g_2^1 = f_2^1 + 0.2 \leq 0.9$. Then, with flow g, player 1 sees costs:

$$\begin{aligned} g_1^1 \cdot C_1(0.6^2) + g_2^1 \cdot C_1(1^2) \;&\leq\; 0.5 \cdot C_1(0.6^2) + 0.9 \cdot C_1(1^2) \\ &\leq\; 12.095 \end{aligned}$$

This is less than at the middle, so the larger player will always want to push the flow away from being evenly balanced. Thus there will be no evenly balance equilibria.

Now, consider the case when the edges are *not* balanced, and all non-atomic users are on the same edge - call it edge 1- which thus has less than 0.8 total flow on it. We'll now show that the larger player would prefer to push extra flow from edge 2 to edge 1. Call the large player B, call s the aggregate of the small players. Call the total flow f. Then by our assumptions, we have:

$$\begin{aligned} f_1 \;&<\; 0.8 < f_2 \\ f_1^s \;&=\; 0.3 \\ f_1^B \;&\leq\; 0.5 \end{aligned}$$

Now, let's construct a better flow g. In this, we'll have player B move enough flow from 2 to 1 so that the volumes on each edge switch. So, we'll have $g_1 = f_2$ and $g_2 = f_1$. Now, edge 2 becomes the cheaper edge. However, player B has more flow on 2 than he had on 1 before, hence her cost is lower. Hence, she prefers g to f. Thus, f cannot be an equilibrium.

Thus, in the game we set up, there can be no pure equilibrium.

5. CONCLUSION

In this paper, we sought to understand when basic results from deterministic, risk-neutral network flow game carry over to the more general stochastic and potentially risk averse case. We've focused on the concavity of the game because this underlies many of the known results for graphs in the risk-neutral setting. We think that understanding this question is fundamental to understanding how robust results about network routing games are to uncertainty and non-linear preferences in the real world.

To model the more complicated setting, we introduced two logical models of player-specific utility over path delay and total cost in atomic splittable flow games. When these games are concave, then we get many results for free, such as pure Nash existence. However, when players have preferences over delays on a path, the game is no longer concave, and thus we do not know if pure Nash always exist. With the advent of larger and larger fractions of internet traffic being used for streaming video from large websites, it is important to understand how such non-linear preferences can affect behavior in network flow games.

5.1. Future Work

We've hinted throughout at one large open problem - whether or not pure Nash equilibria exist when players have per-path cost functions.

Beyond Nash existence, a natural question arises of the uniqueness of equilibria. The results from Bhaskar et al.[2009] should carry over in the total-cost setting, but the per-path model will likely complicate matters.

We'd also like to understand how these preferences affect the social welfare of equilibria and the Price of Anarchy. We can use our formulation to show that robust price of anarchy bounds proven on the deterministic models shown carry over to the stochastic models considered (see this straightforward argument in Section A.1), but we don't have any results for these bounds in the deterministic models.

APPENDIX

A.1. Price Of Anarchy

In this section, we consider the effect of per-player cost functions on price of anarchy results. We consider the framework of (λ,μ)-smoothness, introduced in [Roughgarden 2009]. We show that if the deterministic per-path or total cost flow game is (λ,μ)-(locally) smooth, then the game remains (λ,μ)-(locally) smooth in the stochastic setting. Note however that this is still one step away from a direct reduction from present bounds - we say nothing about how smooth the deterministic per-path or total cost games are.

Recall the definition of (λ,μ)-smoothness:

Definition A.1. [Roughgarden 2009] A cost minimization game is (λ,μ)-smooth if $\sum_i^n c_i(y^i,x^{-i}) \leq \lambda \cdot SC(y) + \mu \cdot SC(x)$, for any pair of actions x, y, where SC indicates the social cost function.

Definition A.2. [Roughgarden and Schoppmann 2011] A cost minimization game is *locally* (λ,μ)-smooth if for any pair of actions x, y:

$$\sum_i^n [c_i(x) + \nabla_i c_i(x)^T (y^i - x^i)] \leq \lambda \cdot SC(y) + \mu \cdot SC(x) \tag{1}$$

Now, we argue that if you prove that the deterministic setting of any of the earlier games is (λ,μ)-(locally) smooth, so too will be a distribution over those games.

LEMMA A.3. *Let $\mathscr{G}$ be a set of (λ,μ)-(locally) smooth, concave, deterministic, n player games, with cost functions $c^G(x^i,x^{-i})$ for each player i, in each game G. Let Q be a stochastic game exhibiting per player costs $c_i(x^i,x^{-i}) = E_{G\sim F}\left[c^G(x^i,x^{-i})\right]\ \ \forall x^i, x^{-i}, G \in \mathscr{G}$. Then Q is (λ,μ)-(locally) smooth.*

PROOF. Follows by straightforward linearity of expectations. First, if it holds for each game in particular for any x and y, then it holds in expectation across all games:

$$\begin{aligned}
E_{G\sim F}\left[\sum_i^n c_i^G(y^i,x^{-i})\right] &\leq E_{G\sim F}\left[(\lambda \cdot SC(y,G) + \mu \cdot SC(x,G))\right] \\
\sum_i^n c_i(y^i,x^{-i}) &\leq E_{G\sim F}\left[(\lambda \cdot SC(y,G) + \mu \cdot SC(x,G))\right] \\
\sum_i^n c_i(y^i,x^{-i}) &\leq \lambda E_{G\sim F}\left[SC(y,G)\right] + \mu E_{G\sim F}\left[SC(x,G)\right] \\
\sum_i^n c_i(y^i,x^{-i}) &\leq \lambda \cdot SC(y,I) + \mu \cdot SC(x,I)
\end{aligned}$$

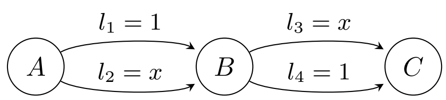

path	f_p	g_p
(1, 3)	1/2	0
(1, 4)	0	1/2
(2, 3)	0	1/2
(2, 4)	1/2	0
Cost:	2.25	2.5

Fig. 4: With per-path function of $C(x) = x^2$, not all path flows that induce the same edge flow are the same - sending half the flow along the top and half along the bottom is better than top-bottom and bottom-top.

□

In the locally smooth case, the gradient of a players cost will simply be the expectation of the gradient, and hence we'll have exactly the same outcome as above.

A.2. Model Differences

In this section, we discuss a few other differences between our models. In particular, whether players see diversification benefits (like in the portfolio optimization problem), and the relationship between optimal edge flows and path flows.

A.2.1. Diversification effects. Our two models have very different behavior when players are splitting flow across multiple paths. In the total-cost function model, players see diversification benefits, whereas they do not in the per-path cost function model.

As an example, consider a network of two parallel links, the first with delay drawn from $U[1, 2]$, the second from $U[1, 2.1]$. Consider a player that can split her flow across these edges. In the deterministic setting, she just chooses the faster path and runs with it. Now, imagine that she controls flow of volume 1, and has a net cost function of $C(x) = x^3$ over the total delay experienced across his paths. Call f_1 the flow he sends on edge 1. Then, she has the following optimization problem:

$$\min_{0 \le f_1 \le 1} \quad E_{l_1 \sim U[1,2], l_2 \sim U[1,2.1]} \left[C(f_1 \cdot l_1 + (1 - f_1) \cdot l_2) \right]$$

$$\min_{0 \le f_1 \le 1} \quad \int_1^2 \int_1^{2.1} \frac{1}{1.1} (f_1 \cdot l_1 + (1 - f_1) \cdot l_2)^3 \, dl_2 \, dl_1$$

Her best action here is to split her flow, 75% on the top edge, 25% on the bottom. So, even though edge 2 is first-order stochastically dominated by edge 1, it still helps our user to diversify her exposure to the uncertainty of edge 1.

A.2.2. Edge and Path Decompositions. When players have per-path cost functions, different path breakdowns can affect the cost the player sees. This was shown in [Nikolova and Stier-Moses 2011] for stochastic settings with mean-risk utility functions - we'll be able to use a few of their following results directly here.

CLAIM 2. *For atomic-splittable, per-path cost function congestion games, not all path decompositions of a flow given as edge delay give the same utility.*

Consider a single player of volume 1 with a per-path cost function $C(x) = x^2$, routing across the graph in Figure 4. Imagine her sending half of his flow along each edge, and consider sending half along edges $(1, 3)$ and $(2, 4)$, vs $(1, 4)$ and $(2, 3)$.

$$\begin{aligned} C(f) &= 1/2(l_1(f_1)+l_3(f_3))^2 + 1/2(l_2(f_2)+l_4(f_4))^2 \\ &= (1+1/2)^2 \\ &= 2.25 \\ C(g) &= 1/2(l_1(g_1)+l_4(g_3))^2 + 1/2(l_2(g_2)+l_3(g_3))^2 \\ &= 1/2(1+1)^2 + 1/2(1/2+1/2)^2 \\ &= 2.5 \end{aligned}$$

Thus, our agent prefers path flow f to path flow g. Now - if we're given an edge flow for a given player and their per-path cost function, can we calculate their best path flow that corresponds to that edge flow? In the nonatomic setting, we have the following result:

THEOREM A.4. *[Nikolova and Stier-Moses 2011] For a nonatomic social optima given as the edge flows $(x_e)_{e \in E}$, there exists a succinct flow decomposition that uses at most $|E| + |K|$ paths.*

Note that although they generally consider a different model of uncertainty, because the path costs are fixed with flows, they only rely on each path having a cost, which is fixed for the purpose of the linear program used. So, for many players in the atomic splittable sense, we can solve the linear program once for each of them, fixing the network flow, giving the following result:

THEOREM A.5. *For an equilibrium of the atomic splittable congestion game with per-path costs, given as the edge flows for each player $((f_e^i)_i)_{e \in E}$, there exists a succinct flow decomposition that uses at most $|E| + |K|$ paths.*

REFERENCES

BHASKAR, U., FLEISCHER, L., HOY, D., AND HUANG, C. 2009. Equilibria of atomic flow games are not unique. In *SODA '09*. 748–757.

COMINETTI, R., CORREA, J., AND STIER-MOSES, N. E. 2009. The Impact of Oligopolistic Competition in Networks. *Operations Research 57,* 6, 1421–1437.

HARKS, T. AND KLIMM, M. 2010. On the existence of pure nash equilibria in weighted congestion games. In *ICALP '10*. 1–20.

MILCHTAICH, I. 1996. Congestion Games with Player-Specific Payoff Functions. *Games and Economic Behavior 13,* 1, 111–124.

MILCHTAICH, I. 2009. Weighted congestion games with separable preferences. *Games and Economic Behavior 67,* 2, 750–757.

NIE, Y. M. 2011. Multi-class percentile user equilibrium with flow-dependent stochasticity. *Transportation Research Part B: Methodological 45,* 10, 1641–1659.

NIKOLOVA, E. AND STIER-MOSES, N. E. 2011. A Mean-Risk Model for the Stochastic Traffic Assignment Problem. Tech. Rep. DRO-2011-03, Columbia Business School.

ORDÓÑEZ, F. AND STIER-MOSES, N. E. 2010. Wardrop Equilibria with Risk-Averse Users. *Transportation Science 44,* 1, 63–86.

ROSEN, J. B. 1965. Existence and Uniqueness of Equilibrium Points for Concave N-Person Games. *The Econometric Society 33,* 3, 520–534.

ROSENTHAL, R. 1973. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory 2,* 1, 65–67.

ROUGHGARDEN, T. 2009. Intrinsic robustness of the price of anarchy. In *STOC '09*. 513.

ROUGHGARDEN, T. AND SCHOPPMANN, F. 2011. Local Smoothness and the Price of Anarchy in Atomic Splittable Congestion Games. In *SODA '11*. 255–267.