

Flow-Based Network Creation Games

Hagen Echzell, Tobias Friedrich, Pascal Lenzner and Anna Melnichenko

Hasso Plattner Institute, University of Potsdam, Germany

Abstract

Network Creation Games (NCGs) model the creation of decentralized communication networks like the Internet. In such games strategic agents corresponding to network nodes selfishly decide with whom to connect to optimize some objective function. Past research intensively analyzed models where the agents strive for a central position in the network. These models agents optimizing the network for low-latency applications like VoIP. However, with today's abundance of streaming services it is important to ensure that the created network can satisfy the increased bandwidth demand. To the best of our knowledge, this natural problem of the decentralized strategic creation of networks with sufficient bandwidth has not yet been studied. We introduce Flow-Based NCGs where the selfish agents focus on bandwidth instead of latency. In essence, budget-constrained agents create network links to maximize their minimum or average network flow value to all other network nodes. Equivalently, this can also be understood as agents who create links to increase their connectivity and thus also the robustness of the network. For this novel type of NCG we prove that pure Nash equilibria exist, we give a simple algorithm for computing optimal networks, we show that the Price of Stability is 1 and we prove an (almost) tight bound of 2 on the Price of Anarchy. Last but not least, we show that our models do not admit a potential function.

1 Introduction

Many of the networks we crucially rely on nowadays have evolved from small centrally designed networks into huge networks created and decentrally controlled by many selfish agents with possibly conflicting goals. For example, the structure of the Internet is essentially the outcome of a repeated strategic interaction by many selfish economic agents [Papadimitriou, 2001; Tardos, 2004], e.g., Internet service providers (ISPs). Hence, such Internet-like networks can be analyzed by considering a strategic game which models the interaction of the involved agents. This insight has inspired researchers to propose game-theoretic models for

the decentralized formation of networks, most prominently the models in [Myerson, 2013; Jackson and Wolinsky, 1996; Bala and Goyal, 2000a; Fabrikant *et al.*, 2003] and many variants of them. In these models selfish agents select strategies to maximize their utility in the created network. Thus, the structure of the created network and corresponding agent utilities depend on the selfishly chosen strategies of all agents.

To the best of our knowledge, the utility of the involved agents in all the existing models depends on the agents' centrality or the size of their connected component and the cost spent for creating links. Striving for centrality, e.g., for low hop-distances to all other nodes in the created network, certainly plays an important role for the involved selfish agents, e.g., ISPs, but it is not the only driving force. What has been completely neglected so far are bandwidth considerations.

In this paper we take the first steps into this largely unexplored area of game-theoretic network formation with bandwidth maximization. In particular, we consider budget-constrained agents that strategically form links to other agents to maximize their flow value towards all other nodes. This models agents that optimize the networks for data-intensive applications like online streaming instead of low-latency applications like VoIP. Interestingly, by the Max Flow Min Cut Theorem [Ford and Fulkerson, 1956], flow maximization corresponds to connectivity maximization which is tightly related to network robustness. While there have been works which consider robustness aspects as side constraints, also the aspect of maximizing the created network's robustness is, to the best of our knowledge, entirely unexplored.

While classical network formation models with focus on centrality yield equilibrium networks that are sparse and have low diameter, our work shows that focusing on bandwidth/connectivity may explain why densely connected sub-networks, like k -core structures, and larger cycles appear in real-world networks: they are essential for its connectivity.

1.1 Model and Notation

We propose the *Flow-Based Network Creation Game* (*flowNCG*), which is a variant of the well-known Network Creation Game [Fabrikant *et al.*, 2003] where the agents create edges to maximize their flow value as defined by the classical Max-Flow Problem [Ford and Fulkerson, 1956]. Given a set of n selfish agents which corresponds to the set of nodes V of a weighted directed network $G = (V, E)$. Each agent

has a uniform budget of $k \in \mathbb{N}$, with $k < n$, to establish connections to other agents. In the created graph G each edge has a capacity $c : E(G) \rightarrow \mathbb{N}$. We define the *degree* of a node v in G as $deg_G(v) = \sum_{(x,v) \in E, x \in V} c(x,v) + \sum_{(v,y) \in E, y \in V} c(v,y)$, i.e., as the sum of the capacities of all incoming and outgoing edges.

The edge set and capacities are defined by the strategies of the agents. In particular, each agent strategically decides how to spend its budget, i.e., which subset of incident edges to buy and for each bought edge its capacity. Therefore, the strategy S_v of an agent v is a set of tuples $(x, c(v,x)) \in V \setminus \{v\} \times [k]$, where $[k] = \{1, \dots, k\}$. Here x denotes the node to which an edge is formed and $c(v,x)$ the capacity of that edge bought by agent v . For nodes $w \in V$ to which no edge is formed, we assume that $c(v,w) = 0$. A strategy S_v is feasible if the total capacity of all edges in S_v does not exceed the given budget k , i.e., $\sum_{x \in V \setminus \{v\}} c(v,x) \leq k$. If $(x, c(v,x)) \in S_v$, we call v the *owner* of the edge (v,x) with capacity $c(v,x)$. The vector of all agents' strategies $\mathbf{s} = (S_{v_1}, \dots, S_{v_n})$ is denoted as a *strategy profile*. Any strategy profile \mathbf{s} uniquely specifies the weighted directed network $G(\mathbf{s}) = (V, E(\mathbf{s}))$ where $(v,x) \in E(\mathbf{s})$ of capacity $c(v,x)$ if and only if agent v buys the edge to x with non-zero capacity $c(v,x)$.

For sending flow in the network $G(\mathbf{s})$, we assume its undirected version $F(\mathbf{s}) = (V, E_F(\mathbf{s}))$ where $\{v,x\} \in E_F(\mathbf{s})$ if $(v,x) \in E(\mathbf{s})$ or $(x,v) \in E(\mathbf{s})$ and the capacity $c(\{v,x\})$ is defined as $c(\{v,x\}) = c(v,x) + c(x,v)$. Thus, for sending flow, any edge of $G(\mathbf{s})$ can be used in both directions. See Figure 1 for an illustration of $G(\mathbf{s})$ and its corresponding undirected version $F(\mathbf{s})$. Note that throughout the paper

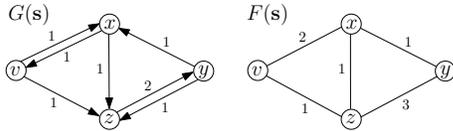


Figure 1: Graph $G(\mathbf{s})$ and its undirected version $F(\mathbf{s})$ for $\mathbf{s} = (S_v, S_x, S_y, S_z)$ with $S_v = \{(x,1), (z,1)\}$, $S_x = \{(v,1), (z,1)\}$, $S_y = \{(x,1), (z,1)\}$, $S_z = \{(y,2)\}$. Agent z has degree 5 and can send a flow of value 3, 3 and 4 to the agents v, x and y respectively.

we only work with $G(\mathbf{s})$ and we will always assume its undirected version $F(\mathbf{s})$ when computing flow values.

We call a directed weighted network G *feasible* if the sum over the weights of all incident outgoing edges for each node is at most k . Thus, we have a bijection between the set of feasible networks and the set of strategy profiles. Therefore, we will further omit \mathbf{s} in $G(\mathbf{s})$ if it is clear from the context and we will use G and \mathbf{s} interchangeably if $G = G(\mathbf{s})$.

Let $\lambda_{G(\mathbf{s})}(v,x) = \sum_{v_1 \in V_1, v_2 \in V_2, \{v_1, v_2\} \in E_F(\mathbf{s})} c(\{v_1, v_2\})$ be the capacity of a minimum v - x -cut (V_1, V_2) disconnecting node v from x in $F(\mathbf{s})$ which we also call the *local edge connectivity* between x and v . The minimum over all pairwise local edge connectivities in $G(\mathbf{s})$ is called the *edge connectivity* of $G(\mathbf{s})$ and is denoted as $\lambda(G(\mathbf{s}))$. A cut (V_1, V_2) in $G(\mathbf{s})$, where $V_1 \cup V_2 = V$ and $V_1 \cap V_2 = \emptyset$, such that $\lambda(G(\mathbf{s})) = \sum_{v_1 \in V_1, v_2 \in V_2, \{v_1, v_2\} \in E_F(\mathbf{s})} c(\{v_1, v_2\})$ is called a *minimum cut (min-cut)* of $G(\mathbf{s})$. Note that cuts in $G(\mathbf{s})$ are

always defined on its undirected version $F(\mathbf{s})$. By the Max Flow Min Cut Theorem [Ford and Fulkerson, 1956], the size of the min- v - x -cut equals the value of the max-flow between nodes v and x . Hence, we use the local edge connectivity of v and x interchangeably with the value of maximum v - x -flow.

We define two versions of the flowNCG. The first is called the *Average Flow NCG (avg-flowNCG)*. For a given strategy profile \mathbf{s} , the utility function of an agent v is defined as the average maximum flow value to each other node in the network: $u(v, \mathbf{s}) = \sum_{i \in V \setminus \{v\}} \frac{\lambda_{G(\mathbf{s})}(v,i)}{n-1}$. The *social utility* of network $G(\mathbf{s})$ is the average agent utility, i.e., $u(\mathbf{s}) = \sum_{v \in V} \frac{u(v, \mathbf{s})}{n}$.

The second variant is called the *Minimum Flow NCG (min-flowNCG)*. The utility of an agent in network $G(\mathbf{s})$ is the size of the min-cut of the network $\lambda(G(\mathbf{s}))$. Note that this value is the same for all agents in $G(\mathbf{s})$. We define a tie-breaking between strategies for an agent which yield the same min-cut of $G(\mathbf{s})$: agents try to maximize the number of nodes with local edge connectivity exceeding $\lambda(G(\mathbf{s}))$, i.e., to increase the number of highly robust connections. In particular, we define the agent's utility as the following vector: $u(v, \mathbf{s}) = (u_1(v, \mathbf{s}), u_2(v, \mathbf{s}))$, where $u_1(v, \mathbf{s}) = \lambda(G(\mathbf{s}))$ and $u_2(v, \mathbf{s}) = |\{i \in V \setminus \{v\} : \lambda_{G(\mathbf{s})}(i,v) > \lambda(G(\mathbf{s}))\}|$. We call the second component of the utility function $u_2(v, \mathbf{s})$ as the number of *well-connected nodes*. We assume that agents aim for maximizing their utility vector lexicographically. We define the social utility of a network in the min-flowNCG as the edge connectivity of the network: $u(\mathbf{s}) = \lambda(G(\mathbf{s}))$. The *social optimum (OPT)* is a network $G(\mathbf{s}^*)$ which maximizes the social utility $u(\mathbf{s}^*)$ over all feasible strategy profiles.

We say agent v *reduces the capacity* of the v - x edge by ℓ if it either decreases the capacity $c(v,x)$ by ℓ (if $c(v,x) > \ell$) or if it deletes the edge (v,x) (if $c(v,x) = \ell$). Analogously, agent v *increases the capacity* of the v - x edge if it either increases the capacity $c(v,x)$ by ℓ (if $(x, c(v,x)) \in S_v$) or if it buys the edge (v,x) with capacity ℓ .

An *improving move* for agent v_i is a strategy change from S_{v_i} to S'_{v_i} such that $u(v_i, (s'_{v_i}, s_{-v_i})) > u(v_i, s)$, where (s'_{v_i}, s_{-v_i}) is a new strategy profile which only differs from s in the strategy of agent v_i . We say that agent v plays its *best response* S_v if there is no improving move for agent v . A strategy change towards a best response is called *best response move*. A sequence of best response moves which starts and ends with the same network is called a *best response cycle*. If every sequence of improving moves is finite, then the game has the *finite improving property (FIP)* or, equivalently, the game is a potential game [Monderer and Shapley, 1996].

We say that a network $G(\mathbf{s})$ is in *pure Nash equilibrium (NE)* if all agents play a best response in $G(\mathbf{s})$. We measure the loss of social utility due to the lack of a central authority and the agents' selfishness with the *Price of Anarchy (PoA)* [Koutsoupias and Papadimitriou, 1999]. Let $\min NE(n, k)$ ($\max NE(n, k)$) be the minimum (maximum) social utility of a NE with n agents with budget k , and let $\text{opt}(n, k)$ be the corresponding utility of the social optimum. For given n and k the PoA is $\frac{\text{opt}(n, k)}{\min NE(n, k)}$ and the *Price of Stability (PoS)* is $\frac{\text{opt}(n, k)}{\max NE(n, k)}$. The latter measures the minimum sacrifice in social utility for obtaining a NE network.

1.2 Related Work

The NCG was proposed in [Fabrikant *et al.*, 2003] and has been an object of intensive study for almost two decades. The model depends on an edge-price parameter α which defines the cost of any edge of the network and the agents' objective function is to minimize their closeness centrality, i.e., their average hop-distance to all other nodes in the created network. A constant PoA was shown for almost all values of α , except for the range where $\alpha \approx n$, where only an upper bound of $o(n^\varepsilon)$, for any $\varepsilon > 0$, is known [Dezmaire *et al.*, 2012]. See [Álvarez and Messegué, 2019; Bilò and Lenzner, 2020] for the latest improvements on the PoA and a more detailed discussion. Moreover, [Kawald and Lenzner, 2013] proved that the NCG does not have the FIP.

Many variants of the NCG have been proposed and analyzed but so far all of them consider agents which strive for centrality or only for creating a connected network. Related to our model are versions where agents can only swap edges [Alon *et al.*, 2013; Mihalák and Schlegel, 2012], bounded-budget versions [Laoutaris *et al.*, ; Ehsani *et al.*, 2015] and two versions which focus on robustness: in [Meirom *et al.*, 2015] agents maintain two vertex-disjoint paths to any other node and in [Chauhan *et al.*, 2016] agents consider their expected centrality with respect to a single random edge failure. Thus agents in these models enforce 2-vertex-connectivity or 2-edge-connectivity in the created network which are much weaker robustness concepts compared to our approach. Another line of research are Network Formation Games [Bala and Goyal, 2000a] where agents only strive for being connected to all other nodes. Several variants with focus on robustness have been introduced: In [Bala and Goyal, 2000b] a model with probabilistic edge failures, in [Kliemann, 2011; Kliemann *et al.*, 2017] adversarial models where a single edge is removed after the network is formed, in [Goyal *et al.*, 2016; Friedrich *et al.*, 2017] a version where a node is attacked with deterministic spread to neighboring nodes, and in [Chen *et al.*, 2019] a model with probabilistic spread. To the best of our knowledge, no related model exists where agents try to maximize their connectivity.

In the realm of classical optimization, maximizing the robustness of a network with a given budget is a frequently studied network augmentation problem, see, e.g., [Watanabe and Nakamura, 1987; Frank, 1994; Nagamochi and Ibaraki, 2008] for surveys. However, due to the centralized optimization view, these problems are very different from our model.

1.3 Our Contribution

By incorporating bandwidth considerations into the classical NCG we shed light on a largely unexplored area of research. In our Flow-Based NCG agents strategically create links under a budget constraint to maximize their average or minimum flow value to all other agents. This is in stark contrast to existing models which focus on agents aiming for centrality or for maximizing the size of their connected component.

For our novel modes, we provide an efficient algorithm to compute a social optimum network and we uncover important structural properties of it. A major part of our research is a rigorous study of the structure and quality of the induced

equilibrium networks. We show that NE networks are guaranteed to be connected and that they contain a subgraph that is at least $(k + 1)$ -edge-connected. Moreover, any NE in the min-flowNCG is always $(k + 1)$ -edge-connected. Most importantly, we prove that the social utility of all NE networks is close to optimum. More precisely, our PoA results for both models guarantee that the social utility of any NE is at least half the optimal utility, i.e., $\text{PoA} \leq 2$, and this bound is tight for the min-flowNCG and almost tight for the avg-flowNCG. Besides this, we prove for both versions that the PoS = 1 and that the FIP does not hold. See Table 1 for an overview.

Due to limited space, some proofs are sketched or omitted.

2 Social Optimum

We analyze the structure of the OPT networks. In particular, we show a generic network construction with maximum social utility $2k$ for both game models.

Algorithm 1: Algorithm for computing the OPT

input : unweighted directed clique $K_n = (V, E)$
output: social optimum network G

- 1 $G \leftarrow (V, \emptyset)$;
- 2 **for** $i \leftarrow 1$ **to** k **do**
- 3 find a directed Hamiltonian cycle C_n in K_n ;
- 4 add the cycle to the output graph $G \leftarrow G \cup C_n$,
 i.e., increase capacity of each (u, v) -connection
 in G by 1 if $(u, v) \in C_n$;

Theorem 1. *Algorithm 1 computes an optimal network with social utility $2k$ for the avg-flowNCG and min-flowNCG in polynomial time.*

Proof. First, we prove a general upper bound on the social utility of OPT. Let G be an optimal network in the min-flowNCG. The social utility of G equals the edge connectivity of G which is at most $\min_{v \in V} \text{deg}_G(v)$. The handshake lemma yields $\sum_{v \in V} \text{deg}_G(v) = 2 \sum_{e \in E(G)} c(e) \leq 2nk$, where the last inequality holds because every agent can build edges of the total capacity of at most k . Since the minimum degree is at most the average degree in the network, we get $u(\mathbf{s}) = \lambda(G(\mathbf{s})) \leq \min_{v \in V} \text{deg}_G(v) \leq 2k$.

In the avg-flowNCG, the social utility is $u(\mathbf{s}) = \frac{1}{n} \sum_{v \in V} u(v, \mathbf{s})$ where $u(v, \mathbf{s})$ is the average local edge connectivity over all pairs $(v, V \setminus \{v\})$. The local edge connectivity of any pair of nodes v, w is at most $\min\{\text{deg}_G(v), \text{deg}_G(w)\}$. Thus, $u(\mathbf{s}) \leq \frac{1}{n} \sum_{v \in V} \text{deg}_G(v) \leq 2k$, again by the handshake lemma.

Now we show that the social utility of the graph G produced by the algorithm matches the upper bound which implies that G is optimal. Since each node in G has an outdegree of k and no self-loops, G is a feasible network for both games. For any two nodes $v, w \in V$ the minimum v - w -cut contains two edges of each Hamiltonian cycle added to G during the algorithm's run. Thus, the local edge connectivity of any pair of nodes equals $2k$ which implies $u(G) = 2k$ in both games.

It is easy to see that the algorithm runs in polynomial time since the Hamiltonian cycle can be constructed greedily. \square

Model	u(OPT)	PoS	PoA	FIP
min-flowNCG	$2k$ (Thm. 1)	1 (Thm. 6)	$\frac{2k}{k+1}$ (Thm. 9)	no (Thm. 15)
			< 2 (Thm. 13)	
avg-flowNCG	$2k$ (Thm. 1)	1 (Thm. 6)	$\geq \frac{2k}{k + \frac{k(k-1)}{n-1}}$ if $k \leq 0.5 + \sqrt{n - 0.75}$ (Thm. 14)	no (Thm. 15)
			$\geq \frac{2k}{k+1}$ if $k > 0.5 + \sqrt{n - 0.75}$ (Thm. 14)	

Table 1: Overview of our results

Theorem 2. *In the min-flowNCG and the avg-flowNCG OPT is a connected $2k$ -regular network.*

3 Structural Properties of Equilibria

In this section we prove structural properties of NE networks, which we will use later on to derive our PoA bounds.

Theorem 3. *In the min-flowNCG and the avg-flowNCG all agents use all of their k budget units in an NE network.*

Proof. Consider a network $G = G(\mathbf{s})$ which is in NE. Assume towards a contradiction that there is an agent $v \in V$ which spends strictly less than k units of its budget in $G(\mathbf{s})$.

For the min-flowNCG, we have, by definition, $u(v, \mathbf{s}) = (u_1(v, \mathbf{s}), u_2(v, \mathbf{s})) = (\lambda(G), u_2(v, \mathbf{s}))$. Let $w \in V$ be an agent such that $\lambda_G(v, w) = \lambda(G)$. Then agent v can improve its strategy by increasing the capacity of the edge (v, w) . Indeed, after the strategy change $(s$ to $s' = (s'_v, s_{-v}))$ either the edge connectivity of the network increases, i.e., $\lambda(G(s')) > \lambda(G(s))$ or the edge connectivity does not change but $u(v, s') = (u_1(v, s), u_2(v, s')) \geq (u_1(v, s), u_2(v, s) + 1) > u(v, s)$. In both cases agent v improves, which contradicts that G is in NE.

For the avg-flowNCG it is easy to see, that building an additional edge towards any agent $w \in V \setminus \{v\}$ increases their local edge connectivity, while not decreasing the local edge connectivity between any pair of agents, and thus this results in a strategy $s' = (s'_v, s_{-v})$ such that $u_v(s') > u_v(s)$. Thus, agent v can improve, which contradicts that G is in NE. \square

Theorem 4. *In the min-flowNCG and the avg-flowNCG any equilibrium network is connected.*

Proof. For both games assume towards a contradiction that there is a disconnected equilibrium network G . Note that $\lambda(G) = 0$. By Theorem 3 all agents use all of their budgets, thus, every vertex has an outgoing edge in G . As the number of nodes is finite, there must be a directed cycle in every weakly connected component of G . Thus, any agent v in such a cycle can remove an existing outgoing cycle-edge without disconnecting the weakly connected component and add the edge to any node of another component in G .

In the min-flowNCG, this move either increases the edge-connectivity of the network, i.e., it increases $u_1(v, G)$, or it increases the number of well-connected agents, i.e., it increases $u_2(v, G)$ and $u_1(v, G)$ remains unchanged. Thus, agent v can improve which contradicts that G is in NE.

For the avg-flowNCG, consider two weakly connected components $X, Y \subseteq V$ in G such that $|X| \leq |Y|$. Thus

there must be an agent v in X which can reduce the capacity of one connection in X by 1 and add an edge of capacity 1 to some other node in Y . This move changes only the local edge-connectivity between v and any other node by at most 1. Thus, its utility changes by at least $-\frac{|X|-1}{n-1} + \frac{|Y|}{n-1} > 0$. Thus, agent v can improve which contradicts that G is in NE. \square

The next property will be about j -clusters, which are subset of nodes $C \subseteq V, |C| \geq 2$, from G such that the C -induced subgraph of G has edge connectivity of at least j .

Theorem 5. *For the min-flowNCG and the avg-flowNCG, if G is in NE, there is a $(k+1)$ -cluster in G .*

Proof. Consider a NE network $G(\mathbf{s})$. We partition $G(\mathbf{s})$, or more precisely its undirected version $F(\mathbf{s})$, into components by the following method: As long as there is a cut (X, Y) of size at most k , remove all edges of (X, Y) from G . At each step of the algorithm either the algorithm terminates, i.e., there is a $(k+1)$ -cluster in G , or the number of connected components increases by at least 1 while the number of edges decreases by at most k . By Theorem 3 the sum of edge weights in the initial graph is nk . Hence, if a $(k+1)$ -cluster was not found, the algorithm produces a network containing n components and a set of edges of the total capacity of at least k . Since self-loops are not allowed, we have a contradiction. Thus, the process must find a $(k+1)$ -cluster. \square

4 Price of Stability

Theorem 6. *The PoS is 1 in the min-flowNCG and the avg-flowNCG.*

Proof. Consider a directed cycle C where every edge has capacity k . We will show that C is a NE in both games.

For the min-flowNCG, consider a node, say v_1 , in $C = G(\mathbf{s})$. Agent v_1 owns one edge, w.l.o.g., (v_1, v_2) of capacity k . The only strategy change that v_1 can perform is to reduce the capacity of the edge (v_1, v_2) and to spend the rest of the budget on new edges to other nodes. This move strictly decreases the degree of the node v_2 . Since the edge connectivity of any network is at most its minimum degree, which then is the degree of v_2 , this implies that the strategy change decreases v_1 's utility. Hence, there is no improving move for agent v_1 , which, by symmetry, implies that every agent plays best response in C which proves that C is in NE.

Analogously, in the avg-flowNCG the same strategy change by any of the agents, say v_1 , decreases the degree of the respective neighbor v_2 . Thus, the local edge connectivity

of this pair of nodes decreases. At the same time, the local edge connectivity of v_1 with other nodes does not change since it does not exceed the minimum degree of the pair, i.e., the degree of node v_1 . This implies a decrease of the average local edge connectivity for agent v_1 , i.e., the decrease of v_1 's utility. Thus, all agents play best response and C is a NE.

By Theorem 1 the social utility of OPT is $2k$. We have proved the existence of a NE network with the same social utility for both models. Thus, $PoS = 1$ for both games. \square

5 Price of Anarchy

In this section we present our main results, which are a tight bound in the PoA for the min-flowNCG and an almost tight PoA bound for the avg-flowNCG. Moreover, whereas all previous results show the similarity of the min-flowNCG and the avg-flowNCG we will now highlight their differences. In particular, we show that NE networks in the min-flowNCG may not be NE networks in the avg-flowNCG and vice versa.

5.1 PoA of the Min-FlowNCG

We start by providing a NE network for the min-flowNCG with minimum social utility.

Theorem 7. *For a given n and k such that $k < n$, there is a NE network G such that $u(G) = k + 1$ in the min-flowNCG.*

Proofsketch. We consider the following generic construction for a given number of agents n and budget $k < n$. Consider a set of nodes $\{v_1, \dots, v_n\} = V(G)$ and a set of edges $E(G)$ where for all $1 \leq i < k$ each node v_i is connected with a node v_{i+1} by an edge (v_i, v_{i+1}) with capacity $c(v_i, v_{i+1}) = k - (i - 1)$ and by an edge (v_{i+1}, v_i) with capacity $c(v_{i+1}, v_i) = i$. Moreover, there is an edge (v_k, v_n) with capacity 1 and a set of sequential edges, for all $k \leq i \leq n - 1$, (v_{i+1}, v_i) with capacity k . See Figure 2 for illustration of the construction. We claim that G is in NE in the min-flowNCG. We omit the proof due to the space restrictions. \square

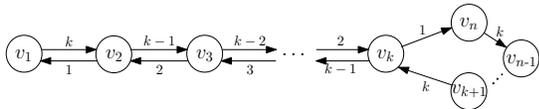


Figure 2: NE network in the min-flowNCG. Note that this network is not a NE in the avg-flowNCG. Indeed, agent v_2 can delete the edge (v_2, v_1) and create the edge (v_2, v_n) with capacity 1.

The next result shows that any NE network in the min-flowNCG is $(k + 1)$ -edge-connected.

Theorem 8. *The social utility of any NE network is at least $k + 1$ in the min-flowNCG.*

Proof. Assume towards a contradiction that there is a NE $G = (V, E)$ such that it contains a cut (A, B) such that $\lambda(G) = \sum_{e \in (A, B)} c(e) \leq k$. By Theorem 5 there is an induced subgraph $K = (C, E(K)) \subseteq G$ such that C is a $(k + 1)$ -cluster in G . W.l.o.g., let $C \subseteq A$. Consider two

nodes $v \in C$ and $w \in B$. Since $k + 1 \leq \lambda(K) \leq \deg_K(v)$, there is a node $x \in K$ which has an edge to v in K .

We will show that agent x can significantly improve on its strategy by reducing the capacity of the edge (x, v) by 1 and creating the edge (x, w) with capacity 1 (or increasing its capacity by 1 if it already exists). We denote the obtained networks as $G' = (V', E')$ and $K' = (C, E'')$. Note that $\lambda_{G'}(x, v) \geq \lambda_{K'}(x, v) + 1 \geq \lambda_K(x, v) - 1 + 1 \geq k + 1$.

If $\lambda(G') > \lambda(G)$, then agent x improved on its strategy since $u_1(x, G') > u_1(x, G)$. Thus, G is not a NE.

If $\lambda(G') < \lambda(G)$, there is a cut (X, Y) , $x \in X$, $v \in Y$ in G' of size strictly less than $\lambda(G) \leq k$. On the other hand, the size of the (X, Y) -cut is at least $\lambda_{G'}(x, v) \geq k + 1 > \lambda(G)$. Hence, we have a contradiction.

Finally, we consider the case $\lambda(G') = \lambda(G)$. We will show that $u_2(x, G') > u_2(x, G)$. First, note that for any node $a \in V \setminus \{x\}$ such that $\lambda_G(x, a) > \lambda(G)$, $\lambda_{G'}(x, a) > \lambda(G')$ holds. Indeed, in case $\lambda_{G'}(x, a) < \lambda_G(x, a)$, we have $\lambda_{G'}(x, a) = \lambda_{G'}(x, v) \geq k + 1 > \lambda(G) = \lambda(G')$. In case $\lambda_{G'}(x, a) \geq \lambda_G(x, a)$, and since $\lambda_G(x, a) > \lambda(G)$, we have $\lambda_{G'}(x, a) > \lambda(G) = \lambda(G')$. Therefore, after the strategy change $u_2(x, G') \geq u_2(x, G)$ holds. Second, $\lambda_G(x, w) = \lambda(G)$, since $x \in A$ and $w \in B$, whereas $\lambda_{G'}(x, w) > \lambda(G')$, i.e., $u_2(x, G') > u_2(x, G)$. Hence, $u(x, G') = (u_1(x, G), u_2(x, G')) > u(x, G)$, i.e., agent x does not play a best response in G . This contradicts the assumption that G is a NE. \square

Now we put the pieces together to derive our tight bound.

Theorem 9. *The PoA in the min-flowNCG is $\frac{2k}{k+1}$.*

Proof. By Theorem 1 and Theorem 8 we get that $\text{PoA} \leq \frac{2k}{k+1}$. By Theorem 7 there is a NE network that meets the social utility lower bound, hence the PoA bound is tight. \square

5.2 PoA of the Avg-FlowNCG

We start with providing constructions for two NE network with almost minimum social utility.

Theorem 10. *For n and k with $2 \leq k < n$, there is a NE network G such that $u(G) = k + \frac{k(k-1)}{n-1}$ in the avg-flowNCG.*

Proofsketch. We construct the following network $G = (V, E)$. We arrange k nodes in a circle such that each node has one edge with capacity k to the subsequent node. The remaining $n - k$ nodes each have a unit-weight edge to each node of the circle. See Figure 3 for an illustration of the construction. We omit the proof due to space constraints. \square

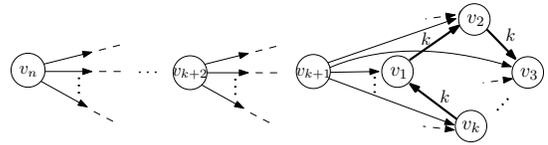


Figure 3: NE network in the avg-flowNCG. Note that this network is not a NE in the min-flowNCG since it is not $(k + 1)$ -edge-connected, e.g., v_1 can improve by reducing the capacity of the edge (v_2, v_3) by 1 and increasing the capacity of the edge (v_2, v_{k+1}) by 1.

Theorem 11. For a given n and k such that $k < n$, there is a NE network G such that $u(G) = k + 1$ in the avg-flowNCG.

Proofsketch. Consider the following star-like network $G = (V, E)$ with the set of nodes $V = \{c\} \cup \{a_1, \dots, a_{k-1}\} \cup \{b_1, \dots, b_{n-k}\}$. Here c is a central node connected with all a_i nodes by an edge with capacity 1. All a_i nodes are connected with the central node by edges with capacity k . Each node $b_i, i = 1, \dots, n-k-1$ has an edge (b_i, c) with capacity $k-1$, and only one node b_{n-k} has an edge (b_{n-k}, c) of capacity k . Also, there is an edge (c, b_1) with capacity 1. Finally, all b_i nodes are sequentially connected by edges (b_i, b_{i+1}) with capacity 1. See Figure 4 for illustration of the construction. The proof of the statement is rather technical, so we omit it due to the space constrains. \square

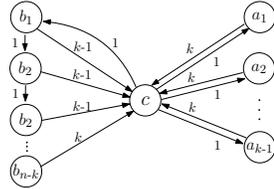


Figure 4: NE network in the avg-flowNCG.

A similar technique as in the proof of Theorem 8 can be applied to prove that the edge connectivity of any NE network in the avg-flowNCG is at least k .

Theorem 12. For $k \geq 2$, the edge connectivity of any NE network in the avg-flowNCG is at least k .

Proofsketch. Assume G is a NE and $\lambda(G) < k$. Consider a min-cut (A, B) of size $\lambda(G)$. By performing the same procedure as in the proof of Theorem 5 in each set A and B , we get two sets $X \subseteq A$ and $Y \subseteq B$ of a maximum size such that for each pair of nodes v_1, v_2 in the same set, $\lambda_G(v_1, v_2) \geq k$. W.l.o.g., $|Y| \geq |X|$. For some edge (v, w) in a k -cluster in X , agent v can reduce its capacity by 1 and create a new edge (v, y) with capacity 1 to $y \in Y$ (or increase its capacity by 1 if it exists). The local edge connectivity decreases by at most 1 only between u and nodes in X , and increases by 1 between u and all nodes in Y . Thus, the utility changes by at least $-\frac{|X|-1}{n-1} + \frac{|Y|}{n-1} > 0$, i.e., it is an improving move. \square

Theorem 13. For the avg-flowNCG we have $PoA < 2$.

Proof. For $k = 1$ a trivial NE of social utility 2 is a directed cycle with all edge capacities equal 1.

For general $k \geq 2$, by Theorem 12, any NE is at least k -edge-connected. On the other hand, by Theorem 5, any NE contains a $(k+1)$ -cluster of at least two nodes. Thus, at least two nodes have their utility of at least $\frac{k+1+k(n-2)}{n-1} > k$ and all other nodes have the utility of at least k . Hence the social utility is strictly greater than $\frac{k(n-2)+2k}{n} = k$. \square

Now we provide an almost matching lower bound on the PoA.

Theorem 14. For the avg-flowNCG, we have $PoA \geq \max\left\{\frac{2k}{k+k(k-1)/(n-1)}, \frac{2k}{k+1}\right\}$.

Proof. By Theorem 10 and Theorem 11 the social utility of a worst-case equilibrium network is at most $\min\{k + \frac{k(k-1)}{n-1}, k+1\}$. Therefore, for $k \leq 0.5 + \sqrt{n-0.75}$ the PoA is at least $2k / \left(k + \frac{k(k-1)}{n-1}\right) = \frac{2(n-1)}{n+k-2}$. Thus, for large n , the PoA tends to 2, i.e., the upper bound. With a high budget, i.e., for $k > 0.5 + \sqrt{n-0.75}$, the PoA is at least $\frac{2k}{k+1}$. \square

6 Game Dynamics

We analyze whether our games are guaranteed to converge to an NE under improving response dynamics. Unfortunately, this is not true for the min-flowNCG and the avg-flowNCG.

Theorem 15. The min-flowNCG and the avg-flowNCG do not admit the FIP.

Proofsketch. We provide an improving response cycle (IRC) for each variant of the game depicted in Figure 5 and Figure 6. Note that the first and the last networks are isomorphic in both sequences, which implies that we indeed have IRCs. \square

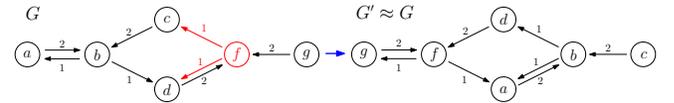


Figure 5: An IRC for the min-flowNCG with $k = 2$. Since G and G' are isomorphic, the sequence of the improving moves recurs to G .

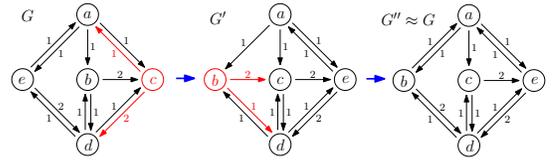


Figure 6: An IRC for the avg-flowNCG with $k = 3$. Since G and G'' are isomorphic, the sequence of the improving moves recurs to G .

7 Conclusion

We made the first steps into the promising direction of investigating network formation games with bandwidth/connectivity considerations. For this, we proposed two versions: the min-flowNCG where agents strive to maximize the connectivity of the entire network; and the avg-flowNCG where agents try to maximize their average flow value towards all other agents. For both versions our main focus was the analysis of equilibrium networks. In particular, we proved that NE networks contain a highly-connected component in the avg-flowNCG, while in the min-flowNCG the entire NE network must be highly-connected. This may explain the occurrence of k -core structures in many real-world networks.

Our main results are (almost) tight bounds of less than 2 on the PoA which shows that NE networks are close to optimum and no external coordination is needed in our games.

An important direction for future work is a generalization of the presented model to non-uniform budgets, non-integer edge capacities and a combination of bandwidth and centrality objectives for the agents.

References

- [Alon *et al.*, 2013] Noga Alon, Erik D Demaine, Mohammad T Hajiaghayi, and Tom Leighton. Basic network creation games. *SIAM Journal on Discrete Mathematics*, 27(2):656–668, 2013.
- [Àlvarez and Messegué, 2019] Carme Àlvarez and Arnau Messegué. On the price of anarchy for high-price links. In *WINE'19*, pages 316–329. Springer, 2019.
- [Bala and Goyal, 2000a] Venkatesh Bala and Sanjeev Goyal. A noncooperative model of network formation. *Econometrica*, 68(5):1181–1229, 2000.
- [Bala and Goyal, 2000b] Venkatesh Bala and Sanjeev Goyal. A strategic analysis of network reliability. *Review of Economic Design*, 5(3):205–228, 2000.
- [Bilò and Lenzner, 2020] Davide Bilò and Pascal Lenzner. On the tree conjecture for the network creation game. *Theory Comput. Syst.*, 64(3):422–443, 2020.
- [Chauhan *et al.*, 2016] Ankit Chauhan, Pascal Lenzner, Anna Melnichenko, and Martin Münn. On selfish creation of robust networks. In *SAGT'16*, pages 141–152, 2016.
- [Chen *et al.*, 2019] Yu Chen, Shahin Jabbari, Michael Kearns, Sanjeev Khanna, and Jamie Morgenstern. Network formation under random attack and probabilistic spread. *arXiv preprint arXiv:1906.00241*, 2019.
- [Demaine *et al.*, 2012] Erik D. Demaine, Mohammad Taghi Hajiaghayi, Hamid Mahini, and Morteza Zadimoghaddam. The price of anarchy in network creation games. *ACM Trans. on Algorithms (TALG)*, 8(2):13, 2012.
- [Ehsani *et al.*, 2015] Shayan Ehsani, Saber Shokat Fadaee, MohammadAmin Fazli, Abbas Mehrabian, Sina Sadeghian Sadeghabad, Mohammadali Safari, and Morteza Saghafian. A bounded budget network creation game. *ACM Trans. on Algorithms (TALG)*, 11(4):34, 2015.
- [Fabrikant *et al.*, 2003] Alex Fabrikant, Ankur Luthra, Elitza Maneva, Christos H Papadimitriou, and Scott Shenker. On a network creation game. In *PODC'03*, pages 347–351, 2003.
- [Ford and Fulkerson, 1956] Lester R. Ford and Delbert R. Fulkerson. Maximal flow through a network. *Canadian Journal of Mathematics*, 8:399–404, 1956.
- [Frank, 1994] András Frank. Connectivity augmentation problems in network design. Technical report, Univ. of Michigan, Ann Arbor, MI (United States), 1994.
- [Friedrich *et al.*, 2017] Tobias Friedrich, Sven Ihde, Christoph Keßler, Pascal Lenzner, Stefan Neubert, and David Schumann. Efficient best response computation for strategic network formation under attack. In *SAGT'17*, pages 199–211, 2017.
- [Goyal *et al.*, 2016] Sanjeev Goyal, Shahin Jabbari, Michael Kearns, Sanjeev Khanna, and Jamie Morgenstern. Strategic network formation with attack and immunization. In *WINE'16*, pages 429–443. Springer, 2016.
- [Jackson and Wolinsky, 1996] Matthew O Jackson and Asher Wolinsky. A strategic model of social and economic networks. *Journal of economic theory*, 71(1):44–74, 1996.
- [Kawald and Lenzner, 2013] Bernd Kawald and Pascal Lenzner. On dynamics in selfish network creation. In *SPAA'13*, pages 83–92. ACM, 2013.
- [Kliemann *et al.*, 2017] Lasse Kliemann, Elmira Shirazi Sheykhdarabadi, and Anand Srivastav. Swap equilibria under link and vertex destruction. *Games*, 8(1):14, 2017.
- [Kliemann, 2011] Lasse Kliemann. The price of anarchy for network formation in an adversary model. *Games*, 2(3):302–332, 2011.
- [Koutsoupias and Papadimitriou, 1999] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *STACS'99*, pages 404–413. Springer, 1999.
- [Laoutaris *et al.*,] Nikolaos Laoutaris, Laura J. Poplawski, Rajmohan Rajaraman, Ravi Sundaram, and Shang-Hua Teng. Bounded budget connection (BBC) games or how to make friends and influence people, on a budget. In *PODC'08*, pages 165–174. ACM.
- [Meirom *et al.*, 2015] Eli A Meirom, Shie Mannor, and Ariel Orda. Formation games of reliable networks. In *INFOCOM '15*, pages 1760–1768. IEEE, 2015.
- [Mihalák and Schlegel, 2012] Matúš Mihalák and Jan Christoph Schlegel. Asymmetric swap-equilibrium: A unifying equilibrium concept for network creation games. In *MFCS'12*, pages 693–704. 2012.
- [Monderer and Shapley, 1996] Dov Monderer and Lloyd S. Shapley. Potential games. *Games and economic behavior*, 14(1):124–143, 1996.
- [Myerson, 2013] Roger B Myerson. *Game theory*. Harvard university press, 2013.
- [Nagamochi and Ibaraki, 2008] Hiroshi Nagamochi and Toshihide Ibaraki. *Algorithmic aspects of graph connectivity*. Cambridge University Press, 2008.
- [Papadimitriou, 2001] Christos H. Papadimitriou. Algorithms, games, and the internet. In *STOC'01*, pages 749–753, 2001.
- [Tardos, 2004] Éva Tardos. Network games. In *STOC'04*, pages 341–342, 2004.
- [Watanabe and Nakamura, 1987] Toshimasa Watanabe and Akira Nakamura. Edge-connectivity augmentation problems. *Journal of Computer and System Sciences*, 35(1):96–144, 1987.