algorithms

MDPI

Article

Network Creation Games with Traceroute-Based Strategies †

Davide Bilò [1], Luciano Gualà [2,*], Stefano Leucci [3] and Guido Proietti [3,4]

[1] Dipartimento di Scienze Umanistiche e Sociali, Università di Sassari, 21, 07100 Sassari SS, Italy; davide.bilo@uniss.it
[2] Dipartimento di Ingegneria dell'Impresa, Università di Roma "Tor Vergata", 50, 00133 Roma (RM), Italy
[3] Dipartimento di Ingegneria e Scienze dell'Informazione e Matematica, Università degli Studi dell'Aquila, 67100 L'Aquila AQ, Italy; stefano.leucci@univaq.it (S.L.); guido.proietti@univaq.it (G.P.)
[4] Istituto di Analisi dei Sistemi ed Informatica, CNR, 19, 00185 Roma (RM), Italy
* Correspondence: guala@mat.uniroma2.it

† A preliminary version of this paper appeared in the proceedings of the 21st International Colloquium on Structural Information and Communication Complexity (SIROCCO'14), Takayama, Japan, 23–25 July 2014.

Abstract: Network creation games have been extensively used as mathematical models to capture the key aspects of the decentralized process that leads to the formation of interconnected communication networks by selfish agents. In these games, each user of the network is identified by a node and selects which link to activate by strategically balancing his/her *building cost* with his/her *usage cost* (which is a function of the distances towards the other player in the network to be built). In these games, a widespread assumption is that players have a *common and complete* information about the evolving network topology. This is only realistic for small-scale networks as, when the network size grows, it quickly becomes impractical for the single users to gather such a global and fine-grained knowledge of the network in which they are embedded. In this work, we weaken this assumption, by only allowing players to have a *partial* view of the network. To this aim, we borrow three popular *traceroute-based* knowledge models used in *network discovery*: (i) *distance vector*, (ii) *shortest-path tree view*, and (iii) *layered view*. We settle many of the classical game theoretic questions in all of the above models. More precisely, we introduce a suitable (and unifying) equilibrium concept which we then use to study the convergence of improving and best response dynamics, the computational complexity of computing a best response, and to provide matching upper and lower bounds to the price of anarchy.

Keywords: network creation games; local-knowledge equilibrium; convergence dynamics; price of anarchy

Citation: Bilò, D.; Gualà, L.; Leucci, S.; Proietti, G. Network Creation Games with Traceroute-Based Strategies. *Algorithms* **2021**, *14*, 35. https://doi.org/10.3390/a14020035

Academic Editor: Angelo Fanelli, Gianpiero Monaco, Luca Moscardelli

Received: 3 December 2020
Accepted: 22 January 2021
Published: 26 January 2021

1. Introduction

The construction of large *communication networks* involves the interaction of many independent and selfish agents with competing interests and, in such a decentralized setting, the problem of understanding the formation process of a network arises naturally. More formally, we model the agents (i.e., *players*) as a set of n vertices in a graph, each controlled by a player that wants to connect itself to all of the other participants of the network. This can happen either directly, by the unilateral and costly activation of a corresponding logical or physical communication channel (i.e., a *link*), or indirectly by routing messages over (a subset of the) existing (bidirectional) links that were activated by other agents. Quite naturally, players prefer low-latency communication paths over longer paths and will therefore always route messages along a shortest path. It is then clear how tension stems from the desire of a player for efficient communication, and his/her interest in minimizing the number of activated links. The study of such a trade-off between a player's *building cost* (which is proportional to the number of links she decides to activate) and her *usage cost* (some function that depends on the distance between the player and the remaining agents in the resulting network) results in the corresponding study of a *communication network creation game*, also known as *network connection game* (simply NCG in the following).

1.1. The Standard Model for NCGs

A first NCG model that received considerable attention in the algorithmic game theory area is due to Fabrikant et al. [1]. In this model, which we will call SUMNCG in the sequel, each player u pays α for each link she activates as his/her *building cost*, and the *sum* of all distances from u to all other vertices in the resulting network as her *usage cost*. More formally, if V is a set of n agents, the set of the possible strategies of player $u \in V$ is $2^{V \setminus \{u\}}$ (Given a set X, 2^X denotes the power set of X). Given a strategy σ_u for each agent $u \in V$, we will call the vector $\sigma = (\sigma_u)_{u \in V}$ a *strategy profile*, and we will denote by $G(\sigma)$ the resulting network, i.e., a graph containing a vertex for each player in V, and such that its edge set $E(\sigma)$ contains an edge (u, v) if and only if $u \in \sigma_v$ or $v \in \sigma_u$. To easy the notation, we will sometimes use G instead of $G(\sigma)$ if no ambiguity arises. Moreover, we will say that $u \in V$ *owns* the edges in $\{u\} \times \sigma_u$ (Notice that, according to the above definition, an edge can be owned by both its endpoints). The goal of each agent $u \in V$ is that of minimizing her cost *cost* $C_u(\sigma)$ in the resulting strategy profile σ, which is defined as

$$C_u(\sigma) = \alpha \cdot |\sigma_u| + \sum_{v \in V} d_{G(\sigma)}(u, v), \tag{1}$$

where $d_{G(\sigma)}(u, v)$ is the distance between player u and player v in $G(\sigma)$. Quite naturally, we define the *social cost* of σ as the sum of the costs $C_u(\sigma)$ incurred by all players $u \in V$.

Notice that in the above description, we are *implicitly* assuming that each player u has a complete knowledge of the current network $G(\sigma)$. This means that u can evaluate both his/her current cost $C_u(\sigma)$, and the cost $C_u(\sigma_{-u}, \sigma'_u)$ that he/she would incur if he/she were to change her strategy from σ_u to σ'_u, as he/she is aware that—following her change—the current network will transit to $G(\sigma_{-u}, \sigma'_u)$ (The notation (σ_{-u}, σ'_u) denotes the vector obtained from σ by replacing component σ_u with σ'_u). Therefore, a (pure strategy) *Nash Equilibrium* (NE) for the game is a strategy profile $\bar{\sigma}$ where, for every player u and every strategy σ_u, we have that $C_u(\bar{\sigma}) \leq C_u(\bar{\sigma}_{-u}, \sigma_u)$. A vast body of literature deals with the problem of characterizing the space of NE in order to bound the *Price of Anarchy* (PoA) as a function of the parameter α. Here, the PoA is defined as the ratio between the social cost of the worst NE (i.e., any NE σ maximizing its social cost $\sum_u C_u(\sigma)$), and the social cost of an *optimal* strategy profile (i.e., a strategy profile σ^* minimizing its social cost $\sum_u C_u(\sigma^*)$) [2–8]. Intuitively, the PoA provides an upper bound on the loss of efficiency (measured in terms of the social cost) between the stable networks that can be reached as a result of the uncoordinated selfish behavior of the agents and that of any optimal (but not necessarily stable) network that minimizes the social cost. Large values of the PoA show that the distributed network formation process might converge to costly networks when compared to a "centralized" design. Conversely, small values of the PoA show that even if the network arises from a combination of selfish strategies, its quality will not be significantly worse than the best possible design. It follows from the best-known results that the PoA of SUMNCG is constant whenever $\alpha \geq n^{1-\varepsilon}$, for any $\varepsilon \geq 1/\log n$ (see in [3,4]) or when $\alpha \leq (1 + \varepsilon')n$ for any constant $\varepsilon' \geq 0$ (see in [6]). For the remaining (small) range of values of α, the PoA is upper bounded by $2^{O(\sqrt{\log n})}$ [7].

In [7], the authors introduced a variant of SUMNCG named MAXNCG. In MAXNCG, the usage cost of a player u is his/her eccentricity in the current network $G(\sigma)$, so that u's cost function becomes

$$C_u(\sigma) = \alpha \cdot |\sigma_u| + \max\{d_{G(\sigma)}(u, v) : v \in V\}. \tag{2}$$

MAXNCG was further considered in [3], where it was shown that the PoA of the game is constant whenever $\alpha \leq 129$ or $\alpha = O(1/\sqrt{n})$.

1.2. Other Models for NCGs

Several variants of these two basic models have been defined, with the goal of better characterizing specific aspects of the network formation process. These variants range from limiting the modifications that players are allowed to perform on their current strategies (see in [9–11]), to budgeting either the number of edges a player can activate or her eccentricity (see in [12–14]), and finally to considering generalized versions of the basic models on edge-weighted graphs (see in [15,16]). Generally speaking, all the aforementioned models exhibit a substantial asymptotic increase in their PoAs compared to that of corresponding basic models introduced above.

Observe that all these models, except for those given in [10,11], share a severe restriction with the basic model, namely, the NP-hardness for a player to select a best-response strategy. Besides that, they also all assume that players have a *common and complete* information about the ongoing network. While this is feasible for small-size instances of the game, this becomes unrealistic for large-size networks. This is rather problematic, given the growing size of the inputs in the practice. Moreover, quite paradoxically, the full-knowledge assumption is not simplifying at all: it makes it computationally unfeasible for a player to select a best-response strategy, as said before, or even to check whether she is actually in a NE! To address these delicate issues, a new model which limits a player's full knowledge of the network structure up to a given radius k from herself has been introduced in [17]. In such a setting, players do not even know the size n of the network (in distributed computing terminology, the system is *uniform*). The authors provided a comprehensive set of upper and lower bounds to the *Price of Anarchy* (PoA) for the entire range of values of k, which have been confirmed by a further experimental evaluation in [18]. Critically, these bounds to the PoA refer to a new suitable equilibrium concept that accounts for the limited knowledge of the players. More precisely, as a player has a partial (defective) view of the network, before changing her strategy she has to evaluate whether such a choice is convenient in *every* realizable network which is compatible with her current view. Then, let σ_u be the strategy played by player u, and define $\Sigma|_{\sigma_u}$ to be the set of strategy profiles $\sigma = (\sigma_{-u}, \sigma_u)$ of the players such that the network $G(\sigma)$ is realizable according to player u's view. Let

$$\Delta(\sigma_u, \sigma'_u) = \max_{\sigma \in \Sigma|_{\sigma_u}} \{C_u(\sigma_{-u}, \sigma'_u) - C_u(\sigma)\} \tag{3}$$

denote the worst possible cost difference u would have in switching from σ_u to σ'_u. This means that if $\Delta(\sigma_u, \sigma'_u) > 0$, then there is some strategy profile σ compatible with u's current view for which switching from σ to (σ_{-u}, σ'_u) would increase u's cost. Conversely, if $\Delta(\sigma_u, \sigma'_u) < 0$, player u can safely change her strategy from σ_u to σ'_u as this will decrease her cost by at least $|\Delta(\sigma_u, \sigma'_u)| > 0$ in all of the conceivable networks $G(\sigma)$ that are compatible with u's view. Then, the *Local Knowledge Equilibrium* (LKE) is defined as a strategy profile $\bar{\sigma}$ such that for every player u and every strategy σ_u, we have $\Delta(\bar{\sigma}_u, \sigma_u) \geq 0$. Notice that our equilibrium concept is actually weaker than the classical NE concept as we have that every NE is also a LKE.

1.3. Our New Local-View Models for NCGs

In this work, we continue the study of NCGs in which players have only a partial information about the current network. Along this direction, we focus on the most prominent *local-knowledge* models used in the area of *network discovery* (see in [19,20]) in which the structure of a large unknown network need to be fully identified using a small number of queries on the network's vertices (where the considered local-knowledge model determines which information are returned by a query). In details, we study the following *traceroute-based* view models, which all find motivations in the practice of probing the topology of a network by tracing the route of packets (According to the spirit of the game, we assume that in all the models, the players initially sit on a connected network):

($\mathcal{M}_1$) *Distance vector*: in addition to his/her incident edges, each player u knows only the distances in $G(\sigma)$ between u and all the agents (this is the minimal knowledge needed by a player in order to compute her current cost).

($\mathcal{M}_2$) *Shortest-Path Tree (SPT) view*: each player u knows the edges of *some* SPT of $G(\sigma)$ rooted at u.

($\mathcal{M}_3$) *Layered view*: each player u knows the set of edges belonging to *at least one* SPT of $G(\sigma)$ rooted at u.

Observe that, in all the above games, each player only has a defective (partial) view of the current network, and therefore the LKE fits perfectly as solution concept. However, unlike the models in [18], our players are always able to compute their current cost $C_u(\sigma)$ as they are aware of their distances towards all the other vertices in the network (i.e., $C_u(\sigma)$ is constant for all $\sigma \in \Sigma|_{\sigma_u}$). For all our models, we study the iterated version of the game and we analyze whether improving and best-response dynamics always converge to a LKE. In these dynamics, we will also assume that the players not only *myopic* but also *oblivious*, i.e., they only compute their responses as a function of their current view, without using any information about their previous views. In addition, we study how the computational difficulty of computing a best response strategy changes according to the partial-knowledge model at hand. Finally, we provide matching upper and lower bounds to the PoA in all of the studied models. See Table 1 for a summary of our results.

Table 1. Summary of our results (and open problems). In the first column, convergence (and divergence) are reported w.r.t. either improving or best-response dynamics (improving response dynamics (IRD) and best response dynamics (BRD), resp.). In the second column, we report the time complexity of selecting a best-response strategy.

	Convergence	Best-Response Complexity	PoA
$\mathcal{M}_1$	SUM: Yes ($\forall$ *improving* response dynamics (IRD)) MAX: Yes ($\forall$ IRD)	SUM: Open MAX: Polynomial	SUM: $\Theta(\min\{1+\alpha, n\})$ MAX: $\Theta(n)$ if $\alpha = \Omega(1)$ $\Theta(1+\alpha n)$ if $\alpha = O(1)$
$\mathcal{M}_2$	SUM: No ($\exists$ *best* response dynamics (BRD) cycle) MAX: No ($\exists$ BRD cycle)	SUM: Polynomial MAX: Polynomial	SUM: $\Theta(\min\{1+\alpha, n\})$ MAX: $\Theta(n)$ if $\alpha = \Omega(1)$ $\Theta(1+\alpha n)$ if $\alpha = O(1)$
$\mathcal{M}_3$	SUM: No ($\exists$ BRD cycle) MAX: No ($\exists$ BRD cycle)	SUM: NP-hard MAX: NP-hard	SUM: $\Theta(\min\{1+\alpha, n\})$ MAX: $\Theta(n)$ if $\alpha = \Omega(1)$ $\Theta(1+\alpha n)$ if $\frac{1}{n-1} \le \alpha = O(1)$

The paper is organized as follows. In Section 2, we focus on convergence issues, while in Section 3 we analyze the computational complexity of finding a best-response. Finally, in Section 4 we study the PoA. All the sections are structured in subsections, according to the various view models. A preliminary version of this work appeared in [21].

2. Convergence

2.1. Model $\mathcal{M}_1$

We start by observing that in $\mathcal{M}_1$ a player u can only infer the existence of the edges (x, y) for which $d_G(u, y) = d_G(u, x) + 1$ and no other vertex $x' \neq x$ satisfies $d_G(u, x) = d_G(u, x')$ (this captures all the edges incident to u as the special case $x = u$).

As a consequence, whenever u has at least two neighbors, *swapping* (i.e., replacing an owned edge with another one) any of his/her edges will never be an improving response as (in u's view) this might cause the resulting network to become disconnected. We summarize this property in the next Lemma.

Lemma 1. *In both* SUMNCG *and* MAXNCG *in* $\mathcal{M}_1$*, any player u that has an improving response involving edge swaps must have degree* 1 *(and u must own its sole incident edge).*

Theorem 1. *In* $\mathcal{M}_1$*, any improving response dynamics for* SUMNCG *must converge to an equilibrium.*

Proof. Consider any strategy profile in which no edge is owned by both endpoints, which can be safely assumed as this can only happen in the starting network.

Notice that, in such a configuration, no edge can be removed by any player (as otherwise the network could become disconnected in the player's view). As the maximum number of edges is upper bounded by $O(n^2)$, it suffices to show that any dynamics in which only swap moves are allowed must converge.

Consider a network G and player u that changes her strategy from σ_u to σ'_u (with $|\sigma_u| = |\sigma'_u|$) in order to improve her (worst-case) cost. As the cost of player u in the resulting network G' must be lower than her corresponding cost in G, and since the building cost of u is unchanged, we know that the usage cost of u decreases when moving from G to G'.

In the rest of the proof we will argue that the function $\Phi(G) = \sum_{x \in V} \sum_{v \in V} d_G(x, v)$ is a potential function for this game, i.e., the value of Φ decreases whenever any improving response is played. The claim will immediately follow by noticing that Φ has a finite domain. As the degree of u in G is 1, as shown by Lemma 1, we can write

$$\begin{aligned}\Phi(G) - \Phi(G') &= \sum_{x \in V} \sum_{v \in V} d_G(x, v) - \sum_{x \in V} \sum_{v \in V} d_{G'}(x, v) \\ &= \sum_{x \in V \setminus \{u\}} \sum_{v \in V \setminus \{u\}} d_G(x, v) + 2 \sum_{v \in V} d_G(u, v) - \sum_{x \in V \setminus \{u\}} \sum_{v \in V \setminus \{u\}} d_{G'}(x, v) - 2 \sum_{v \in V} d_{G'}(u, v) \\ &= 2 \left(\sum_{v \in V} d_G(u, v) - \sum_{v \in V} d_{G'}(u, v) \right) > 0.\end{aligned}$$

where the fact that u has degree 1 in both G and G' implies that $d_G(x, v) = d_{G'}(x, v)$ for every $x, v \in V \setminus \{u\}$. □

A similar result can be shown for MAXNCG in model $\mathcal{M}_1$:

Theorem 2. *In* $\mathcal{M}_1$*, any improving response dynamics for* MAXNCG *must converge to an equilibrium.*

Proof. Similarly to the proof of Theorem 1, it will suffice to prove that every dynamics involving only swap moves converges to an equilibrium. We will show that this holds true using similar argument to those in [22]. Consider a graph G corresponding to some profile of strategies and look at the n-dimensional vector ε_G in which the i-th entry is the *eccentricity* of the i-th vertex of G. Now, consider the sorted version of ε_G in which entries appear in non-increasing order. We will assume that comparisons among n-dimensional vectors be performed in lexicographic order and we will consider the function that maps each strategy profile the corresponding sorted vector ε_G. We will prove that the value of such a function can only decrease following an improving response.

As show in Lemma 1, all players that are able to change their strategy to an improving response must have degree 1 in the current network G. Consider a player u changing her strategy to an improving response and let G' be the resulting network and denote by $\varepsilon_H(x)$ the eccentricity of vertex $x \in V(H)$ in H. As the number edges of edges owned by u cannot decrease, we must have $\varepsilon_G(u) > \varepsilon_{G'}(u)$.

We now prove that if the eccentricity a player $x \neq u$ increases between G and G' (as a consequence of u's strategy change), then $\varepsilon_{G'}(x) < \varepsilon_{G'}(u)$. Consider a vertex x such that $\varepsilon_G(x) < \varepsilon_{G'}(x)$ and notice that, as u has degree 1 in both G and G', it must hold that $\varepsilon_{G'}(x) = d_{G'}(x, u) \leq \varepsilon_{G'}(u)$, which concludes the proof. □

2.2. Models $\mathcal{M}_2$ and $\mathcal{M}_3$

Here, we show that, by contrast to model $\mathcal{M}_1$, a best-response dynamics is not guaranteed to converge in models in $\mathcal{M}_2$ and $\mathcal{M}_3$, i.e., *best responses* cycles can arise. Even thought we are dealing with undirected graphs (i.e., each edge can be traversed in both directions no matter which of its endpoints is the buyer), we will draw some edges as arrows directed away from the buyer in order to stress edge ownership. The next vertex to change her strategy to a best response is shown in dark gray and the set of edges that belong to her current view are depicted using bold lines/arrows.

We start by considering $\mathcal{M}_2$:

Theorem 3. *Best-response cycles are possible in both* SUMNCG *and* MAXNCG *in $\mathcal{M}_2$.*

Proof. As far as MAXNCG is concerned, a best-response cycle for $\alpha \geq 6$ is depicted in Figure 1. The current player in graphs (a) to (d) always replaces his/her current strategy with an best response decreasing her cost. The graph in Figure 1d is isomorphic to that in Figure 1a, where the roles of the vertices labeled a and b are swapped. By mirroring the previous best-responses moves, the configuration in graph (a) will be reached again, thus completing the cycle for MAXNCG.

Figure 2 shows a best-response cycle for SUMNCG when $\alpha \geq 6$, which concludes the proof. □

We also have similar results in model $\mathcal{M}_3$.

Theorem 4. *Best-response cycles are possible in both* SUMNCG *and* MAXNCG *in $\mathcal{M}_3$.*

Proof. A best-response cycle for SUMNCG in $\mathcal{M}_3$ when $\alpha = 15$ is shown in Figure 3. Graphs (a) and (d) are isomorphic and, in particular, the roles of vertices a and c are exchanged.

Figure 4 shows a best-response cycle for MAXNCG when $\alpha = 2 - \epsilon$, for a sufficiently small value of $\epsilon > 0$. □

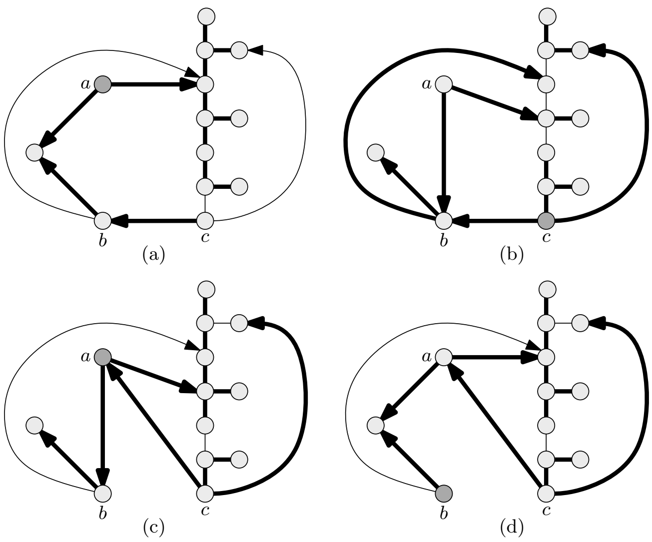

Figure 1. Cycle of best responses for MAXNCG in $\mathcal{M}_2$ for $\alpha \geq 6$. The cycle is obtained by first traversing the configurations (**a**), (**b**), (**c**), and (**d**), in order, and then traversing the configurations obtained from (**b**), (**c**), and (**d**) by exchanging the roles of vertices a and b.

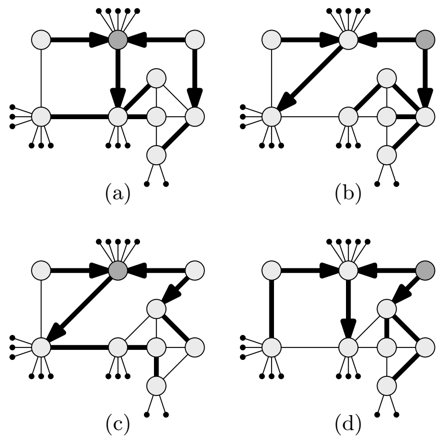

Figure 2. Cycle of best responses for SUMNCG in $\mathcal{M}_2$ for $\alpha \geq 6$. The cycle traverses configurations (**a**), (**b**), (**c**), (**d**), and (**a**), in this order.

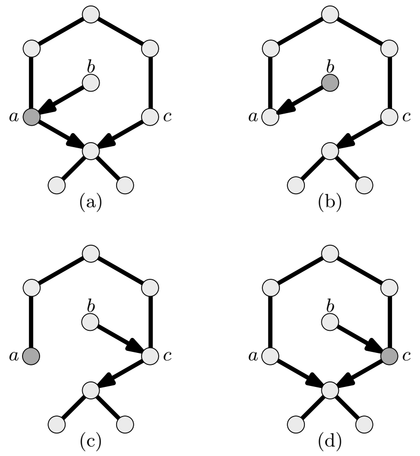

Figure 3. Cycle of best responses for SUMNCG in $\mathcal{M}_3$ for $\alpha = 15$. The cycle is obtained by first traversing the configurations (**a**), (**b**), (**c**), and (**d**), in order, and then traversing the configurations obtained from (**b**), (**c**), and (**d**) by exchanging the roles of vertices a and c.

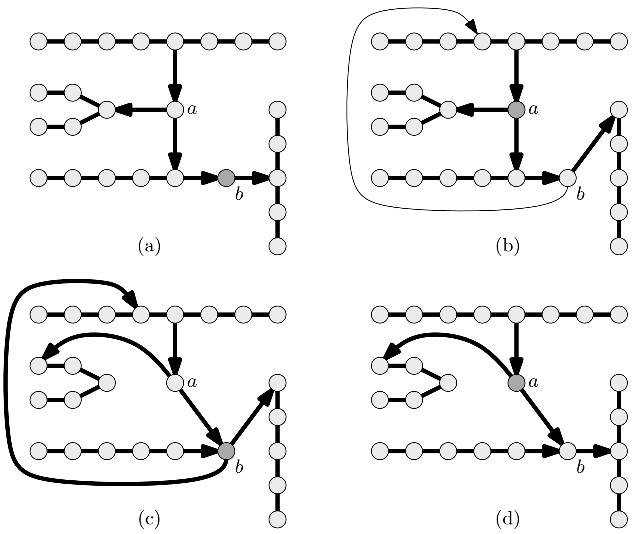

Figure 4. Cycle of best responses for MAXNCG in $\mathcal{M}_3$ for $\alpha = 2 - \epsilon$. The cycle traverses configurations (**a**), (**b**), (**c**), (**d**), and (**a**), in this order.

3. Complexity of Computing a Best Response

We now consider the computational complexity of the problem of computing a best response strategy in our models. For $\mathcal{M}_1$ we will show that it is possible to compute a best response strategy for MAXNCG in polynomial time. The corresponding problem form SUMNCG remains open, although we conjecture that a polynomial-time algorithm exists for this game as well.

Concerning $\mathcal{M}_2$, we prove that a best response strategy can be efficiently computed for SUMNCG as well as MAXNCG. If we consider $\mathcal{M}_3$ instead, then the problem becomes NP-hard for both games.

3.1. Model $\mathcal{M}_1$

In this section, we design a polynomial-time dynamic programming algorithm that computes a best response of a player for MAXNCG in $\mathcal{M}_1$. Fix a strategy profile σ and a player u, and let ℓ be the eccentricity of u in $G(\sigma)$. We will use $\mathcal{G}$ to refer to the set of all graphs that are compatible with u's view in σ. Furthermore, we will use the set $\mathcal{G}'$ consisting of all graphs that can be obtained by removing the edges incident to u from some graph in $\mathcal{G}$. For $i = 0, \ldots, \ell$, let L_i be the set of vertices at distance exactly i from u in $G(\sigma)$. Finally, for $X, Y \subseteq V$ and a graph G having V as its vertex set, we define $d_G(X, Y) = \max_{y \in Y} \min_{x \in X} d_G(x, y)$ and $\delta(X, Y) = \max_{G \in \mathcal{G}'} d_G(X, Y)$. With a slight abuse of notation we will write $\delta(X, y)$ in place of $\delta(X, \{y\})$ whenever $y \in V$.

Our dynamic programming algorithm will exploit the structural properties of best-response strategies in the worst-case graph that are summarized in the following lemmata.

Lemma 2. *Let $\emptyset \subsetneq L'_i \subsetneq L_i$ and $y \in L_j$, with $i, j > 0$. If there exists an index $0 < t \leq \min\{i, j\}$ such that $|L_t| = 1$, then let h be the maximum such index t; otherwise, let $h = \perp$. Then,*

$$\delta(L'_i, y) = \begin{cases} j + i - 2h & \text{if } h \neq \perp; \\ +\infty & \text{if } h = \perp. \end{cases}$$

Proof. We start by considering the case $h \neq \perp$. Let the sole vertex in L_h be z. It is not hard to see that $\delta(L'_i, y) \leq j + i - 2h$ as any $G \in \mathcal{G}'$ must contain (i) some path of length $i - h$ from z to each of the vertices in L_i, and (ii) a path from z to y having a length of $j - h$. Consider now any tree $T \in \mathcal{G}$ in which the lowest common ancestor between y and each vertex in L'_i is z (notice that it is always possible to find such a tree T). For every $x \in L'_i$ we have $d_T(x, y) = j + i - 2h$, implying $\delta(x, y) \geq d_T(x, y) = i + i - 2h$.

We now turn our attention to the case $h = \perp$. Pick any graph $G \in \mathcal{G}'$ such that (i) G is disconnected, (ii) all the vertices in L_i belong to the same connected component $\mathcal{C}$ of G, and (iii) y is not among the vertices in $\mathcal{C}$ (again, it is always possible to find that such a graph G). Then, $\delta(L'_i, y) = \infty$. □

Lemma 3. *Let $y \in L_j$, with $j > 0$. If there exists an index $0 < t \leq j$ such that $|L_t| = 1$, then let h be the maximum such index; otherwise, let $h = \perp$. Then, the following holds $\forall i = 1, \ldots, \ell$,*

$$\delta(L_i, y) = \begin{cases} j - i & \text{if } i \leq j; \\ j + i - 2h & \text{if } h \neq \perp \text{ and } i > j; \\ +\infty & \text{if } h = \perp \text{ and } i > j; \end{cases}$$

Proof. The case $h = \perp$ and $i > j$ has already been handled in the proof of Lemma 2, we therefore focus on the two remaining cases.

When $i \leq j$, we have that $d_G(L_i, y) = j - i$ for all graphs $G \in \mathcal{G}$, thus $\delta(L_i, y) = j - i$. When $i > j$ and $h \neq \perp$, we can consider any tree $T \in \mathcal{G}$ in which the lowest common ancestor between $x \in L_i$ and y is the solve vertex z in L_h. Then we have $d_T(x, y) = j + i - 2h$ for every $x \in L_i$, which implies that $\delta(x, y) = j + i - 2h$. □

Define $S = \{v \in L_1 : u \in \sigma_v\}$ as the set of players v that have bought the edge (v, u) towards u in the strategy profile σ. We can show the following.

Lemma 4. *There exists a best-response strategy σ_u^* for u such that, for every $i = 2, \ldots, \ell$, either $L_i \cap \sigma_u^* = \emptyset$ or $L_i \subseteq \sigma_u^*$. Moreover, at least one of the following two conditions holds: $|L_1| = 1$ or $L_1 \subseteq (\sigma_u^* \cup S)$.*

Proof. We start by noticing that when $|L_1| \geq 2$ and $L_1 \not\subseteq (\sigma_u^* \cup S)$ we can find a network $G \in \mathcal{G}$ that becomes disconnected when player u changes her strategy from σ to σ^*. This implies that σ_u^* cannot be a best response for u and proves the second part of our claim.

Given a possible strategy σ_u' for u, define $f(\sigma_u') = |\{i \in \{2, \ldots, \ell\} \text{ s.t. } \sigma_u' \cap L_i \neq \emptyset \text{ and } L_i \not\subseteq \sigma_u'\}|$. In other words, $f(\sigma_u')$ is the number of levels L_i with $i \geq 2$ for which u is buying some but not all edges towards the vertices in L_i, according to strategy σ_u'.

We prove that if σ_u' is a strategy for which $f(\sigma_u') > 0$, then there exists another strategy σ_u'' for u such that $f(\sigma_u'') < f(\sigma_u')$ and $\Delta(\sigma_u, \sigma_u') \geq \Delta(\sigma_u, \sigma_u'')$. As $f(\sigma_u') > 0$ we are always able to find an index i for which neither $L_i \cap \sigma_u' = \emptyset$ nor $L_i \subseteq \sigma_u'$. Define $L_i' = L_i \cap \sigma_u'$ and let $j \in \{1, 2, \ldots, i-1\}$ be the maximum index for which $|L_j| = 1$, if any. If such an index exists, we let $\sigma_u'' = (\sigma_u' \setminus L_i') \cup L_j$; otherwise (if j does not exist), we let $\sigma_u'' = \sigma_u' \setminus L_i'$. By our choice of σ_u'', we must have both $f(\sigma_u'') < f(\sigma_u')$ and $|\sigma_u''| \leq |\sigma_u'|$. Moreover, invoking Lemmas 2 and 3, we obtain $\delta(L_j, V) \leq \delta(L_i', V)$. It follows that $\delta(\sigma_u'', V) \leq \delta(\sigma_u', V)$ and hence, for every $G \in \mathcal{G}$:

$$C_u((\sigma_{-u}, \sigma_u''), G) = \alpha \cdot |\sigma_u''| + \delta(\sigma_u'', V) \leq \alpha \cdot |\sigma_u'| + \delta(\sigma_u', V) = C_u((\sigma_{-u}, \sigma_u'), G),$$

thus showing that $\Delta(\sigma_u, \sigma_u') \geq \Delta(\sigma_u, \sigma_u'')$. □

We can now design a polynomial-time algorithm to compute a best-response in $\mathcal{M}_1$:

Theorem 5. *In $\mathcal{M}_1$ the best response of a player can be computed in polynomial time for* MAXNCG.

Proof. We will design a dynamic programming algorithm that computes a best-response strategy satisfying the conditions of Lemma 4.

For $i = 0, \ldots, \ell$ and $\eta = 0, \ldots, n-2$, we denote by $A[i, \eta]$ the minimum criminality of a set of vertices X such that $L_i \subseteq X$ and $\delta\big(X, \bigcup_{h=1,\ldots,i} L_h\big) \leq \eta - 2$. The idea is that $A[i, \eta]$ will be (proportional to) the minimum building cost needed to ensure that the distance between u and all the vertices in $\bigcup_{h=1,\ldots,i} L_h$ will not exceed $\eta + 1$ with the restriction that u can only buy edges towards sets of vertices X for which $L_i \subseteq X \subseteq \bigcup_{h=1,\ldots,i} L_h$.

For technical simplicity (and with a small abuse of notation) we will consider L_0 to be the empty set. Moreover, we define $g(i, \eta)$ as the smallest index for which $0 \leq g(i, \eta) \leq i$ and $\delta(L_{g(i,\eta)} \cup L_i, \bigcup_{h=g(i,\eta),\ldots,i} L_h) \leq \eta$. According to the above discussion, the cost of a best-response strategy will be equal to:

$$1 + \min_{\eta=0,\ldots,n-2} \{\eta + \alpha \cdot \min_{\max\{1, g(\ell,\eta)\} \leq j \leq \ell} A[j, \eta]\}.$$

For every $\eta = 0, \ldots, n-1$, we have $A[0, \eta] = 0$ and $A[1, \eta] = |\sigma_u \cap L_1|$ as each edge (u, v) between u and $v \in L_1 \setminus \sigma_u$ is bought by v. Moreover, for $i = 1, \ldots, \ell$ and $\eta = 0, \ldots, n-2$, we can efficiently compute $A[i, \eta]$ using the following identity,

$$A[i, \eta] = |L_i| + \min_{g(i,\eta) \leq j < i} A[j, \eta].$$

After the cost of a best-response strategy σ_u^* for u has been found, one can easily find the strategy σ_u^* itself by retracing the dynamic-programming choices backwards. □

3.2. Model $\mathcal{M}_2$

In model $\mathcal{M}_2$, we are able to design polynomial-time algorithms to compute a best-response strategy for a player for both SUMNCG and MAXNCG, as the following theorem shows.

Theorem 6. *In $\mathcal{M}_2$, the best response of a player can be computed in polynomial time for both* SUMNCG *and* MAXNCG.

Proof. Let σ be a strategy profile, where σ_u denotes the current strategy of some player u, and let $T(u)$ be the shortest-path tree (rooted in u) corresponding to u's view of the network $G(\sigma)$. We now design a dynamic programming algorithm computing a best response σ_u^x for u (in polynomial-time). For SUMNCG and MAXNCG, our algorithm will respectively solve a variant of the *k-median* and the *k-center* problem on trees. The input of the k-median (resp., k-center) problem is a graph H, and the goal is that of finding a set S of exactly k vertices in H in order to minimize the sum of distances (resp. the *maximum* distances) between each node in H and the closest node in S (Actually, the k-median problem asks to minimize the *average* distance between the vertices of H and the closest vertex in S. However, it is easy to see that this is equivalent to minimizing the sum of the above distances).

By deleting u from $T(u)$ we are left with a some forest F. Let $T_1, \dots, T_h$ be the trees in F and let r_i be the root of T_i. Let $z = |\sigma_u|$. We can assume (w.l.o.g.) that the trees T_i with $i \leq z$ (for some value of z) are those for which u owns the edge (u, r_i) and that all edges (u, r_j) for $j > z$ are owned by r_j.

For $i \leq z$ and $0 \leq k < n$, we let $A[i,k]$ be the cost of an optimal solution the instance of the k-median (resp. k-center) problem in which $H = T_i$. Similarly, for $z < i \leq h$ and $0 \leq k < n$, we let $A[i,k]$ be the cost of an optimal solution to the instance of a *constrained version* of the k-median (resp. k-center) problem in which $H = T$ and we additionally require the solution to contain vertex r_i. Whenever there is no feasible solution to the above instances (e.g., when $k \geq |V(T_i)|$) we let $A[i,k] = +\infty$. It is well known both the k-median and the k-center problem can be solved in polynomial time when the input graph H is a trees (and it is easy to see that this is still true for their constrained variants).

We now consider the first i trees and, for each value of j, we compute the minimum *aggregate cost* $B[i,j]$ that can be obtained by solving i (possibly constrained) instances of k-median (resp. k-center) on the trees $T_1, \dots, T_i$ while selecting exactly j vertices *across all such instances*. More formally, for the k-median problem, we define

$$B[i,j] = \min_{\substack{j_1, j_2, \dots, j_i \in \mathbb{N} \\ j_1 + j_2 + \dots + j_i = j}} \sum_{\ell=1}^{i} A[\ell, j_\ell],$$

while, for the k-center problem we define

$$B[i,j] = \min_{\substack{j_1, j_2, \dots, j_i \in \mathbb{N} \\ j_1 + j_2 + \dots + j_i = j}} \max \left\{ A[\ell, j_\ell] \;:\; 1 \leq \ell \leq i \right\}.$$

Notice that, by definition, $B[1,j] = A[1,j]$ for both problems. For $i > 1$, we can efficiently compute $B[i,j]$ using the following recursive formulas. For the k-median problem we have

$$B[i,j] = \min_{1 \leq t \leq j-i+1} \{ A[i,t] + B[i-1, j-t] \},$$

while, for the k-center problem,

$$B[i,j] = \min_{1 \leq t \leq j-i+1} \max\{ A[i,t], B[i-1, j-t] \}.$$

To find the cost of u's best response σ_u^* in SUMNCG, we observe that $n-1+B[h,j]$ is exactly the smallest usage cost incurred by u if she decides to buy exactly $j-(h-z)$ edges (as $h-z$ edges are already owned by the vertices r_i with $z+1 \le i \le h$). We therefore know that, in SUMNCG,

$$C_u((\sigma_{-u}, \sigma_u^*), G) = n - 1 + \min_{h-z \le j < n} \{B[h,j] + \alpha(j - h + z)\}.$$

Similarly, $1 + B[h,j]$ is the smallest usage cost for u when he/she constructs exactly $j-(h-z)$ edges in MAXNCG, therefore

$$C_u((\sigma_{-u}, \sigma_u^*), G) = 1 + \min_{h-z \le j < n} \{B[h,j] + \alpha(j - h + z)\}.$$

In both cases, the number of edges $|\sigma_u^*|$ that u buys in a best response is $\arg\min_{h-z \le j < n} \{B[h,j] + \alpha(j-h+z)\}$. Retracing the dynamic programming choices backwards, it is possible to compute (in polynomial time) the optimal number j_i^* of edges to be bought towards the vertices in each tree T_i. The actual edges to activate in SUMNCG (resp. MAXNCG) can be found by inspecting the optimal solutions for the (possibly constrained) k-median (reps. k-center) instances associated with each $A[i, j_i^*]$. □

3.3. Model $\mathcal{M}_3$

As far as model $\mathcal{M}_3$ is concerned, the problem of computing a best response is NP-hard for both SUMNCG and MAXNCG, as shown by the following results:

Theorem 7. *Computing a best response for a player in $\mathcal{M}_3$ for* SUMNCG *is* NP*-hard.*

Proof. We probe the claim by providing a polynomial-time reduction from the (decision version of) the *Minimum Dominating Set problem* (MDSP) on bipartite graphs to the problem of computing a best response in $\mathcal{M}_3$, similarly to that used in [1].

An instance of MDSP is a pair $\langle G, k \rangle$ where G is a graph and k is a positive integer. The problem is that of deciding whether there exists a subset D of vertices of H such that (i) each vertex $v \in V'$ is either in D or is adjacent to a vertex in D, and (ii) $|D| \le k$. It is well-known that MDSP is hard even when H is bipartite [23].

Let $\langle H, k \rangle$ be an instance of MDSP where $H = (U' \cup V', E')$ is a bipartite graph. Fix any value of $\alpha \in (1,2)$ and let G be the graph obtained from H by adding an additional vertex (player) u along with all the edges in $\{u\} \times U'$. Consider any strategy profile σ such that $G(\sigma) = G$ and all edges incident to u are bought solely by u. Notice that each edge $e \in E'$ belongs to at least one shortest path from u in G and therefore the view of u coincides with G itself.

Let σ_u^* be a best response strategy for u and notice that the eccentricity of u in $G(\sigma_{-u}, \sigma_u^*)$ must be at most 2. Indeed, if a vertex v was at a distance of at least 3 from u in $G(\sigma_{-u}, \sigma_u^*)$, then u could improve her cost by playing $\sigma_u^* \cup \{u\}$ instead of σ_u^* (thus decreasing the distance to v by at least 2 while increasing her building cost by $\alpha < 2$). It follows that σ_u^* must be a dominating set of H.

We now argue that σ_u^* must be a minimum dominating set of H. Indeed, given any strategy σ_u' of player u for which the eccentricity of u in $G(\sigma_{-u}, \sigma_u)$ is at most 2, we have that

$$\begin{aligned} C_u(\sigma_{-u}, \sigma_u') &= \alpha|\sigma_u'| + \sum_{v \in V'} d_{G((\sigma_{-u}, \sigma_u'))}(u,v) \\ &= \alpha|\sigma_u'| + \sum_{v \in \sigma_u'} d_{G((\sigma_{-u}, \sigma_u'))}(u,v) + \sum_{v \in V' \setminus \sigma_u'} d_{G((\sigma_{-u}, \sigma_u'))}(u,v) \\ &= \alpha|\sigma_u'| + |\sigma_u'| + 2(|V'| - |\sigma_u'|) = (\alpha - 1)|\sigma_u'| + 2|V'|, \end{aligned}$$

which is minimized when $|\sigma_u'|$ is minimized.

The above discussion implies that H admits a dominating set D with $|D| \leq k$ vertices if $|\sigma^*| \leq k$ and concludes the proof. □

Theorem 8. *Computing a best response for a player in $\mathcal{M}_3$ for* MAXNCG *is* NP*-hard.*

Proof. As in the proof of Theorem 7, we prove the claim through a reduction (similar to than in [3]), from the *Minimum Dominating Set problem* (MDSP) on bipartite graphs [23]. Given an instance $\langle H, k \rangle$ of (the decision version) of MDSP on bipartite graphs, we further assume w.l.o.g. that $H = (U' \cup V', E')$ admits a dominating set D^* of size smaller than $\frac{|U' \cup V'|}{2}$ (If this is not the case, then it suffices to consider the graph $\overline{H}$ obtained from the disjoint union of H with a star on $|U' \cup V'| + 2$ vertices centered in some vertex v. H' has $2|U' \cup V'| + 2$ vertices and admits a minimum dominating set of cardinality at most $|U' \cup V'| + 1$. Moreover, $D \subseteq U' \cup V'$ is a dominating set for H iff $D \cup \{v\}$ is a dominating set of $\overline{H}$).

Pick $\alpha = \frac{2}{|U' \cup V'|}$ and let G be the graph obtained from H by adding an additional vertex (player) u along with all the edges in $\{u\} \times U'$. Consider any strategy profile σ such that $G(\sigma) = G$ and all edges incident to u are bought solely by u, so that the view of u coincides with G itself.

Let σ_u^* be a best response strategy for u and let $\mathcal{E}$ be the eccentricity of u in $G(\sigma_{-u}, \sigma_u^*)$. Clearly, $\sigma^* \neq \emptyset$. If $\mathcal{E} > 2$, then u could buy an edge towards every vertex in G increasing her building cost by at most $\alpha|(U' \cup V') \setminus \sigma^*| < 2$ and reducing her usage cost by $\mathcal{E} - 1 \geq 2$. If $\mathcal{E} = 1$, then $|\sigma^*| = |U' \cup V'|$ and u could delete all the edges towards vertices that are not in D^*, thus saving at least $\alpha \frac{|U' \cup V'|}{2} = 2$ on his/her building cost, while increasing the usage cost by 1. We can therefore conclude that $\mathcal{E} = 2$, i.e., σ^* must be a dominating set of H.

As each dominating sets D of H is also a strategy for u such that u has cost $\alpha|D| + 2$ in $G(\sigma_{-u}, D)$, we have that H admits a dominating set D with $|D| \leq k$ vertices iff $|\sigma^*| \leq k$. □

4. Price of Anarchy

Models $\mathcal{M}_1$, $\mathcal{M}_2$ and $\mathcal{M}_3$ are equivalent with respect to the PoA, as shown by the following two Theorems.

Theorem 9. *The PoA for* SUMNCG *in $\mathcal{M}_1$, $\mathcal{M}_2$, and $\mathcal{M}_3$ is $\Theta(\min\{1 + \alpha, n\})$ for every α.*

Proof. Notice that each network that is compatible with the view of u in model $\mathcal{M}_2$ (resp. $\mathcal{M}_3$) must also be compatible with the view of u in $\mathcal{M}_1$ (resp. $\mathcal{M}_2$). This shows that the PoA can only decrease when we consider $\mathcal{M}_1$, $\mathcal{M}_2$, and $\mathcal{M}_3$ in this order. We will prove the claim by showing that $O(\min\{1 + \alpha, n\})$ is an upper bound for $\mathcal{M}_1$ and $\Omega(\min\{1 + \alpha, n\})$ a lower bound in $\mathcal{M}_3$.

We start by proving our lower bound in $\mathcal{M}_3$. To this aim, consider the complete graph and notice that it is an equilibrium since the only SPT rooted in each vertex (player) is a star, and no edge can be removed. The corresponding social cost is $\Omega(\alpha n^2 + n^2)$, while the cost of an optimal network is upper bounded by that of a star, i.e., $O(\alpha n + n^2)$. It follows that the PoA must be $\Omega(\min\{n, 1 + \alpha\})$.

We now turn our attention to the upper bound in $\mathcal{M}_1$. We start by showing that the diameter of an equilibrium graph G can be at most $O(1 + \sqrt{\alpha})$. Let u and v be two players such that the eccentricity D of u in G is equal to $d(u, v) \geq 4$ (we can assume that such a pair exists as otherwise the diameter of G would be at most 3). If player u were to buy the edge (u, v), he/she would decrease her distances towards at least $\frac{D}{4}$ vertices, namely, those on the shortest path between u and v, by at least $\frac{3}{4}D - \frac{1}{4}D - 1 = \frac{1}{2}D - 1$. This would decrease her usage cost by at least $\Omega(D^2)$. Notice that this reasoning is compatible with the local view of u in model $\mathcal{M}_1$, since u only needs to know $d(u, v)$ (while no knowledge of the internal vertices of π is necessary). As the edge (u, v) is not in G, we can infer that $\alpha = \Omega(D^2)$, which implies $D = O(1 + \sqrt{\alpha})$.

The PoA can then be upper bounded as follows,

$$O\left(\frac{\alpha n^2 + n^2(1 + \sqrt{\alpha})}{\alpha n + n^2}\right) = O(\min\{n, 1 + \alpha\}). \qquad \square$$

Theorem 10. *The PoA for* MAXNCG *in* $\mathcal{M}_1$, $\mathcal{M}_2$, *and* $\mathcal{M}_3$ *is* $\Theta(n)$ *for* $\alpha = \Omega(1)$, *and* $\Theta(1 + \alpha n)$ *if* $\alpha = O(1)$.

Proof. By similar arguments as the ones used in the proof of Theorem 9, we only need to show an upper bound to the PoA in model $\mathcal{M}_1$ and a lower bound to the PoA in model $\mathcal{M}_3$.

Once again, we can consider a complete graph H to obtain our lower bound. The same discussion in the proof of Theorem 9 also applies to MAXNCG, and shows that H is an equilibrium network in $\mathcal{M}_1$. The social cost of H is $\Omega(\alpha n^2 + n)$, while that of a star is $O(\alpha n + n)$. As a consequence, we obtain a lower bound on the PoA of $\Omega\left(\frac{\alpha n^2 + n}{\alpha n + n}\right)$, which reduces to $\Omega(n)$ for $\alpha = \Omega(1)$, and to $\Omega(1 + \alpha n)$ for $\alpha = O(1)$.

To prove our upper bound on $\mathcal{M}_3$ we show that, in any equilibrium graph G, the diameter of G must be at most $O(1 + \alpha n)$. Let u be a player having eccentricity $D \geq 4$ in G, define X as the set of vertices at distance $\lfloor D/2 \rfloor$ from u in G. If u were to buy all edges in $\{u\} \times X$, his/her eccentricity would decrease by $\Omega(D)$. As G is an equilibrium, we know that the building cost must exceed $\Omega(D)$, in formulas: $\alpha n \geq \alpha \cdot |X| = \Omega(D)$, i.e., $D = O(1 + \alpha n)$. This shows that the social cost of G must be at most $O(\alpha n^2 + n + n^2)$ and we can upper bound the PoA as $\frac{O(\alpha n^2 + n + n^2)}{\Omega(\alpha n + n)}$, which yields the sought bounds when instantiated with $\alpha = \Omega(1)$ and $\alpha = O(1)$. $\square$

Author Contributions: Writing— original draft, D.B., L.G., S.L. and G.P. All authors have read and agreed to the published version of the manuscript.

Funding: This work was partially supported by the “Fondo di Ateneo per la Ricerca 2019” grant of the University of Sassari and by the project E89C20000620005 “ALgorithmic aspects of BLOckchain TECHnology” (ALBLOTECH).

Data Availability Statement: Data sharing not applicable.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Fabrikant, A.; Luthra, A.; Maneva, E.; Papadimitriou, C.H.; Shenker, S. On a network creation game. In Proceedings of the ACM 22nd Symposium on Principles of Distributed Computing (PODC), Boston, MA, USA, 13–16 July 2003; pp. 347–351.
2. Albers, S.; Eilts, S.; Even-Dar, E.; Mansour, Y.; Roditty, L. On Nash Equilibria for a Network Creation Game. *ACM Trans. Econ. Comput.* **2014**, *2*, 2:1–2:27. [CrossRef]
3. Mihalák, M.; Schlegel, J.C. The price of anarchy in network creation games is (mostly) constant. *Theory Comput. Syst.* **2013**, *53*, 53–72. [CrossRef]
4. Mamageishvili, A.; Mihalák, M.; Müller, D. Tree Nash equilibria in the network creation game. *Internet Math.* **2015**, *11*, 472–486. [CrossRef]
5. Àlvarez, C.; Messegué, A. Network Creation Games: Structure vs Anarchy. *arXiv* **2017**, arXiv:1706.09132.
6. Àlvarez, C.; Messegué, A. On the Price of Anarchy for High-Price Links. In Proceedings of the Web and Internet Economics—15th International Conference (WINE 2019), New York, NY, USA, 10–12 December 2019; Caragiannis, I., Mirrokni, V.S., Nikolova, E., Eds.; Lecture Notes in Computer Science; Springer: Berlin/Heidelberg, Germany, 2019; Volume 11920, pp. 316–329. [CrossRef]
7. Demaine, E.D.; Hajiaghayi, M.T.; Mahini, H.; Zadimoghaddam, M. The Price of Anarchy in Network Creation Games. *ACM Trans. Algorithms (TALG)* **2012**, *8*, 13. [CrossRef]
8. Bilò, D.; Lenzner, P. On the Tree Conjecture for the Network Creation Game. *Theory Comput. Syst.* **2020**, *64*, 422–443. [CrossRef]
9. Alon, N.; Demaine, E.D.; Hajiaghayi, M.T.; Leighton, T. Basic network creation games. *SIAM J. Discret. Math.* **2013**, *27*, 656–668. [CrossRef]

10. Lenzner, P. Greedy selfish network creation. In *International Workshop on Internet and Network Economics (WINE)*; Springer: Berlin/Heidelberg, Germany, 2012; pp. 142–155.
11. Mihalák, M.; Schlegel, J.C. Asymmetric swap-equilibrium: A unifying equilibrium concept for network creation games. In *International Symposium on Mathematical Foundations of Computer Science (MFCS)*; Springer: Berlin/Heidelberg, Germany, 2012; pp. 693–704.
12. Laoutaris, N.; Poplawski, L.J.; Rajaraman, R.; Sundaram, R.; Teng, S. Bounded Budget Connection (BBC) games or how to make friends and influence people, on a budget. *J. Comput. Syst. Sci.* **2014**, *80*, 1266–1284. [CrossRef]
13. Ehsani, S.; Fadaee, S.S.; Fazli, M.; Mehrabian, A.; Sadeghabad, S.S.; Safari, M.; Saghafian, M. A bounded budget network creation game. *ACM Trans. Algorithms (TALG)* **2015**, *11*, 34. [CrossRef]
14. Bilò, D.; Gualà, L.; Proietti, G. Bounded-distance network creation games. *ACM Trans. Econ. Comput. (TEAC)* **2015**, *3*, 16. [CrossRef]
15. Bilò, D.; Gualà, L.; Leucci, S.; Proietti, G. The max-distance network creation game on general host graphs. *Theor. Comput. Sci.* **2015**, *573*, 43–53. [CrossRef]
16. Bilò, D.; Friedrich, T.; Lenzner, P.; Melnichenko, A. Geometric Network Creation Games. In Proceedings of the 31st ACM on Symposium on Parallelism in Algorithms and Architectures (SPAA 2019), Phoenix, AZ, USA, 22–24 June 2019; Scheideler, C., Berenbrink, P., Eds.; ACM: New York, NY, USA, 2019; pp. 323–332. [CrossRef]
17. Bilò, D.; Gualà, L.; Leucci, S.; Proietti, G. Locality-based network creation games. In Proceedings of the 26th ACM on Symposium on Parallelism in Algorithms and Architectures (SPAA 2014), Prague, Czech Republic, 23–25 June 2014; Blelloch, G.E., Sanders, P., Eds.; ACM: New York, NY, USA, 2014; pp. 277–286. [CrossRef]
18. Bilò, D.; Gualà, L.; Leucci, S.; Proietti, G. Locality-Based Network Creation Games. *ACM Trans. Parallel Comput.* **2016**, *3*, 6:1–6:26. [CrossRef]
19. Bampas, E.; Bilò, D.; Drovandi, G.; Gualà, L.; Klasing, R.; Proietti, G. Network verification via routing table queries. *J. Comput. Syst. Sci.* **2015**, *81*, 234–248. [CrossRef]
20. Beerliova, Z.; Eberhard, F.; Erlebach, T.; Hall, A.; Hoffmann, M.; Mihalák, M.; Ram, L.S. Network Discovery and Verification. *IEEE J. Sel. Areas Commun.* **2006**, *24*, 2168–2181. [CrossRef]
21. Bilò, D.; Gualà, L.; Leucci, S.; Proietti, G. Network creation games with traceroute-based strategies. In *International Colloquium on Structural Information and Communication Complexity (SIROCCO)*; Springer: Berlin/Heidelberg, Germany, 2014; pp. 210–223.
22. Kawald, B.; Lenzner, P. On dynamics in selfish network creation. In Proceedings of the 25th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), Montreal, QC, Canada, 23–25 July 2013; pp. 83–92.
23. Chlebík, M.; Chlebíkova, J. Approximation hardness of dominating set problems. In *European Symposium on Algorithms*; Springer: Berlin/Heidelberg, Germany, 2004; pp. 192–203.