Asymmetric Swap-Equilibrium: A Unifying Equilibrium Concept for Network Creation Games

Matúš Mihalák and Jan Christoph Schlegel

Institute of Theoretical Computer Science, ETH Zurich, Switzerland
matus.mihalak@inf.ethz.ch, jansc@student.ethz.ch

Abstract. We introduce and study the concept of an *asymmetric swap-equilibrium* for network creation games. A graph where every edge is *owned* by one of its endpoints is called to be in *asymmetric swap-equilibrium*, if no vertex v can delete its own edge $\{v, w\}$ and add a new edge $\{v, w'\}$ and thereby decrease the sum of distances from v to all other vertices. This equilibrium concept generalizes and unifies some of the previous equilibrium concepts for network creation games. While the structure and the quality of equilibrium networks is still not fully understood, we provide further (partial) insights for this open problem. As the two main results, we show that (1) every asymmetric swap-equilibrium has at most one (non-trivial) 2-edge-connected component, and (2) we show a logarithmic upper bound on the diameter of an asymmetric swap-equilibrium for the case that the minimum degree of the unique 2-edge-connected component is at least n^{ε}, for $\varepsilon > \frac{4 \lg 3}{\lg n}$. Due to the generalizing property of asymmetric swap equilibria, these results hold for several equilibrium concepts that were previously studied. Along the way, we introduce a node-weighted version of the network creation games, which is of independent interest for further studies of network creation games.

1 Introduction

Many communication networks (such as the Internet) are planned, maintained and built locally by individual entities (such as the autonomous systems). This new phenomenon contrasts with centrally planned and built networks. Network creation games study the quality and structure of communication networks that are created in this non-central way.

The first and arguably the most prominent game-theoretic consideration in this field is the *(original) network creation game* [5]) – a strategic game parameterized by a value $\alpha > 0$ in which the players buy adjacent edges at the price α each, and aim to minimize the cost expressed as the usage cost plus the cost for the bought edges, where a usage cost is the sum of all distances from the respective player. A series of papers [5,1,3,7] studied the *structure* and the *price of anarchy* of (Nash) equilibria in network creation games. The price of anarchy is expressed as the social cost of the worst equilibrium network divided by the

B. Rovan, V. Sassone, and P. Widmayer (Eds.): MFCS 2012, LNCS 7464, pp. 693–704, 2012.

social cost of an optimum (centrally-planned) network (where, a social cost is the sum of all individual costs of the players). It is believed that the price of anarchy is constant for all values of α. This has been shown for almost all values of α – with the exception of the range $n^{1-\varepsilon} \leq \alpha \leq 273 \cdot n$ (for $\varepsilon \geq 1/\lg n$). It remains a major open problem to prove/disprove the conjecture. In each of these papers, a completely new proof-technique for giving an upper bound on the price of anarchy has been developed, largely depending on the parameter α. To overcome this rather unnatural "dependency" on α, Alon et al. [2] defined and studied a new equilibrium concept: a graph is called to be in *swap equilibrium*, if no vertex v (a player) can delete an existing edge $\{v, w\}$ and add a new edge $\{v, w'\}$ and thereby decrease the usage cost of v. Swap equilibria indeed do not depend on α, but they do fail to generalize Nash equilibria. The authors conjecture that swap equilibria have at most polylogarithmic diameter. We give a partial affirmation of this.

Another approach to "get rid of" α (although having a different motivation than Alon et al. [2]) has been presented by Ehsani et al. [4] which studied the so called *bounded-budget network creation game* – a variant of the original network creation game in which the players' only goal is to minimize their usage cost (as opposed to minimizing the usage cost plus the creation cost in the original game) given that every player can buy at most a given number of edges. While conceptually different, these two approaches enjoy several similarities, which shall become obvious with our work.

Driven by the desire to answer the aforementioned open problem, and following the original idea of Alon et al. to "get rid of" of α, we define and study the *asymmetric swap-equilibrium*, a natural modification of the swap equilibrium: a graph where every edge is *owned* by one of its endpoints is called to be in *asymmetric swap-equilibrium*, if no vertex v can delete its own edge $\{v, w\}$ and add a new edge $\{v, w'\}$ and thereby decrease the usage cost of v. Asymmetric swap-equilibria have, on top of the inherent properties of swap equilibria (such as that best responses can be calculated efficiently), the interesting property that they generalize swap equilibria and they also generalize Nash equilibria in both the original and the bounded-budget network creation games. Thus, any quality and structural "upper bounds" that one proves for asymmetric swap-equilibrium immediately hold for these equilibrium concepts as well.

As solving the main open problem seems to be difficult (as evidenced by the many papers on the topic that only partially solve it), we are also interested in partial results towards this direction. Besides the quality of asymmetric swap-equilibria, we thus also study their structure, which helps understanding equilibrium graphs. In fact, analyzing the diameter of equilibrium networks is another main open problem.

Definition of the Problem and Related Concepts. For every (undirected) graph G we use the following notation. We denote the vertex set of G by $V(G)$ and its edge set by $E(G)$. For $u, v \in V(G)$ we denote by $d_G(u, v)$ the length of a shortest u-v-path in G. If G is not connected we define $d_G(u, v) := \infty$. We denote the diameter of G by $\text{diam}(G)$ and the radius by $\text{rad}(G)$. We sometimes omit writing G and write simply $d(u, v)$, diam or rad if the underlying graph G is clear from the context. Recall that a *2-edge-connected graph* is a graph of size at least 2 that does not contain a *bridge*, i.e. an edge whose removal makes the graph disconnected, and that a *2-edge-connected component* of a graph G is a maximal 2-edge-connected subgraph of G.

The (original) *network creation game* (as defined by Fabrikant et al. [5]) is a strategic game parameterized by a positive real number α called the *edge price.* The game is played by n players $[n]$ representing nodes in a graph, where a strategy s_i of a player $i \in [n]$ is a set of adjacent edges that it *buys* (or *builds*). The (played) strategies $s = (s_1, s_2, \dots, s_n)$ of the players naturally define the graph $G(s) = ([n], \bigcup s_i)$. The *creation cost* of player i in the game is $\alpha|s_i|$, i.e., the amount it pays for the edges s_i, and the *usage cost* of player i in the game is the sum of distances from i to all other nodes in the graph $G(s)$. The *cost function* $c_i(s)$ of player i is expressed as the creation cost plus the usage cost. A *Nash equilibrium* (NE for short) of the network creation game are strategies $s = (s_1, \dots, s_n)$ of the players such that no player can lower her current cost $c_i(s)$ by changing its chosen strategy s_i to a different one. It is easy to see that for every finite α every NE induces a connected graph. We call a graph induced by a Nash equilibrium a *stable graph* or simply, by abusing the notation a bit, a *Nash equilibrium.*

The *bounded-budget network creation game* (as defined and studied by Ehsani et al. [4]) is a network creation game without the parameter α where every player i has a *budget* b_i on the number of edges it can buy. The set of strategies S_i contains only sets of adjacent edges of i of cardinality no more than b_i. The cost function c_i of player i is then just the usage cost of the original network creation game. We will see that several of the results for the bounded-budget network creation game carry over to our model studied in this paper.

The *basic network creation game* (as defined and studied by Alon et al. [2]) is not a strategic game. Rather, it is an equilibrium concept for graphs. Analogously to network creation games, each vertex possesses a cost function (which is the usage cost of the original network creation game). A graph G is called to be in *swap equilibrium* if no vertex $v \in V(G)$ can improve its cost by deleting an adjacent edge $\{v, w\} \in E(G)$ and creating a new adjacent edge $\{v, w'\}$.

In this paper we define the *asymmetric swap-equilibrium* based on the *ownership* of edges. An *ownership* of a graph G is a function $o : E(G) \to V(G)$ that assigns to every edge $\{u, v\} \in E(G)$ either u or v. If $o(\{u, v\}) = u$ we say that u *owns* the edge $\{u, v\}$ and that the edge $\{u, v\}$ is owned by u. Again, every vertex (a player) has a cost – the usage cost of the original network creation game. A graph G with an ownership o is called to be in *asymmetric swap-equilibrium* if no vertex $v \in V(G)$ can improve its cost by deleting its *own* adjacent edge $\{v, w\} \in E(G)$ and creating a new adjacent edge $\{v, w'\}$. Such a modification is called a *swap* (of an edge), and results in a modified graph and modified ownership in that the newly created edge is owned by the vertex v.

Asymmetric swap equilibria generalize these equilibrium concepts in the following sense. Every stable graph $G(s)$ of the (original) network creation game induces an ownership o in which each edge is owned by the player which bought it in the Nash equilibrium s (and observe that in a Nash equilibrium no edge is bought by two players). It is easy to see that for this ownership o, the graph $G(s)$ is in asymmetric swap equilibrium. One can make similar arguments about Nash equilibria of the bounded-budget network creation games. Furthermore, a swap equilibria of the basic network creation game is an asymmetric swap-equilibrium for any ownership o. Thus, we have:

Proposition 1. *Every stable graph of the original network creation game, every stable graph of the bounded-budget network creation game, and every swap equilibrium graph is an asymmetric swap-equilibrium graph.*

Another motivation to study (asymmetric) swap equilibria is that computing a best swap of a player is easy (i.e., polynomial), while computing the best strategy s_i of a player i (given the strategies of all other players) in the original/bounded-budget network creation game is an NP-hard problem [5,4]. We note that this is also true for the basic network creation game if an arbitrary number of swaps is permitted [6].

2 The Structure of Asymmetric Swap-Equilibria

In the following we state and prove the main result of this paper.

Theorem 1. *Every graph in asymmetric swap-equilibrium has at most one* 2*-edge-connected component.*

Proof. Let G be an asymmetric swap-equilibrium and assume for contradiction that there are two 2-edge-connected components $H_1, H_2 \subset G$. Assign every vertex $v \in V(G)$ whose shortest v-H_1-path passes through H_2 to H_2 and every vertex $v \in V$ whose shortest v-H_2-path passes through H_1 to H_1. Denote by $\tilde{H}_1$ the vertices assigned to H_1 and by $\tilde{H}_2$ the vertices assigned to H_2. Without loss of generality $\tilde{H}_1$ is the smallest of the two, i.e., $|\tilde{H}_1| \leq |\tilde{H}_2|$. Let $\{x_1, x_2\} \in E(G)$ with $x_1 \in H_1$ be the first edge in the (unique) shortest H_1-H_2-path in G (and observe that $\{x_1, x_2\}$ is a bridge of G). A schematic illustration of the situation is depicted in Figure 1.

The main idea of the proof is the following. Observe that every vertex u of H_1, with the exception of x_1, decreases its distance to every vertex of $\tilde{H}_2$, if an edge $\{u, x_2\}$ is added to G. We will show that there always exists a vertex u in H_1 that owns an edge $\{u, w\}$ such that the deletion of the edge does not increase the usage cost too much so that this vertex v can swap $\{u, w\}$ for $\{u, x_2\}$ and improve its usage cost – a contradiction with the assumption that the graph is in asymmetric swap-equilibrium. We consider three cases: Either there is a vertex u which owns an edge $\{u, v\} \in E(H_1)$ such that v is closer to x_1 than u, or this is not the case but there is a vertex u which owns an edge $\{u, v\} \in E(H_1)$ such that u and v are at the same distance from x_1, or neither is the case and therefore every edge $\{u, v\} \in E(H_1)$ has both its vertices at different distances

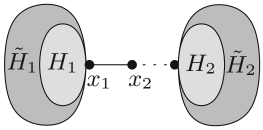

Fig. 1. 2-edge-connected components

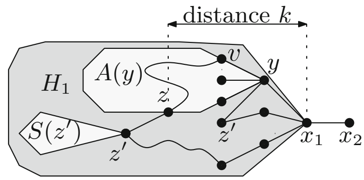

Fig. 2. Case 3 of the proof of Theorem 1

from x_1 and the vertex which owns the edge is closer to x_1 than the other vertex of the edge.

Case 1. In the first case, consider the swap of the edge $\{u, v\}$ with the edge $\{u, x_2\}$. The distance of u to every vertex of $\tilde{H}_2$ decreases by $d(u, x_2) - 1 = d(u, x_1)$ and the distance of u to every vertex of $\tilde{H}_1 \setminus \{u\}$ increases by at most $d(v, x_2) = d(u, x_1)$. The distance of u to any other vertex does not increase. Hence, as $|\tilde{H}_1 \setminus \{u\}| < |\tilde{H}_2|$, u could improve by swapping, a contradiction.

Extra Notation. We will consider all vertices of $\tilde{H}_1$ aligned in *layers* according to the distance from x_1; all vertices at distance k from x_1 will be referred to as *layer* k. We will also call an edge a *layer-edge* if both its endpoints lie in the same layer. Let p be a path connecting two vertices a and b in $\tilde{H}_1$ and $\{u, v\}$ an edge in it, where u is from layer k and v from layer $k+1$ for some k. We can consider the path as oriented from a to b or vice versa. For a considered orientation, we call $\{u, v\}$ a *forward-edge* in p if u precedes v on p. Similarly, we call $\{u, v\}$ a *backward-edge* in p if v precedes u on p. Thus, "forward"/"backward" reflects on the progression of the path away from x_1.

Case 2. Consider now the second case where a vertex u owns an edge $\{u, v\}$ such that u and v are at the same distance k from x_1, i.e., $\{u, v\}$ is a layer-edge. Consider the swap of edge $\{u, v\}$ with edge $\{u, x_2\}$. Recall that the endpoints of the edge $\{u, v\}$ lie in the same layer, and therefore the distance from x_1 to every vertex of $\tilde{H}_1$ cannot increase after the swap. Let us investigate how much the distances from u to $\tilde{H}_1$ could increase. Let us first consider the simple case when $k = 1$. Then the increase of the distance from u to vertices in $\tilde{H}_1 \setminus \{u\}$ is at most $d(u, x_1) + d(x_1, v) - 1 = 1$. Vertex u decreases its distance to all vertices in $\tilde{H}_2$ by $d(u, x_2) - 1 = 1$, and as $|\tilde{H}_2| > |\tilde{H}_1 \setminus \{u\}|$, u improves its usage cost by the swap, a contradiction. We therefore assume that $k > 1$. The distances from u to every vertex of $\tilde{H}_2$ decrease by $d(u, x_2) - 1 = k$. Let us now consider the increase of the distances from u. In general, the length of a shortest path from u to a vertex $w \in \tilde{H}_1$ could change after the swap, but the length is upper bounded by the length of the u-w-path that uses the new edge $\{u, x_2\}$ and goes via x_1. Obviously, after the swap, u can increase its distance only to vertices w for which a shortest u-w-path p_w uses the edge $\{u, v\}$. We classify the vertices w according to the shortest u-w-path p_w before the swap: (i) path p_w contains, besides $\{u, v\}$, only forward-edges, (ii) p_w contains, besides $\{u, v\}$, exactly one layer-edge and forward-edges, (iii) p_w is none of the first two. Let $S(i)$, $S(ii)$ denote, respectively, the vertices w for which p_w satisfies (i) and (ii). By the swap, the distances of u to vertices in $S(i)$ increase by at most

$1+1+d(x_1,v)-1=k+1$. The distances of u to vertices in $S(ii)$ increase by at most $2+d(x_1,w)-d(u,w)\le k$. The distances of u to other vertices w increase by at most $k-1$. Comparing the total increase and decrease of the distances from u, the total decrease of the usage cost of vertex u is at least

$$k\cdot|\tilde{H}_2|-(k+1)\cdot|S(i)|-k\cdot|S(ii)|-(k-1)\cdot(|\tilde{H}_1\setminus\{u\}|-|S(i)|-|S(ii)|) > |\tilde{H}_1|-2|S(i)|-|S(ii)|.$$

If $|\tilde{H}_1|-2|S(i)|-|S(ii)|\ge 0$ then, by the above equation, vertex u improves its usage cost by the swap, a contradiction. Assume therefore that $|\tilde{H}_1|-|S(i)|-|S(ii)|<|S(i)|$. But then $S(i)$ contains a lot of vertices and vertex x_1 can swap the incident edge $\{x_1,y\}$ in a shortest x_1-v-path (which x_1 owns as we are not in case 1) with the edge $\{x_1,v\}$ and improve its usage cost: observe that such a swap decreases the distances of x_1 to $S(i)$ by $k-1$, does not increase the distances to $S(ii)$, and increases the distances to $\tilde{H}_1\setminus\{S(i)\cup S(ii)\}$ by at most $1+d(v,y)-1=k-1$. This is a contradiction.

Further Extra Notation. For the analysis of the third case we introduce the following notation. For every vertex $v\in H_1$ we define $S(v)\subset\tilde{H}_1$ to be the set of vertices $u\in\tilde{H}_1$ such that some (but not necessarily every) shortest u-x_1-path passes through v, and we define $A(v)\subset S(v)$ to be the set of vertices $u\in\tilde{H}_1$ such that all shortest u-x_1-paths pass through v. Note that by definition $v\in A(v)$.

Case 3. Consider now the third case, i.e., the case where every edge $\{u,v\}\in E(H_1)$ has the vertices u, v at different distances from x_1 and the "owner" of the edge is closer to x_1 than the other vertex. Observe that, as there is no layer-edge, and because H_1 is 2-edge-connected, there are at least three layers in H_1 (including the 0-th layer consisting of x_1).

Note that, as H_1 is 2-edge-connected, x_1 has at least 2 neighbors in H_1, and therefore it has a neighbor $y\in V(H_1)$ such that $|A(y)|<|\tilde{H}_1|/2$. Because H_1 is 2-edge-connected, every vertex $v\in A(y)$ has an alternative v-x_1-path in H_1 that does not go via y. Moreover, for every $v\in A(y)$, there is such an alternative v-x_1-path of the following type: starting from v, the first (nonempty) part of the path is using only vertices from $A(y)$, and the second (nonempty) part is using only vertices from $V(H_1)\setminus A(y)$ and is a shortest path to x_1 (see Figure 2). Let us consider the edge $\{z,z'\}$ where such an alternative v-x_1-path (for any vertex v) leaves the first part of the path, i.e., where z is from $A(y)$ and z' is not from $A(y)$ anymore. Obviously, z and z' are from different layers (as we assume there is no layer-edge). Moreover, z has to be closer to x_1 than z' is, as otherwise there would be a shortest path from z to x_1 that does not go via y, a contradiction with the assumption that $z\in A(y)$. Let k denote the distance of z from x_1, i.e., $k=d(x_1,z)$.

First, if $k=1$, i.e., $z=y$, and $\{z,z'\}=\{y,z'\}$, we consider the swap of $\{y,z'\}$ with $\{y,x_2\}$ by vertex y. The distance of y to $\tilde{H}_2$ decreases by $d(y,x_2)-1=1$, the distance of y to $S(z')$ increases by at most $d(y,x_1)+d(x_1,z')-1=2$ (recall that $z'\notin A(y)$ and therefore the swap cannot increase the distance of z' from x_1), and no other distances increase. Therefore, as we assume y cannot improve

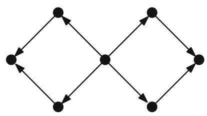

Fig. 3. An example of a graph in asymmetric swap-equilibrium that has two 2-vertex-connected components. The ownership of an edge is depicted by the arrows – the owner is the tail of the respective edge.

by this swap, we have $2|S(z')| \geq |\tilde{H}_2|$ $(\geq |\tilde{H}_1|)$. But then vertex x_1 can improve its usage cost by swapping $\{x_1, y\}$ with $\{x_1, z'\}$: the distance to $S(z')$ decreases by 1, the distance to $A(y)$ increases by 1, and no other distance increases; as $|A(y)| < |\tilde{H}_1|/2$ and $|S(z')| \geq |\tilde{H}_1|/2$, vertex x_1 improves its usage cost. This is a contradiction.

Assume now that $k \geq 2$. Consider the vertex z and the swap of the edge $\{z, z'\}$ with the edge $\{z, x_2\}$. After the swap, vertex z decreases its distance to the vertices in $\tilde{H}_2$ by $d(z, x_2) - 1 = k$. It may increase its distance to the vertices w for which the shortest z-w-path p_w used the edge $\{z, z'\}$. Let us classify such vertices w according to the number of backward-edges in p_w: let $S(i)$ denote the set of vertices $w \neq z$ for which p_w uses only forward-edges (i.e., $S(i)$ is equal to $S(z)$); and let $S(ii)$ denote the set of vertices $w \neq z$ for which p_w uses exactly one backward-edge. Then, after the swap, z increases its distances to vertices in $S(i)$ by at most $1 + 1 + d(x_1, z') - 1 = k + 2$ (z can now get to w via the direct connection to x_2) and it increases its distances to vertices in $S(ii)$ by at most k, and it increases its distances to any other vertex of $\tilde{H}_1 \setminus z$ by at most $k - 2$. (This follows because the shortest z-w path p_w to any other vertex w either does not use $\{z, z'\}$, or it uses at least 2 backward-edges). Therefore, the total decrease of the usage cost of z after the swap is at least

$$k \cdot |\tilde{H}_2| - (k+2) \cdot |S(i)| - k \cdot |S(ii)| - (k-2) \cdot (|\tilde{H}_1 \setminus \{z\}| - |S(i)| - |S(ii)|) > 2 \cdot |\tilde{H}_1| - 4 \cdot |S(i)| - 2 \cdot |S(ii)|.$$

The graph is in asymmetric swap-equilibrium and therefore z cannot improve by this swap. Thus, $2 \cdot |\tilde{H}_1| - 4 \cdot |S(i)| - 2 \cdot |S(ii)| < 0$, or equivalently, $|\tilde{H}_1| - |S(i)| - |S(ii)| < |S(i)|$. But in this case $S(i)$ is very large and x_1 would benefit from an edge to z': Consider a swap of $\{x_1, y\}$ with $\{x_1, z'\}$; x_1 decreases its distances to $S(i)$ by $d(x_1, z') - 1 = k$; it decreases its distances to $S(ii)$ by $k - 2$ (recall that $k \geq 2$; thus, x_1 cannot increase its distance to $S(ii)$); it increases its distances to any other vertex in $\tilde{H}_1$ by at most $1 + d(z', y) - 1 = k$; thus, x_1 improves by the swap, a contradiction. □

We note that the theorem cannot be made stronger in that there are asymmetric swap-equilibria that have more than one 2-(vertex)-connected components – see Figure 3. Not many constructions of bridge-less (Nash or swap) equilibrium graphs are known (a small cycle or a complete graph are the firm favorites) and they all have diameter ≤ 3 (to the best of our knowledge). Figure 3 gives a simple example of a bridge-less asymmetric swap-equilibrium with diameter 4. Finding less trivial examples is definitely an interesting quest.

Besides the 2-edge-connected component H, an asymmetric swap-equilibrium also contains trees. If H is an empty graph, then the equilibrium is a tree. Ehsani et al. [4] analysed trees in the bounded-budget network creation games. A careful inspection of their proofs shows that their analysis can be taken 1-to-1 to argue about asymmetric swap equilibria, too:

Theorem 2 ([4]). *A tree in asymmetric swap-equilibrium has diameter* $\mathcal{O}(\log n)$*. Moreover, a complete binary tree where every node owns all adjacent edges to its children is in asymmetric swap-equilibrium.*

2.1 A Vertex-Weighted Version of the Game

An asymmetric swap-equilibrium G thus consists of a nontrivial 2-edge-connected component H with trees attached to the vertices of H. This suggests the following natural variant of the game. Consider a node-weighted graph $\bar{G}$, i.e., a graph with a weight $c(v) \in \mathbb{N}$ for every vertex $v \in V(\bar{G})$. For such a graph, consider the modified usage cost where every distance is "weighted" by the corresponding weight of the vertex: $\sum_{v \in V(\bar{G})} c(v) \cdot d(u,v)$. We further define $c(\bar{G}) := \sum_{v \in \bar{G}} c(v)$ for any graph $\bar{G}$. Instead of studying G, we may study the node-weighted graph $\bar{G}$ where $\bar{G}$ is the 2-edge-connected component H of G, and the weight $c(v)$ of vertex $v \in H$ is the number of vertices in the attached trees. For node-weighted graphs, we are interested in both the swap equilibria and the asymmetric swap-equilibria. Adapting the proofs and results for the non-weighted setting (the missing proofs can be found in [8]), we obtain:

Proposition 2. *Let* T *be a tree and* $c : V(T) \longrightarrow \mathbb{N}$ *an arbitrary weight function. If* T *is in asymmetric swap-equilibrium in the weighted version of the game, then* $diam(T) = \mathcal{O}(\log c(T))$*.*

Corollary 1. *Let* G *be a non-tree graph in asymmetric swap-equilibrium for the unweighted version of the game and* H *be its unique* 2*-edge-connected component. Then* $diam(G) = diam(H) + 2\log n$*.*

Proposition 3. *Let* T *be a tree and* $c : V(T) \longrightarrow \mathbb{N}$ *a weight function. If* T *is a swap equilibrium in the weighted version of the game then* $diam(T) \leq 2$*.*

Corollary 2. *Let* G *be a non-tree swap equilibrium and* H *its 2-edge-connected component in the unweighted version of the game. Then* $diam(G) \leq diam(H) + 4$*.*

3 Diameter of Non-tree Asymmetric Swap-Equilibria

In this section we consider the problem of bounding the diameter of asymmetric swap-equilibria. Ehsani et al. [4] showed that the diameter of any Nash equilibrium in the bounded-budget network creation game is at most $2^{O(\sqrt{\log n})}$. They proved this for a more general concept of equilibrium graphs, which, it turns out, is equivalent to the asymmetric swap-equilibrium. Therefore, they proved:

Theorem 3 ([4]). *The diameter of a graph in asymmetric swap-equilibrium is at most* $2^{O(\sqrt{\log n})}$*.*

It is believed, however, that the diameter of equilibrium graphs is much smaller. Similarly to the original network creation game, where for various values of α different techniques have been applied to show constant price of anarchy, we believe that studying various classes of asymmetric swap-equilibria can have a similar effect. In the following we focus on a special case where the unique 2-edge-connected component of an equilibrium graph has a large minimum degree. Along the way we present a more general approach that can possibly be applied to more general (asymmetric) swap-equilibria. Let G be a non-tree asymmetric swap-equilibrium on n vertices. By Theorem 1 we know that G has a unique 2-edge-connected component H. We will show that in our special case, the diameter of H is a constant, and hence, by Corollary 1, a $O(\log n)$ upper bound on the diameter of G, or respectively, by Corollary 2, a constant upper bound on the diameter of G if G is a swap equilibrium. This problem can also be seen as a problem to show a constant (with respect to $c(H)$) bound on the diameter of a bridge-less asymmetric swap-equilibrium H in the weighted version of the game (with appropriately chosen weight function $c : V(H) \longrightarrow \mathbb{N}$) – in the following we use this approach.

We define and use the following notation. For every vertex $v \in V(H)$ let $T(v)$ denote the set of vertices $u \in V(G)$ for which a shortest u-H-path ends in v. Note that $v \in T(v)$, $T(v)$ induces a tree in G, and V is a disjoint union of $T(v)$, $v \in V(H)$. We define the weight function $c : V(H) \longrightarrow \mathbb{N}$ for vertices in H by setting $c(v) := |T(v)|$, and introduce the notation $c(H') := \sum_{v \in H'} c(v)$ for any $H' \subset H$. We note that $c(H) = n$. From now on we only consider H as a stand-alone, vertex-weighted, bridge-less graph. For $k \in \mathbb{N}$ and $u \in V(H)$ we define $B_k(u) := \{v \in V(H) : d_H(u, v) \leq k\}$ to be the *ball* of radius k and center u in H, and we define $S_k(u) := \{v \in V(H) : d_H(u, v) = k\}$ to be the *sphere* of radius k and center u in H. We further define $C_k := \min_{u \in V(H)} c(B_k(u))$ to be the minimum weighted size of any ball of radius k in H. For a vertex $u \in V(H)$ we denote its *eccentricity* in H by $D(u)$ (and recall that $D(u) = \max_{v \in V(H)} d(u, v)$).

We will need the following lemma, which shows that in asymmetric swap-equilibria a large (linear in n) number of vertices is far away from any given vertex u (the proof can be found in [8]).

Lemma 1. *For every vertex* $u \in V(H)$ *and* $k < \frac{D(u)-1}{2}$ *we have* $c(B_k(u)) < \frac{D(u)+1}{2(D(u)-k)-1} \cdot n$.

Corollary 3. *If* $r := rad(H) > 14$ *then for every vertex* $u \in V(H)$ *we have* $c(B_{r/4}(u)) < \frac{3}{4}n$.

In the following discussion, we bound the diameter of H using the "region-growing" technique of Demaine et al. [3] for the original network creation game (which showed an upper bound $2^{O(\sqrt{\log n})}$ on the price of anarchy). The details, however, differ significantly from the proofs in [3] due to the different definition of the games (players are only allowed to swap and not to buy new edges) and our new structural insights from Section 2.

Lemma 2. *If H has minimum degree $d(H) \geq n^\varepsilon$ for $\frac{4\lg 3}{\lg n} < \varepsilon < 1$, then for every vertex $u \in V(H)$ there is an edge $\{x, y\}$ induced by $B_{8/\varepsilon}(u)$ owned by x whose deletion increases x's usage cost by at most $\frac{20}{\varepsilon} n^{1-\varepsilon/2}$.*

Proof. We first show that there exists a vertex $v \in B_{2/\varepsilon}(u)$ which owns at least $\frac{n^{\varepsilon/2}}{2}$ edges in H. The main argument goes along the line "if there are m' edges among n' vertices, then there is a vertex that owns at least m'/n' edges (pigeon-hole principle)". First we consider the case when $|B_{k+1}(u)| \geq n^{\varepsilon/2}|B_k(u)|$ for every $k < 2/\varepsilon$. Clearly in this case, $|B_k(u)|$ is at least $n^{k\varepsilon/2}$ and therefore $|B_{2/\varepsilon}(u)| = |H|$. As H contains at least $\frac{|H|n^\varepsilon}{2}$ edges, there is a vertex $v \in B_{2/\varepsilon}(u) = H$ which owns at least $\frac{|H|n^\varepsilon}{2|H|} = n^\varepsilon/2$ edges.
Assume now that $|B_{k+1}(u)| < n^{\varepsilon/2}|B_k(u)|$ for some $k < 2/\varepsilon$. The ball $B_{k+1}(u)$ contains at least $\frac{n^\varepsilon}{2}|B_k(u)|$ edges (as all edges adjacent to vertices in $B_k(u)$ need to lie within $B_{k+1}(u)$). Therefore, there is a vertex in $B_{k+1}(u)$ which owns at least $\frac{n^\varepsilon}{2}|B_k(u)|/|B_{k+1}(u)| > \frac{n^\varepsilon}{2}|B_k(u)|/(n^{\varepsilon/2}|B_k(u)|) = \frac{n^{\varepsilon/2}}{2}$ edges.

We now investigate whether among the at least $\frac{n^{\varepsilon/2}}{2}$ edges which vertex $v \in B_{2/\varepsilon}(u)$ owns there is an edge whose deletion increases the usage cost by the claimed amount. For every edge $\{v, w\}$ which v owns, let $A(w)$ be the vertices x such that *every* shortest path from x to v goes via w. Observe that, as v owns at least $\frac{n^{\varepsilon/2}}{2}$ edges, there is an edge $\{v, w\}$ such that $c(A(w)) \leq n/(\frac{n^{\varepsilon/2}}{2}) = 2n^{1-\varepsilon/2}$. If this edge $\{v, w\}$ lies in a "short" cycle of length l, then deleting this edge would increase the usage cost of v by at most $l \cdot c(A(w))$. Consider vertices of $A(w)$ ordered in layers according to the distance to v. Let $k \in \mathbb{N}$ be the smallest index for which an edge between $S_k(v) \cap A(w)$ and $V(H) \setminus A(w)$ exists (such an edge indeed exists for some k as otherwise $\{v, w\}$ would be a bridge). If k is "small", then the edge lies in a "short" cycle of length $2k$ and we can bound the increase of the usage cost as suggested. Define for every $j \leq k$ the set $B_j := B_j(v) \cap A(w)$. In the case that $|B_{j+1}| \geq n^{\varepsilon/4}|B_j|$ for all $j < k$, we have $|B_k| \geq n^{(k\varepsilon)/4}$. But as $|B_k|$ is clearly upper bounded by n, we get $k < 4/\varepsilon$. Thus, k is a constant and deleting the edge $\{v, w\}$ increases the usage cost of v by at most $2k \cdot c(A(w)) < \frac{16n^{1-\varepsilon/2}}{\varepsilon}$. In the other case, let $j < k$ such that $|B_{j+1}| < n^{\varepsilon/4}|B_j|$. Note that, as $|B_j| < n$, $j < 4/\varepsilon$. In this case we do not consider deleting $\{v, w\}$ but instead we find another edge within B_j that can be deleted and which does not increase the usage cost of the owner of the edge too much. There are at least $\frac{n^\varepsilon}{2}|B_j|$ edges that are incident to vertices of B_j. If we subtract from these edges the edges that form a breadth-first search tree of B_{j+1} (there are at most $|B_{j+1}| \leq n^{\varepsilon/4}|B_j|$ of these), we obtain at least $\frac{n^\varepsilon - 2n^{\varepsilon/4}}{2}|B_j|$ edges that are part of a "short" cycle of length no more than $(2j + 2) \leq \frac{8}{\varepsilon} + 2$. Observe that $n^\varepsilon - 2n^{\varepsilon/4} = n^{3\varepsilon/4}(n^{\varepsilon/4} - 2n^{-2\varepsilon/4}) \geq n^{3\varepsilon/4}(n^{\varepsilon/4} - 2) > n^{3\varepsilon/4}(n^{\lg 3/\lg n} - 2) \geq n^{3\varepsilon/4}$. Therefore, there are at least $\frac{n^{3\varepsilon/4}}{2}|B_j|$ edges within B_{j+1} that are part of a "short cycle". There has to be a vertex $w \in B_{j+1}$ which owns at least $\frac{n^{3\varepsilon/4}/2|B_j|}{|B_{j+1}|} \geq \frac{n^{\varepsilon/2}}{2}$ of these edges. By the pigeon-hole principle, among these edges of w, there has to be one edge whose deletion increases the usage cost of w by at most $2(j+1)\frac{2n}{n^{\varepsilon/2}} < (\frac{16}{\varepsilon}+4)n^{1-\varepsilon/2} < \frac{20}{\varepsilon}n^{1-\varepsilon/2}$. □

Theorem 4. *If H has minimum degree $d(H) \geq n^{\varepsilon}$ for $\frac{4\lg 3}{\lg n} < \varepsilon < 1$, then there is a constant $C(\varepsilon) > 0$ (depending on ε) such that diam$(H) \leq C(\varepsilon)$.*

Proof. We will show that for every $k \leq \frac{r}{8} - 1$, where $r := \text{rad}(H)$: $C_{3k+3+8/\varepsilon} > \frac{\varepsilon \cdot n^{\varepsilon/2}}{40} C_k$. Assuming this, the result follows immediately: Let $u \in V(H)$ be a vertex with eccentricity $D(u) = \text{diam}(H)$ and let $\tilde{C} > 0$ be a constant such that $\tilde{C}k \geq 3k+3+8/\varepsilon$. We have $c(B_{\tilde{C}k}(u)) \geq C_{3k+3+8/\varepsilon} > \frac{\varepsilon \cdot n^{\varepsilon/2}}{40} C_k$ for $k \leq \frac{r}{8} - 1$. Now, in the trivial case when n is at most the constant threshold $\left(\frac{40}{\varepsilon}\right)^{4/\varepsilon}$, then of course the diameter of H is at most this value (a constant). For the general case when n is larger than the threshold, we must have $\frac{r}{8} - 1 < \tilde{C}^{4/\varepsilon}$ as otherwise $c(B_{\frac{r}{8}-1}(u)) \geq c(B_{\tilde{C}^{4/\varepsilon}}(u)) \geq \left(\frac{\varepsilon n^{\varepsilon/2}}{40}\right)^{4/\varepsilon} C_1 \geq n$. Thus, in this case, the diameter of H is at most $2r < 16\tilde{C}^{4/\varepsilon} + 16$, a constant.

We now prove that $C_{3k+3+8/\varepsilon} > \frac{\varepsilon \cdot n^{\varepsilon/2}}{40} C_k$ for every $k \leq \frac{r-1}{4}$. Consider an arbitrary vertex $u \in V(H)$. By Lemma 2 there is an edge $\{x, y\}$ within $E(B_{8/\varepsilon}(u))$ owned by x whose deletion increases x's usage cost by at most $\frac{20n^{1-\varepsilon/2}}{\varepsilon}$. We select a maximal subset $\{x_1, \ldots, x_l\} \subset S_{2k+3}(x)$ subject to the condition $d(x_i, x_j) \geq 2k+1$ for every $i \neq j$. We assign every vertex in $S_{2k+3}(x)$ to the closest x_i, breaking ties arbitrarily. Let $S_{2k+3}(x) = \bigcup_{i=1}^{l} A_i$ be the obtained partition. We now prove that $l \geq \frac{\varepsilon \cdot n^{\varepsilon/2}}{40}$. We extend the partition and also assign each vertex $z \in V(H) \setminus B_{2k+2}(x)$ to one of the x_i, as follows: Pick any shortest path from z to x, and assign z to the same x_i as the (unique) vertex $w \in S_{2k+3}(x)$ which is contained in the path. After this step, $V(H) \setminus B_{2k+2}(x) = \bigcup_{i=1}^{l} A_i$ is the resulting partition. Consider vertex x and the swap of the edge $\{x, y\}$ with $\{x, x_i\}$ for arbitrary i: the distance of x to x_i decreases by $2k+2$ and hence, by the construction of A_i, the distance of x to the vertices of A_i decreases by at least 2. On the other hand, by Lemma 2, the swap increases x's usage cost by at most $\frac{20n^{1-\varepsilon/2}}{\varepsilon}$. Hence, as x cannot improve its usage cost by the swap (we are considering an asymmetric swap-equilibrium), $\frac{20n^{1-\varepsilon/2}}{\varepsilon} \geq 2c(A_i)$. As i was arbitrary, we have $l \cdot \frac{20n^{1-\varepsilon/2}}{\varepsilon} \geq 2\sum_{i=1}^{l} c(A_i)$. On the other hand, as $2k+2 \leq r/4$, we have by Corollary 3: $c(B_{2k+2}(x)) < 3n/4$, so $\sum_{i=1}^{l} c(A_i) = c(V(H) \setminus B_{2k+2}(x)) \geq n/4$. Therefore $l \cdot \frac{20n^{1-\varepsilon/2}}{\varepsilon} \geq 2\sum_{i=1}^{l} c(A_i) \geq n/2$ and hence $l \geq \frac{\varepsilon \cdot n^{\varepsilon/2}}{40}$.

By definition, $B_k(x_i) \cap B_k(x_j) = \emptyset$ for every $i \neq j$. Hence $c(\bigcup_{i=1}^{l} B_k(x_i)) = \sum_{i=1}^{l} c(B_k(x_i)) \geq l \cdot C_k$. Furthermore, for every $1 \leq i \leq l$, we have $d(u, x_i) \leq d(u, x) + d(x, x_i) \leq 2k+3+8/\varepsilon$, so vertex u has a path of length at most $3k+3+8/\varepsilon$ to every vertex in $B_k(x_i)$. Therefore $c(B_{3k+3+8/\varepsilon}(u)) \geq \sum_{i=1}^{l} c(B_k(x_i)) > l \cdot C_k \geq \frac{\varepsilon n^{\varepsilon/2}}{40} C_k$. Hence, as $u \in V(H)$ was chosen arbitrarily, $C_{3k+3+8/\varepsilon} > \frac{\varepsilon n^{\varepsilon/2}}{40} C_k$. □

Together with Corollary 1 resp. Corollary 2 this implies a logarithmic bound on the diameter of asymmetric swap-equilibria, respectively a constant bound on the diameter of swap equilibria:

Corollary 4. *For any ϵ and any non-tree asymmetric swap equilibrium G, the following holds. If the unique 2-edge-connected component H of G has minimum degree $d(H) \geq n^{\varepsilon}$ and $\frac{4 \lg 3}{\lg n} < \varepsilon < 1$ then diam$(G) = \mathcal{O}(\lg n)$.*

Corollary 5. *For every constant $0 < \varepsilon < 1$ there is a constant $C(\varepsilon)$ such that the following holds. The diameter of any graph G, where $d(H) \geq n^{\varepsilon}$ and $\frac{4 \lg 3}{\lg n} < \varepsilon$ (and H is the unique 2-edge-connected component of G), is at most $C(\varepsilon)$.*

Inspired by Theorem 4 and by the fact that one could not find any bridge-less equilibrium graph of diameter greater than 4, we conjecture:

Conjecture 1. There exists a constant $C > 0$ such that the diameter of the unique 2-edge connected component of any asymmetric swap-equilibrium is smaller than C. In particular, every asymmetric swap-equilibrium has diameter $\mathcal{O}(\log n)$ and every swap equilibrium has diameter $\leq C + 4$.

References

1. Albers, S., Eilts, S., Even-Dar, E., Mansour, Y., Roditty, L.: On Nash equilibria for a network creation game. In: Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, pp. 89–98 (2006)
2. Alon, N., Demaine, E.D., Hajiaghayi, M., Leighton, T.: Basic network creation games. In: Proceedings of the 22nd ACM Symposium on Parallelism in Algorithms and Architectures, SPAA, pp. 106–113 (2010)
3. Demaine, E.D., Hajiaghayi, M., Mahini, H., Zadimoghaddam, M.: The price of anarchy in network creation games. In: Proceedings of the Twenty-Sixth Annual ACM Symposium on Principles of Distributed Computing, PODC, pp. 292–298 (2007)
4. Ehsani, S., Fazli, M., Mehrabian, A., Sadeghian Sadeghabad, S., Safari, M., Saghafian, M., ShokatFadaee, S.: On a bounded budget network creation game. In: Proceedings of the 23rd ACM Symposium on Parallelism in Algorithms and Architectures, SPAA, pp. 207–214 (2011)
5. Fabrikant, A., Luthra, A., Maneva, E., Papadimitriou, C.H., Shenker, S.: On a network creation game. In: Proceedings of the Twenty-Second Annual Symposium on Principles of Distributed Computing, PODC, pp. 347–351 (2003)
6. Lenzner, P.: On Dynamics in Basic Network Creation Games. In: Persiano, G. (ed.) SAGT 2011. LNCS, vol. 6982, pp. 254–265. Springer, Heidelberg (2011)
7. Mihalák, M., Schlegel, J.C.: The Price of Anarchy in Network Creation Games Is (Mostly) Constant. In: Kontogiannis, S., Koutsoupias, E., Spirakis, P.G. (eds.) SAGT 2010. LNCS, vol. 6386, pp. 276–287. Springer, Heidelberg (2010)
8. Mihalák, M., Schlegel, J.C.: Asymmetric swap-equilibrium: A unifying equilibrium concept for network creation games. Tech. Rep. 764, Department of Computer Science, ETH Zurich (June 2012), `https://www1.ethz.ch/inf/research/disstechreps/techreports`