Article

# The Price of Anarchy for Network Formation in an Adversary Model 

Lasse Kliemann

Department of Computer Science, Christian-Albrechts-Universität zu Kiel, Christian-Albrechts-Platz 4, Kiel 24098, Germany; E-Mail: lki@informatik.uni-kiel.de; Tel.: +49-431-880-7454;
Fax: +49-431-880-1725

Received: 27 May 2011; in revised form: 18 July 2011 / Accepted: 8 August 2011 /
Published: 23 August 2011


#### Abstract

We study network formation with $n$ players and link cost $\alpha>0$. After the network is built, an adversary randomly deletes one link according to a certain probability distribution. Cost for player $v$ incorporates the expected number of players to which $v$ will become disconnected. We focus on unilateral link formation and Nash equilibrium. We show existence of Nash equilibria and a price of stability of $1+o(1)$ under moderate assumptions on the adversary and $n \geq 9$. We prove bounds on the price of anarchy for two special adversaries: one removes a link chosen uniformly at random, while the other removes a link that causes a maximum number of player pairs to be separated. We show an $O(1)$ bound on the price of anarchy for both adversaries, the constant being bounded by $15+o(1)$ and $9+o(1)$, respectively.


Keywords: network formation; equilibrium; price of anarchy; unilateral link formation; adversary model; network robustness

## 1. Network Formation

In network formation, a multitude of individuals, called players, form a network in such a way that each player decides for herself to which other players she would like to connect. So players can be considered vertices [1] of a (to-be-built) network. Any outcome of this, i.e., any network, can be evaluated from the point of view of each player via an individual cost. Individual cost comprises building cost, proportional to the number of links [2] built by the player, and indirect cost, which expresses properties of
the network. Social cost is the sum of individual cost over all players. There are two parameters: $n$ is the number of players and $\alpha>0$ is the cost of a link. Another crucial feature is how links can be formed: unilaterally or bilaterally. Under unilateral link formation a player can connect to any other player and is charged the amount of $\alpha$ for each link. Under bilateral link formation, the consent of both endpoints is required and if they both agree, then they pay $\alpha$ each. When the network is so that no player sees a way to improve her individual cost, we speak of an equilibrium. The finer facets of the equilibrium concept have to be chosen according to the link formation rule: Nash equilibrium is well suited for unilateral link formation, whereas for bilateral link formation, pairwise Nash equilibrium or pairwise stability are better suited. When the social cost is minimal, we speak of an optimum. [3] The price of anarchy [4,5] measures overall performance loss due to distributed operation, compared to when a central authority would enforce an optimum: the price of anarchy is the worst-case ratio of the social cost of an equilibrium to that of an optimum. One is interested in bounds on the price of anarchy, especially upper bounds. The price of anarchy is a static measure in the sense that it does not try to assess how a network might evolve over time. It instead builds upon the assumption that equilibrium networks will emerge from evolutionary processes. A related concept is price of stability, i.e., the best-case ratio of the social cost of an equilibrium to that of an optimum. This work's focus is on upper bounds on the price of anarchy, although we prove tight bounds on the price of stability and some structural results along the way.

Our Contribution. We study the price of anarchy in an adversary model. After the network is built, an adversary deletes exactly one link from it. The adversary is modeled by a random experiment; hence in general there is an uncertainty which link will be deleted, but players know the probability distribution according to which the adversary chooses the link to destroy. Indirect cost of a player $v$ is defined as the expected number of players to which $v$ will lose connection when the adversary strikes. Formally, an adversary is a mapping from networks to probability distributions on the edges of the particular network.

Although it appears limiting that the adversary can only destroy one link, this model already is challenging to analyze. It is a contribution to the understanding of how networks are formed when it is important that every vertex can reach every other vertex, for example for data transmission or delivery of goods.

After preparations and discussion of related work (Sections 2 to 5) we start out with a simple $O\left(1+\frac{n}{\alpha}\right)$ bound on the price of anarchy for any adversary and independent of the link formation rule and the equilibrium concept, but under the assumption that equilibria only have a linear (in $n$ ) number of edges (Section 6). This assumption will later be shown to be valid for the two special adversaries under consideration and unilateral link formation. In the three remaining sections (Sections 7 to 9), we consider unilateral link formation. We constructively show existence of Nash equilibria for all adversaries (under some moderate additional assumptions), including a co-existence of two very different topologies for the same range of parameters. A $1+o(1)$ bound on the price of stability follows from our existence results. Then for two specific adversaries we prove an $O(1)$ bound on the price of anarchy for all ranges of $\alpha$ and $n \geq 9$. These two adversaries are chosen to represent extreme cases: the first one, called simple-minded, chooses one link uniformly at random. The second one, called smart, chooses uniformly at random from the set of those links whose removal causes a maximum number of vertex pairs to be separated. The proof techniques for the simple-minded adversary are roughly similar to what has been used for other models
before, e.g., [6], namely we relate to the diameter of equilibria. For the smart adversary, a new approach has to be taken; it works by an appropriate decomposition of the graph.

Open Problems and Future Work. Tight bounds on the price of anarchy for other adversaries, or for a general adversary are left for future work. As one of the most intriguing open problems, we leave the case of an adversary removing more than one link. Since our proofs rely heavily on the restriction of only one link being removed, this is expected to be a new challenge.

This work's focus is on unilateral link formation and Nash equilibrium. Bilateral link formation with its appropriate equilibrium concepts is planned to be studied in a separate publication. Currently, for bilateral link formation we can show a bound of $O(1+\sqrt{n / \alpha})$ for the simple-minded adversary, if $\alpha>\frac{1}{2}$. However, it is not known whether this bound is tight. For the smart adversary, we can show a tight $\Theta(n)$ bound, if $\alpha>2$ considered constant and $n \geq 10$.

## 2. Model Framework

We give a rigorous description of the model framework that will be used in the following. Let $n \geq 3$ (we will later refine this to $n \geq 9$ ) and $V$ a set of $n$ vertices, say $V=[n]:=\{1, \ldots, n\}$. Each vertex represents an individual, called a player. Each player names a list of other players to which she would like to build an edge. The decisions of player $v$ are collected in a vector $S_{v} \in\{0,1\}^{n}$, with $S_{v w}=1$ meaning that $v$ would like to have the edge $\{v, w\}$ in the network. Such an $S_{v}$ is called a strategy for player $v$. A vector of strategies $S=\left(S_{1}, \ldots, S_{n}\right)$, one for each player, is called a strategy profile. A strategy profile can be written as a matrix $\{0,1\}^{n \times n}$ and interpreted as the adjacency matrix of a directed graph $\vec{G}(S)=(V, \vec{E}(S))$. Then $(v, w) \in \vec{E}(S)$ if and only if player $v$ would like to have the edge $\{v, w\}$. We will often work with this representation of strategy profiles. Denote $\mathcal{S}(n)$ the set of all strategy profiles for $n$ players. We use sets $F \subseteq V \times V$ to denote strategy changes. Define $S+F$ and $S-F$ by setting for all $x, y \in V$

$$
(S+F)_{x y}:=\left\{\begin{array}{ll}
1 & \text { if }(x, y) \in F \\
S_{x y} & \text { otherwise }
\end{array} \quad \text { and } \quad(S-F)_{x y}:= \begin{cases}0 & \text { if }(x, y) \in F \\
S_{x y} & \text { otherwise }\end{cases}\right.
$$

If $F=\{(v, w)\}$, we write $S+(v, w)$ and $S-(v, w)$. For instance, $S+(v, w)$ means that we add to $S$ the request of player $v$ for the edge $\{v, w\}$.

The graph which is actually built is called the final graph and denoted $G(S)=(V, E(S))$, where

$$
E(S):=\left\{\{v, w\} ; S_{v w}=1 \vee S_{w v}=1\right\}
$$

So the wish of one endpoint, either $v$ or $w$, is enough to have $\{v, w\}$ in the final graph. We also call this way of forming the final graph unilateral link formation (ULF).

For bilateral link formation (BLF), $E(S):=\left\{\{v, w\} ; S_{v w}=1 \wedge S_{w v}=1\right\}$ are the edges of the final graph. So both endpoints, $v$ and $w$, have to agree on having $\{v, w\}$ in the final graph. Otherwise, it will not be built. BLF will be considered in future work, and we only consider ULF here (save some results that hold independently of the link formation rule and the equilibrium concept, cf. Section 6).

We speak of selling (or deleting, removing) an edge $e$ if a player changes her strategy so that $e$ is no longer part of the final graph. We speak of buying (or adding, building) an edge $e$ if a player changes her strategy so that $e$ is then part of the final graph.

Fix parameters $n$ and $\alpha>0$. Given a strategy profile $S \in \mathcal{S}(n)$ each player experiences a cost $C_{v}(S)$, her individual cost. It is comprised of building cost plus indirect cost. Building cost is computed by counting $\alpha$ for each edge that $v$ requested. Indirect cost can be defined in many different ways and usually captures properties of the final graph, we denote it $I_{v}(G(S))$ and sometimes just $I_{v}(S)$ for a streamlined notation. Denoting $\left|S_{v}\right|:=\sum_{w \in V} S_{v w}$, we can write out the individual cost $C_{v}(S):=\left|S_{v}\right| \alpha+I_{v}(G(S))$. An equivalent concept found in the literature is payoff: properties of the final graph are expressed by income, and payoff is income minus building cost.

The indirect cost $I_{v}(\cdot)$ is a placeholder to be filled in in order to have a concrete model. For example, the model in [6] uses $I_{v}(G)=\sum_{w \in V} \operatorname{dist}_{G}(v, w)$, where the $\operatorname{distance} \operatorname{dist}(v, w)$ is the length of a shortest path between $v$ and $w$ and equals $\infty$ if there is no such path. We call this the sum-distance model. The price of anarchy in the sum-distance model is particularly well-studied. We will introduce our definition of indirect cost in Section 3.

We call indirect cost anonymous if for each final graph $G=(V, E)$ and each graph automorphism $\phi: V \longrightarrow V$ of $G$, we have $I_{v}(G)=I_{\phi(v)}(G)$ for all $v \in V$. In other words, anonymity of indirect cost means that $I_{v}(G)$ does not depend on $v$ 's identity, but only on $v$ 's position in the final graph $G$. This is of importance in particular if $G$ has symmetry. For instance, if $G$ is a cycle, then all vertices experience the same indirect cost. If $G$ is a path, then both endpoints experience the same indirect cost.

The social cost of $S$ is $C(S):=\sum_{v \in V} C_{v}(S)$. When we sum up the building cost over all players, we also speak of total building cost; when we sum up the indirect cost over all players, we speak of total indirect cost. Hence social cost is total building cost plus total indirect cost. A strategy profile $S^{*}$ is called optimal if $C\left(S^{*}\right)=\min _{S \in \mathcal{S}(n)} C(S)$. This is with respect to fixed $\alpha$; denote $\operatorname{OPT}(n, \alpha)$ the social cost of an optimum for given $n$ and $\alpha$. An undirected graph $G$ is called optimal if $G=G\left(S^{*}\right)$ for an optimal $S^{*}$.

A strategy profile $S$ is called essential [7] if for all $v, w \in V$ the following implication holds:

$$
S_{v w}=1 \Longrightarrow S_{w v}=0
$$

In other words, an essential strategy profile does not contain unnecessary requests. In an essential strategy profile, building cost deserves its name in the following sense: players pay only for edges that would not be in the final graph without this payment. This means that each edge in the final graph is paid for $\alpha$ by exactly one of its endpoints, namely the one who requested it. Hence if $S$ is essential then $C(S)=|E(S)| \alpha+\sum_{v \in V} I_{v}(G(S))$. Social cost then only depends on the final graph.

For each strategy profile $S$, dropping all unnecessary requests [8] results in an essential strategy profile $S^{\prime}$ with the same final graph and with the same or a smaller individual cost for each player. Moreover, it is easy to see that if $S$ is a Nash equilibrium (introduced below), then $S^{\prime}$ is a Nash equilibrium. It is hence reasonable to restrict to essential strategy profiles, and we will do so in the following.

Recall that we can specify strategy profiles as directed graphs. Furthermore, since the social cost is fully determined by the final graph (since we restrict to essential strategy profiles), it suffices to consider the final graph (which is an undirected graph) in places where only the social cost is relevant.

A strategy profile $S$ is called a Nash equilibrium (NE) if no player can strictly improve her individual cost by changing her strategy given the strategies of the other players, i.e.,

$$
C_{v}(S+A-D) \geq C_{v}(S) \quad \forall A, D \subseteq\{v\} \times V \quad \forall v \in V
$$

Denote the set of all NE for given $n$ and $\alpha$ by $\mathrm{NE}(n, \alpha)$. An undirected graph $G$ is called a NE if there exists a NE $S$ such that $G=G(S)$. The price of anarchy (with respect to NE) is the social cost of a worst-case NE divided by the social cost of an optimum, i.e.,

$$
\frac{\max _{S \in \mathrm{NE}(n, \alpha)} C(S)}{\operatorname{OPT}(n, \alpha)}
$$

When we replace "max" for "min", we speak of price of stability. Both notions are meant relative to a given $n$ and $\alpha$.

We call a NE $S$ a maximal Nash equilibrium (MaxNE) if $C_{v}\left(S+\left(v, w_{1}\right)+\ldots+\left(v, w_{k}\right)\right)>C_{v}(S)$ for all $\left\{v, w_{1}\right\}, \ldots,\left\{v, w_{k}\right\} \notin E(S)$. That is, we exclude the possibility that a player can buy additional links so that the gain in her indirect cost and the additional building cost nullify each other. This notion is useful to carry over results for ULF to BLF, and we treat it here for future reference.
2.1 Remark. A NE $S$ is maximal if indirect cost $I_{v}(S)$ has its minimum possible value for all players $v$ (which is 0 for most models). A NE is also maximal, if there exists $\varepsilon>0$ such that it is still a NE for link $\operatorname{cost} \alpha-\varepsilon$ instead of $\alpha$. Hence, if $S$ is a NE for all $\alpha \geq f(n)$, for some function $f$, this implies that $S$ is a MaxNE for all $\alpha>f(n)$.

We require some basic graph-theoretic notions. Let an undirected graph $G=(V, E)$ be given, that is, $V$ is a finite set and $E \subseteq\binom{V}{2}$. A walk of length $\ell$ is a sequence of vertices $W=\left(v_{0}, \ldots, v_{\ell}\right)$ such that $\left\{v_{i-1}, v_{i}\right\} \in E$ for all $i \in[\ell]$. Denote $V(W):=\left\{v_{0}, \ldots, v_{\ell}\right\}$ its vertices and $E(W):=\left\{\left\{v_{i-1}, v_{i}\right\} ; i \in[\ell]\right\}$ its edges. The walk is called a path if all its vertices are distinct, that is, if $|V(W)|=\ell+1$. The walk is called a cycle if all its vertices except the last are distinct (i.e., $\left|\left\{v_{0}, \ldots, v_{\ell-1}\right\}\right|=\ell$ ) and the walk is closed (i.e., $v_{0}=v_{\ell}$ ). Sometimes we use a notation that gives names to the edges in the walk, like $\left(v_{0}, e_{1}, v_{1}, e_{2}, \ldots, e_{\ell}, v_{\ell}\right)$. If $C$ is a cycle and $e=\{u, w\}$ is an edge with $u, w \in V(C)$ but $e \notin E(C)$, we call $e$ a chord. For a subset $W \subseteq V$ denote $G[W]:=\left(W,\binom{W}{2} \cap E\right)$ the induced subgraph of $W$, or the subgraph induced by $W$. If $G$ is a graph, then $V(G)$ denotes its set of vertices and $E(G)$ its set of edges; this is useful when $G$ was not introduced writing " $G=(V, E)$ ". More graph-theoretic notions will be introduced along the way as we need them.

One might suggest using multigraphs instead of graphs, since in our adversary model, connectivity under removal of edges is relevant. However, none of our results becomes false when we allow multigraphs. Where not obvious, a remark on this is made. So we can stick to the simpler notion of graphs.

In order to not have to introduce names for all occurring constants, we use " $O$ " and " $\Omega$ " notation. For our results, we use this notation in the following understanding (it does not necessarily apply to all cited results). We write " $x=O(y)$ " if there exists a constant $c>0$ such that $x \leq c y$. The constant may only depend on other constants and is in particular independent of the non-constant quantities that constitute $x$ and $y$, e.g., parameters $n$ and $\alpha$. We do not implicitly require that some quantities, e.g., $n$, have to be large. Analogously, we write " $x=\Omega(y)$ " if there exists a constant $c>0$ such that $x \geq c y$. Note that " $O$ " indicates an upper bound, making no statement about a lower bound; while " $\Omega$ " indicates a lower bound, making no statement about an upper bound. We write $x=\Theta(y)$ if $x=O(y)$ and $x=\Omega(y)$; the constants used in the " $O$ " and the " $\Omega$ " statement may be different, of course.

The " $o$ " notation is only used in one form, namely $o(1)$ substituting a quantity that tends to 0 when $n$ tends to infinity, regardless whether other parameters are fixed or not. Whenever we write " $o(1)$ " in an expression, it is meant as an upper bound, making no statement about a lower bound.

## 3. Adversary Model

An adversary $\mathcal{A}$ is a mapping assigning to each graph $G=(V, E)$ a probability measure $\operatorname{Pr}_{G}^{\mathcal{A}}$ on the edges $E$ of $G$. Given a connected graph $G$, the relevance of an edge $e$ for a player $v$ is the number of vertices that can, starting at $v$, only be reached via $e$. We denote the relevance by $\operatorname{rel}_{G}(e, v)$ and the sum of all relevances for a player by $R_{G}(v):=\sum_{e \in E} \operatorname{rel}_{G}(e, v)$. An edge of a connected graph is called a bridge if its removal destroys connectivity, or equivalently, if it is no part of any cycle. The relevance $\operatorname{rel}_{G}(e, v)$ is 0 iff $e$ is not a bridge. Given a strategy profile $S$ where $G(S)$ is connected, we define the individual cost of a player $v$ by

$$
C_{v}(S):=\left|S_{v}\right| \alpha+\sum_{e \in E(S)} \operatorname{rel}_{G(S)}(e, v) \operatorname{Pr}_{G(S)}^{\mathcal{A}}(\{e\})
$$

The indirect cost is the expected number of vertices to which $v$ will lose connection when exactly one edge is removed from $G(S)$ randomly and according to the probability measure given by the adversary $\mathcal{A}$. For this indirect cost, we use the term disconnection cost in the following instead of "indirect cost". We define the indirect cost to be $\infty$ when $G(S)$ is not connected, so we can concentrate on connected graphs in our study of optima and equilibria. We usually omit the " $G(S)$ " subscripts and also the " $\mathcal{A}$ " superscript; we also write " $E$ " instead of " $E(S)$ " and " $m$ " for the number of edges, i.e., $m=|E(S)|$.

Remark. Since $\infty$ is assigned to disconnected final graphs, optima and NE are connected.
Proof. It is clear for optima. For NE note that since a connected graph has finite indirect cost, a player would always choose to build enough links in order to make the graph connected.

The separation of an edge $e$, denoted $\operatorname{sep}(e)$, is the number of ordered vertex pairs that will be separated by the removal of $e$. For a bridge $e$, denote $\nu(e)$ the number of vertices in the component of $G-e$ that has a minimum number of vertices; we have $\nu(e) \leq\left\lfloor\frac{n}{2}\right\rfloor$. If $e$ is not a bridge, we define $\nu(e):=0$. Then $\operatorname{sep}(e)=2 \nu(e)(n-\nu(e))$ and also $\operatorname{sep}(e)=\sum_{v \in V} \operatorname{rel}(e, v)$. If $e$ is a bridge, then $\operatorname{sep}(e) \geq 2(n-1)$. We can express the social cost now:

$$
C(S):=\sum_{v \in V} C_{v}(S)=m \alpha+\sum_{v \in V} \sum_{e \in E} \operatorname{rel}(e, v) \operatorname{Pr}(\{e\})=m \alpha+\sum_{e \in E} \operatorname{sep}(e) \operatorname{Pr}(\{e\})
$$

We call an adversary symmetric if for a fixed graph, the probability of an edge only depends on its separation, i.e., $\operatorname{sep}(e)=\operatorname{sep}\left(e^{\prime}\right)$ implies $\operatorname{Pr}(\{e\})=\operatorname{Pr}\left(\left\{e^{\prime}\right\}\right)$ for all $e, e^{\prime}$. The following proposition is proved straightforwardly.

### 3.1 Proposition. A symmetric adversary induces anonymous disconnection cost.

Proof. Let $G=(V, E)$ be connected and $\phi: V \longrightarrow V$ a graph automorphism of $G$. If $e=\{v, w\}$ is a non-bridge, then $\phi(e):=\{\phi(v), \phi(w)\}$ is a non-bridge as well, and so $\operatorname{sep}(e)=0=\operatorname{sep}(\phi(e))$ and $\operatorname{rel}(e, v)=0=\operatorname{rel}(\phi(e), \phi(v))$ for all $v \in V$. Let $e=\{v, w\}$ be a bridge and $G_{1}, G_{2}$ be the two components of $G-e$. Then $\phi(e)$ is a bridge as well. Let $G_{1}^{\prime}, G_{2}^{\prime}$ be the two components of $G-\phi(e)$. Then $\phi\left(V\left(G_{i}\right)\right)=V\left(G_{i}^{\prime}\right)$ for all $i \in\{1,2\}$, or $\phi\left(V\left(G_{i}\right)\right)=V\left(G_{j}^{\prime}\right)$ for all $i, j \in\{1,2\}, i \neq j$. In either
case, $\operatorname{sep}(e)=\operatorname{sep}(\phi(e))$, and also $\operatorname{rel}(e, v)=\operatorname{rel}(\phi(e), \phi(v))$ for all $v \in V$. In total, we have for all $e \in E$ and all $v \in V$ :

$$
\begin{align*}
\operatorname{sep}(e) & =\operatorname{sep}(\phi(e))  \tag{1}\\
\operatorname{rel}(e, v) & =\operatorname{rel}(\phi(e), \phi(v)) \tag{2}
\end{align*}
$$

Let $v \in V$. Then:

$$
\begin{array}{rlr}
I_{\phi(v)}(G) & =\sum_{e \in E} \operatorname{rel}(e, \phi(v)) \operatorname{Pr}(\{e\}) \\
& =\sum_{e \in E} \operatorname{rel}(\phi(e), \phi(v)) \operatorname{Pr}(\{\phi(e)\}) & \quad \phi \text { is bijective } \\
& =\sum_{e \in E} \operatorname{rel}(e, v) \operatorname{Pr}(\{\phi(e)\}) & \text { by (2) }  \tag{2}\\
& =\sum_{e \in E} \operatorname{rel}(e, v) \operatorname{Pr}(\{e\}) & \text { by (1) and symmetric adversary } \\
& =I_{v}(G) &
\end{array}
$$

The converse of Proposition 3.1 does not hold, as shown in Figure 1 on this page. Proposition 3.1 is useful since symmetry can be recognized directly from the definitions of the two special adversaries studied later, and so we know that they induce anonymous disconnection cost.

Figure 1. Let this be the final graph $G=(V, E)$ and $G_{1}$ the subgraph to the left starting with $u$ (i.e., the cycle on 5 vertices), and $G_{2}$ the subgraph to the right starting with $w$. Let the adversary assign probabilities, $\operatorname{say} \operatorname{Pr}(\{\{u, v\}\}):=2 / 3$ and $\operatorname{Pr}(\{\{v, w\}\}):=1 / 3$ and 0 to all other edges. Then all players in $G_{1}$ experience the same disconnection cost, and the same holds for $G_{2}$. (Precisely, we have $I_{x}(G)=2 / 3 \cdot 6+1 / 3 \cdot 5=4+5 / 3$ for all $x \in V\left(G_{1}\right)$ and $I_{y}(G)=1 / 3 \cdot 6+2 / 3 \cdot 5=2+10 / 3$ for all $y \in V\left(G_{2}\right)$ and $I_{v}(G)=2 / 3 \cdot 5+1 / 3 \cdot 5=$ 5.) The adversary is not symmetric, since $\operatorname{sep}(\{u, v\})=5 \cdot 6=\operatorname{sep}(\{v, w\})$. However, disconnection cost is anonymous, since an automorphism can only permute players within $G_{1}$ and $G_{2}$, respectively.


## 4. Related Work

There is a vast body of literature on game-theoretic network formation, by far not limited to studies of the price of anarchy. A good starting point is the survey by Jackson [12] from 2004. We cite several publications below with a bias towards studies of the price of anarchy. In a separate subsection on page 311, we give a detailed comparison of our model with work being particularly related to it, namely [13-18].

Bilateral link formation (BLF) follows a concept given by Myerson [19] (p. 228) in a different context. Jackson and Wolinsky [14] in 1996 introduced the symmetric connections model, using BLF, and the equilibrium concept of pairwise stability. The symmetric connections model is best described using the notions of income and payoff. The income for player $v$ is $\sum_{\substack{w \in V \\ w \neq v}} \delta^{\operatorname{dist}_{G(S)}(v, w)}$, where $\delta \in(0,1)$ is a parameter. Her payoff is income minus building cost. Note that we have an exponential dependence on distance. This models to some extent that each link has a probability of $1-\delta$ for failure. We will elaborate on this later.

Pairwise stability (PS) is an equilibrium concept suited for BLF. Essentially, it introduces a minimum of cooperation between players: a link not being in the final graph requires the additional justification that building the link would be an impairment for at least one of its endpoints. On the other hand, PS is only concerned with single-link deviations. Jackson and Wolinsky discussed several variations of PS, including what would later be known as pairwise Nash equilibrium (PNE), a strengthening of PS.

Watts [20] in 2001 studied the symmetric connections model with an extended equilibrium concept: a graph is considered stable if no player wishes to sell any link and if no two players wish to establish an additional link while deleting any number of their links. Calvó-Armengol and Ilkiliç [21] and Corbo and Parkes [22] in 2005 discussed different equilibrium concepts and their relations: PNE, PS, and proper equilibrium [19].

Bloch and Jackson [23] in 2007 introduced a model with transfers: each player $v$ decides how much she is willing to pay for a link $\{v, w\}$ or how much she would demand the other endpoint $w$ to pay for the link. If $v$ offers at least as much as $w$ demands, or vice versa, the link $\{v, w\}$ is established in the final graph. Appropriate equilibrium concepts were introduced and discussed. Bloch and Jackson also compared PS, PNE, and their transfer model in a separate publication [24].

Bala and Goyal [18] in 2000 and in a unilateral setting studied a model where players wish to be connected by a path to as many other players as possible, but path lengths are unimportant. They also considered a unilateral version of the symmetric connections model. In another publication [9] in the same year, they extended the first model by allowing each link to fail with a probability $1-p$. Haller and Sarangi $[16,17]$ in 2003 extended this model again by allowing each link $\{v, w\}$ to fail with its own probability $1-p_{v w}$. We will elaborate on this later.

Anshelevich, Dasgupta, Tardos, and Wexler [25] in 2003 studied the price of anarchy and algorithmic aspects of a model in which each player has a set of terminals and aims to construct a network which connects her terminals. For a related model, Anshelevich, Dasgupta, Kleinberg, Tardos, Wexler, and Roughgarden [26] in 2004 studied the price of stability. Also in 2004, Christin and Chuang [27] studied a model for network formation with an extended cost function modeling peer-to-peer networks, and Christin, Grossklags, and Chuang [28] looked at it under the aspect of different game-theoretic principles.

Chun, Fonseca, Stoica, and Kubiatowicz [13] in 2004 experimentally studied an extended version of the sum-distance model.

Johari, Mannor, and Tsitsiklis [29] in 2006 studied a model in which each vertex wishes to send a given amount of traffic to some of the other vertices, and only cares whether the traffic eventually arrives at the destination. There is a handling cost at each vertex, which is proportional to the amount of traffic through that vertex.

The work of Fabrikant, Luthra, Maneva, Papadimitriou, and Shenker [6] from 2003 is to the best of the author's knowledge the first quantitative study of the price of anarchy in a model that fits into the framework considered here, as per Section 2. They considered the unilateral sum-distance model and proved a bound of $\max \{1, O(\sqrt{\alpha})\}$ on the price of anarchy in general, and an $O(1)$ bound for $\alpha>\frac{(n-1) n}{2}$. They conjectured that for $\alpha=\Omega(1)$, all non-transient NE were trees-the Tree Conjecture. A NE is called transient when there exists a sequence of strategy changes in which each player changing her strategy maintains her individual cost, and finally a strategy profile is reached which is not a NE anymore. The Tree Conjecture was based on the observation that all NE constructed so far at that time, for $\alpha>2$, were trees or transient ones (namely the Petersen graph for $\alpha \leq 4$ ). The Tree Conjecture was later, in 2006, disproved by Albers, Eilts, Even-Dar, Mansour, and Roditty [30] by showing that for each $n_{0}$, there exists a non-transient NE on $n \geq n_{0}$ vertices containing cycles, for any $1<\alpha \leq \sqrt{n / 2}$.

Corbo and Parkes [22] in 2005 considered the bilateral version of the sum-distance model. They showed an $O(\sqrt{\alpha})$ bound for $1 \leq \alpha<n^{2}$ on the price of anarchy. As noticed later in 2007 by Demaine et al. [31], the proof in fact yields $O(\min \{\sqrt{\alpha}, n / \sqrt{\alpha}\})$.

Albers et al. [30] in 2006 not only disproved the Tree Conjecture, but also improved the bounds on the price of anarchy for the unilateral sum-distance model: they gave constant upper bounds for $\alpha=O(\sqrt{n})$ and $\alpha \geq 12 n\lceil\log n\rceil$, as well as an upper bound for any $\alpha$ of

$$
15\left(1+\left(\min \left\{\alpha^{2} / n, n^{2} / \alpha\right\}\right)^{1 / 3}\right)
$$

An $O(1)$ upper bound for $\alpha=O(\sqrt{n})$ was also independently proved by Lin [32]. These bounds were again improved by Demaine, Hajiaghayi, Mahini, and Zadimoghaddam [31] in 2007. They showed a bound of $2^{O(\sqrt{\log n})}$ for any $\alpha$ and a constant bound for $\alpha=O\left(n^{1-\varepsilon}\right)$ for any constant $\varepsilon>0$. For the bilateral version, they proved the $O(\min \{\sqrt{\alpha}, n / \sqrt{\alpha}\})$ bound of Corbo and Parkes tight. Recently, in 2010, Mihalák and Schlegel [33] proved that for the unilateral sum-distance model and $\alpha \geq 273 n$, all equilibria are trees, which implies a constant bound on the price of anarchy in that range of $\alpha$.

Moscibroda, Schmid, and Wattenhofer [34] in 2006 studied the price of anarchy in a variation of the sum-distance model where the distance between two vertices is generalized, that is, it may be given by any metric. The cost function uses the stretch, that is the actual distance in the constructed graph divided by the distance that a direct connection would provide. Halevi and Mansour [35] in 2007 studied the price of anarchy in the sum-distance model under the generalization that each player has a list of "friends", that is, a list of other vertices and she is only interested in her distance to those. Demaine et al. in [31] in 2007 also considered the max-distance model: indirect cost for $v$ is $\max _{w \in V} \operatorname{dist}(v, w)$. Upper bounds were shown for ULF and tight bounds for BLF. For ULF, improved bounds were recently shown in [33].

Brandes, Hoefer, and Nick [36] in 2008 studied a variant of the sum-distance model assigning a finite distance to pairs of disconnected players, allowing for disconnected equilibria. They proved structural properties and bounds on the price of anarchy. Laoutaris, Poplawski, Rajaraman, Sundaram, and Teng [37] in 2008 considered bounded budget connection games, a variant of the sum-distance model with player-dependent link costs, lengths, and preferences $w(u, v)$ expressing the importance for player $u$ of having a good connection to player $v$, and finally a budget for each player limiting the number of links that this player can build. They considered existence of equilibria and proved bounds on the price of anarchy and stability. An important special case is the uniform version, which has link costs, link lengths,
and preferences all equal, and also all players have the same limit on their budget. Recently, this uniform version was also studied by Demaine and Zadimoghaddam [38]. They proved a tight upper bound and, more importantly, showed how to induce equilibria with small social cost. They used a technique called public service advertising, previously studied for different games by Balcan, Blum, and Mansour [39].

Baumann and Stiller [15] in 2008 considered the price of anarchy in the symmetric connections model. Demaine et al. [40] in 2009 studied the price of anarchy in a cooperative variant of the sum-distance model. They also looked at the case that links can only be formed for certain pairs of vertices, that is, the underlying "host" graph needs not to be a complete one.

### 4.1. Comparison of Our Model with Related Work

Our adversary model addresses robustness in a way that, to the best of the author's knowledge, has not been studied theoretically before. We compare our approach to previous work that also addresses robustness.

Chun, Fonseca, Stoica, and Kubiatowicz [13] experimentally studied an extended version of the sum-distance model and considered robustness. To simulate failures, they removed some vertices randomly. To simulate attacks, they removed vertices starting with those having highest degree.

The symmetric connections model of Jackson and Wolinsky [14] can also be interpreted from a robustness point-of-view. Recall that in the symmetric connections model there is a parameter $\delta \in(0,1)$, and payoff $\pi_{v}(S)$ for player $v$ under strategy profile $S$ is defined

$$
\pi_{v}(S):=\sum_{\substack{w \in V \\ w \neq v}} \delta^{\operatorname{dist}_{G(S)}(v, w)}-\left|S_{v}\right| \alpha
$$

An interpretation is that $v$ receives one unit of income from each other vertex $w$ along a shortest path between $v$ and $w$. However, each link has a probability $1-\delta$ of failure, so the expected income from $w$ is the probability that none of the $\operatorname{dist}_{G(S)}(v, w)$ links fails, which is $\delta^{\operatorname{dist}_{G(S)}(v, w)}$ if we assume stochastic independence of failures. For BLF, Baumann and Stiller [15] gave an expression for the exact price of anarchy for $\alpha \in\left(\delta-\delta^{2}, \delta-\delta^{3}\right)$, which implies an $O(1)$ bound (the constant being bounded by $\frac{4}{1+2 \delta}$ ). The price of anarchy is 1 for $\alpha<\delta-\delta^{2}$, following from [14]. The price of anarchy in the range $\alpha>\delta-\delta^{3}$ is not fully understood yet.

The symmetric connections model is different from ours in many respects:

- All links have the same probability of failure. In our model, links can have different probabilities, and these may even depend on the final graph. [41]
- The failure of a link $e$ and the failure of a link $f$ are independent events for $e \neq f$, at least along the concerned paths. In our model, the failures of $e$ and $f$ are mutually exclusive events.
- Alternative paths are not considered; it is assumed that routing happens along a specific shortest path that is fixed before the random experiment that models the link failures is conducted. In our model, all paths are considered. However, we do not consider path lengths.

Bala and Goyal [9] studied a variation of the symmetric connections model, which is closer to ours. In their model, each vertex receives an amount of 1 from each vertex it is connected to via some path. Each
link has a probability $1-p$ of failure, $p \in[0,1]$ being the same for all links and independent of the final graph. Failures of two distinct links are stochastically independent. Income of a vertex $v$ is the expected number of vertices to which $v$ is connected via a path. Unilateral link formation is used. They considered structural properties of optima and NE, in particular pointing out cases where NE are "super-connected", i.e., connected and not containing bridges. They also showed that for some regions of parameters, there exist NE that are also optima (i.e., they show a price of stability of 1 for these regions).

Haller and Sarangi $[16,17]$ studied an extension of the model of Bala and Goyal [9]. In their model, each link $\{v, w\}$ may fail with its $o w n$ probability $1-p_{v w}$. They also considered structural properties of optima and NE as well as relations of optima and NE, including the price of stability similar to [9]. Like the symmetric connections model, their model shows several differences to ours:

- The failure probability of each link $\{v, w\}$ is $1-p_{v w}$, independent of the final graph [42]. In our model, failure probabilities depend on the final graph.
- Failures of two different links are stochastically independent. In our model, they are mutually exclusive events. (This difference is exactly as between the symmetric connections model and ours.)

Generally, independent link failures model the unavailability of links due to, e.g., deterioration, maintenance times, or influences affecting the whole infrastructure or large parts of it (e.g., natural disasters). Our adversary model, on the other hand, models the situation when faced with an entity that is malicious but only has limited means so that it can only destroy a limited number of links (we limit this number to 1 in this work).

## 5. The Bridge Tree

We conduct some preparations for the analysis of equilibria in our adversary model, which will be useful regardless of the link formation rule and the equilibrium concept. In the end, in Lemma 5.2, we will have developed a simple method to bound the sum of relevances $R(v)$ for each player $v$, which will later help to bound the disconnection cost. It will be helpful in several places to consider a variation of the block graph [44], which we call the bridge tree. Its definition requires some preparation. If $W \subseteq V$ is maximal under the condition that the induced subgraph $G[W]$ is connected and does not contain any bridges of $G[W]$, we call $W$ a bridgeless connected component, abbreviated "BCC". The proof of the following proposition is straightforward.
5.1 Proposition. A set of vertices $W \subseteq V$ is a BCC if and only if $W$ is maximal under the condition that the induced subgraph $G[W]$ is connected and does not contain any bridges of $G$.

Proof. Let $W$ be maximal under the condition of $G[W]$ being connected and not containing any bridges of $G[W]$, i.e., we follow the original definition given above. Clearly, $G[W]$ does not contain any bridges of $G$, since if removal of some edge disconnects $G$, then it also disconnects $G[W]$ if the endpoints of this edge are in $W$. We choose $U \supseteq W$ maximal under the condition that $G[U]$ is connected and $G[U]$ does not contain any bridges of $G$. Suppose $U \neq W$. Then $G[U]$ contains a bridge $e$ of $G[U]$. Since this is no bridge of $G$, it is located on a cycle $C$. Then $V(C) \nsubseteq U$, since $e$ is a bridge of $G[U]$. But $G[U \cup V(C)]$ would still be connected and would contain no bridge of $G$. This contradicts the maximality of $U$.

Now let $W$ be maximal under the condition of $G[W]$ being connected and not containing any bridges of $G$. If $G[W]$ contained a bridge $e$ of $G[W]$ (but not of $G$ ), we could use the cycle-argument from before to augment $W$ and have a contradiction to its maximality. Suppose there is $U \supsetneq W$ such that $G[U]$ is connected and $G[U]$ does not contain any bridges of $G[U]$. Then $G[U]$ contains a bridge of $G$. As noted earlier, this is also a bridge of $G[U]$, a contradiction.

What we call "BCC" is sometimes called "block" in the literature, and what we call "bridge tree" is then called "bridge-block tree". We refrain from using the term "block" here, since it usually is related to vertex-connectivity; see [43] (p.55).

Every vertex is contained in exactly one BCC. If $W$ is a BCC, we have to remove at least 2 edges from $G[W]$ in order to make it disconnected. A graph from which we have to remove at least 2 edges to make it disconnected is also called being " 2 -edge-connected" in common terminology, provided that it has more than 1 vertices; see [43] (p. 12).

Now we introduce the bridge tree. It is the graph $\widetilde{G}=(\widetilde{V}, \widetilde{E})$ defined by:

$$
\begin{aligned}
& \widetilde{V}:=\{B \subseteq V ; B \text { is a BCC }\} \\
& \widetilde{E}:=\left\{\left\{B, B^{\prime}\right\} ; B, B^{\prime} \in \widetilde{V}, \exists v \in B, w \in B^{\prime}:\{v, w\} \in E\right\}
\end{aligned}
$$

Then $\widetilde{G}$ is a tree (assuming $G$ is connected). By Proposition 5.1 there is a $1: 1$ mapping between the edges of $\widetilde{G}$ and the bridges of $G$. We make the following special convention concerning the bridge tree:

Convention. Whenever we speak of the number of vertices in a subgraph $T$ of the bridge tree, we count $|B|$ for each vertex $B \in V(T)$.

In other words, we count the vertices that would be there if we expanded $T$ back to its corresponding subgraph of $G$. Figure 2 shows an example. Since each vertex of $G$ is in exactly one BCC, counting in this way for $\widetilde{V}$ yields the number of vertices in $G$, i.e., $n$.

Figure 2. Bridge tree construction. Vertices representing BCCs of more than 1 vertices have their number of vertices attached, here 4,7 , and 3 , respectively.
(a) A graph $G$.

(b) The corresponding bridge tree $\widetilde{G}$.


On several occasions, when considering the effect of building additional edges, we treat vertices of the bridge tree as players. This is justified since edges inside BCCs have relevance 0 . Hence for a strategy profile $S$ and $B, B^{\prime} \in \widetilde{V}$ the effect in disconnection cost of a new edge between a player from $B$ and a player of $B^{\prime}$ is specific to the pair $\left\{B, B^{\prime}\right\}$ and not to the particular players.

For a path $P$ in $G$, let $\widetilde{P}$ be its contracted counterpart in $\widetilde{G}$, i.e., we replace in $P$ each maximal sequence of vertices from the same BCC $B \in \widetilde{V}$ with $B$. Then the length $|\widetilde{P}|$ of $\widetilde{P}$ is the number of bridges in $P$. For each pair $v, w \in V$ denote $P(v, w)$ an arbitrary shortest path from $v$ to $w$; and $\mathcal{P}(v):=\{P(v, w) ; w \in V\}$. The bridge tree helps bounding the disconnection cost. We conclude this section with a preparation for this. For each $v \in V$ and $e \in E$ we easily observe:

$$
\operatorname{rel}(e, v)= \begin{cases}0 & \text { if } e \text { is a non-bridge }  \tag{3}\\ |\{P \in \mathcal{P}(v) ; e \in E(P)\}| & \text { if } e \text { is a bridge }\end{cases}
$$

We use this to prove the following lemma, which we will apply in Section 8. The lemma relates relevance to path length, and so to diameter. This is possible since it is the same to count for each edge the number of paths that cross this edge (establishing the connection to relevance) as to count for each path the number of its edges (establishing the connection to path length).
5.2 Lemma. For each $v \in V$ we have $R(v) \leq(n-1) \operatorname{diam}(\widetilde{G})$.

Proof. Fix $v \in V$. We have

$$
\begin{aligned}
R(v) & \stackrel{\text { def }}{=} \sum_{e \in E} \operatorname{rel}(e, v) \stackrel{(3)}{=} \sum_{\substack{e \in E \\
e \text { is a bridge }}}|\{P \in \mathcal{P}(v) ; e \in E(P)\}| \\
& =\sum_{P \in \mathcal{P}(v)} \mid\{e \in E(P) ; e \text { is a bridge }\}\left|=\sum_{P \in \mathcal{P}(v)}\right| \widetilde{P} \mid \leq(n-1) \operatorname{diam}(\widetilde{G})
\end{aligned}
$$

The last estimation is true since the bridge tree is a tree and so every path is a shortest path.

## 6. A Simple Bound on the Price of Anarchy

We give an upper bound on the price of anarchy for a general adversary. It holds independently of the link formation rule and the equilibrium concept, provided that equilibria have few edges.
6.1 Proposition. Let $S$ be any strategy profile and (as usual) $m=|E(S)|$.
(i) If $m=O(n)$, then $\frac{C(S)}{\operatorname{OPT}(n, \alpha)}=O\left(1+\frac{n}{\alpha}\right)$.
(ii) If $m=O(n)$ and $\alpha=\Omega(n)$, then $\frac{C(S)}{\operatorname{OPT}(n, \alpha)}=O(1)$.

Proof. Since $\operatorname{sep}(e)=O\left(n^{2}\right)$ for all $e$, we have

$$
C(S)=O\left(m \alpha+n^{2} \sum_{e \in E} \operatorname{Pr}(\{e\})\right)=O\left(n \alpha+n^{2}\right)
$$

Since an optimum is connected, the optimal social cost is $\Omega(n \alpha)$. Dividing by this yields (i). Assertion (ii) follows from (i).

The following corollary is obvious.
6.2 Corollary. Fix any link formation rule and equilibrium concept.
(i) If the number of edges in each equilibrium is $O(n)$, then the price of anarchy is $O\left(1+\frac{n}{\alpha}\right)$.
(ii) If the number of edges in each equilibrium is $O(n)$ and moreover $\alpha=\Omega(n)$, then the price of anarchy is $O(1)$.

A remark on the meaning of $O$ and $\Omega$ is in order. Recall that we use this notation to avoid having to introduce all occurring constants explicitly, and that a constant is required to be independent of all game parameters, strategy profiles, etc. For example, the constant hidden in the premise " $m=O(n)$ " in Proposition 6.1(i) is required to be independent of $n, \alpha$, and $S$, while the constant hidden in the conclusion " $\frac{C(S)}{\operatorname{OPT}(n, \alpha)}=O\left(1+\frac{n}{\alpha}\right)$ " is guaranteed to have the same independence. The proof reflects that this is true. In Corollary 6.2(i), it is required that there exists a constant $c>0$ such that for each equilibrium $S$ we have $|E(S)| \leq c n$. Here as well, $c$ is required to be independent of $n, \alpha$, and $S$.

The main goal of Sections 8 and 9 is to show a bound of $O(1)$ for ULF and NE and all [45] $n \geq 9$ and $\alpha$, but restricted to two special adversaries, which are chosen to mark extreme cases. We will there proceed in showing the $O(n)$ bound on the number of edges in a NE first. In Section 8, we will then bound disconnection cost of NE by $O(n \alpha)$. In Section 9, we achieve the same bound under the condition that $\alpha<c n$ for a constant specified there. If $\alpha \geq c n$, then we are done by Corollary 6.2(ii).

The $O(1)$ bound would follow trivially if we could show an $O(n \alpha)$ bound for the social cost of any strategy profile. However, later Proposition 8.1 shows that there is no hope for this, and hence we will have to exploit characteristics of NE in order to prove our bounds.

## 7. Optima, Nash Equilibria, Price of Stability

We stick to ULF and NE for the rest of this work. The aim of this section is to construct optima and NE, and finally to show how a bound on the price of stability follows easily. The adversaries considered are a general one, i.e., without any additional assumptions, and one inducing anonymous disconnection cost.
7.1 Proposition. An optimum has social cost $\Theta(n \alpha)$. More precisely:
(i) If $\alpha \leq 2(n-1)$, the cycle is an optimum; it has social cost $n \alpha$.
(ii) If $\alpha \geq 2(n-1)$, a star is an optimum; it has social cost $(n-1)(\alpha+2)$.

Proof. An optimum can only be the cycle or a tree, because any graph containing a cycle has already the building cost $n \alpha$ of the cycle, and the cycle has optimal disconnection cost. So an optimum is either the cycle, or it is cycle-free. Let $T$ be any tree. We have its indirect cost:

$$
\sum_{e \in E(T)} \operatorname{sep}(e) \operatorname{Pr}(\{e\})=2 \sum_{e \in E(T)} \nu(e)(n-\nu(e)) \operatorname{Pr}(\{e\}) \geq 2 \cdot 1(n-1) \sum_{e \in E(T)} \operatorname{Pr}(\{e\})=2(n-1)
$$

Hence the social cost of a tree is at least

$$
(n-1) \alpha+2(n-1)=(n-1)(\alpha+2)
$$

The social cost of the cycle is $n \alpha$. So if $\alpha \leq 2(n-1)$, the cycle is better or as good as any tree, hence it is an optimum. If $\alpha>2(n-1)$, then we look for a good tree. A star has social cost $(n-1)(\alpha+2)$, which matches the lower bound given above, and is hence optimal (and better than the cycle). However, for $\alpha=2(n-1)$, both cycle and star are optimal with social cost $2 n(n-1)$.

The following simple remark later will help establishing concrete bounds on the price of anarchy.
7.2 Remark. Assume there are constants $c_{0}, c_{1}>0$ such that the social cost of all NE is bounded by $\left(c_{1} n+c_{0}\right) \alpha$. Then the price of anarchy is bounded by $c_{1}+\frac{c_{1}+c_{0}}{n-1}$.
Proof. If the optimum is $n \alpha$, we have the ratio $\frac{\left(c_{1} n+c_{0}\right) \alpha}{n \alpha}=c_{1}+\frac{c_{0}}{n} \leq c_{1}+\frac{c_{1}+c_{0}}{n-1}$. Otherwise, if the optimum is $(n-1)(\alpha+2)$, we have the ratio $\frac{\left(c_{1} n+c_{0}\right) \alpha}{(n-1)(\alpha+2)}<\frac{c_{1} n+c_{0}}{n-1}=\frac{c_{1}(n-1)+c_{1}+c_{0}}{n-1}=c_{1}+\frac{c_{1}+c_{0}}{n-1}$.

The following two propositions can be proved by appropriate cost-benefit analysis.
7.3 Proposition. Let $S$ be a star with edges pointing outward.
(i) If $\alpha \geq n-1$, then $S$ is a $N E$.
(ii) If $\alpha \geq 2-\frac{1}{n-1}$, then $S$ is a NE if disconnection cost is anonymous.

In both cases, strict inequality implies a MaxNE.
Proof. (i) Since all edges point outwards, the center player is the only one that could sell edges, but this would make the graph disconnected. Exchanges of edges by the center cannot lead to a different strategy profile. The maximum disconnection cost is experienced by a leaf vertex when the probability measure is concentrated on the one edge that connects it to the rest. The disconnection cost is then $n-1$. Since this is at most $\alpha$, there is no incentive to buy additional edges. Hence no player can strictly improve her individual cost by changing her strategy. (If multiple edges between the same players were allowed, the center could build additional edges. However, since $\alpha>1$, this would increase her cost.)
(ii) Disconnection cost of the center is 1 . By anonymity, all leafs experience the same disconnection cost. It follows easily from this that all edges have the same probability, namely $\frac{1}{n-1}$. Disconnection cost of a leaf hence is

$$
\frac{(n-1)+(n-2)}{n-1}=\frac{(n-1)+(n-1)-1}{n-1}=2-\frac{1}{n-1}
$$

Now we apply the same arguments as for part (i). Maximality is clear in both cases by Remark 2.1.
7.4 Proposition. Let $S$ be a cycle with all edges pointing in the same direction (either all clockwise or all counter-clockwise).
(i) If $\alpha \leq 1$, then $S$ is a MaxNE.
(ii) If $\alpha \leq \frac{1}{2}\left\lfloor\frac{n-1}{2}\right\rfloor$, then $S$ is a MaxNE if disconnection cost is anonymous.

Proof. Maximality in both cases is due to the cycle having minimum disconnection cost, namely 0 . Buying or exchanging edges is also not beneficial since the cycle already has minimum disconnection cost. We only have to check whether it is beneficial for a player $v$ to sell her one edge. Selling the edge yields a path with $v$ at one of its ends. This increases disconnection cost for $v$ to at least 1 , since the removal of any edge disconnects $v$ from at least one other vertex. This proves (i).

To prove (ii), we have to establish a better lower bound on the new disconnection cost for $v$. Anonymity of indirect cost allows us to do so. Let the path be $\left(v_{1}, e_{1}, v_{2}, \ldots, e_{n-1}, v_{n}\right)$ with $v=v_{1}$. We claim that

$$
\begin{equation*}
\operatorname{Pr}\left(\left\{e_{i}\right\}\right)=\operatorname{Pr}\left(\left\{e_{n-i}\right\}\right) \quad \text { for all } i \in\{1, \ldots, n-1\}, \tag{4}
\end{equation*}
$$

i.e., the adversary behaves like a symmetric one on this graph. From (4) it follows $\sum_{i=1}^{[n-1) / 2\rceil} \operatorname{Pr}\left(\left\{e_{i}\right\}\right) \geq \frac{1}{2}$. Since each of the edges $e_{1}, \ldots, e_{\lceil(n-1) / 2\rceil}$ has relevance at least $\lfloor(n-1) / 2\rfloor$ for $v$, the proposition follows.

We are left with proving (4). For each $i \in[n]$ we write $I_{i}$ for the indirect cost of vertex $v_{i}$, and moreover define its left indirect cost by $I_{i}^{\text {left }}:=\sum_{j=1}^{i-1} j \operatorname{Pr}\left(\left\{e_{j}\right\}\right)$ and its right indirect cost by $I_{i}^{\text {right }}:=$ $\sum_{j=1}^{n-i} j \operatorname{Pr}\left(\left\{e_{n-j}\right\}\right)$. Then clearly $I_{i}=I_{i}^{\text {left }}+I_{i}^{\text {right }}$. It suffices to show (4) for $i<\frac{n}{2}$. We have (in fact even for $i \leq n-1$ ) on the one hand:

$$
\begin{aligned}
I_{i} & =I_{i}^{\text {left }}+(n-i) \operatorname{Pr}\left(\left\{e_{i}\right\}\right)+I_{i+1}^{\text {right }} \\
& =I_{i}^{\text {left }}+(n-i) \operatorname{Pr}\left(\left\{e_{i}\right\}\right)+I_{i+1}-I_{i+1}^{\text {left }} \\
& =I_{i}^{\text {left }}+(n-i) \operatorname{Pr}\left(\left\{e_{i}\right\}\right)+I_{i+1}-\left(I_{i}^{\text {left }}+i \operatorname{Pr}\left(\left\{e_{i}\right\}\right)\right) \\
& =(n-2 i) \operatorname{Pr}\left(\left\{e_{i}\right\}\right)+I_{i+1}
\end{aligned}
$$

On the other hand, along the same lines we prove $I_{n-i+1}=(n-2 i) \operatorname{Pr}\left(\left\{e_{n-i}\right\}\right)+I_{n-i}$. By anonymity, $I_{i}=I_{n-i+1}$ and $I_{i+1}=I_{n-i}$. It follows $(n-2 i) \operatorname{Pr}\left(\left\{e_{i}\right\}\right)=(n-2 i) \operatorname{Pr}\left(\left\{e_{n-i}\right\}\right)$, and since $i<\frac{n}{2}$, this means $\operatorname{Pr}\left(\left\{e_{i}\right\}\right)=\operatorname{Pr}\left(\left\{e_{n-i}\right\}\right)$.

For anonymous disconnection cost, this proves existence of NE for all ranges of $\alpha$ provided that $n \geq 9$, since then $2-\frac{1}{n-1} \leq 2 \leq \frac{1}{2}\left\lfloor\frac{n-1}{2}\right\rfloor$. In the range $2-\frac{1}{n-1} \leq \alpha \leq \frac{1}{2}\left\lfloor\frac{n-1}{2}\right\rfloor$ two very different topologies-namely cycle and star-co-exist as NE.

Convention. In the following, we assume $n \geq 9$ where necessary. In particular, when speaking of the price of anarchy or stability, we require existence of equilibria, so $n \geq 9$ will be assumed in all those places. We then also use $n \geq 9$ in the calculation of bounds, where appropriate.

The following is a consequence of Propositions 7.1, 7.3, and 7.4.
7.5 Theorem. For anonymous disconnection cost the price of stability is $1+o(1)$, more precisely it is bounded by $1+\frac{8}{n-2}$.

Proof. For $\alpha \leq \frac{1}{2}\left\lfloor\frac{n-1}{2}\right\rfloor$ the cycle is a NE as well as an optimum, and so the price of stability is 1. For $\alpha \geq 2(n-1)$ a star is a NE as well as an optimum, and so the price of stability is 1.

For $\frac{1}{2}\left\lfloor\frac{n-1}{2}\right\rfloor \leq \alpha \leq 2(n-1)$, the star is a NE and the cycle is an optimum. The price of stability so is upper-bounded by

$$
\frac{(n-1)(\alpha+2)}{n \alpha} \leq 1+\frac{2}{\alpha} \leq 1+\frac{4}{\left\lfloor\frac{n-1}{2}\right\rfloor} \leq 1+\frac{8}{n-2}
$$

## 8. Simple-Minded Adversary

The simple-minded adversary picks an edge uniformly at random, that is, $\operatorname{Pr}(\{e\})=\frac{1}{m}$ for all $e \in E$. Then we have individual and social cost:

$$
\begin{aligned}
C_{v}(S) & =\left|S_{v}\right| \alpha+\frac{1}{m} \sum_{e \in E} \operatorname{rel}(e, v)=\left|S_{v}\right| \alpha+\frac{1}{m} R(v) \quad \text { for } v \in V \\
C(S) & =m \alpha+\frac{1}{m} \sum_{v \in V} R(v)
\end{aligned}
$$

Clearly, this is a symmetric adversary and hence disconnection cost is anonymous. All results in this section are for the simple-minded adversary. As promised earlier, we give an example for a non-linear (in $n$ ) social cost.
8.1 Proposition. Social cost of a path is $(n-1) \alpha+\frac{1}{3} n(n+1)=\Theta\left(n \alpha+n^{2}\right)$.

Proof. We have the social cost of a path:

$$
\begin{aligned}
& m \alpha+\frac{1}{m} \sum_{v \in V} \sum_{e \in E} \operatorname{rel}(e, v)=(n-1) \alpha+\frac{1}{n-1} \sum_{e \in E} \sum_{v \in V} \operatorname{rel}(e, v) \\
& =(n-1) \alpha+\frac{1}{n-1} \sum_{e \in E} \operatorname{sep}(e) \\
& =(n-1) \alpha+\frac{1}{n-1} \sum_{e \in E} 2 \nu(e)(n-\nu(e)) \\
& =(n-1) \alpha+\frac{2}{n-1} \sum_{k=1}^{n-1} k(n-k) \\
& =(n-1) \alpha+\frac{2}{n-1}\left(n \frac{(n-1) n}{2}-\frac{(n-1) n(2 n-1)}{6}\right) \\
& =(n-1) \alpha+\frac{1}{3} n(n+1)=\Theta\left(n \alpha+n^{2}\right)
\end{aligned}
$$

### 8.1. Bounding Cost Changes and Cycle Length

We estimate the benefit for a player of building or selling a particular edge. This will become useful in several places. It moreover immediately leads to a structural result on the length of cycles. The following remark is purely graph-theoretic.
8.2 Remark. Let $G=(V, E)$ be a graph.
(i) Let $e=\{v, w\} \notin E$ and $C$ be any cycle in $G+e$ with $e \in E(C)$. Then all bridges in $G$ that are non-bridges in $G+e$ are located on $C$.
(ii) Let $e=\{v, w\} \in E$ be a non-bridge and $C$ be any cycle with $e \in E(C)$. Then all bridges of $G-e$ that are non-bridges in $G$, are in $E(C)$.
Proof. (i) The additional edge $e$ creates exactly one cycle $\widetilde{C}$ in the bridge tree. All bridges in $G$ that are non-bridges in $G+e$ correspond to edges on $\widetilde{C}$, and all those in turn correspond to edges on $C$.
(ii) Let $f$ be a non-bridge in $G$ and a bridge in $G-e$. Then $G-e$ consists of two subgraphs $G_{1}$ and $G_{2}$ that are connected only by $f$. Since $f$ was no bridge before $e$ was removed, $e$ must also connect $G_{1}$ with $G_{2}$. Moreover, there are no other edges between $G_{1}$ and $G_{2}$. It follows that any cycle that contains $e$ also contains $f$.
8.3 Proposition. For each player $v$ we have $R(v) \leq \frac{n(n-1)}{2}$.

Proof. We repeat the counting argument from the proof of Lemma 5.2:

$$
\begin{aligned}
R(v) & \stackrel{\text { def }}{=} \sum_{e \in E} \operatorname{rel}(e, v) \stackrel{(3)}{=} \sum_{\substack{e \in E \\
e \text { is a bridge }}}|\{P \in \mathcal{P}(v) ; e \in E(P)\}| \\
& =\sum_{P \in \mathcal{P}(v)} \mid\{e \in E(P) ; e \text { is a bridge }\}\left|\leq \sum_{P \in \mathcal{P}(v)}\right| E(P) \mid=\sum_{w \in V} \operatorname{dist}(v, w)
\end{aligned}
$$

This is maximal if $G$ is a path with $v$ at its end; then $R(v)=\frac{n(n-1)}{2}$.
Fix a player $v$. Let $R:=R(v)$ and let $R^{\prime}$ be the same quantity when an additional edge $e$ is built by $v$. By the previous proposition, we have $R, R^{\prime} \leq \frac{n(n-1)}{2}$. The benefit in disconnection cost of building this edge for player $v$ is $\frac{1}{m} R-\frac{1}{m+1} R^{\prime}$. Due to the change in denominators from " $m$ " to " $m+1$ " this expression looks somewhat unhandy. Yet, we can give good bounds incorporating the change in relevances, $\Delta R:=R-R^{\prime} \geq 0$, with one denominator. We can do something similar for the case when the player sells an edge, where we put $\Delta R:=R^{\prime}-R \geq 0$.

### 8.4 Proposition.

(i) If a player builds an additional edge and the sum of her relevances drops from $R$ to $R^{\prime}$ by $\Delta R:=R-R^{\prime}$, then her improvement in disconnection cost is at least $\frac{1}{m+1} \Delta R$ and at most $\frac{1}{2}+\frac{1}{m+1} \Delta R \leq \frac{n}{2}$.
(ii) If a player sells a non-bridge and the sum of her relevances increases from $R$ to $R^{\prime}$ by $\Delta R:=R^{\prime}-R$, then her impairment in disconnection cost is at least $\frac{1}{m} \Delta R$ and at most $\frac{1}{2}+\frac{1}{m} \Delta R \leq \frac{n}{2}$.

Proof. (i) We have

$$
\begin{aligned}
& \frac{1}{m} R-\frac{1}{m+1} R^{\prime}=\frac{1}{m} R-\frac{1}{m+1}\left(R+\left(R^{\prime}-R\right)\right)=\left(\frac{1}{m}-\frac{1}{m+1}\right) R+\frac{1}{m+1} \Delta R \\
= & \frac{1}{m(m+1)} R+\frac{1}{m+1} \Delta R \quad\left\{\begin{array}{l}
\leq \frac{1}{2}+\frac{1}{m+1} \Delta R \leq \frac{n}{2} \\
\geq \frac{1}{m+1} \Delta R
\end{array}\right.
\end{aligned}
$$

We used $\Delta R \leq \frac{n(n-1)}{2}$ and $n-1 \leq m$ for the upper bound. (ii) is proved alike, using $n \leq m$ since the graph contains a cycle.

### 8.5 Proposition.

(i) If a player builds an edge creating a cycle of length $\ell$, the improvement in disconnection cost is at most $\frac{1}{2}+\frac{1}{m+1}(\ell-1)\left(n-\frac{\ell}{2}\right)$. (The graph is allowed to already contain other cycles.)
(ii) If a player sells an edge destroying a cycle of length $\ell$, the impairment in disconnection cost is at most $\frac{1}{2}+\frac{1}{m}(\ell-1)\left(n-\frac{\ell}{2}\right)$.

Proof. (i) Let $C=\left(v, e_{1}, v_{1}, \ldots, v_{\ell-1}, e_{\ell}, v\right)$ be any new cycle, created by the new edge $e_{\ell}$ bought by $v$. By Remark 8.2(i), all edges for which a change in relevance occurs by adding $e_{\ell}$, i.e., all edges that were bridges and become non-bridges due to the new edge, are located on this cycle. In the best case, i.e., in case of maximal improvement,

- all $\ell$ edges were bridges before and became non-bridges now, and
- without the additional edge, $n-1$ vertices are reached from $v$ only through $e_{1}, n-2$ through the next edge, and so on; edge $e_{\ell-1}$ is relevant for $(n-(\ell-1))$ vertices.
It follows $\Delta R \leq \sum_{k=1}^{\ell-1}(n-k)=(\ell-1) n-\sum_{k=1}^{\ell-1} k=(\ell-1) n-\frac{(\ell-1) \ell}{2}=(\ell-1)\left(n-\frac{\ell}{2}\right)$. The statement follows with Proposition 8.4(i).
(ii) By Remark 8.2(ii), we may consider any cycle that is destroyed. The rest is the same calculation as for (i).
8.6 Proposition. Let $\ell<\alpha+\frac{1}{2}$.
(i) If a player builds an edge creating a cycle of length $\ell$, she suffers an impairment in her cost.
(ii) If a player sells an edge destroying a cycle of length $\ell$, she experiences an improvement in her cost.

Proof. (i) By Proposition 8.5(i), the player suffers an impairment in her cost if

$$
\alpha>\frac{1}{2}+\frac{1}{m+1}(\ell-1)\left(n-\frac{\ell}{2}\right)
$$

Since $m \geq n-1$, this is the case if $\alpha>\frac{1}{2}+\frac{1}{n}(\ell-1)\left(n-\frac{\ell}{2}\right)$, which is the same as $n\left(\alpha+\frac{1}{2}\right)>\ell\left(n-\frac{\ell}{2}+\frac{1}{2}\right)$. Since $\ell \geq 3 \geq 1$, this is the case if $n\left(\alpha+\frac{1}{2}\right)>\ell n$.

We show (ii) in almost exactly the same way, using Proposition 8.5 (ii) and that $m \geq n$, since the original graph contains a cycle.

It follows the structural result:
8.7 Corollary. No NE contains cycles shorter than $\alpha+\frac{1}{2}$.

## Bounding the Price of Anarchy

The following observation is the key to showing that a NE does not have many more edges than a tree.

### 8.8 Proposition. A NE is chord-free.

Proof. Selling a chord $e=\{v, w\}$ from a cycle $C=(v, \ldots, w, \ldots, v)$ does not increase the relevance of any edge for any player. To see this, let $C^{\prime}$ be a cycle which contains $e$. Then $C^{\prime}-e$ also forms a cycle with a part of $C$, say $(v, \ldots, w)$. Hence, if the graph is bridgeless, removing a chord would decrease the player's building cost without increasing the disconnection cost. Now let the graph contain a bridge $e^{\prime}$. Due to the decrease in the denominator of the disconnection cost, removing a chord impairs the disconnection cost. However, the player owning the chord, say $v$, would rather remove the chord and instead build an edge to form a new cycle containing $e^{\prime}$. The only case where this is impossible is when $v$ is one endpoint of the bridge $e^{\prime}=\{v, u\}$, and $u$ is a leaf vertex. Then, a double-edge between $v$ and $u$ would be needed, which is not allowed unless we use a multigraph.

We consider this case now and show that we in fact do not need a multigraph. By selling the chord, the disconnection cost for $v$ increases by $\frac{1}{m(m-1)} R(v)$. If this increase is strictly smaller than $\alpha$, we are done. Hence assume $\frac{1}{m(m-1)} R(v) \geq \alpha$ now. Edge $\{v, u\}$ has relevance $n-1$ for $u$. Due to the positions of $v$ and $u$, we have $R(u)=R(v)+(n-1)-1$. If $u$ builds an edge to any other vertex, save $v$, edge $e^{\prime}$ is put on a cycle. The improvement in disconnection cost for $u$ by building such an edge is at least

$$
\begin{aligned}
& \frac{1}{m} R(u)-\frac{1}{m+1}(R(u)-(n-1))=\frac{1}{m}(R(v)+n-2)-\frac{1}{m+1}(R(v)-1) \\
= & \left(\frac{1}{m}-\frac{1}{m+1}\right) R(v)+\frac{n-2}{m}+\frac{1}{m+1}=\frac{1}{m(m+1)} R(v)+\frac{n-2}{m}+\frac{1}{m+1} \\
= & \left(\frac{1}{m(m-1)}+\frac{1}{m(m+1)}-\frac{1}{m(m-1)}\right) R(v)+\frac{n-2}{m}+\frac{1}{m+1} \\
\geq & \alpha-\frac{1}{m}\left(\frac{1}{m-1}-\frac{1}{m+1}\right) R(v)+\frac{n-2}{m}+\frac{1}{m+1} \\
\geq & \alpha-\frac{1}{m}\left(\frac{1}{m-1}-\frac{1}{m+1}\right) \frac{n(n-1)}{2}+\frac{n-2}{m}+\frac{1}{m+1} \\
\geq & \alpha-\left(\frac{1}{m-1}-\frac{1}{m+1}\right) \frac{n-1}{2}+\frac{n-2}{m}+\frac{1}{m+1} \\
= & \alpha-\frac{2}{(m-1)(m+1)} \frac{n-1}{2}+\frac{n-2}{m}+\frac{1}{m+1} \\
\geq & \alpha-\frac{1}{m+1}+\frac{n-2}{m}+\frac{1}{m+1}>\alpha
\end{aligned}
$$

So $u$ has an incentive to buy an additional edge, a contradiction to NE.
The next two are graph-theoretic results. The first is a straightforward adaption of a result (and its proof) on vertex-connectivity to edge-connectivity; see [43] (Prop.3.1.3) for the version for vertex-connectivity.
8.9 Proposition. Any bridgeless connected graph can be constructed from a cycle by successively adding paths or cycles of the form $\left(u, e_{1}, v_{1}, \ldots, v_{k}, e_{k+1}, w\right)$, where $u$, $w$ are vertices of the already constructed graph and $v_{1}, \ldots, v_{k}$ are zero or more new vertices.

Proof. Clearly, any graph that was constructed in this manner is connected and bridgeless. Now let $G$ be connected and bridgeless and $H$ a subgraph of $G$ that is constructible in this manner, chosen such that it has a maximum number of edges among all such subgraphs. Since $G$ contains a cycle, $H$ is not empty. Also, $H$ is an induced subgraph since $H+e$ is also constructible for any edge $e$. If $H \neq G$, then since $G$ is connected, there is an edge $e=\{v, w\}$ with $v \notin V(H)$ and $w \in V(H)$. Since $G$ is bridgeless, this edge is on a cycle $C=\left(w, e, v=v_{1}, \ldots, v_{k}=w\right)$. Let $v_{i}$ be the first vertex with $v_{i} \in V(H)$. Then $P:=\left(w, \ldots, v_{i}\right)$ is a path or cycle of the form used in the construction, and so $H+P$ is constructible and has more edges than $H$, a contradiction.
8.10 Proposition. A chord-free graph on $n$ vertices contains no more than $3 n=O(n)$ edges.

Proof. Let $G$ be a chord-free graph, w.l.o.g. being connected. We first consider the case that $G$ is bridgeless. By the previous proposition, $G$ can be constructed from a cycle on, say, $N_{0}$ vertices, by successively adding paths of the form $\left(u, e_{1}, v_{1}, \ldots, v_{k}, e_{k+1}, w\right)$, where $u, w$ are vertices of the already constructed graph and $v_{1}, \ldots, v_{k}, k \in \mathbb{N}_{0}$, are zero or more new vertices. For any two vertices $u, w$ in the already constructed graph, there is a cycle $C$ with $u, w \in V(C)$. Since $G$ is chord-free, we may not add a path $\left(u, e_{1}, w\right)$. Hence $k \geq 1$ in each step, i.e., at least one new vertex is added. It follows that there are at most $t \leq n-N_{0}=: N_{1}$ steps in this construction. Let $n_{i}$ and $m_{i}$ be the number of new
vertices and edges, respectively, inserted in step $i$. Then $m_{i}=n_{i}+1$ for each $i \in[t]$ and so we add $\sum_{i=1}^{t} m_{i}=\sum_{i=1}^{t}\left(n_{i}+1\right)=N_{1}+t \leq 2 N_{1}$ edges to the initial cycle. It follows that $G$ has at most $N_{0}+2 N_{1} \leq 2 N_{0}+2 N_{1}=2 n$ edges.

If $G$ is not bridgeless, we consider each of its BCCs; these correspond to vertices of the bridge tree. Altogether, they cannot contribute more than $2 n$ edges. In addition, there are at most $n-1$ edges, namely bridges of $G$. So we have a bound of $2 n+n-1 \leq 3 n$.
8.11 Corollary. A NE has at most $3 n=O(n)$ edges.

Proof. Follows from Prop. 8.8 and 8.10.
Now we know that the total building cost in a NE is $O(n \alpha)$, hence it is of the same order as the optimal social cost. In order to bound the price of anarchy, we are left with bounding the disconnection cost. To this end, we make use of the bridge tree. The following is a corollary to Lemma 5.2.
8.12 Corollary. The disconnection cost is bounded by $n \operatorname{diam}(\widetilde{G})$.

Proof. We have by Lemma 5.2:

$$
\begin{aligned}
\frac{1}{m} \sum_{v \in V} \sum_{e \in E} \operatorname{rel}(e, v) & =\frac{1}{m} \sum_{v \in V} R(v) \leq \frac{1}{m} \sum_{v \in V}(n-1) \operatorname{diam}(\widetilde{G}) \\
& =\frac{n}{m}(n-1) \operatorname{diam}(\widetilde{G}) \leq n \operatorname{diam}(\widetilde{G})
\end{aligned}
$$

A bound on the diameter of the bridge tree holding for all NE will hence yield a bound on the price of anarchy. This is accomplished by the following Lemma.
8.13 Lemma. The bridge tree of a NE has its diameter bounded by $\frac{12 n+4}{n} \alpha=O(\alpha)$.

Proof. Let $G$ be a NE. Let $\widetilde{P}=\left(v_{0}, e_{1}, v_{1}, \ldots, e_{\ell}, v_{\ell}\right)$ be a path in the bridge tree $\widetilde{G}$ connecting two leaves $v_{0}$ and $v_{\ell}$. Let $\bar{\ell}:=\left\lceil\frac{\ell}{2}\right\rceil \geq 1$. Then at least one of the following is true (recall the convention on page 313 regarding vertex-counting in the bridge tree):

- At least $\left\lceil\frac{n}{2}\right\rceil$ vertices lie beyond $e_{\bar{\ell}}$ from the view of $v_{0}$.
- At least $\left\lceil\frac{n}{2}\right\rceil$ vertices lie beyond $e_{\bar{\ell}}$ from the view of $v_{\ell}$.

Let us assume the first; the other case can be treated alike. Let $v:=v_{0}$ and $w:=v_{\ell}$ and recall that we may treat vertices of the bridge tree $\widetilde{G}$ as single players with respect to building of new links. Then $e_{1}, \ldots, e_{\bar{\ell}}$ for $v$ have relevance at least $\left\lceil\frac{n}{2}\right\rceil$ each. So $\sum_{i=1}^{\bar{\ell}} \operatorname{rel}\left(e_{i}, v\right) \geq \bar{\ell} \frac{n}{2} \geq \frac{\ell}{2} \frac{n}{2}=\Omega(\ell n)$. By building $\{v, w\}$, player $v$ would have a benefit in disconnection cost of at least $\frac{1}{m+1} \frac{\ell n}{4} \geq \frac{1}{3 n+1} \frac{\ell n}{4}=\Omega(\ell)$, using the bound $m \leq 3 n$ from Corollary 8.11. Since the edge is not built, $\alpha$ is larger than this benefit, so $\ell \leq(3 n+1) \frac{4}{n} \alpha=O(\alpha)$.
8.14 Corollary. The disconnection cost in a NE is bounded by $(12 n+4) \alpha=O(n \alpha)$.

Proof. Follows from Corollary 8.12 and Lemma 8.13.
8.15 Theorem. The price of anarchy with a simple-minded adversary is bounded by $O(1)$.

Proof. The building cost and the disconnection cost in a NE are both $O(n \alpha)$ by Corollary 8.11 and 8.14. The theorem follows with Proposition 7.1, which states that the optimum social cost is $\Theta(n \alpha)$.

A closer look at Lemma 8.13 and its proof reveals that there exists a constant $c>0$ such that if $m=O(n)$ then there are players who can improve their disconnection cost by $c$ through the building of new links, as long as the graph contains bridges. It follows that for $\alpha<c$, all NE are bridgeless, i.e., they have disconnection cost 0 and the adversary cannot harm them. This does not rule out, however, that they may contain an unnecessarily high number of links, compared to an optimum. On the other hand, the ratio cannot be more than $O(1)$ by Theorem 8.15.
8.16 Remark. The constant in Theorem 8.15 is $15+o(1)$, more precisely it is bounded by $15+\frac{19}{n-1} \leq 18$.

Proof. Building cost of a NE is bounded by $3 n \alpha$ by Corollary 8.11. Disconnection cost of a NE is bounded by $(12 n+4) \alpha$ by Corollary 8.14. In total, social cost of a NE is bounded by $(15 n+4) \alpha$. Using Remark 7.2 with $c_{1}:=15$ and $c_{0}:=4$, and finally $n \geq 9$ proves the claim.

## 9. Smart Adversary

We consider an adversary that destroys an edge which separates a maximum number of vertex pairs. If there are several such edges, one of them is chosen uniformly at random. In other words, we replace the uniform probability distribution on the edges for one that is concentrated on the edges which cause maximum overall damage. Recall that $\operatorname{sep}(e)$ is the number of separated vertex pairs when edge $e$ is deleted. Let $\operatorname{sep}_{\max }:=\max _{e \in E} \operatorname{sep}(e)$ and $E_{\max }:=\left\{e \in E ; \operatorname{sep}(e)=\operatorname{sep}_{\max }\right\}$ and $m_{\max }:=\left|E_{\max }\right|$. These are the edges of which each causes a maximum number of separated vertex pairs when it is deleted. We call them the critical edges. The adversary chooses one of those uniformly at random. Clearly, this yields a symmetric adversary, and so disconnection cost is anonymous. We have the individual and social cost:

$$
\begin{array}{rlr}
C_{v}(S) & =\left|S_{v}\right| \alpha+\frac{1}{m_{\max }} \sum_{e \in E_{\max }} \operatorname{rel}(e, v) & \text { for } v \in V \\
C(S) & =m \alpha+\frac{1}{m_{\max }} \sum_{e \in E_{\max }} \operatorname{sep}_{\max }=m \alpha+\operatorname{sep}_{\max } &
\end{array}
$$

If $\operatorname{sep}_{\text {max }}=0$, then the graph is bridgeless and all edges are critical—however, their removal does not separate any vertex pairs. If $\operatorname{sep}_{\max }>0$, then there are one or more critical edges, and each of them is a bridge. Recall that if $e$ is a bridge, $\nu(e)$ denotes the number of vertices in the smaller component of $G-e$, or $\frac{n}{2}$ if both are of equal size. If $e$ is no bridge, then $\nu(e)=0$. We have $\operatorname{sep}(e)=2 \nu(e)(n-\nu(e))$ for all edges. The function $x \mapsto 2 x(n-x)$ is strictly increasing on $\left[0, \frac{n}{2}\right]$, so $\nu(e)=\nu\left(e^{\prime}\right)$ follows from $\operatorname{sep}(e)=\operatorname{sep}\left(e^{\prime}\right)$. Hence $\nu(e)=\nu\left(e^{\prime}\right)$ for all critical edges $e, e^{\prime} \in E_{\max }$.
9.1 Proposition. If $\operatorname{sep}_{\max }>0$ and if there are more than one critical edges, they form a subgraph that is a star in the bridge tree $\widetilde{G}$.

Proof. Let $\operatorname{sep}_{\max }>0$. For any two distinct bridges $e$ and $e^{\prime}$, one component of $G-e$ is strictly contained in one component of $G-e^{\prime}$. Therefore, with multiple critical edges, $\nu(e)<\frac{n}{2}$ for all $e \in E_{\max }$, and so also for all other bridges (since they have smaller $\nu(\cdot)$ value). In other words, there is always a small and a large component of $G-e$, with $e$ being a bridge.

Let $P=\left(v_{0}, e_{1}, v_{1}, \ldots, v_{\ell-1}, e_{\ell}, v_{\ell}\right)$ be a path in the bridge tree $\widetilde{G}$ with $e_{1}$ and $e_{\ell}$ being distinct critical edges. First assume that $v_{\ell}$ is in the larger component of $G-e_{\ell}$. Then $v_{0}$ is in the smaller component of $G-e_{1}$. Then the smaller component of $G-e_{2}$ cannot contain $v_{0}$, since otherwise $\nu\left(e_{1}\right)<\nu\left(e_{2}\right)$, and $e_{1}$ would not be critical. So the component of $G-e_{2}$ containing $v_{0}$ is the larger one, and then the same holds for the component of $G-e_{\ell}$ containing $v_{0}$. This contradicts that $v_{\ell}$ is in the larger component of $G-e_{\ell}$. We can carry out the same argument with $v_{0}$ and $e_{1}$. Summarizing, now we know that the smaller component of $G-e_{1}$ is located "before" $P$ and that the smaller component of $G-e_{\ell}$ is located "beyond" $P$.

If $\ell \geq 3$, then there is an edge $f$ between $e_{1}$ and $e_{\ell}$ on $P$. The smaller component of $G-f$ strictly contains either the smaller component of $G-e_{1}$ or $G-e_{\ell}$. Since $\nu\left(e_{1}\right)=\nu\left(e_{2}\right)$, we have thus in particular, $\nu(f)>\nu\left(e_{1}\right)$, a contradiction that $e_{1}$ is critical. Hence there is no such edge $f$, and so $\ell=2$. Since this holds for all pairs $\left(e_{1}, e_{\ell}\right)$ of critical edges, the set of all critical edges forms a star (in the bridge tree).

The smart adversary admits a new NE topology:
9.2 Proposition. If $\alpha \geq \frac{n}{2}$, then a path with all edges pointing to the nearest end (in case of even $n$, the middle edge having arbitrary orientation) is a NE with social $\operatorname{cost} \Theta\left(n \alpha+n^{2}\right)$. If $\alpha>\frac{n}{2}$, it is a MaxNE.

Proof. The social cost of the path is $(n-1) \alpha+\left\lfloor\frac{n}{2}\right\rfloor\left\lceil\frac{n}{2}\right\rceil=\Theta\left(n \alpha+n^{2}\right)$. The adversary removes the one middle edge if $n$ is even, or one of the two middle edges if $n$ is odd. The disconnection cost for each player is $\frac{n}{2} \leq \alpha$ if $n$ is even and at most $\frac{1}{2}\left(\left\lfloor\frac{n}{2}\right\rfloor+\left\lceil\frac{n}{2}\right\rceil\right)=\frac{n}{2} \leq \alpha$ if $n$ is odd. Hence, there is no incentive for any player to build more edges than she currently owns, even after exchanging the currently built edges for others.

Now consider that a player $v$ sells one (or two) of her edges and buys one (or two) different ones instead. First consider that one edge is exchanged. Since all edges point outwards, the part of the path that becomes disconnected from $v$ does not contain the critical edge(s). So, after reconnecting it with $v$ via a new edge, there are as many vertices on both sides of the formerly critical edge(s) as before the exchange. No separation value increases. Hence the formerly critical edges remain critical. They also maintain their relevance for $v$. With the same argument, if the exchanged edge itself was critical, the new one will be critical as well, also with the same relevance for $v$.

When two edges are exchanged, $v$ is the center vertex, and in particular all critical edge(s) are among the exchanged ones, see Figure 3. This again means that the disconnected parts do not contain critical edges, and so the exchange cannot change that each of the two edges has $\left\lfloor\frac{n}{2}\right\rfloor$ vertices on the one and $\left\lceil\frac{n}{2}\right\rceil$ vertices on the other side, so they remain critical. Also, their relevance for $v$ does not change. The MaxNE property is clear by Remark 2.1.

Figure 3. NE if $\alpha \geq \frac{n}{2}$. Critical edges are drawn dashed. Disconnection cost for the center vertex is $\left\lfloor\frac{n}{2}\right\rfloor$ in both cases, which is 4 here.
(a) A path with all edges pointing to the nearest end.

(b) Center vertex cannot improve by exchanging her edges.


### 9.1. Bounding the Price of Anarchy

The proof of the following is even easier than previously:
9.3 Remark. A NE is chord-free.

Proof. Removing a chord does not change the relevance of any edge, nor does it change sep max $_{\text {max }}$, hence it does not change $E_{\text {max }}$. Selling a chord so is always beneficial.

With Proposition 8.10, it follows immediately:
9.4 Corollary. A NE has at most $3 n=O(n)$ edges.

We are again left with bounding the disconnection cost of NE. This requires some effort and is accomplished in the following remark and two lemmas.
9.5 Remark. If there are $k \geq 2$ critical edges, say $E_{\max }=\left\{e_{1}, \ldots, e_{k}\right\}$, and $e_{1}$ is put on a cycle by an additional edge, but not $e_{2}, \ldots, e_{k}$, then the new critical edges are $e_{2}, \ldots, e_{k}$. If $k \geq 3$ and the additional edge puts $e_{1}$ and $e_{2}$ on a cycle, but not $e_{3}, \ldots, e_{k}$, then the new critical edges are $e_{3}, \ldots, e_{k}$.

Proof. An additional edge $e$ only changes the $\nu(\cdot)$ value of those edges which are put on a cycle by $e$, namely it reduces them to 0 . Hence, none of the edges in $\left\{e_{2}, \ldots, e_{k}\right\}$ (or $\left\{e_{3}, \ldots, e_{k}\right\}$ ) becomes less attractive for the adversary when $e$ is added. Also no other edge becomes more attractive by the addition of $e$, since no $\nu(\cdot)$ value increases.
9.6 Lemma. Let $\alpha \leq$ cn for a constant $c>0$ and fix a NE with $m_{\max } \geq 3$. Then we have $\operatorname{sep}_{\max } \leq$ $2(1+9 c) n \alpha=O(n \alpha)$.

Proof. Fix two critical edges $e_{1}$ and $e_{2}$, and set $n_{0}:=\nu\left(e_{1}\right)$. For each $i \in\{1,2\}$ fix a player $v_{i}$ in the smaller component of $G-e_{i}$. Then for each $i \in\{1,2\}$ we have $\operatorname{rel}\left(e_{i}, v_{i}\right)=n-n_{0}$ and $\operatorname{rel}\left(e, v_{i}\right)=n_{0}$ for all critical edges $e \neq e_{i}$; recall that all critical edges have the same $\nu(\cdot)$ value. Building $\left\{v_{1}, v_{2}\right\}$
puts $e_{1}$ and $e_{2}$ on a cycle and leaves the other $m_{\max }-2$ critical edges critical by Remark 9.5. For each $i \in\{1,2\}$ [46], player $v_{i}$ has her disconnection cost decreased by:

$$
\begin{aligned}
& \frac{1}{m_{\max }} \sum_{e \in E_{\max }} \operatorname{rel}\left(e, v_{i}\right)-\frac{1}{m_{\max }-2} \sum_{\substack{\left.e \in E_{\max } \\
e \notin e_{1}, e_{2}\right\}}} \operatorname{rel}\left(e, v_{i}\right) \\
& =\frac{1}{m_{\max }}\left(\left(m_{\max }-1\right) n_{0}+n-n_{0}\right)-\frac{1}{m_{\max }-2}\left(m_{\max }-2\right) n_{0} \\
& =\frac{1}{m_{\max }}\left(\left(m_{\max }-2\right) n_{0}+n\right)-n_{0}=\frac{1}{m_{\max }}\left(n-2 n_{0}\right)
\end{aligned}
$$

Since we are in a NE, this is at most $\alpha$. Since $n \geq m_{\max } n_{0}$, we have $n-2 n_{0} \geq\left(m_{\max }-2\right) n_{0}$, and so it follows $\alpha \geq\left(1-\frac{2}{m_{\max }}\right) n_{0} \geq \frac{1}{3} n_{0}$. Moreover, it follows $m_{\max } \alpha+n_{0} \geq n-n_{0}$. With these two inequalities at hand, we can bound $\operatorname{sep}_{\max }$. We have

$$
\begin{aligned}
\operatorname{sep}_{\max } & =2 n_{0}\left(n-n_{0}\right) \leq 2 n_{0}\left(m_{\max } \alpha+n_{0}\right) \leq 2\left(n_{0} m_{\max } \alpha+9 \alpha^{2}\right) \\
& \leq 2\left(n \alpha+9 \alpha^{2}\right) \stackrel{\alpha \leq c n}{\leq} 2(n \alpha+9 c n \alpha)=2(1+9 c) n \alpha
\end{aligned}
$$

9.7 Lemma. Fix a NE with $m_{\max } \in\{1,2\}$. Then

- we have $\operatorname{sep}_{\max } \leq 4 n \alpha=O(n \alpha)$
- or we have $\alpha \geq \frac{1}{6} n=\Omega(n)$.

Proof. First we consider the case $m_{\max }=2$. A player can make the two critical edges part of a cycle by building an additional edge. The difficulty lies in that new critical edges, with a smaller separation value, can emerge. We will have to put some more effort into estimating the improvement in disconnection cost that a player is able to achieve by building another edge. Consider the bridge tree. There are two subtrees $T_{1}$ and $T_{2}$ that are connected to the rest by the two critical edges $e_{1}$ and $e_{2}$, respectively. They both have $n_{0}:=\nu\left(e_{1}\right)=\nu\left(e_{2}\right)$ vertices. There may be more subtrees $T_{3}, \ldots, T_{N}$ connected by $e_{3}, \ldots, e_{N}$ to the center vertex. To streamline notation, we often write $T_{k}$ instead of $V\left(T_{k}\right), k \in[N]$, when we refer to the set of vertices of a tree. Figure 4 on the next page depicts the situation. Figure 5 shows how a new edge would put $e_{1}$ and $e_{2}$ on a cycle.

First assume that we can arrange $v_{1} \in T_{1}$ and $v_{2} \in T_{2}$ such that after building $\left\{v_{1}, v_{2}\right\}$, there are no critical edges in $T_{1}$ nor in $T_{2}$. If there are no subtrees except $T_{1}$ and $T_{2}$, i.e., if $N=2$, this means that we can make the graph bridgeless by the additional edge. The improvement in disconnection cost for $v_{1}$ (and also for $v_{2}$ ) of building $\left\{v_{1}, v_{2}\right\}$ is hence their original disconnection cost, i.e., $\frac{1}{2}\left(n-n_{0}+n_{0}\right)=\frac{1}{2} n \geq \frac{1}{6} n$. If $N \geq 3$, then critical edges emerge in one or more of the $T_{3}+e_{3}, \ldots, T_{N}+e_{N}$ after building. Fix $k \in\{3, \ldots, N\}$. Since $e_{k}$ is not critical without the new edge, we have $\left|T_{k}\right|<n_{0}$ or $\left|T_{k}\right| \geq\left\lceil\frac{n}{2}\right\rceil$. The latter can be excluded, since it would imply that the smaller (or equally sized) component of $G-e_{k}$ includes $T_{1}$ and the center vertex, and so $\nu\left(e_{k}\right)>n_{0}=\nu\left(e_{1}\right)$, in which case $e_{1}$ would not be critical. Moreover, we have $\left|T_{k}\right| \leq n-2 n_{0}<n-2\left|T_{k}\right|$, so $\left|T_{k}\right|<\frac{1}{3} n$. For a player in $T_{1}$ (or $T_{2}$ ), a critical edge in $T_{k}+e_{k}$ can have relevance at most $\left|T_{k}\right|$ and so no more than $\frac{1}{3} n$. The improvement in disconnection $\operatorname{cost}$ for $v_{1}$ (and also for $v_{2}$ ) gained by building $\left\{v_{1}, v_{2}\right\}$ is hence at least the original disconnection cost minus $\frac{1}{3} n$, i.e., $\frac{1}{2}\left(n-n_{0}+n_{0}\right)-\frac{1}{3} n=\left(\frac{1}{2}-\frac{1}{3}\right) n=\frac{1}{6} n$.

Figure 4. Schematic view of the bridge tree with two critical edges $e_{1}$ and $e_{2}$, drawn dashed. Subtrees are represented by triangles.


Figure 5. How $e_{1}$ and $e_{2}$ are put on a cycle by a new edge $\left\{v_{1}, v_{2}\right\}$. Paths that are part of the new cycle and located inside $T_{1}$ and $T_{2}$ are depicted as zigzag paths. New critical edges can emerge, e.g., $e_{3}$ can become critical.


Now consider that for all choices of $v_{1} \in T_{1}$ and $v_{2} \in T_{2}$, building $\left\{v_{1}, v_{2}\right\}$ induces a critical edge in at least one of $T_{1}$ or $T_{2}$. For each $i \in\{1,2\}$ we can do the following. Let $u_{i}$ be the vertex where $T_{i}$ is connected to the rest of the graph and consider $T_{i}$ being rooted at $u_{i}$. Let $P_{i}$ be a path starting at $u_{i}$ and ending at one of the leaves of $T_{i}$, say $w_{i}$, such that the path always descends into a subtree that has a maximum number of vertices, as shown in Figure 6 on the following page. If we choose $v_{i}:=w_{i}$, $i=1,2$, then each $P_{i}$ does not contain a critical edge when we build $\left\{v_{1}, v_{2}\right\}$, since these paths then both are located on a cycle. However, by assumption, there is a critical edge $f$ in, say $T_{1}$. By construction of $P_{1}$, we have $\nu(f) \leq \frac{n_{0}}{2}$. So, player $v_{1}$ (and also $v_{2}$ ) can reduce her disconnection cost to no more than $\frac{n_{0}}{2}$. It follows that the improvement in disconnection cost is at least $\frac{1}{2}\left(n-n_{0}+n_{0}\right)-\frac{1}{2} n_{0}=\frac{1}{2}\left(n-n_{0}\right)$, which is at most $\alpha$, since we are in a NE. It follows $\operatorname{sep}_{\max }=2 n_{0}\left(n-n_{0}\right) \leq 2 n_{0} \cdot 2 \alpha \leq 4 n \alpha$.

Figure 6. Detailed view of $T_{i}$ for one $i \in\{1,2\}$. Path $P_{i}$ is highlighted. Recall that vertices of the bridge tree are counted according to the size of the respective BCCs. Here, in this example, we assume that each vertex counts 1 . The path is drawn accordingly, i.e., always descending into a subtree with a maximum number of vertices.


The case of $m_{\max }=1$ can be treated similarly. Let $e_{1}$ be the critical edge and $T_{1}$ the subtree with $n_{0}:=\nu\left(e_{1}\right)$ vertices. There are zero or more additional subtrees, say $T_{2}, \ldots, T_{N}$. If there are zero such trees, define $T_{2}:=\widetilde{G}-T_{1}$, which consists of just one vertex in the bridge tree then (but can consist of multiple vertices in $G$ ). Let the ordering be such that $\left|T_{2}\right| \geq\left|T_{k}\right|$ for all $k \in\{3, \ldots, N\}$. Then we argue similar to before with $T_{1}$ and $T_{2}$ in the roles of the former subtrees of the same name. Assume first that we can find $v_{1} \in T_{1}$ and $v_{2} \in T_{2}$ such that building $\left\{v_{1}, v_{2}\right\}$ does not induce any critical edges in $T_{1}$ nor $T_{2}$. If $N \leq 2$, then we can make the graph bridgeless and this means an improvement for $v_{1}$ of at least $n-n_{0}$, and so $\operatorname{sep}_{\max }=2 n_{0}\left(n-n_{0}\right) \leq 2 n_{0} \alpha \leq 4 n \alpha$. If $N \geq 3$, then fix $k \in\{3, \ldots, N\}$. We have $\left|T_{k}\right|<n_{0}$. Moreover, we have $\left|T_{k}\right| \leq n-\left(\left|T_{2}\right|+n_{0}\right) \leq n-2\left|T_{k}\right|$, and so $\left|T_{k}\right| \leq \frac{1}{3} n$. Then building $\left\{v_{1}, v_{2}\right\}$ reduces the disconnection cost of $v_{1}$ to no more than $\frac{1}{3} n$. This means an improvement for $v_{1}$ of at least $n-n_{0}-\frac{1}{3} n \geq \frac{2}{3} n-\frac{1}{2} n=\frac{1}{6} n$.

If each choice of $v_{1}$ and $v_{2}$ induces a critical edge in $T_{1}$ or $T_{2}$, we can, as before, show that by a careful choice of these vertices, building $\left\{v_{1}, v_{2}\right\}$ reduces the disconnection cost for $v_{1}$ (and $v_{2}$ ) to at most $\frac{n_{0}}{2}$. Player $v_{1}$ originally has disconnection cost $n-n_{0} \geq n_{0}$, so she experiences an improvement of at least $\frac{n_{0}}{2}$. (Player $v_{2}$ originally has disconnection cost $n_{0}$, so she as well experiences an improvement of at least $\frac{n_{0}}{2}$.) It follows $n_{0} \leq 2 \alpha$ and so $\operatorname{sep}_{\max }=2 n_{0}\left(n-n_{0}\right) \leq 4 \alpha n$.
9.8 Theorem. The price of anarchy with a smart adversary is $O(1)$.

Proof. Let $c>0$ be the constant from the " $\alpha=\Omega(n)$ " statement of Lemma 9.7, e.g., we may choose $c:=\frac{1}{6}$. Consider first $\alpha<c n$. We use $c$ as the constant in the premise in Lemma 9.6. So if $m_{\max } \geq 3$, then Lemma 9.6 gives $\operatorname{sep}_{\max }=O(n \alpha)$. Otherwise, if $m_{\max } \in\{1,2\}$, Lemma 9.7 gives the same, since $\alpha \geq c n$ is ruled out. Since $\operatorname{sep}_{\max }$ is the total disconnection cost, it so has a ratio of $O(1)$ to the optimum. Remark 9.3 ensures that the same holds for the building cost.

If $\alpha \geq c n$, then this and Remark 9.3 allow us to invoke Corollary 6.2(ii).
9.9 Remark. The constant in Theorem 9.8 is $9+o(1)$, more precisely it is bounded by $9+\frac{9}{n-1} \leq 11$.

Proof. We proceed as in the proof of the theorem, but do more detailed calculations. Let $c:=\frac{1}{6}$. We start again with $\alpha<c n$. If $m_{\max } \geq 3$, then Lemma 9.6 gives $\operatorname{sep}_{\max } \leq 2(1+9 c) n \alpha=5 n \alpha$. Otherwise, if $m_{\text {max }} \in\{1,2\}$, Lemma 9.7 gives sep max $\leq 4 n \alpha$.

Next consider $\alpha \geq c n$, i.e., $\frac{\alpha}{c} \geq n$. Using the trivial bound $n^{2}$ on the disconnection cost yields the $\frac{\alpha}{c} n=6 n \alpha$ bound on it.

In all cases disconnection cost is bounded by $6 n \alpha$. Building cost is bounded by $3 n \alpha$ by Remark 9.3 . Finally invoking Remark 7.2 with $c_{1}:=9$ and $c_{0}:=0$ yields a bound on the price of anarchy of $9+\frac{9}{n-1} \leq 11$, using $n \geq 9$.

## Bibliographic Information

This work is based on parts of Chapter 5 of my dissertation [11]. It has also been presented in form of a brief announcement at PODC 2010 [47].

## Notes Added in Proof

The argument in the beginning of the proof of Proposition 8.8 is flawed, since $C$ could share more vertices with $C^{\prime}$ than just $v$ and $w$. A correct argument is as follows.

We show that the bridge tree does not change by removal of $e$. Assume for contradiction that there exists an edge $e^{\prime}$ which is a bridge in $G^{\prime}:=G-e$, but which is no bridge in $G$. Then $G^{\prime}-e^{\prime}$ consists of two components $G_{1}$ and $G_{2}$. Since $e^{\prime}$ is no bridge in $G$, the edge $e$ connects $G_{1}$ and $G_{2}$. But then, due to the existence of $C$, in addition to $e$ there are two more edges between $G_{1}$ and $G_{2}$. Hence removal of the single edge $e^{\prime}$ from $G^{\prime}$ cannot disconnect $G_{1}$ from $G_{2}$.

## Acknowledgements

I thank the German Research Foundation (Deutsche Forschungsgemeinschaft (DFG)) for financial support through Priority Program 1307 "Algorithm Engineering" (Grant SR7/12-2). I thank the two anonymous referees for their very helpful and constructive comments. I thank everyone at Games for their support and the open access policy.

## References

1. We use "player" and "vertex" synonymously.
2. We use "link" and "edge" synonymously.
3. Optimal networks are also called "efficient" in the literature.
4. Koutsoupias, E.; Papadimitriou, C.H. Worst-case equilibria. In Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science, Trier, Germany, March 1999; pp. 404-413.
5. Papadimitriou, C.H. Algorithms, Games, and the Internet. In Proceedings of the 33rd Annual ACM Symposium on Theory of Computing, Crete, Greece, July 2001; pp. 749-753.
6. Fabrikant, A.; Luthra, A.; Maneva, E.; Papadimitriou, C.H.; Shenker, S. On a network creation game. In Proceedings of the 22nd Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing, Boston, MA, USA, July 2003; pp. 347-351.
7. The "essential" term was used in [9] in the context of ULF. In [10], the concept is called "non-superfluous" in the context of BLF. In [11], "clean" was used the way we use "essential" here.
8. This means if $S_{v w}=1$ and $S_{w v}=1$, we arbitrarily switch to either $S-(v, w)$ or to $S-(w, v)$, so we drop exactly one of the two requests, not both.
9. Bala, V.; Goyal, S. A strategic analysis of network reliability. Rev. Econ. Des. 2000, 5, 205-228.
10. Chakrabarti, S.; Gilles, R.P. Network potentials. Rev. Econ. Des. 11, 13-52.
11. Kliemann, L. Networks and Distributed Operation: The Price of Anarchy in Non-Atomic Routing and Network Formation. Ph.D. thesis, Institut für Informatik, Christian-Albrechts-Universität zu Kiel, Kiel, Germany, 2010.
12. Jackson, M.O. A Survey of Models of Network Formation: Stability and Efficiency. In Group Formation in Economics; Networks, Clubs and Coalitions; Demange, G., Wooders, M., Eds.; Cambridge University Press: Cambridge, UK, 2004; Chapter 1.
13. Chun, B.G.; Fonseca, R.; Stoica, I.; Kubiatowicz, J. Characterizing selfishly constructed overlay routing networks. In Proceedings of the 23rd IEEE Conference on Computer Communications, Hong Kong, China, March 2004.
14. Jackson, M.O.; Wolinsky, A. A strategic model of social and economic networks. J. Econ. Theor. 1996, 71, 44-74.
15. Baumann, N.; Stiller, S. The price of anarchy of a network creation game with exponential payoff. In Proceedings of the 1st Annual ACM-SIAM Symposium on Algorithmic Game Theory, Paderborn, Germany, April-May 2008; pp. 218-229.
16. Sarangi, S.; Haller, H. Nash Networks with Heterogeneous Agents. Departmental Working Papers 2003-06. Department of Economics, Louisiana State University: Baton Rouge, LA, USA, 2003.
17. Haller, H.; Sarangi, S. Nash networks with heterogeneous links. Math. Soc. Sci. 2005, 50, 181-201.
18. Bala, V.; Goyal, S. A noncooperative model of network formation. Econometrica 2000, 68, 1181-1229.
19. Myerson, R.B. Game Theory: Analysis of Conflict; Harvard University Press: Cambridge, MA, USA, 2002.
20. Watts, A. A dynamic model of network formation. Games Econ. Behav. 2001, 34, 331-341.
21. Calvó-Armengol, A.; İlkiliç, R. Pairwise Stability and Nash Equilibria in Network Formation. Technical Report 34.2005. Fondazione Eni Enrico Mattei: Milano, Italy, 2005.
22. Corbo, J.; Parkes, D. The price of selfish behavior in bilateral network formation. In Proceedings of the 24th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing, Las Vegas, NV, USA, July 2005; pp. 99-107.
23. Bloch, F.; Jackson, M.O. The formation of networks with transfers among players. J. Econ. Theor. 2007, 133, 83-110.
24. Bloch, F.; Jackson, M.O. Definitions of equilibrium in network formation games. Int. J. Game Theor. 2006, 34, 305-318.
25. Anshelevich, E.; Dasgupta, A.; Tardos, É.; Wexler, T. Near-optimal network design with selfish agents. Theor. Comput. 2008, 4, 77-109.
26. Anshelevich, E.; Dasgupta, A.; Kleinberg, J.; Tardos, É.; Wexler, T.; Roughgarden, T. The price of stability for network design with fair cost allocation. In Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science, Rome, Italy, October 2004; pp. 295-304.
27. Christin, N.; Chuang, J. On the cost of participating in a peer-to-peer network. In Proceedings of Revised Selected Papers of the 3rd International Workshop on Peer-to-Peer Systems, La Jolla, CA, USA, February 2004; pp. 22-32.
28. Christin, N.; Grossklags, J.; Chuang, J. Near rationality and competitive equilibria in networked systems. In Proceedings of the ACM SIGCOMM Workshop on Practice and Theory of Incentives in Networked Systems, Portland, OR, USA, August-September 2004.
29. Johari, R.; Mannor, S.; Tsitsiklis, J.N. A contract-based model for directed network formation. Games Econ. Behav. 2006, 56, 201-224.
30. Albers, S.; Eilts, S.; Even-Dar, E.; Mansour, Y.; Roditty, L. On Nash Equilibria for a network creation game. In Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms, Miami, FL, USA, January 2006; pp. 89-98.
31. Demaine, E.D.; Hajiaghayi, M.T.; Mahini, H.; Zadimoghaddam, M. The price of anarchy in network creation games. In Proceedings of the 26th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing, Portland, OR, USA, August 2007; pp. 292-298.
32. Lin, H. On the price of anarchy of a network creation game. Unpublished Manuscript, 4 pages, 2003.
33. Mihalák, M.; Schlegel, J.C. The price of anarchy in network creation games is (mostly) constant. In Proceedings of the 3rd Annual ACM-SIAM Symposium on Algorithmic Game Theory, Athens, Greece, October 2010; pp. 276-287.
34. Moscibroda, T.; Schmid, S.; Wattenhofer, R. On the topologies formed by selfish peers. In Proceedings of the 25th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing, Denver, CO, USA, July 2006; pp. 133-142.
35. Halevi, Y.; Mansour, Y. A network creation game with nonuniform interests. In Proceedings of the 3rd International Workshop on Internet and Network Economics, San Diego, CA, USA, December 2007; pp. 287-292.
36. Brandes, U.; Hoefer, M.; Nick, B. Network creation games with disconnected equilibria. In Proceedings of the 4th International Workshop on Internet and Network Economics, Shanghai, China, December 2008; pp. 394-401.
37. Laoutaris, N.; Poplawski, L.J.; Rajaraman, R.; Sundaram, R.; Teng, S.H. Bounded budget connection (BBC) games or how to make friends and influence people, on a budget. In Proceedings of the 27th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing, Toronto, ON, Canada, August 2008; pp. 165-174.
38. Demaine, E.D.; Zadimoghaddam, M. Constant price of anarchy in network creation games via public service advertising. In Proceedings of the 7th International Workshop on Algorithms and Models for the Web-Graph, Stanford, CA, USA, December 2010; pp. 122-131.
39. Balcan, M.F.; Blum, A.; Mansour, Y. Improved equilibria via public service advertising. In Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms, New York, NY, USA, January 2009; pp. 728-737.
40. Demaine, E.D.; Hajiaghayi, M.T.; Mahini, H.; Zadimoghaddam, M. The price of anarchy in cooperative network creation games. In Proceedings of the 26th Annual Symposium on Theoretical Aspects of Computer Science, Freiburg, Germany, February 2009; pp. 301-312.
41. However, large parts of our analysis will be restricted to two specific cases: One in which the adversary picks a link uniformly at random (simple-minded adversary) and another in which he picks a link that causes maximum overall damage (smart adversary).
42. Haller and Sarangi also briefly discussed failure probabilities depending on the final graph. They considered an example where for non-increasing functions $f_{v}(\cdot)$ and parameters $P_{v w}$ the probabilities are defined $p_{v w}(S):=f_{v}\left(\operatorname{deg}_{G(S)}(v)\right) f_{w}\left(\operatorname{deg}_{G(S)}(w)\right) P_{v w}$ if $v$ and $w$ have a link between them, and 0 otherwise.
43. Diestel, R. Graph Theory, 3rd ed.; Springer-Verlag: Heidelberg, Germany, 2005.
44. See [43] (p. 56) for the definition of the block graph.
45. The restriction to $n \geq 9$ stems from the need for existence guarantees for equilibria. This is established in Section 7.
46. It would suffice to restrict to $i=1$ or $i=2$. However, here and in the proof of the following Lemma 9.7, we point out arguments that are symmetric in the sense that both endpoints would like to build the edge. This is interesting in order to realize which parts of the proof could be carried over to BLF.
47. Kliemann, L. Brief Announcement: The price of anarchy for distributed network formation in an adversary model. In Proceedings of the 29th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing, Zurich, Switzerland, July 2010; pp. 229-230.
(c) 2011 by the author; licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution license (http://creativecommons.org/licenses/by/3.0/.)
