Tjalling C. Koopmans Research Institute

Discussion Paper Series nr: 10-13

Strategic Network Disruption and Defense

Britta Hoyer
Kris De Jaegher

Universiteit Utrecht

Utrecht School of Economics

Tjalling C. Koopmans Research Institute
Utrecht School of Economics
Utrecht University

Janskerkhof 12
3512 BL Utrecht
The Netherlands
telephone +31 30 253 9800
fax +31 30 253 7373
website www.koopmansinstitute.uu.nl

The Tjalling C. Koopmans Institute is the research institute and research school of Utrecht School of Economics.
It was founded in 2003, and named after Professor Tjalling C. Koopmans, Dutch-born Nobel Prize laureate in economics of 1975.

In the discussion papers series the Koopmans Institute publishes results of ongoing research for early dissemination of research results, and to enhance discussion with colleagues.

Please send any comments and suggestions on the Koopmans institute, or this series to J.M.vanDort@uu.nl

ontwerp voorblad: WRIK Utrecht

How to reach the authors

Please direct all correspondence to the first author.

Britta Hoyer
Kris De Jaegher
Utrecht University
Utrecht School of Economics
Janskerkhof 12
3512 BL Utrecht
The Netherlands.
E-mail: b.hoyer@uu.nl
k.dejaegher@uu.nl

This paper can be downloaded at: http:// www.uu.nl/rebo/economie/discussionpapers

Utrecht School of Economics
Tjalling C. Koopmans Research Institute
Discussion Paper Series 10-13

Strategic Network Disruption and Defense

Britta Hoyer
Kris De Jaegher

Utrecht School of Economics
Utrecht University

June 2010

Abstract
Networks are one of the essential building blocks of society. Not only do firms cooperate in R&D networks, but firms themselves may be seen as networks of information-exchanging workers. Social movements increasingly make use of networks to exchange information, just as on the negative side criminal and terrorist networks use them. However, the literature on networks has mainly focused on the cooperative side of networks and has so far neglected the competition side of networks. Networks themselves may face competition from actors with opposing interests to theirs. Several R&D networks may compete with one another. The firm as a network of employees obviously faces competition. In particular, given the importance of connectivity for networks, competing networks may try to disrupt each other, by trying to convince key players in competing networks to defect, or to stop sponsoring key links (strategic network disruption). In response, networks that face competition will adapt their structure, and will avoid vulnerable network structures. Such network competition is what our paper is concerned with.

Keywords: Strategic Network Disruption, Strategic Network Design, Non-cooperative Network Games

JEL classification: C72, D85

1. Introduction

In The Strategy of Conflict, Schelling (1960) criticizes that economics, and game theory in particular, puts too much emphasis on games with conflict of interests and too little emphasis on games with at least partially common interests. A large and recent network literature in economics (for an overview, see Goyal (2009)[1]) does put emphasis on common interests. For instance, one application of the models of this literature lies in the formation of cooperative networks of firms, which e.g. exchange R&D information. Given individual firms' incentives on obtaining the most possible information, the question is then how they connect to existing networks, with a further question of what network structures arise from this individual behavior in equilibrium. The starting point of our paper is that this approach perhaps puts too much emphasis on common interests. Staying with the example of firm networks, there may be several competing firm networks, rather than one network, and such competition in turn may have important consequences for the structure of networks. In particular, given the importance of connectivity for networks, competing networks may try to disrupt each other, by trying to convince key players in competing networks to defect, or to stop sponsoring key links (strategic network disruption). In response, networks that face competition will adapt their structure, and will avoid having vulnerable network structures.

To capture this intuition, our paper models a sequential zero-sum game between a network designer and a network disruptor. We model network structure as being determined by a network designer, because in the first instance we want to gain insight into what is an efficient defensive strategy for the network as a whole. The same applies to the disruptive strategy of the competing network, so that we can simply model the game as being played between a network designer and a network disruptor. At stage 1 of the game, the network designer uses costly links to connect nodes. At stage 2 of the game, in two different versions of our model, the network disruptor either has the opportunity to delete links from the designer's network (link deletion), or nodes (node deletion), where deleting is costly to the disruptor. The network designer's preferences are modeled such that he is better off, the larger the largest component in the post-disruption network is; the disruptor is better off, the smaller the largest post-disruption component is.

We start by treating several recent papers from the game-theoretic and economic literature, which are related to our work. In Ballester et al. (2006), the network is exogenously given, and a game with strategic substitutes or complements is played on the network. An example of the game played is coordinating criminal activity. Because the network is exogenously given, the focus is only on optimal network disruption, and not on network defense as in our paper. In particular, the disruptor tries to find the "key player" which is shown to be the node with the highest degree of Bonacich centrality (a centrality measure used in social network analysis). Our focus is on network defense as well as on network disruption. The optimal network defense strategy in our game constitutes at avoiding that there is any key player or key link. In Hong (2009) terrorists try to carry an explosive through an exogenously given transport network, modeled as a directed flow network. Security services try to stop the explosive from reaching its destination by shutting down a minimal number of links in the flow network. We instead focus on undirected networks, and in our model network defense consists of adding links, not deleting them.

Enders and Su (2007), Enders and Jindapon (2010), and Baccara and Bar-Isaac (2008) take a game-theoretic approach similar to ours, where more ties allow more information to be produced in the network. However, the network adversary's purpose is to learn as much of this information as possible, whereas the purpose of the network itself is to keep this information as secret as possible from the network adversary. In our model, on the contrary, the network adversary tries to block to the maximal extent the sharing of information in the network. Goyal and Vigier (2009)'s model has a similar focus on network defense against a disruptor as in our model. However, in their approach, network defense takes the form of the creation of a "firewall" around key nodes; in our approach on the contrary, a network is defended by adding links that would otherwise be redundant. A similar approach to the one of Goyal and Vigier is found in Hong (2008). In Billand et al. (2009) model of corporate espionage, firms do not only compete by price competition, but by finding out each other's information. Espionage is modeled as unilateral link formation. In Bier et al. (2006), a defender needs to decide how to allocate defensive efforts over two targets for attack. Just as is

[1]For a good introduction on the literature and research on social and economic networks see Jackson (2008).

the case in our model, it may be optimal to defend the locations in an asymmetric way, leaving weak spots. Yet, their results are not treated in the context of the design of a network.

In non game-theoretic/non-economic literature related to our paper, the following papers providing related intuitions are worth mentioning. An influential paper is Albert et al. (2000), which treats a stochastic network generation process that yields networks with properties that are often observed in real-world networks (namely preferential attachment). It is shown that these networks are robust against random attacks, but vulnerable to targeted attacks.[2] By the same intuition, stars do badly in our analysis under node deletion. In the context of vulnerability of road networks, Taylor et al. (2006) treat the adding of links as a mechanism of network protection. They are interested, however, in the effect that this has on several vulnerability measures, whereas our focus is on network structure. The non game-theoretic paper most related to our work is Dekker and Colbert (2004), who study the node (link) connectivity of networks, which is the smallest number of nodes (links) which upon deletion results in a disconnected graph. They say that a graph is optimally connected if its node respectively link connectivity is equal to the minimal degree in the network. Finally, they show that networks that have certain symmetry properties are optimally connected. While the authors do not consider linking to be costly, and do not model strategic disruption, their treatment can be seen as corresponding to our treatment of low linking costs. Still, we show that not all best-response defensive networks are symmetric.

In Section 2, we start by listing some potential applications of the model. Section 3 treats the model, and some of the basic results. In Section In Sections 4, 5, and 6, we treat the differing results of the model according to whether the designer's costs of linking nodes are low, high, or intermediate. We end with a discussion of the results in Section 7.

2. Applications

As already suggested by the game-theoretic literature above, the most straightforward applications of our model lie in the military and security field (for a comprehensive overview with references outside the field of game theory, see Lipsey (2006)). The potentially disrupted network may be a physical network or a network of people. An army may try to use minimal means to incapacitate to the maximal extent the transport (e.g. bombing bridges) or communication network of an enemy country (e.g. by jamming). [3] Terrorist groups or hackers with limited means may similarly try to disrupt transport, communication or computer networks to the maximal extent. A dictatorial regime may try to disrupt a dissident or resistance movement to the maximal extent, police forces may try to do the same with terrorist or criminal networks, and with online peer-to-peer networks sharing illegal material. Each time, the potentially disrupted network may adapt its structure to be better protected against disruption.[4]

A recent paper (Posner et al., 2009) provides a useful general framework to think about these and other applications. Players of a collective action game may face an adversary, who tries to block collective action to the maximal extent, in a so-called divide and conquer, named after the Roman strategy of creating discord. Obviously, collective action requires the participation of several players, but it also requires communication between the players, as has been empirically shown. Posner et al. conclude that the adversary of collective action can disrupt collective action in two manners, namely by stopping players from participating in collective action, or by blocking communication opportunities. E.g., an employer may block trade union activity either by stopping the activities of trade union members (e.g. by firing them), or by blocking

[2]For a more strictly mathematical treatment of such models see Bollobàs and Riordan (2003).

[3]That such a threat actually exists can be seen from the efforts taken by homeland security and other government agencies to find a disruption tolerant network. Raytheon BBN Technologies reportedly "was awarded a $81 million contract to create a collaborative technology alliance in network science" (Baburajan, 2010) and in 2010 demonstrated a field experiment of a disruption tolerant military network (Baburajan, 2010).

[4]An additional problem in the disruption of covert networks as resistance movements, and criminal and terrorist networks is that the network structure may not be known to the network disruptor. This is an aspect that our basic model does not deal with. Before such research is undertaken, it is interesting to know the equilibrium when network structure is observable.

communication between trade union members (e.g. by spreading trade union members over separated departments). [5] Posner et al.'s analysis, is based on a simple monotonously increasing effect, of the number of participants in collective action and the number of communication possibilities on the success of collective action. Yet, depending on the communication network that links the participants in collective action, the effect of taking out different players and communication links may be widely different. Still, Posner et al.'s framework can be directly interpreted in terms of network analysis, where the disruptor may either take nodes (what we call node deletion) or links (what we call link deletion) away from the disrupted network.

As Posner et al.'s analysis suggests, the applications of strategic network disruption and defence are not confined to military and security applications. Historically, firms have not only competed with one another through price and quality competition, but also by disrupting each other's operations. 16th century English and Dutch commercial fleets sunk each other's boats (Francois, 2006). Brevoord and Marvel (2004) report that in the beginning of the 20th century, NCR had a competition department the aim of which was to literally sabotage the entry of any competitor in the market for cash registers. Yet even in times where law enforcement prevents such practices, competition can still take disruptive forms. In industrial organization, this goes under the common denominator of raising rivals' costs (Salop and Scheffman, 1983).

In the context of this paper, where the importance of network structure is emphasized, raising rivals' costs may take the following form. In the standard model of competition, where competition is between firms, internally, every firm can be considered as a network of employees. In standard production functions, production is a simple function of the number of employees. However, consider two employees i and j who are identical except for their position inside the network of employees, where employee i has the function of a bridge between different groups of employees. Then the loss of employee i may have a hugely different impact than the loss of employee j. For this reason, a competitor of the firm may engage in predatory hiring practices (Kim, 2007), and offer specifically employee i a high wage to defect (node deletion). Because of this fact, firms may restructure in such a manner that they are not vulnerable to such practices.

Furthermore, staying within the model of competition between firms, some markets are characterized by network externalities (Katz and Shapiro, 1985), where the utility of a particular good to consumers is an increasing function of the number of consumers who use the good. E.g., a particular software program is more useful to the individual consumer the more other consumers use it. In standard models, the utility of a good is then a simple increasing function of the number of its consumers. Yet, these consumers may be connected in a particular network. In this case, consumers with a central position in a consumer network may be of special importance to the firm, as they can lead other consumers to adopt a certain good. A firm who depends on such a central consumer may be very vulnerable to disruption by a competitor, and firms may therefore aim to provide for consumers who are not linked in a vulnerable network. In another application, the good produced by a firm may literally consist of a network, such as a network of airports served by an airline, or a network of motels spread over a country (Hendricks et al., 1999). Such physical networks are characterized by the fact that there may locally be only a limited space available to construct a node in the network (such as a motel along a particular highway), or to construct a link in the network (such as a flight connection between two cities). Again, the individual seller of such a networked good wants to avoid that its network has vulnerable nodes or links. Otherwise, by pre-emptively occupying a certain physical spot, a competitor may deprive the firm of nodes or links that are crucial to the operation of its network.[6]

Recent models of competition also model competition taking place between competing networks of firms. We can distinguish here between vertical alliances, such as buyer-seller networks (Kranton and Minehart, 2001) of firms at different stages of the commodity chain, and horizontal networks, such as R&D or standard-forming alliances (for an overview, see Bloch (2002)). Again, the contribution of a single firm

[5]That such techniques, although largely illegal, are still being used by employers, can be seen by the frequent complaints of trade union members. In 2009 a report about practices used by the German company Deutsche Telekom to prevent workers in their US subsidiary from organizing in a Union was published. According to the report, such practices included hiring known "union avoidance companies" (Logan, 2009) to write manuals for managers to explicitly prevent workers from organizing in a union, distributing anti-union adds and union member intimidation (Logan, 2009).

[6]For a different approach to competition between transport networks, see van der Leij (2003).

to an alliance may depend on its position in the network. A vertical network of buyers and sellers is vulnerable if there are a few firms, or relations between firms, that are crucial to the operation of the entire network. Through exclusive dealings (Aghion and Bolton, 1987), equivalent to link deletion, or through vertical integration (Salinger, 1988), equivalent to node deletion, a competing network may then disrupt the functioning of the vertical network. A horizontal network of firms in e.g. an R&D alliance again faces the risk that a crucial firm is lured away by a competing network.

3. Model

The two-player full information game is played by a network designer and a network disruptor. At stage 1, the network designer has a set of $N = \{1, ..., n\}$ nodes available. The network designer uses the nodes in N to build a pre-disruption network g^1. If two nodes i and j are directly linked, we say that $g^1_{ij} = 1$. If they are not linked to one another $g^1_{ij} = 0$. Given this notation, the pre-disruption network g^1 is the set of g^1_{ij} such that $g^1_{ij} = 1$ holds.[7] Links are undirected so that $g^1_{ij} = g^1_{ji} = 1$ always holds. Nodes are indirectly linked to each other if a path exists between them. We assume that there exists a path between two nodes i and j if there exists a sequence of nodes $[i_1, ..., i_k]$ such that $g^1_{ii_1} = g^1_{i_1 i_2} = ... = g^1_{i_{k-1} i_k} = g^1_{i_k j} = 1$. We denote g^1_{-i} as network g^1 with node i removed, and g^x_{-ij}, as network g with link ij removed.

At stage 2, the network disruptor observes the network and can then choose to disrupt it. We consider two models of network disruption, reflecting simply the constituting parts of any network. In the link deletion model, the disruptor can decide to delete a number of links D_l from the pre-disruption network. This leads to a post-disruption network g^2 consisting of all the links for which $g^2_{ij} = 1$. In the node deletion model, the disruptor can decide to delete a subset v of nodes from the pre-disruption network, where D_v denotes the cardinality of this set. The post-disruption network g^2 then consists of all links such that $g^2_{ij} = 1, i, j \notin v$. At stage 3, both players obtain their payoffs.

If we use symbol g without superscript, this may refer to both the pre-disruption and the post-disruption network. Define as $N_i(g)$ the set of nodes with whom node i maintains a (direct or indirect) link. Given a network g, a set $C \subset N$ is called a component of g if for every distinct pair of nodes i and j in C we have $j \in N_i(g)$, and there is no strict superset C' of C for which this is true. The degree of connection of each node $\eta_i(g)$ is defined as the number of direct links the node has, so in effect the number of nodes it is directly linked with.

At stage 3 the players obtain their payoffs. Each of the identical nodes ex ante has one unit of information, but obtains information $|N_i(g^2)|$, i.e. obtains the total information of all nodes it is connected to in the post-disruption network. The value of a node i equals $u_i[|N_i(g^2)|]$, where $u'(.) > 0$, i.e. the value of node i is an increasing function of the amount of information obtained. As is clear from the above, contrary to what is the case in e.g. Jackson and Wolinsky (1996) and Bala and Goyal (2000), the nodes in our networks are not individual decision makers, and the linking decisions result from the design decisions of the designer at stage 1, and the disruption decisions of the disruptor at stage 2. The network designer's payoff is a function of the sum of the value of each node in the post-disruption network, i.e. the designer's payoff equals the sum of the values of all nodes, $\sum_i u_i[|N_i(g^2)|]$. The costs of the network designer are a function $c_{DES}(.)$, with $c'_{DES}(.) > 0$, of the number of links used in the pre-disruption network g^1. This number of links used is the sum of the degrees of each node divided by 2, or $[\sum_{i \in N} \eta_i(g^1)/2]$ and the costs therefore equal $[c_{DES}\left(\sum_{i \in N} \eta_i(g^1)/2\right)]$. The network designer plays a zero-sum game against a network disruptor, therefore the disruptor's benefit equals $-\sum_i u_i[|N_i(g^2)|]$. The disruptor's cost function is an increasing function $c_{DIS}(.)$, with $c_{DIS'}(.) > 0$, of the number of links D_l or nodes D_v taken out from the pre-disruption

[7] We thereby exclude empty networks, however, empty networks are irrelevant for our analysis.

network g^1. We consciously assume information decay (see Jackson and Wolinsky (1996)) and heterogeneous values of the nodes (see Galeotti et al. (2006)) away. In this way, in the absence of network disruption ($D_l = 0$ or $D_v = 0$), any minimally connected pre-disruption architecture [8], which uses $(n-1)$ links, is a best response to the designer. This means that, if there is a network disruptor, any restrictions that we obtain on the set of best-response architectures, or any non-minimal links in best-response architectures, are strictly better due to attempts to prevent disruption. Our model thus allows us to isolate the pure effect of defense against network disruption.

As will become clear from the analysis, even with continuous functions v, c_{DES} and c_{DIS}, the designer's and disruptor's maximands are discontinuous. It may be the case that a number of links added to g^1 does not increase the designer's payoff at all, in that the designer cannot protect the network in a better way by means of these links; but it may also be the case that one single link added has a huge effect on the designer's payoff, in making a huge difference for network protection. Similarly, it may be the case that a number of extra links or nodes deleted has little effect on the disruptor's payoff, in that they do not cause any reduction in the designer's payoff; but it may also be that one extra link or node added has a huge effect, in enabling substantial disruption of the designer's network. In view of this, we cannot find the equilibria by means of maximization of continuous maximands.

Instead, first, we simplify the benefit functions of designer and disruptor, in assuming that the payoff of the designer (disruptor) is a positive (negative) function only of the order[9] of the largest component in g^2. This reflects the increasing marginal benefits of the information generated within a single component due to the non-excludable nature of information in our model. Let every node have one unit of information. Consider now a component of order x. The information generated by this component equals $(x)^2$. Everyone benefits from the information of an added node, and so the added benefit of increasing the order of a component are larger the more nodes this component already has. Second, we assume that the disruptor has a fixed link deletion budget D_l or node deletion budget D_v , so has already decided on how many links or nodes to delete in g^1. This is without loss of generality. In any equilibrium, the disruptor decides on a certain number of links or nodes to be deleted. Given such a number of nodes or links, it must be the case that the disruptor chooses a best response in the form of an optimal network disruption strategy.

Third, we model the designer's decision in either of two simplified manners, depending pragmatically on which manner allows us best to obtain results. In one approach, we assume that the designer has a linking budget B; given such a budget, the designer then maximizes the value of the network, in maximizing the order of the largest component in the post-disruption network. In the other approach, we assume that the designer decides on achieving a fixed order for the largest component in g^2, and so a fixed benefit; given such an order, the designer aims to achieve this with a minimal number of links in g^1. Note that this reflects what takes place at an optimal solution: given the order of the largest component achieved there, this is achieved with a minimal number of links; given the number of links used, a maximal order for the largest component is aimed at. This is illustrated in Figure 1. On the X- and Y-axis the two elements of the designer's payoff are represented; namely the number of links used in the pre-disruption network, and the order of the largest remaining component in the post-disruption network. To each combination of these two measures corresponds a maximal payoff that can be achieved on the Z-axis, depending on an appropriate pre-disruption network structure. The payoff of the designer is represented in Figure 1 by means of a hill (even though the real payoff function is likely to be discontinuous). The curves represent the two approaches treated. In the first approach, in planes parallel to the Y-axis, we look for the largest achievable remaining component for a given number of links B. In the second approach, in planes parallel to the X-axis, we look how a fixed order of the largest remaining component can be achieved with a minimal number of links.

The second approach deserves further attention. We start by defining the concept of proofness of a pre-disruption network, which defines the levels along the Y-axis in Figure 1.

[8] For definitions and proofs concerning the graph theoretic terms used in this paper see the Appendix.

[9] The order or the network, is the size of the network in terms of the number of nodes it includes. This term is used rather than size, so as to avoid confusion, as it is commonly used in graph theory.

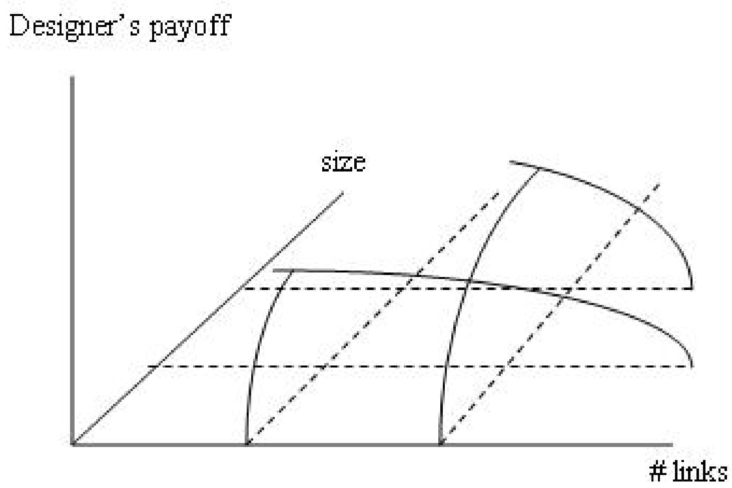

Figure 1: *Designer's Payoffs*

Definition 1. *A pre-disruption network g^1 is said to be $(n-x)$ proof against a link (node) deletion budget $D_l(D_v)$ if the largest remaining post-disruption component upon strategic link (node) deletion contains exactly $(n-x)$ nodes.*

This definition leads to two observations, in the form of the following Lemmata.

Lemma 1. *For $x \geq 0$, a pre-disruption network is $(n-x)$ proof against a link deletion budget D_l if*

(i) every set of y nodes, with $y > x$, has at least $(D_l + 1)$ links to nodes outside the set y, and

(ii) for any non-empty set of x nodes, there is at least one such set that has D_l or less links to nodes outside the set of x nodes.

We then denote the set of all networks that are $(n-x)$-proof against a link deletion budget of D_l as $\Gamma_{l,D_l}^{(n-x)}$. It is these networks that are represented in the planes parallel to the Y-axis in Figure 1. They differ according to the number of links that they use, as represented by the curves in the planes parallel to the Y-axis.

Lemma 2. *For $x \geq D_v$, a pre-disruption network is $(n-x)$ proof against a node deletion budget D_v if*

(i) every set of y nodes, for $y > x$, has at least $(D_v + 1)$ links to neighbors connecting it to nodes outside the set y, and

(ii) for any non-empty set of $(x - D_v)$ nodes, there is at least on such set that has D_v or less neighbors connecting it to nodes outside of the set of $(x - D_v)$ nodes.

We then denote the set of all networks that are $(n-x)$-proof against a node deletion budget of D_v as $\Gamma_{v,D_v}^{(n-x)}$. Again, these networks are represented in the planes parallel to the Y-axis in Figure 1. If we then look at the two cases of link- and node deletion, we can see that it is straightforward to deduce how large the post-disruption network can maximally be in either case.

Lemma 3. *Under link deletion with a deletion budget of D_l, a pre-disruption network can be at most n-proof. Under node deletion with a deletion budget of D_v, a pre-disruption network can be at most $(n - D_v)$-proof.*

Proof The result for link deletion follows from the fact that it is not possible for a post-disruption network to contain more nodes than the pre-disruption network contained. The result for node deletion follows simply from the fact that at least D_v nodes will be removed from the network, due to the definition of node deletion. □

In order to make it easier to compare link deletion and node deletion, we define as a *max*-proof network a network that achieves the highest achievable proofness under either link deletion or node deletion. As follows from Lemma 3, *max*-proofness then means *n*-proofness for the case of link deletion and $(n - D_v)$-proofness for the case of node deletion. In the same manner, we define as a $(max - x)$-proof network, a network where the largest remaining component after strategic disruption has x less nodes than the *max*-proof network.

A first result that we can then immediately state is that node deletion restricts network structure to a higher extent than link deletion. Thus when facing an attack on its nodes, the network designer is more restricted in the way he builds up his network than when the attack is directed at the links of his network.

Proposition 1. *Compare node and link deletion when $D_v = D_l$. Then any network that is $(max - x)$-proof under node deletion is also $(max - x)$-proof under link deletion. But not all networks that are $(max - x)$-proof under link deletion are also $(max - x)$-proof under node deletion. Put otherwise, for $D_v = D_l$, the set $\Gamma_{v,D_v}^{(max-x)}$ is a strict subset of the set $\Gamma_{l,D_l}^{(max-x)}$.*

Proof Any set of nodes x that has l neighbors connecting it to nodes outside of the set x necessarily has l links to nodes outside of the set x. Thus by taking out l links, no more than l nodes can be separated from the network. However, any set of x nodes that has l links connecting it to nodes outside of the set of x nodes does not necessarily have l neighbors connecting it to nodes outside of the set of x nodes. For instance, two of these links may be to the same neighbor, in which case the network is $(max - x)$-proof under link deletion, but not $(max - x)$-proof under node deletion. □

As linking is costly, a network designer that aims at achieving $(max - x)$-proofness will do this with a minimal number of links. This leads us to the definition of minimal $(max - x)$-proofness.

Definition 2. *A network g is said to be* ***minimal*** *$(max - x)$-proof against a node (link) deletion budget $D_v(D_l)$, if no network exists that achieves $(max - x)$-proofness using less links.*

We denote the set of all networks that are minimal $(max - x)$-proof against a node (link) deletion budget of $D_v(D_l)$ as $\Gamma_{v,D_v}^{(max-x),min}$ ($\Gamma_{l,D_l}^{(max-x),min}$). It should be noted that the fact that for $D_v = D_l$, the set $\Gamma_{v,D_v}^{(max-x)}$ is a strict subset of the set $\Gamma_{l,D_l}^{(max-x)}$ (see proposition 1) does not imply that the set $\Gamma_{v,D_v}^{(max-x),min}$ is a strict subset of the set $\Gamma_{l,D_l}^{(max-x),min}$.
We now treat the cases of low, intermediate and high linking costs. Some of the basic graph-theoretic definitions and results that are used in these sections are treated in the Appendix.

4. Low Linking Costs

In this section, we assume that linking costs are low enough to assure that the designer always wants to construct a *max*-proof network, so that only a minimum number of nodes can be removed from the network. Yet, because links continue to be costly, the designer will aim at *max*-proofness with a minimal number of links. We thus apply the approach where we look for the minimal *max*-proof networks. To find such minimal *max*-proof networks, is a straightforward task. We know from Lemmata 1 and 2 that a necessary condition for *max*-proofness under a disruption budget of $D_v = (r - 1)$ or $D_l = (r - 1)$ is that each node receives at least r links. Good candidates for minimal *max*-proof networks are therefore networks in which each node has exactly degree r, because then each link is crucial in assuring *max*-proofness. Such networks are known as r-regular networks.

Definition 3. *An r-regular network is a network in which each node is connected exactly of degree r.*

For any given n and r, any r-regular network has the same number of links [10]. Thus, if there is an r-regular network that is *max*-proof, it is necessarily also minimal.

Lemma 4. *If an r-regular network with $r = (D_v + 1)$ (respectively $r = (D_l + 1)$) exists that is max-proof given $D_v(D_l)$, then this network is also minimal max-proof. It is then the case that the sets $\Gamma^{(n-x),min}_{v,D_v}$ and respectively $\Gamma^{(n-x),min}_{l,D_l}$ only contain r-regular networks.*

Proof By Lemmata 1 and 2, any *max*-proof network must have nodes that each have at least r links. Any network where each node has exactly r links has exactly the same number of links [11]. It follows that any r-regular network that is *max*-proof is also minimal *max*-proof. □

Lemma 4 is lacking in two aspects: first, it does not prove that such minimal *max*-proof networks actually exist. Second, even if these do exist, it does not characterize what form they take. The rest of the section deals with these two problems. They are easiest to solve for the simple case where $D_v = D_l = 1$, as in this case the circle is the unique minimal *max*-proof architecture:

Corollary 1. *For $D_v = D_l = 1$, the unique minimal* max-*proof architecture is the circle containing all n nodes.*

Proof Graph-theoretically, the only connected 2-regular network is the circle. In the circle, every set of connected nodes has two links and two neighbors connecting the set to the other nodes.[12] Every set of nodes including nodes unconnected to one another has more than two links and two neighbors connecting it to the rest of the nodes. Given Lemma 4, the circle is minimal *max*-proof. □

The circle has several graph-theoretic properties (see e.g. Chartrand (1977) or Diestel (2005)) that we will generalize, and use to show the general existence of r-regular networks that are minimal *max*-proof under a link or node deletion budget of $(r - 1)$. We now define these properties.[13].

Definition 4. *A **hamiltonian** network is a network that contains a circle spanning all nodes in the network.*

Definition 5. *Denote by d_{max} the maximal distance between any two nodes in a network. A network is symmetric if every node in the network has exactly N_x nodes at distance x, with $x = 1, 2, ..., d_{max}$.*

As any network where each node has exactly N_1 nodes at distance 1 is N_1-regular, it follows that every symmetric network is also a regular network. However, oppositely, not every regular network is also symmetric. The circle network is clearly a regular, symmetric network. A further property of the circle is that there are no multiple links between the same nodes, i.e. it is what graph theorists call a simple network.

Definition 6. *A **simple** network is an undirected network containing no multiple links between two nodes and no links beginning and ending at the same node (commonly referred to as loops).*

As the circle is a Hamiltonian, symmetric, simple network that is minimal *max*-proof for the smallest deletion budget, this suggests that a network with these properties is minimal *max*-proof for general disruption budgets. We now prove that this to be indeed the case, thus showing the non-emptiness of the set of minimal *max*-proof networks.

Lemma 5. *For n and/or r even, with $r > 2$, a simple, Hamiltonian, symmetric r-regular network exists that is minimal max-proof under link and node deletion budget $(r - 1)$.*

[10]For an existence proof for r-regular networks see the Appendix.
[11]For a the necessary conditions for the existence of r-regular networks see Appendix Lemma A.7.
[12]For a proof see Lemma A.1 in the Appendix.
[13]There are more graph-theoretic concepts that we took into consideration, such as Menger's proposition (Menger, 1927) or the concept of *k-connectivity* (Frank, 1995), however, while they are close to what we are doing, we did not find them useful in further characterizing networks more stringently.

Proof Step 1 shows that an r-regular network can be constructed whenever n and/or r is even. Step 2 shows that the constructed networks are *max*-proof under a disruption budget *(r - 1)*.
Step 1. Label each node i with a label $l_i \in \{1, 2, ..., n\}$. For each $\delta \leq r/2$, define a network $g_\delta = \{ij : l_i = l_j + \delta, i_x j_x \not\equiv i_y j_y (\text{mod} n)\}$. For n odd, and for n even and $\delta \neq r/2$, all such networks take the form of a circle, or if (n/δ) is an integer, take the form of a set of δ circle components of order (n/δ). For n even and $\delta = r/2$, $g_{r/2}$ consists of *(n/2)* components consisting of a single link. An r-regular symmetric network when n is odd, or when n is even and r is even, can be constructed as follows: $g_1 \cup g_2 \cup ... \cup g_{(r/2)-1}$. An r-regular symmetric network when n is even and r is odd can be constructed as follows: $g_{r/2} g_1 \cup g_2 \cup ... \cup g_{(r-1)/2-1}$. Note that each of these networks is a Hamiltonian network, as g_1 is an n circle. By construction, each of these networks is also simple.
Step 2. In the constructed simple r-regular networks, every node has r links and r neighbors. Further, it can be checked that in networks constructed in Step 1, every set of nodes also has at least r links and r neighbors connecting the set to the rest of the nodes. □

Lemma 5 does not imply, however, that every minimal *max*-proof network is a Hamiltonian, symmetric simple graph. We can find graphs that do not have all of these graph-theoretic properties but are still minimal *max*-proof. A well-known example of such a graph is the so-called Petersen graph, which is the graph on the left in Figure 2. This can be checked to be minimal *max*-proof for $D_v = 2$ or $D_l = 2$, but it is non-Hamiltonian. The middle graph in Figure 2 is also minimal *max*-proof but is not not a simple network. The right graph, while being minimal *max*-proof, is non-symmetric.

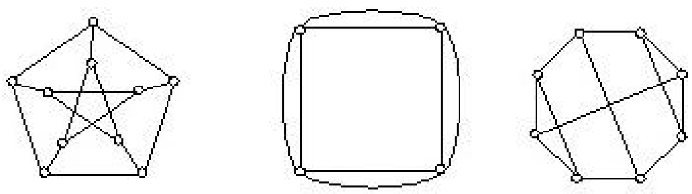

(a) Non-Hamiltonian (b) Non-Simple (c) Non-Symmetric

Figure 2: Minimal Max-proof Networks

We conclude from these examples that the characterization of minimal *max*-proof networks must be formulated in wider way than suggested by Lemma 5. We simply note that minimal *max*-proof networks do not contain any bridge link sets, or bridge node sets.[14]

In effect, this means that in networks that do not contain any bridge link sets or bridge node sets, not only all single nodes need to be linked of a certain degree to the rest of the network, but also all sets of nodes are linked to the rest of the network with a certain number of links. Thus, we can come to the following proposition.

Proposition 2. *Let n and/or* $(D_l + 1)$ *be even. Then* $\Gamma^{max,min}_{v,D_v}$ *(and respectively* $\Gamma^{max,min}_{l,D_l}$*) is the set of all networks* g *with the following characteristics:*

(i) g *is a connected* $(D_v + 1)$ *(respectively* $(D_l + 1)$*) regular network;*

(ii) g *does not contain any bridge link sets (bridge node sets) of order* $x \leq r$.

Proof We prove each part of the proposition independently.

[14]See the Appendix for a definition of these concepts.

(i) Under link deletion, any minimal *max*-proof pre-disruption network must be connected, since otherwise the post-disruption network is not connected no matter which links are deleted. Under node deletion, any minimal *max*-proof pre-disruption network must be connected, since otherwise the disruptor can remove at least D_v nodes in the largest pre-disruption component. The rest of the proof of this part follows by noting that all r-regular networks use exactly the same number of links and from Lemma 4.

(ii) This follows directly from Lemmata 1 and 2. □

An illustration of Proposition 2 is given in Figure 3, for the case where $n = 16$ and $r = 3$. Whereas the only connected 2-regular architecture is the circle, a myriad of architectures is connected 3-regular. It can be checked that all represented graphs in the figure use exactly 24 links. We investigate what happens when they are facing a network disruptor with a disruption budget of $D_v = D_l = 2$. The largest remaining

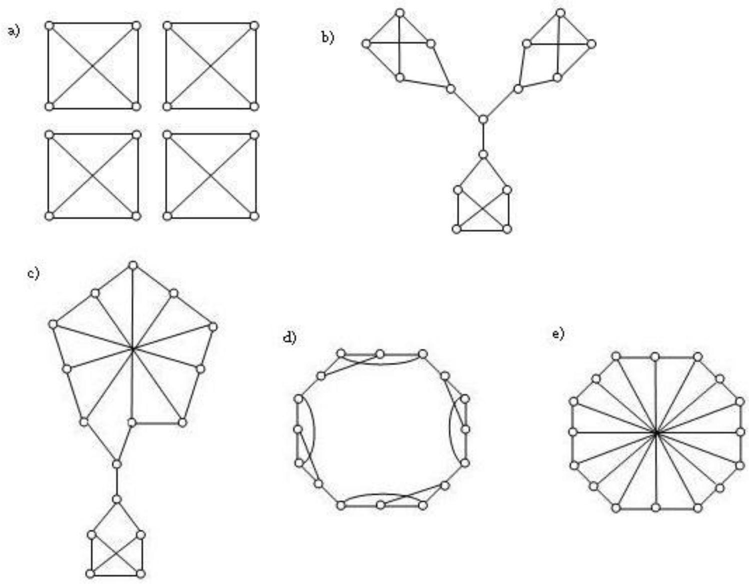

Figure 3: *Max- and Non-Max-proof 3-regular Networks*

connected component in the non-connected network (a) in the graph, consists of only 4 nodes. Network (b), which has been taken from Dekker and Colbert (2004), contains a bridge node as well as a bridge link. Under node deletion the largest remaining post-disruption component contains only 5 nodes and 6 nodes under link deletion. Network (c) also contains a bridge node and a bridge link. However, under node deletion the largest remaining component includes 8 nodes, and under link deletion 10 nodes. Network (d) contains a 2 bridge link set and a 2 node bridge set. Therefore the largest remaining component under node deletion contains only 7 nodes and under link deletion 8 nodes. Network (e) is *max*-proof.

We can thus conclude, that minimal *max*-proof networks are connected regular graphs that do not contain small bridge node sets or small bridge link sets. Intuitively, this means that networks that are *max*-proof should not contain clusters of highly connected local cliques, with few links between the cliques. Clearly, such networks are easy to disconnect.

We have already shown in Proposition 1 that network structure is more restrictive for node deletion than for link deletion, when building *max*-proof networks. We now extend this to show that this also holds for *minimal max*-proof networks (note that this does not imply that this is a general result for minimal *(max - x)*-proof networks).

Proposition 3. *Let $D_v = D_l$, and let n and/or D_v be even.The set $\Gamma^{max,min}_{v,D_v}$ is a subset of the set $\Gamma^{max,min}_{l,D_l}$.*

Proof *Step 1*. In this step, we show that any network that is minimal *max*-proof under node deletion is also minimal *max*-proof under link deletion. In any network that is minimal *max*-proof under node deletion, by Proposition 2, every node and set of nodes has at least D_v neighbors connecting it to the rest of the network, and therefore at least D_l links connecting it to the rest of the network, making it *max*-proof. By Lemmata 4 and 5, since any such network is $(D_l + 1)$ regular, it is also minimal *max*-proof.
Step 2. The set $\Gamma^{(n-x),min}_{v,D_v}$ contains only simple graphs. When there are multiple links between two nodes, removing one of these links does not make any difference for the proofness of the network under node deletion.
Step 3. To show that graphs may exist that are minimal *max*-proof under link deletion, but not under node deletion, consider a non-simple graph of the following form. For $r \geq 4$ and being an even number, construct a circle, and double ($r = 4$), triple ($r = 6$), etc. each link. Such a graph is *r*-regular and minimal *max*-proof under link deletion, however, by step 1, not under node deletion. □

An illustration of Proposition 3 is in Figure 4, representing the case for $n = 16$ and $r = 4$, which is minimal *max*-proof under link deletion but not *max*-proof under node deletion.

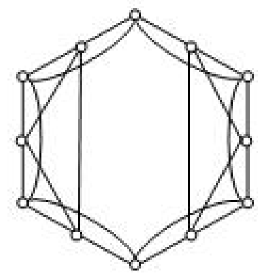

Figure 4: *Max*-proof under Link Deletion but not under Node Deletion

It is easy to see that keeping a network *max*-proof comes at a high cost to the network designer, as for every additional link (node) in the network disruptor's budget, he needs to increase his defense budget by $n/2$ links. Therefore, if linking costs are relatively high, it is unlikely that the network designer will construct a *max*-proof network. In the rest of the paper, we look at the case of higher costs.

5. High Linking Costs

Now looking at the second approach introduced in Figure 1, we investigate what happens if the network designer has a fixed linking budget. In this case here, we are looking at prohibitively high linking costs, which leads to a defense budget of 0 links. Thus adding additional links next to the original linking budget of $B = (n - 1)$, is prohibitively expensive; leaving the network designer to structure his network as robustly as possible with this limited amount of links.

As in the previous case, the structure of the network now becomes decisive again. However, while in the case of low linking costs, the same type of architectures are best responses (with the case of link deletion more restrictive), here it quickly becomes apparent that fundamentally different architectures are best responses under link deletion compared to node deletion. This is obvious by simply looking at minimally connected networks. While the network disruptor can completely disconnect the star network, by taking out the central node, the maximum damage he can cause to the chain network is to separate two additional nodes from the network, thus leaving a largest remaining component of $n/2$. For link deletion, however, the star network seems like a good option, since the maximal damage the network disruptor can cause by taking out a single link, is to disconnect one node from the network, whereas the chain network can be cut in half. Thus, here again, we will start by analyzing the link deletion case, and then turning to the node deletion case.

5.1. Link Deletion - High Linking Costs

For high linking costs, we assume that the network designer has a linking budget of $B = (n-1)$ but no defense budget at all. The network designer cannot afford to use any more links than are necessary to build a minimally connected network. He puts up with the fact that some nodes will be disconnected from the network, and prefers this to adding additional links, because such links are simply too expensive. In the modeling section, we have discussed that if there is no threat of an attack on the links of the network, all minimally connected networks are best responses. However, now taking into account that there is a network disruptor with a disruption budget $D_l > 0$, not all minimally connected networks are best responses. We will show in the following that the star network is the network structure that weakly dominates all other structures in case the network designer faces a network disruptor with a disruption budget of $D_l \geq 0$.

In the star network, the maximal damage a network disruptor with any positive disruption budget can cause can be straightforwardly calculated as follows.

Lemma 6. *The maximal damage a network disruptor with a disruption budget of D_l can cause in a star network is to disconnect exactly D_l nodes.*

Proof The star consists of exactly $(n-1)$ links, which connect one node with degree $(n-1)$ with $(n-1)$ nodes with degree 1. With each link that the network disruptor can delete, he can therefore only separate one of these nodes with degree 1 from the central node. Therefore, with a disruption budget of D_l, he can separate exactly D_l nodes. □

To show that the star is the most robust network topology for a linking budget of $B = (n-1)$ and a disruption budget of D_l, we need to consider then what happens to other minimally connected networks.

Lemma 7. *In every non-star minimally connected network, at least $(D_l + 1)$ nodes can be removed.*

Proof Every non-star minimally connected network has a diameter of 3 or larger. To see why, take any minimally connected network and let node k be an end-node (thus connected of degree $\eta_k = 1$). As is shown in the Appendix (Lemma A.5), there need to be at least 2 end-nodes in the network, since by definition any minimally connected network does not contain a circle. By definition node k must have a link to a node i in the network (otherwise the network would not be connected). Since the network is not a star, node i cannot receive another $(n-2)$ links, next to link g_{ki}. Thus there needs to be at least one more node h, which is only linked to i indirectly through node j. The distance between node k and h, is then 3 and no matter how else the network looks like it has at least diameter 3. It thus follows that if we delete link g_{ij}, at least 2 nodes are separated from the largest remaining component. This is because either nodes i and k are separated from the largest remaining component, or nodes j and h (possibly along with further nodes to which they are connected). It follows that removal of link g_{ij} results in the separation of two nodes from the largest component. The $(D_l - 1)$ remaining links that are deleted each time result in the separation of at least one node, as every link in a minimally connected network is a bridge link (see Appendix Lemma A.4). □

Thus the minimal damage that can be done in any minimally connected network that is not the star, is always larger than the damage that can be done to the star for the same disruption budget, as can be seen from Lemmata 6 and 7. Therefore the star network needs to at least weakly dominate all other minimally connected networks in this case.

What we have not considered so far is, however, the case when the network designer might build a smaller but stronger network by not including all nodes in the component, but using links instead to make the network more robust against attack. We now turn to such cases. We know from the graph-theoretic Appendix that in any minimally connected network there is no circle and in any network that is not minimally connected there is at least one circle. To build a circle we know we need exactly n links. Since we have a linking budget of $B = (n-1)$, the network with a circle can therefore only include $(n-1)$ nodes. We also know that any end node can be disconnected by disrupting one link only. Therefore, if we do build a network including less nodes, we should not leave any end nodes in the network because they are natural weak spots. Therefore the network designer, should include all $(n-1)$ nodes in the circle. For any additional link the network designer adds to the circle, he has to leave one additional node out of the connected component. Therefore, we can calculate exactly how large the network in question will be, depending on the number of links added to the basic circle.

Lemma 8. *Consider a network designer who builds a pre-disruption network consisting of a connected component, and a set of isolated nodes. The network designer with a linking budget of $B = (n-1)$, needs to leave one additional node out of the connected component, as compared to a minimally connected network, for each additional non-minimal link he wants to add to the connected component.*

Proof For a circle this is trivially true, as we know from the graph-theoretic Appendix that any circle needs exactly n links. *Step 1*. If the designer constructs a connected component consisting of $(n-1-x)$ connected nodes and x isolated nodes, it will need exactly $(n-1-x)$ links to assure that the connected component is connected. Given that the linking budget is $(n-1)$, x non-minimal links can now be added.
Step 2. If the designer puts x non-minimal links in the connected component, it uses $(\gamma - 1 + x)$ links, where γ is the number of nodes in the connected component. If the designer uses up all links, $(\gamma - 1 + x) = (n-1)$. It follows that $\gamma = (n-1-x)$. □

Knowing this, the question then is, how many nodes remain in the post-disruption network g^2, because the component does not only get smaller but also stronger. Thee answer to this question is easy for the case $D_l = 1$. In this case, the designer can leave out at most one node to construct a smaller connected component. The optimal architecture of this component of $(n-1)$ nodes is the circle, as any end node in the component means that one node can be removed from the connected component. It follows that in the $(n-1)$ circle, the largest remaining component has order $(n-1)$. This is the same result as for the star with $(n-1)$ nodes. We next look at the general case.

For r-regular networks we know that we need exactly $(n*r)/2$ links. Thus if we rearrange this equation, we get that for a linking budget of $(n-1)$, we can build a (D_l+1)-regular network with exactly $(n-1)*2/(D_l+1)$ nodes. Then, if the connected component is (D_l+1)-regular and meets the conditions of Proposition 2, no nodes can be separated from it. Therefore, whenever $(n-1)*2/(D_l+1) < (n-D_l)$, the star network is a better choice than the (D_l+1) regular component. For $D_l = 1$ and $D_l = (n-2)$, the two sides in the latter inequality are equal; for levels of D_l in between, the inequality is valid. The case $D_l = 1$ simply was already treated above. For $D_l = (n-2)$, the structure of the pre-disruption network is irrelevant, as in any case a single connected pair remains in the post-disruption network. It follows that only the intermediate levels are relevant, and for these the inequality is valid.[15] However, there still is the possibility of making a larger connected component, from which nodes can then still be separated. The following Proposition 4 shows that the designer cannot do better in this way.[16]

Proposition 4. *When facing a network disruptor with a disruption budget of D_l, with $1 \leq D_l \leq (n-2)$, the weak best response of a designer with linking budget B = (n - 1) is to build a star network.*

[15]Thus in all interesting cases because since $D_l = r - 1$, otherwise we end up with only one link in the post disruption network in any case.

[16]Making several smaller components also will never be better than the star network, as the same logic applies that they can be taken apart again by the network disruptor. Additionally already at least one link is used in making a second connected non minimal component. Therefore, the star strictly dominates this case.

Proof We prove this proposition in three steps. Step 1 shows that non-star minimally connected networks never do better than the star. Steps 2 and 3 show that the designer is never better off when constructing a smaller connected component.
Step 1. By Lemma A.4, in every minimally connected graph, at least D_l nodes can be removed. Further, we know from the Appendix that every non-star minimally connected graph contains at least two nodes with degree larger than 1. From this fact, it follows that at least one extra node can be separated from the network, on top of the D_l nodes removed.
Step 2. Denote by x the number of nodes with zero degree in the pre-disruption network. Consider the case where $D_l \leq x$. This means that at least D_l nodes are already separated from the pre-disruption network. Clearly then, this cannot do better than the star.
Step 3. Consider the case where $x < D_l$. By Lemma 8, the connected component than has x non-minimal links. Suppose that the disruptor first deletes these non-minimal links. Then at best, a minimally connected component remains. From this component, the disruptor can now still remove $(D_l - x)$ nodes. Since the connected component in g^1 has $(n - x)$ nodes, and at least $(D_l - x)$ nodes are removed from it, the largest component in g^2 has at most $(n - D_l)$ nodes, which is the same as in the star network. □

So while all minimally connected networks are equally good responses if there is no threat of an attack, the star is the only best-response minimally connected network in case of an impending attack for a linking budget of $B = (n - 1)$. The star network is also always at least as good as any other possible network made up of $(n - 1)$ links. It appears that it is only in special cases, such as when there is a very small or very large disruption budget, that there is an alternative best response to the star.

5.2. *Node Deletion - High Linking Costs*

For node deletion we can immediately show that for the case of $D_v = 1$, the best the network designer with a linking budget of $B = (n - 1)$ can do is to build a circle containing $(n - 1)$ nodes. This suggests that in general, the network designer should build a smaller, stronger component, an intuition that we will indeed confirm in this section.

Proposition 5. *For $D_v = 1$, and a linking budget of $B = (n - 1)$ nodes the designer's best-response network architecture is the circle containing $(n - 1)$ nodes.*

Proof *Step 1.* In any minimally connected network, every node is a bridge node (see Appendix). We already know that the in the star network, the largest remaining post-disruption component has order 1. Any non-star connected network contains at least one path of links *ij*, *jk*, *kl*. Given that every node is a bridge node, by removing *j* or *k*, the disruptor can disconnect at least three nodes from the network. *Step 2.* Every component that links $(n - 1)$ nodes using $(n - 1)$ links, and is not a circle of $(n - 1)$ nodes, has at least one end node. It follows that, in this component, the disruptor can delete one extra node on top of the deleted node. Together with the node that was not connected in the pre-disruption network, this means that a largest post-disruption component of at most $(n - 3)$ nodes remains. *Step 3.* Every network that links less than $(n - 1)$ nodes has a largest post-disruption component that is smaller than $(n - 2)$. *Step 4.* In the circle that links $(n - 1)$ nodes using $(n - 1)$ links, if one node is deleted, a post-disruption component connecting $(n - 2)$ nodes remains. □

For larger disruption networks, we cannot give a full characterization of the designer's best response pre-disruption network. Still, we can show that an essential feature of any best-response pre-disruption network is that this network does not connect all nodes to one another. In order to show that it is not a best response for the designer to construct a minimally connected pre-disruption network, we must first know the order of the largest remaining post-disruption component given a minimally connected pre-disruption network. We start by deriving this for the simple case of a deletion budget $D_v = 1$. We denote the smallest integer larger than a number x as $\lceil x \rceil$.

Lemma 9. *In any minimally connected network, the largest remaining component after an attack by a network disruptor with a disruption budget of $D_v = 1$, will be maximally of order $\lceil (n - 1)/2 \rceil$.*

Proof *Step 1.* By Lemma A.1, in any minimally connected network, every node is a bridge node, so that any minimal connected network can always be separated into two parts by removing one node.
Step 2. A disruptor can always do better than to take out a node such that the largest remaining component is of an order larger than $\lceil (n-1)/2 \rceil$. Let the disruptor take out node x with degree $\eta_i(g) = d$, which has links to nodes $y_1, y_2, \ldots, y_d$. Given that by Lemma A.1, each node in a minimally connected graph is a bridge node, taking out node x leads to d separated components, which we can denote as $g_1, g_2, \ldots, g_d$. Let the node labeled y_d and the corresponding component g_d have contain $s$nodes, with $s > \lceil (n-1)/2 \rceil$. Then it follows that $g - g_d$, meaning the network obtained when component g_d is removed from the network, is of an order smaller than $\lceil (n-1)/2 \rceil$. By instead removing y_d in g_d, the disruptor can assure that the component g_{z_d} which includes nodes connected to a neighbor z_d of y_d is of an order of at most $(s-1)$, so that the order of this component is smaller than the order of g_d. At the same time, we have already seen that the order of component $g - g_d$ is smaller than $\lceil (n-1)/2 \rceil$. It follows that the disruptor is better off by disrupting y_d. Therefore, a disruption strategy where a largest component larger than $\lceil (n-1)/2 \rceil$ is left can never be optimal for the disruptor.
Step 3. Given that by Step 1 a network disruptor never leaves a largest component of an order larger than $\lceil (n-1)/2 \rceil$, the best that the designer can possibly do is to leave a largest component of an order of exactly $\lceil (n-1)/2 \rceil$. □

The *chain* architecture shows that the network designer can actually achieve the maximal order of the post-disruption network suggested by Lemma 9. The following Lemma now generalizes the result of Lemma 9 to generic disruption budgets.

Lemma 10. *In any minimally connected network, the largest remaining component after an attack by a network disruptor on the nodes of the network with a disruption budget D_v, will be maximally of an order $\lceil (n - D_v)/[D_v + 1] \rceil$.*

Proof We prove this by induction. Step 1 is the base step, step 2 the inductive step.
Step 1. The largest remaining component for $D_v = 1$ has order $\lceil (n-1)/2 \rceil$, as is shown by Lemma 9.
Step 2. We show in this step that, if it is true that the largest remaining post-disruption component in a pre-disruption minimal connected network of order γ, given $D'_v = (D_v - 1)$, is $\lceil (\gamma - D'_v)/(D'_v + 1) \rceil$, then it follows that the largest remaining component for a minimal connect network, given D_v, is $\lceil (n - D_v)/(D_v + 1) \rceil$.
Consider a link ij in a minimal connected g^1, where node j has the following properties. Among the components in g^1_{-j} not connected to i, the largest component has an order of at most $\lceil (n - D_v)/[D_v + 1] \rceil$; in g^1_{-i}, the component connected to j has an order larger than $\lceil (n - D_v)/[D_v + 1] \rceil$. Every minimal connected network contains at least one link ij where node j has these properties. This is because in a minimal connected network, every node lies on a path between two end nodes. Every deleted node cuts the network in at least two components. As one deletes consecutive nodes along this path, the maximal order of one component becomes smaller, while the maximal order of the other component gets larger. By continuity, one must meet a node j with the properties above.
Suppose that the disruptor deletes node j. Then i is part of a connected component C of an order of at most $[n - \lceil (n - D_v)/(D_v + 1) \rceil - 1]$ remains, in which the disruptor can delete a further $D'_v = (D_v - 1)$ nodes. By the assumption at the start of this step, the largest remaining component in C after the disruptor has taken out a further D'_v nodes from C has an order of at most $\lceil (n - D_v)/[D_v + 1] \rceil$. It follows that this is also the largest component that the designer can keep when the disruptor deletes D_v nodes from the entire network. □

As we now show, the network designer can do better than with a minimally connected network by constructing a circle of $(n-1)$ nodes. By Proposition 5, we already know that this result holds for $D_v = 1$, where in fact the circle of order $(n-1)$ is the unique best-response architecture. We now show that the fact that the circle of $(n-1)$ nodes is a better response holds for any relevant disruption budget. We start by deriving the order of the largest component that can remain after disruption.

Lemma 11. *In a circle network, the network disruptor with a disruption budget of D_v will cause maximal damage by cutting the network into D_v separate components, each maximally of order $\lceil (n-1-D_v)/D_v \rceil$.*

Proof As the circle is completely symmetric, any disruption strategy by the disruptor can be seen as the deletion of one random node in the circle, and $(D_v - 1)$ further nodes. After the deletion of this random node, the remaining network takes the form of a chain, i.e. a minimal connected component of order $(n-1)$, in which the disruptor can delete $(D_v - 1)$ nodes. It follows directly from Lemma 9 that the largest remaining post-disruption component has an order of $\lceil ([(n-2)-(D_v-1)]/[(D_v-1)+1] \rceil$. □

We are now ready to show that the minimal connected network is never better than the circle that excludes one node. Moreover, for the range of disruption budgets for which a post-disruption largest component of order larger than 1 remains, and for a range of relatively large n, the circle is strictly better.

Proposition 6. *When facing a network disruptor with a node disruption budget larger than or equal to $(n-1)/2$, the line architecture with order n and the circle architecture with $(n-1)$ are equally good responses. When facing a network disruptor with a disruption budget smaller than $(n-1)/2$, the line architecture with order n is never better than the circle architecture with $(n-1)$ nodes; for relatively large n, the circle of order $(n-1)$ is strictly better.*

Proof By Lemma 9 we know that the largest remaining post-disruption component with a pre-disruption line of n nodes has order $\lceil (n-D_v)/(D_v+1) \rceil$. By Lemma 11 we know that the largest remaining post-disruption component in a circle of order $(n-1)$ has order $\lceil (n-D_v-1)/D_v \rceil$. It is easy to calculate that $(n-D_v)/(D_v+1) < (n-D_v-1)/D_v \iff D_v < (n-1)/2$. It is also easy to calculate that $[(n-D_v-1)/D_v] - [(n-D_v)/(D_v+1)]$ is an increasing function of n and is larger than 1 for a range of n above a certain threshold.
For $D_v \geq (n-1)/2$, in both the mentioned circle and line, the disruptor can reduce the post-disruption network to a set of isolated components, so that both architectures are equivalent in this extreme case. □

This of course does not imply that the circle network is the best possible network for the network designer to build. What it does show, however, is that it is never good to build a network including all nodes, if you have a limited linking budget. This suggests that there is a tradeoff between, on the one hand, constructing a large pre-disruption component, which still leaves the possibility of a large post-disruption component, but is vulnerable to disruption; and, on the other hand, constructing a small pre-disruption component, which immediately decreases the possible number of nodes in the post-disruption network, but which is more vulnerable to disruption. In the analysis so far, we study this tradeoff only for the decision to leave a single node unconnected. But better architectures may exist where more nodes are left out, enabling the designer to construct a stronger component. Since these cases are hard to characterize, we will only hint at what such networks can possibly look like by means of an example.

We have already shown in the case of the low cost links, that to make his network maximally robust against any attack, the network designer would have to build an r-regular network with $r = D_v + 1 (r = D + l + 1)$ that meets the conditions of Proposition 2. However, due to the limited linking budget in the present case, constructing an r-regular network is only possible when leaving out a considerable number of nodes. For a disruption budget of $D_v = 2$, this would then lead to building a connected, symmetric 3-regular network. Knowing that the linking budget of $B = (n-1)$, and that the network designer needs exactly $3/2 * n$ links to make a network 3-regular, leads to the conclusion that the network designer is forced to leave $1/3 * (n-1)$ nodes out of the connected network to build a 3-regular network. Now taking as an example a network with $n = 25$ nodes and a linking budget of $B = 24$ links that is build in 2 different ways in Figure 5.

The 3-regular network in Figure 5(a), only uses $2/3 * B$ nodes, as has been calculated above. In the network in Figure 5(b), there is exactly one node between any two highly connected nodes. Thus we know that the post-disruption network will consist of 3 nodes less than the pre-disruption network. However, the second factor that we need to take into account is the order of the pre-disruption network. The network in Figure 5(b) uses 20 nodes and the network in Figure 5(a) uses only 16 nodes. Therefore, although less

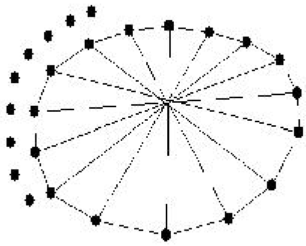

(a) 3-regular

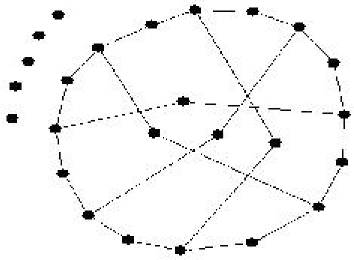

(b) bipartite

Figure 5: Different 25 node networks

damage in terms of disconnected nodes can be done to the network in Figure 5(a), as compared to the network in Figure 5(b), the post-disruption network might still turn out to be smaller than in the larger network. In this particular example, we can see, that the network in Figure 5(b) has the larger post-disruption network, since it contains a connected component of 17 nodes, as compared to 14 in the network in Figure 5(a). Even in this small example with a limited number of nodes only, we can see that there is a definite trade-off between including more nodes in the pre-disruption network and making a stronger network. It can be seen that the 3-regular network has a smaller post-disruption network than the larger and less strong network in Figure 5(b). However, when comparing this to the circle network, we have shown in Lemma 11 above that the largest remaining component will only consist of 11 nodes, which is even smaller than the post-disruption network in the 3-regular network. Thus, a circle network seems to be a too weak pre-disruption network.

Knowing from the example above that the circle can be too weak a pre-disruption network in certain cases, we can also look at the other end of the scale and see if it would make sense for the network designer to build a *max*-proof network instead, thereby basically leaving a maximal amount of nodes out of the component to build a network that then again is fully proof against disruption. From the previous discussion, we know that in a *max*-proof network, we need exactly $(D_v + 1) * n/2$ links. Now if we look at a linking budget of $B = (n - 1)$, we know that for a given D_v we can build a *max*-proof network consisting of $y = 2 * (n - 1)/(D_v + 1)$ nodes. After the network disruption, the largest remaining connected component then includes $\lceil 2 * (n - 1)/(D_v + 1) - D_v \rceil$ nodes. Now if we compare this to the circle network, we can see that the *max*-proof network is actually better than the circle network for large n or small a small disruption budget D_v.

Lemma 12. *For any linking budget $B = (n - 1)$ links, and any disruption budget D_v, a max-proof network weakly dominates the circle network for all $(n - 1) \geq D_v(D_v + 1)$.*

Proof The largest remaining component in a circle network after an attack by a network disruptor with a disruption budget of D_v is $\lceil [(n - 1) - D_v]/D_v \rceil$. The largest remaining component in a *max*-proof network after disruption with a disruption budget of D_v is $\lceil 2 * (n - 1)/(D_v + 1) - D_v \rceil$.

$[2 * (n - 1)/(D_v + 1)] - D_v \geq [(n - 1) - D_v]/D_v$

$\Leftrightarrow [2 * (n - 1)/(D_v + 1)] - [(n - 1)]/D_v \geq D_v - 1$

$\Leftrightarrow [(n - 1)(D_v - 1)]/[D_v(D_v + 1)] \geq D_v - 1$

$\Leftrightarrow (n - 1) \geq D_v(D_v + 1)$ □

It follows that there are cases where, if the only options are the circle or the *max*-proof network, the *max*-proof networks are the better choice. However, the result also suggests that for different cases, the *max*-proof network might be too small as a pre-disruption network. This again implies that there is some tradeoff involved for the network designer between building smaller but stronger and weaker but larger pre-disruption networks. But the circle and the *max*-proof networks are two extreme cases, and it may be better to construct a network that leaves more than one node unconnected, and at the same time does not contain a *max*-proof component. That this indeed is the case in some cases is shown in Appendix 2. We treat this separately in an appendix, because we use some concepts there that are introduced only in Section 6. We have seen in this section that while there is a tradeoff between larger and stronger networks for the case of node deletion, no such tradeoff takes place in the case of link deletion. Therefore, the most robust network topology is quite different for the two cases. In the link deletion case, it is always optimal to include all nodes in the pre-disruption component, whereas in the node deletion case, the network designer is better off leaving out a number of nodes to build a smaller but stronger pre-disruption network. In general it seems that nodes are harder to protect in a network than links, not only because in the node deletion case nodes will be disrupted by definition, but also because all links attached to a node may be removed from the network once the node has been deleted. Therefore it is much harder to keep nodes safe from disruption than to keep links safe.

6. Intermediate Linking Costs

We have so far treated the extreme cases where linking is cheap enough for the network designer to build a *max*-proof network, and where linking is so expensive that the network designer does not want to add any links above the minimum needed to connect all nodes. In this section, we explore the cases in between, where linking costs are intermediate, so that the designer is willing to add defensive links, but not to the extent that the network is protected to the highest level achievable. We limit ourselves to the case where the network designer builds a $(max - 1)$-proof network, meaning that the designer tolerates that the largest component in the post-disruption network has one node less than is possible with maximal network defense. This case is analytically tractable for the following reason. We know from Section 4 which linking budget B is minimally needed to achieve *max*-proofness. Suppose that we take a linking budget smaller than B. Then we know that *max*-proofness is not achievable for this linking budget, so that the best that can be achieved is $(max - 1)$-proofness. Thus, if we find networks for such linking budgets smaller than B that achieve $(max - 1)$-proofness, then these networks do the best possible with this linking budget.

From the previous section, we already know that for a linking budget $B = (n - 1)$, if $D_l = 1$, the network designer achieves $(max - 1)$-proofness in the star architecture; if $D_v = 1$, the network designer achieves $(max - 1)$-proofness in the circle of order $(n - 1)$. These results suggest that in general, under link deletion, the network designer should construct a star or star-like architecture in which the network has a set of one or more strong (i.e. high-degree) nodes which the disruptor cannot remove, and a set of weak (i.e. low-degree) nodes, where the disruptor is able to remove only one of the weak nodes. Under node deletion, the basic results suggest that, in general, the network designer should avoid having nodes with higher degree, as these then become targets to the disruptor; instead, the disruptor should assure that all nodes have the same degree. Our analysis below indeed confirms this intuition.

We focus on linking budgets that are exactly large enough to allow the network designer to build what we now define as a pair r-regular network, where either $r = (D_v + 1)$ or $r = (D_l + 1)$.

Definition 7. *In any pair r-regular network, each pair of neighboring nodes has exactly r links to the rest of the nodes.*

We show below that, for fixed linking budgets that are exactly large enough to build specific pair r-regular networks, such networks are the best that the network designer can do, since they achieve $(max - 1)$-proofness. Unfortunately, we cannot show that these pair r-regular networks are *minimal* $(max - 1)$-proof, so that it remains possible that there are other networks which achieve $(max - 1)$-proofness with less links. Still, each $(max - 1)$-proof pair r-regular is a at least a local optimum to the network designer in the following sense. In any $(max - 1)$-proof network, each connected pair of nodes should have at least $(D_v + 1) = r$ neighbors (respectively $(D_l + 1) = r$ links), connecting it to the rest of the nodes under node deletion (link deletion); otherwise, the disruptor is able to remove the connected pair from the network, so that the network is not $(max - 1)$-proof. In a pair r-regular $(max - 1)$-proof network, this goal is just achieved for every connected pair, so that the network designer is worse off both when adding links to such a pair r-regular network, and when deleting links form a pair r-regular network. It should be noted in this respect that the star is pair $(n-2)$-regular, and that the circle is pair 2-regular.

A further attractive feature of pair r-regular networks is that they allow us to analyze in a tractable manner the following tradeoff faced by the network designer, which is relevant for the case of link deletion. Either the designer can use few links and leave many weak spots in the network, which means many nodes can be deleted from it. An extreme case where this takes place is the star network, which is pair $(n-2)$-regular. Instead, the designer can use many links and leave no weak spots in the network, in the sense that each node is an equally likely target. As we will show below, in any pair r-regular network, each node has one out of a set of at most two degrees, where high-degree nodes can be interpreted as strong spots, and low-degree nodes as weak spots. Moreover, for any r, several pair r-regular networks may exist, varying according to the extent to which these two degrees are different from each other, according to the relative number of strong and weak spots in the network, and according to the number of links used overall. In this manner, we are still able to identify in the set of pair r-regular networks those that achieve $(max-1)$-proofness with a minimal number of links, and are able to interpret this result.

We start by showing that any pair r-regular network is a bipartite network, in which each node has one degree out of a set of at most two degrees.

Lemma 13. *In any connected pair r-regular network, each node has either degree r_1, or degree r_2, where $(r_1+r_2-2)=r$, $r_1 \geq 1$, $r_2 \geq 1$, $r \geq 2$, and nodes with degree r_1 are only linked to nodes with degree r_2.*

Proof Note first that any pair 0-regular network consists of separated components of 2 connected nodes, and any pair 1-regular network consists of separated minimal connected components of 3 connected nodes. It follows that $r \geq 2$. Note further that nodes in a connected pair cannot have degree zero. Consider now a pair r-regular network where in a single connected pair x_1x_2, node x_1 has degree r_1 and node x_2 has degree r_2, such that $(r_1+r_2-2)=r$. Note that the 2 in the latter expression accounts for the direct link of the two nodes, so that r is indeed their number of links to the rest of the nodes. The (r_1-1) links of node x_1 to nodes other than x_2 each form a pair. Take one such node y_1 connected to x_1. In the pair x_1y_1, x_1 has degree r_1, so that y_1 must necessarily have degree r_2. In the same manner, every node connected to y_1 must again have degree r_1. And so on. Similarly, each neighbor of x_2 must have degree r_1, the neighbors of this neighbor must again have degree r_2, and so on. □

Knowing that any connected pair r-regular network is bipartite, we can label the two groups of nodes. By n_1 we denote the number of nodes with degree r_1, and by n_2 the number of nodes with degree r_2, where we label the type-1 and type-2 nodes such that $r_1 \geq r_2$, where $r_2 \geq 1$ (note that for $r_2 = 1$, the only pair r-regular network is the star, which is pair (n - 2) regular. We then call the type-1 nodes *high-degree nodes*, and the type-2 nodes *low-degree nodes*. As shown in the following Lemma, we can now exactly calculate how many high-degree and low-degree nodes there are in a pair r-regular network, and how many links it uses.

Lemma 14. *In any connected pair r-regular network, we have $n_1 = n*[r_2/(r_1+r_2)] = n*[r_2/(r+2)]$ and $n_2 = n*[r_1/(r_1+r_2)] = n*[r_1/(r+2)]$, and the network has exactly $n*[r_1r_2/(r_1+r_2)] = n*[r_1(r+2-r_1)/(r+2)]$ links.*

Proof By Lemma 13, any pair r-regular network is a bipartite graph, with nodes with only two kinds of degrees. It follows that for L, the number of links used in the network, it is the case that $L = n_1r_1 = n_2r_2$. Combining this with the fact that $(n_1+n_2)=n$, and using the fact that $(r_1+r_2-2)=r$, the given expressions for n_1 and n_2 are obtained. These expressions, and the fact that $L=n_1r_1=n_2r_2$ again allow us to calculate that $L = n*[r_1r_2/(r1+r2)] = n*[r_2(r+2-r_2)/(r+2)]$. □

Note that it follows from Lemma 14 that the number of high-degree nodes is smaller the less links are given to the low-degree nodes. Intuitively, by making the distribution of degrees over any pair of nodes more unequal, so that r_2 becomes smaller and r_1 bigger, we will have fewer high-degree nodes and more low-degree nodes. This can be seen in the extreme case of the star network, where we have only one core node and $(n-1)$ peripheral nodes with $r_2 = 1$ and $r_1 = (n-1)$.

Given that we have obtained from Lemma 14 numbers for the high- and low-degree nodes we should have in any pair r-regular network, we know that these should be integer numbers. We focus only on n such that these are indeed integer numbers. We are then ready to show the existence of pair r-regular networks.

Lemma 15. *Consider a number $r \geq (n-2)$, and consider two numbers $r_1 \geq 1$, $r_2 \geq 1$, $r_1 \geq r_2$ such that $(r_1 + r_2 - 2) = r$. Let $n_1 = n * [r_2/(r_1 + r_2)]$ and $n_2 = n * [r_1/(r_1 + r_2)]$ be integer numbers. Then the set of pair r-regular networks is not empty.*

Proof Construct a circle of $2 * n_1$ nodes, namely n_1 type 1 nodes and n_2 type 2 nodes. In the circle every type 1 node is connected to type 2 nodes, and vice versa. Next, add the appropriate number of links to each pair on the circle, namely $(r_1 - 2)$ links for each of the type 1s, and $(r_2 - 2)$ for each of the type 2 nodes. These nodes added to the circle make a total of $n_1(r_1 + r_2 - 4)$ links connected on one side of the link to nodes on the circle, and the other side of which still needs to be connected to other nodes, on or off the circle. Also, $(n_1 + n_2 - 2n_1) = (n_2 - n_1)$ nodes still need to be connected to the network. These remaining nodes are all type 2 nodes. All the $n_1(r_2 - 2)$ links of the type 2 nodes on the circle need to be to a type 1 node on the circle, as type 1 nodes by assumption only lie on the circle. We can let each of the n_1 nodes on the circle receive $n_1 * [(r_2 - 2)/n_1] = (r_2 - 2)$ such links from type 2 nodes on the circle. For $r_1 > r_2$, all of the remaining $[n_1 * (r_1 - 2) - n_1 * (r_2 - 2)] = n_1(r_1 - r_2)$ links of the type 1 nodes on the circle need to be to type 2 nodes not on the circle; each such type 2 node not on the circle needs to receive links from exactly r_2 of the type 1 nodes, each type 1 node needs to have $(r_1 - r_2)$ such links. This means that there must be exactly $n_1(r_1 - r_2)/r_2$ type 2 nodes that do not lie on the circle. We therefore have $n_2 = n_1 + n_1 * [(r_1 - r_2)/r_2] = n_1 r_1/r_2$ type 2 nodes, and n_1 type 1 nodes. Also, $n_2 = n_1 r_1/r_2 \Leftrightarrow n = n_1(r_1 + r_2)/r_2 \Leftrightarrow n_1 = n * [r_2/(r_1 + r_2)]$, in accordance with Lemma 14. □

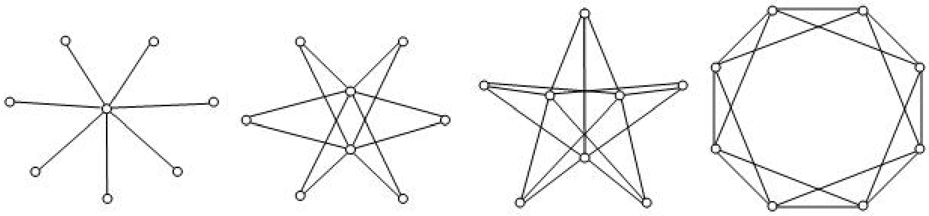

(a) 1 core node (b) 2 core nodes (c) 3 core nodes (d) 4 regular

Figure 6: Pair 6-regular Networks

Figure 6 represents four pair 6-regular networks for the case where $n = 8$, $r = 6$. We next show that the budgets described above that are just sufficient to build pair r-regular, do not allow *max*-proofness to be achieved for the case where $D_l = (r-1)$, respectively $D_v = (r-1)$. It follows that, if we can show that a pair r-regular network is $(max-1)$-proof, then for the number of links that is used in the given network, the network is the best that the network designer can do. The simple case where $r = 2$ and $D_l = 1$, respectively $D_v = 1$ is an exception, as the pair 2-regular network (namely the circle) is also *max*-proof in this case - this is why we exclude it. Further, we show that for sufficiently large n, the network designer cannot build an r-regular component connecting $(n-1)$ nodes. It follows then that any best response pre-disruption network for these linking budgets must be connected. Note that for the case in Figure 6, the connected *max*-proof network uses 24 links, and the *max*-proof component connecting 7 nodes uses 21 links. It can be checked that this is more than any of the number of links used in the four represented networks.

Lemma 16. *Let it not be the case that both $r_1 = 2$ and $r_2 = 2$. Then with any number of links $L = n * [r_1 r_2/(r_1 + r_2)]$ that exactly allows one to build a pair r-regular network, one cannot build a connected r-regular network. Moreover, a critical n_C exists such that, for all $n > n_C$, the linking budget L is also too small to construct an r-regular component connecting $(n-1)$ nodes.*

Proof The r-regular connected network uses more links than the pair r-regular network iff $n * r/2 > n * [r_1 r_2/(r_1 + r_2)] \Leftrightarrow (r_1 + r_2 - 2)(r_1 + r_2) > 2 * r_1 r_2 \Leftrightarrow r_1^2 + r_2^2 > 2(r_1 + r_2)$. The latter is true for the specified r. The r-regular component connecting $(n-1)$ nodes uses more links than the pair r-regular network iff $(n-1) * r/2 > n * [r_1 r_2/(r_1 + r_2)]$. For large n, this is true by the same calculations. □

Given Lemma 16, we know that for each pair r-regular network, if the linking budget needed to construct it is available to the designer, there is a potential case where it is optimal to construct a $(max - 1)$-proof network. Moreover, given Lemma 14, we know that for each r, there may be several pair r-regular networks, differing according to how the r links that connect any connected pair are distributed over the two nodes in that pair. As they each potentially achieve the same maximally attainable goal of $(max - 1)$-proofness, we are interested in which networks achieve this with the smallest number of links. We now show that the more asymmetric the distribution of links in any given connected pair is, the less links the network uses overall.

Lemma 17. *Consider the set of connected pair r-regular networks, where n is assumed such that for all $r_1 \geq r_2$ (with $r_1 + r_2 - 2 = r$), it is the case that $n * [r_2/(r_1 + r_2)]$ and $n * [r_1/(r_1 + r_2)]$ are integer numbers. Then in this set, pair r-regular networks have less links the smaller their r_2, and the pair r-regular networks with $r_2 = 2$ have the smallest number of links in this set.*

Proof A pair r-regular network with $r_2 = 1, r_1 = (r+1)$ is only possible with the star architecture, and is pair (n - 2)- regular. This case is here excluded given that $r < (n-2)$. The smallest possible r_2 is then $r_2 = 2$. In the expression $L = n * [r_2(r + 2 - r_2)/(r+2)]$ derived in the proof of Lemma 14, the number of links used is smaller the smaller r_2. The result follows. □

Lemma 17 suggests that the network designer who decides to construct a pair r-regular network with the purpose of achieving $(max - 1)$-proofness, in order to save as much as possible on links, should distribute links across pairs as unequally as possible. In Figure 6 (case where $n = 8$, $r = 6$), network (a) uses 7 links, network (b) 12 links , network (c) 15 links, and network (d) 16 links. However, as we now show, a distribution that is too unequal makes it possible for the disruptor to remove multiple low-degree nodes from the network. This puts a cap on how unequal the distribution may be.

Lemma 18. *A necessary condition for a pair r-regular network to be $(max - 1)$ proof under link deletion with $D_l = (r-1)$, respectively node deletion with $D_v = (r-1)$, is that $r_2 > (r-1)/2 \Leftrightarrow r_2 > (r_1 - 3)$.*

Proof If $r_2 \leq (r-1)/2$, the disruptor is able to delete several nodes with degree r_2, by either deleting all their links in the case of link deletion, or all their neighbors in case of node deletion. □

Illustrating Lemma 18, in network (b) in Figure 6, if either $D_l = 5$ or $D_v = 5$, the disruptor is able to take out at least two low-degree nodes. Moreover, for node deletion, an additional danger of an unequal link distribution is that the disruptor can do a lot of damage by targeting the high-degree nodes. By the same intuition, when the linking budget is equal to $(n-1)$, the star does very bad under node deletion. For the case of link deletion see the section on high linking costs.

Lemma 19. *A necessary condition for a pair r-regular network to be $(n - D_v - 1)$-proof under node deletion with $D_v = (r-1)$ is that $r_2 > (r_1 - 2)$.*

Proof Consider two neighbors of a type-1 node x_1. By definition, these two neighbors are type-2 nodes. Each of them have $(r_2 - 1)$ type-1 neighbors other than x_1. If $[2 * (r_2 - 1) + 1] \leq D_v = (r-1)$, then by taking all the $[2(r_2 - 1) + 1]$ type-1 neighbors of the two mentioned type-2 nodes out, the disruptor can take out two extra nodes. □

Illustrating Lemma 19, in network (c) in Figure 6, under node deletion, by deleting the three high-degree nodes, the disruptor can make sure that all nodes are isolated.

We are now ready to summarize our results about pair r-regular networks that achieve $(max-1)$-proofness, and do so with a minimal number of links. Our main intuition is that under link deletion, links can to a limited extent be unequally distributed over the two sides of any connected pair, saving links. Under node deletion, links should not be unequally distributed. For this reason, we focus on the case where r is even; for r is odd, there is a divisibility problem in that links simply can not be equally distributed over the two sides of a connected pair.

Proposition 7. *Let* r *be even, where* $r \geq 4$. *For any* $r_1 \geq 1, r_2 \geq 1$ *such that* $r = (r_1 + r_2 - 2)$, *let* $n * [r_2/(r_1+r_2)]$ *and* $n * [r_1/(r_1+r_2)]$ *be integer numbers. Then under link deletion,* $(max-1)$*-proofness can be achieved with a minimal number of links when* $r_1 = r/2+2, r_2 = r/2$, *so that the type-1 nodes have two extra links compared to type-1 nodes. Under node deletion,* $(max-1)$*-proofness can be achieved with a minimal number of links when* $r_1 = r_2 = (r+2)/2$, *so that all nodes have the same degree, meaning that we have a* $(r+2)/2$*-regular network.*

Proof This follows directly from Lemmata 17, Lemma 18 and Lemma 19. □

Proposition 7 is illustrated by Figure 6. Networks (a) and (b) are not $(max-1)$-proof under link deletion. Both networks (c) and (d) are $(max-1)$-proof, but network (d) uses more links. Note that in network (c), the disruptor will always target a low-degree node, as the disruptor then even does not need to use up his entire disruption budget. In this sense, low-degree nodes can be considered as weak spots. The way in which network (c) is represented in Figure 6 emphasizes that this is a star-like architecture. In the simple case where $D_l = 1$, the star is minimal $(max-1)$ proof. This is achieved by making every non-central node a possible target for removal from the network. This saves links, and at the same time the disruptor is only able to remove a single node. The same intuition applies to network (c). There are multiple weak spots, so that the designer saves links, and at the same time the disruptor is only able to remove a single weak spot. Under node deletion, only network (d) is $(max-1)$-proof. It can be checked that the largest post-disruption component in network (d) has order $(n - D_v - 1) = (8-5-1) = 2$. This in turn is a circle-like structure, in that it is completely symmetric. Under node deletion, the designer should not make central, high-degree nodes as these will then be targeted.

7. Conclusion

Summarizing our results, when linking costs are low, the network designer protects his network by constructing a regular network, where all nodes are equally well protected. Multiple architectures meet this requirement, but the designer should take care to avoid that the network has local cliques connected by few links, as such networks are easy to disrupt. This requirement is more restrictive for node deletion than for link deletion. The set of best-response networks under node deletion is a subset of the set of best-response networks under link deletion. A class of easily-described best-response networks is the class of symmetric networks, consisting of a circle encompassing all nodes, plus added links. Intuitively, circles plus added links assure that there are alternative paths between nodes, so that if one path is disrupted, players are still connected by an alternative path.

When linking costs are high, contrary to what is the case for low linking costs, the best-response architectures under link and node deletion look fundamentally different. Under link deletion, it is a best response to connect all nodes in a star network. In such a network, the disruptor can only take out one node for each disrupted link. Intuitively, the star network keeps all nodes as close as possible to one another, in such a manner that the disruption of one link cannot disconnect several nodes. Under node deletion, the star network is on the contrary a very bad network, since deletion of the central node disconnects all nodes. In general, any minimal connected network is a bad response, as it can easily be cut up into components. Instead, it is a best response to leave some nodes out of the network, and build a smaller and stronger component. Such components again should not involve local cliques connected by few links, and examples of good components are again symmetric, consisting of a circle with added links. It follows that, under node deletion, network disruption causes a further inefficiency, where the designer connects less links than is optimal in the absence of network disruption.

For intermediate linking costs, the network designer finds it too expensive to fully protect all nodes. Our analysis suggests that, both under link and node deletion, the network designer constructs connected networks, i.e. does not leave nodes unconnected to construct smaller and stronger networks. But otherwise, the intuitions for high-linking costs are confirmed. Our analysis suggests that under link deletion, star-like networks should be constructed, consisting of low-degree, weak, nodes, and high-degree, strong, nodes. Just as in the star, there are multiple weak nodes, but only one can be disrupted. Also, the diameter of the network is kept small so that deletion of links cannot cause large parts of the network do be disconnected. Under node deletion, whenever possible, all nodes get the same degree. This is because high-degree nodes are not strong as is the case under link deletion, but would on the contrary be likely targets for disruption. This means again that the network designer does well by constructing a symmetric network.

We end by exploring possibilities for future research, where the key question is to extend our present approach of a network designer to a multi-player game, where the nodes in the network are actual players. Let us start by looking at agent's incentives to form links, independently form the presence of a disruptor. In a more realistic model, there may be information decay, where information is worth less the larger the distance it traveled in the network (Jackson and Wolinsky (1996) and Bala and Goyal (2000)). From the perspective of information sharing, it is efficient for nodes to be as close to one another as possible, as is the case in the star; in equilibrium, players also have the tendency to connect to a central node, such that the star is likely to arise. As our analysis shows, at least for high linking costs, the star is also efficient under link deletion. However, it is a bad network under node deletion. Further, players' incentives to link to certain nodes may not only depend on the information obtained from those nodes, but may also depend on players preferences. A well-known phenomenon in sociology is homophily, where in networks, birds of a feather flock together (McPherson et al., 2001). As shown in our analysis, such preferences are in direct conflict with efficient defense against network disruption, as a deletion of a few links or nodes may then cause great damage to the network. Further, players may also directly take into account network defense and network disruption when deciding on which links to form. In one type of network formation extension of our model, we could assume that players dislike being removed from a network. An example would be a member of in an illegal network who does not want to be arrested. In some of our results, it is efficient to leave weak spots in the network, which are then more likely to be removed from the network. But individual players may not want to be at such weak spots then. In another type of network extension of our model, players may on the contrary like to be at vulnerable positions in a network. If a firm defects to a competing alliance, then this need not make the firm worse off. In its present network, each firm may try to manoeuvre itself in a crucial position, in order to have larger bargaining power in its network. This follows Burt (1992) argument that individual may gain advantage by bridging structural holes in networks, thus assuring that information exchange takes place between different groups (for recent game-theoretic translations of this argument, see Goyal and Vega-Redondo (2007) and Kleinberg et al. (2008)). In our argument, this strong position as such does not relate to the bridging function, but to the fact that a disruptor is willing to give the such a bridge player a large payment for defecting, creating a very good outside option for the player, and increasing his bargaining power in his present network. This does not mean, however, that in equilibrium such bridging positions may ever arise, as every players seeks to obtain them.

Graph-Theoretic Appendix

Graph Theoretic Lemmata and Definitions:

Definition I. *A* star network *has a central node i, such that $g_{ij} = 1$ for all $j \in N \setminus i$ and no other links.*

Definition II. *A* chain graph *is an alternating sequence of nodes and links, which begins and ends with a node and each link is incident with exactly two nodes.*

Definition III. *A graph where each node is connected exactly of degree 2, is called a* circle graph.

Definition IV. *An* end node *is a node that is connected exactly of degree $\eta_i(g) = 1$.*

Lemma A.1. *The only connected 2-regular architecture is the circle.*

Proof A 2-regular network cannot contain any end nodes, as these have degree 1. Alternatively, construct a 2-regular network step by step. Start with one node. This node should have two neighbors. These two neighbors should each have an extra neighbor. Continue this procedure until one node remains to be added, this leads to the construction of a line of $(n-1)$ nodes. The network can only be made 2-regular by connecting the 2 extreme nodes of this line, resulting in a circle. □

Definition V. *A graph G is called* bipartite*, if it is possible to divide the node set into two sets, N_1 and N_2, where each link connects a node of subset N_1 with a node of subset N_2 and no two nodes of the same set are directly linked.*

Lemma A.2. *Every connected graph contains at least $(n-1)$ links.*

Proof Build up a network step by step. Start with one node. Connect another node to it, and so on (note that any network can be constructed in such a manner). For each node you connect, you need at least one link. □

Definition VI. *Each graph that uses exactly $(n-1)$ links to connect n nodes, is called* minimally connected.

Lemma A.3. *Every graph that contains a circle uses at least n links.*

Proof Suppose that a network contains a circle of x nodes. Consider this circle in isolation. It uses at least x links. Now add the further $(n-x)$ nodes, one by one. We need at least $(n-x)$ links to do this. So to build a graph that contains all n nodes and a circle, at least n links are needed. □

Corollary A.1. *Minimally connected graphs do not contain any circles.*

Proof This follows from Definition VI and Lemma A.3. □

Definition VII. *A link (ij) in a connected graph G, is called a* bridge link*, if G_{-ij} is disconnected. Similarly, every set of links L is called a bridge link set, if G_{-L} is disconnected.*

Definition VIII. *A node i in a connected graph G, is called a* bridge node*, if G_{-i} is disconnected. Similarly, every set of nodes V is called a bridge node set, if G_{-V} is disconnected.*

Lemma A.4. *In every minimally connected graph, every link is a bridge link, and every node is a bridge node.*

Proof By Lemma A.2, we know that every minimally connected graph does not contain a circle. Thus the removal of any link or node will disconnect the graph. □

Lemma A.5. *In every minimally connected graph, there are at least two end-nodes.*

Proof We know by Lemma A.2 that any minimally connected network does not contain a circle. Knowing that no circle is contained in any minimally connected network, there have to be at least 2 end-nodes in the network, since without a circle at least two nodes need to be connected of degree $\eta_i(g) = 1$, which are end-nodes by definition. □

Lemma A.6. *Every minimally connected graph that is not a star has at least two nodes with degree larger than 1.*

Proof Any star has $(n-1)$ nodes with degree 1. In every minimally connected graph, each node has degree at least 1. It follows that in every non-star minimally connected network, at most two nodes have degree 1. It follows that at least two nodes have degree larger than 1. □

Necessary conditions on the existence of an r-regular network are:

Lemma A.7. *A necessary condition for existence of an r-regular network is that n and/or r is an even number, where the r-regular network then has exactly* $(n*r)/2$ *links.*

Proof As each node receives exactly r links, and since each link is shared by exactly two nodes, the total number of links in any r-regular network is $(n*r)/2$. It follows that an r-regular network only exists if n and/or r is even. □

Appendix 2

Addition to Section on Node Deletion with High Linking Costs

The difference between the circle network and the *max*-proof network, is quite extreme. In Section 6, it is shown how a *(max* - 1)-proof network can be constructed in the form of a pair r-regular network. Such an architecture can also be used to make a smaller but stronger component in the case of node deletion and high linking costs. As shown in Lemma 16, for a number of nodes γ, a pair r-regular network $L = \gamma * [r_1 r_2/(r_1+r_2)]$ links, where by Proposition 7, $r_1 = r_2 = (r+2)/2$. It follows that $L = (\gamma*(r+2))/4$. Given that the linking budget is $(n-1)$, it follows that the pair r-regular network has a number of nodes $\gamma = [4*(n-1)]/(r+2) = [4*(n-1)]/(Dv+3)$. We can now show that building a *max*-proof network is never the best option, and is either dominated by the circle network or by the *(max* - 1)-proof network.

Lemma A.8. *For a linking budget of* $B = (n-1)$ *links, where* $(n-1) \geq D_v(D_v+1)$ *and a disruption budget of* D_v*, the* $(max-1)$*-proof network strictly dominates the max-proof network.*

Proof The largest remaining component in a *max*-proof network, after an attack by a network disruptor with a disruption budget of D_v will be $\lceil 2*(n-1)/(D_v+1) - D_v \rceil$. The largest remaining component in a *max-1*-proof network, after an attack by a network disruptor with a disruption budget of D_v will be $\lceil 4*(n-1)/(D_v+3) - (D_v+1) \rceil$ [17].
$4*(n-1)/(D_v+3) - (D_v+1) \geq 2*(n-1)/(D_v+1) - D_v$
$\Leftrightarrow 4*(n-1)/(D_v+3) - 2*(n-1)/(D_v+1) \geq 1$
$\Leftrightarrow (n-1) \geq (D_v+3)(D_v+1)/2(D_v-1)$
Thus what remains to be shown is that this is indeed larger than $D_v(D_v+1)$.
$(D_v+3)(D_v+1)/2(D_v-1) > D_v(D_v+1)$
$\Leftrightarrow (D_v+3)/2(D_v-1) > D_v$
$\Leftrightarrow (D_v+3)/2 > D_v^2 - D_v$
$\Leftrightarrow (D_v+1)/D_v^2 > 2/3$ This holds for all $D_v > 0$, since D_v by definition is a natural number. □

[17] We show this here explicitly for an even r, however, the same holds for odd r

Bibliography

References

Aghion, P. and Bolton, P. (1987). Contracts as a barrier to entry. *The American Economic Review*, 77(3):388–401.

Albert, R., Jeong, H., and Barabasi, A.-L. (2000). Error and attack tolerance of complex networks. *Nature*, 406(6794):378–382. M3: 10.1038/35019019; 10.1038/35019019.

Baburajan, R. (2010). Raytheon bbn technologies demos disruption-tolerant network. *TMCnet.com*, 05 Feb 2010:04 June 2010.

Baccara, M. and Bar-Isaac, H. (2008). How to organize crime. *The Review of Economic Studies*, 75(4):1039–1067.

Bala, V. and Goyal, S. (2000). A noncooperative model of network formation. *Econometrica*, 68(5):1181–1229.

Ballester, C., Calvo-Armengol, A., and Zenou, Y. (2006). Who's who in networks. wanted: The key player. *Econometrica*, 74(5):1403–1417.

Bier, V., Oliveros, S., and Samuelson, L. (2006). Choosing what to protect: Strategic defensive allocation against an unknown attacker. *Journal of Public Economic Theory*, 9:1–25.

Billand, P., Bravard, C., Chakrabarti, S., and Sarangi, S. (2009). Corporate espionage. *Discussion Paper 854, DIW Berlin*.

Bloch, F. (2002). Coalitions and networks in industrial organization. *The Manchester School*, 70(1):36–55.

Bollobàs, B. and Riordan, O. (2003). Robustness and vulnerability of scale-free random graphs. *Internet Mathematics*, 1(1):1–35.

Brevoord, K. and Marvel, H. P. (2004). *Successful Monopolization Through Predation: The National Cash Register Company*, pages 85–126. Antitrust Law and Economics, Research in Law and Economics. Elsevier, Amsterdam, vol. 21 edition.

Burt, R. S. (1992). *Structural Holes*. Harvard University Press, Cambridge.

Chartrand, G. (1977). *Introductory Graph Theory*. Dover Publication, New York.

Dekker, A. and Colbert, B. (2004). Network robustness and graph topology. In Estivill-Castro, V., editor, *27th Australasian Computer Science Conference*, volume 26, pages 359–368.

Diestel, R. (2005). *Graph Theory*. Springer Verlag, Heidelberg.

Enders, W. and Jindapon, P. (2010). Network externalities and the structure of terror networks. *Journal of Conflict Resolution*, 54(2):262–280.

Enders, W. and Su, X. (2007). Rational terrorists and optimal network structure. *Journal of Conflict Resolution*, 51(1):33–57.

Francois, J. (2006). Firms behaving badly: International regulations of anticompetitive behavior. *Working Paper, Tinbergen Institute*.

Frank, A. (1995). *Connectivity and Network Flows*, pages 111–178. Handbook of Combinatorics. Elsevier Science B.V., Amsterdam, 1 edition.

Galeotti, A., Goyal, S., and Kamphorst, J. (2006). Network formation with heterogeneous players. *Games and Economic Behavior*, 54(2):353–372.

Goyal, S. (2009). *Connections: An Introduction to the Economics of Networks*. Princeton University Press, Princeton.

Goyal, S. and Vega-Redondo, F. (2007). Structural holes in social networks. *Journal of Economic Theory*, 137(1):460–492.

Goyal, S. and Vigier, A. (2009). Robust networks. *Working Paper, Cambridge University*.

Hendricks, K., Piccione, M., and Tan, G. (1999). Equilibria in networks. *Econometrica*, 67(6):1407–1434.

Hong, S. (2008). Hacking-proofness and stability in a model of information security networks. *Working Paper, Vanderbilt University*.

Hong, S. (2009). Enhancing transportation security against terrorist attacks. *Working Paper, Vanderbilt University*.

Jackson, M. O. (2008). *Social and Economic Networks*. Princeton University Press, Princeton.

Jackson, M. O. and Wolinsky, A. (1996). A strategic model of social and economic networks. *Journal of Economic Theory*, 71(1):44–74.

Katz, M. L. and Shapiro, C. (1985). Network externalities, competition, and compatibility. *The American Economic Review*, 75(3):424–440.

Kim, J.-H. (2007). Employee poaching, predatory hiring, and covenants not to compete.

Kleinberg, J., Suri, S., Tardos, E., and Wexler, T. (2008). Strategic network formation with structural holes. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 284–293.

Kranton, R. E. and Minehart, D. F. (2001). A theory of buyer-seller networks. *The American Economic Review*, 91(3):485–508.

Lipsey, R. A. (2006). Network warfare operations: Unleashing the potential. *Mimeo Center for Strategy and Technology, Air War College, Air University*.

Logan, J. (2009). Lowering the bar or setting the standard? deutsche telekom's u.s. labor practices. Technical report, American Rights at Work Education Fund.

McPherson, M., Smith-Lovin, L., and Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27:415–444.

Menger, K. (1927). Zur allgemeinen kurventheorie. *Fundamenta Mathematicae*, 10:96–115.

Posner, E. A., Spier, K. E., and Vermeule, A. (2009). Divide and conquer. *Working Paper*.

Salinger, M. A. (1988). Vertical mergers and market foreclosure. *The Quarterly Journal of Economics*, 103(2):345–356.

Salop, S. C. and Scheffman, D. T. (1983). Raising rivals' costs. *The American Economic Review*, 73(2, Papers and Proceedings of the Ninety-Fifth Annual Meeting of the American Economic Association):267–271.

Schelling, T. (1960). *The Strategy of Conflict*. Cambridge.

Taylor, M., Sekhar, S., and D'Este, G. (2006). Application of accessibility based methods for vulnerability analysis of strategic road networks. *Networks and Spatial Economics*, 6(3):267–291. M3: 10.1007/s11067-006-9284-9.

van der Leij, M. (2003). Competing transport networks. *Discussion Paper, Tinbergen Institute*.