Budapest University of Technology and Economics
Department of Telecommunications and Media Informatics

The Network Forming Effects of Routing Policy

Dávid Szabó

PhD Dissertation

Advisor:

András Gulyás, Ph.D. (Assistant Professor)

*Dept. of Telecommunications and Media Informatics and
Hungarian Academy of Science Future Internet Research Group,
Budapest University of Technology and Economics*

Budapest, Hungary
2016.

...to my parents.

Abstract

Knowing more about the topology formation of complex networks – like social networks, biological networks or the Internet – is of great importance. Such knowledge can help us gain insight into the decision making process of individuals in the network and understand the principles that govern the emerging topological features. However, despite the large existing literature from multidisciplinary research areas, our knowledge is still limited and incomplete. In general, we have valuable information about network dynamics from multiple aspects, but still the *topological footprints inflicted by routing policies, which is one of the most determining factors on the resulting network, is an unacquainted aspect.*

I aim to fill this gap by investigating the topology formation effects of different routing policies connected to complex networks. In order to do this I design analytical game theoretical models in which the policy rules exist in their pure and realistic form, carefully separated from other factors that could confuse the analytic inference. Throughout the analysis I identify the equilibrium states that emerge as a direct consequence of the used routing policies.

In this dissertation two fundamental routing policies are discussed. First I focus on the *Border Gateway Protocol* that encompasses the inter-domain routing policies on the Internet autonomous system level topology, and within that the two most important rules called *valley-free* and *highest local preference* rule. After that I broaden the context from the Internet specific routing policies to investigate *greedy navigation*, the most accepted routing policy of general complex networks. Finally, I give an overview about the applicability of the presented results.

Kivonat

A komplex hálózatok – mint a közösségi hálózatok, biológiai hálózatok vagy az Internet – topológia formálódásával kapcsolatos ismeretek nagy fontossággal bírnak. Az ilyen jellegű tudás betekintést enged a hálózaton belüli szereplők döntéshozási folyamataiba, továbbá segít megérteni azokat az elveket, amelyek a kialakuló topologikus tuljadonságokat megszabják. Azonban a hatalmas - multidiszciplináris kutatási területekről érkező - szakirodalom dacára jelenlegi tudásunk továbbra is korlátozott, nem teljes. Általánosságban elmondható, hogy értékes eredmények születtek a hálózatok dinamikájáról, méghozzá különböző aspektusok mentén, viszont az útvonalválasztási irányelvek topológikus lenyomatainak vizsgálata (amely a hálózatformálódás egyik legmeghatározóbb tényezője) továbbra is meglehetősen hiányos terület.

Célom ennek a hiánynak a pótlása, azáltal, hogy különböző útvonalválasztási irányelvek topológia formálódásra gyakorolt hatását visgálom. Ehhez olyan analitikus játékelméleti modelleket definiálok, melyek az irányelveket, szabályokat tisztán, realisztikus formában tartalmazzák, óvatosan elkülönítve azokat olyan más tényezőktől, amelyek megzavarhatja az analízist. A vizsgálatok során azonosítom az egyensúlyi állapotokat, amelyek az egyes útvonalválasztási irányelvek közvetlen következményeként foghatók fel.

A disszertációban két alapvető útvonalválasztási irányelv kerül tárgyalásra: először a Border Gateway Protocol-ra fókuszálok, amely az Internet autonóm hálózatai közötti útvonalválasztási elveket tartalmazza, azon belül pedig a két legfontosabb szabályra a valley-free és a highest local preference szabályra. Ezt követően szélesítem a kontextust és az Internet-specifikus útvonalválasztási irányelvekről áttérek az ún. mohó navigációra, ami jelenleg a komplex hálózatok legáltalánosabban elfogadott útvonalválasztási irányelve. Végül egy általános áttekintést adok a disszertációban tárgyalt eredmények alkalmazhatóságáról.

Acknowledgements

First of all, I would like to thank my supervisor, András Gulyás for all the time and energy he has put into supervising and guiding my research. His patience, encouragement and support were indispensable to becoming a researcher in the field of telecommunications. He taught me a lot about the ways of thinking, developing and presenting. He always gave me directions and thoughts when I felt myself lost.

I am also indebted to Zalán Heszberger, my supervisor during my MSc. years and who first advised me to start a PhD. Moreover, I must thank all of my co-authors, from whom I have learnt how to turn mere ideas into scientific results. I owe Attila Kőrösi a debt of gratitude for his selfless assistance with the mathematical formulations throughout this dissertation.

My work was carried out in the Department of Telecommunications and Media Informatics at Budapest University of Technology and Economics and in the High Speed Networks Laboratory (HSNLab), therefore I would like to thank to József Bíró, Erzsébet Győri, Attila Vidács and Gábor Magyar for helping through the administrative process and providing sufficient funding during my PhD years. I also want to thank my roommates at the department for the good atmosphere.

Last but not least, I would like to thank my parents for all the support they have provided me over the years and made it possible to me to concentrate on my studies. Without my great family, this accomplishment would not have been possible.

Contents

Chapter 1

Introduction

The spreading of computers and communication networks drove the network research towards new directions in the last decades. The new infrastructure enabled to gather and analyze data on a scale far larger than previously was possible. This change shifted the focus from the analysis of small networks - and the properties of individual connections within them - to considerating large-scale statistical properties and placed complex networks in the center of interest. This means even though studies used to consider networks with just up to hundreds of nodes, nowadays it is not uncommon to see networks with millions of nodes.

Knowing more about real world networks has great importance as a wide range of networks belong to this group such as the Internet, social networks or biological networks. However, the investigation of these networks requires not only to change the set of reasonable questions but also the used methods. Nowadays research can be described along three main activities: (*i*) **find and highlight macroscopical statistical properties** such as degree distribution, path length or clustering that characterize the structure of the networks and suggest ways to measure these properties, (*ii*) **create models** that enables us to understand the meaning of these properties and how they influence each other, and (*iii*) **predict the behavior** of the networks based on these metrics.

The thing these approaches have in common is that all are tightly connected to the topology of the network. This is not a surprise at all as the topology directly affects how a network sustains to its function, which ultimately breaks down to how effectively the nodes in the network can communicate with each other. A trivial

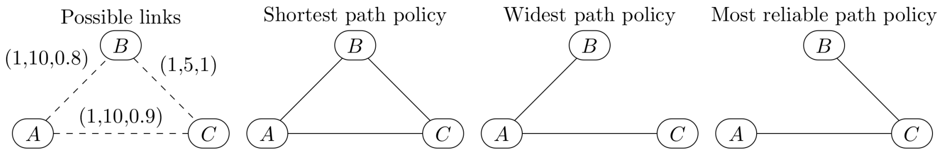

Figure 1.1: Policy drives topology: We have three nodes (A, B, C) over which we decide to build a network. The characteristics of the possible edges (A, B), (A, C), (B, C) are defined by a triplet: latency, bandwidth and reliability (l, b, r). Now we have to decide what topology should be built. Clearly, if the routing policy is the *shortest path* policy, then all the links have to be used. If the policy is the *widest path*, meaning that the highest bandwidth is preferred between the nodes, then the third topology from the left is the appropriate choice. Finally, if the policy is the *most reliable path* policy then the corresponding network will be the rightmost topology.

example for this topological influence is that information can flow only along the existing connections among nodes. On the other hand, if we consider the network as a constantly altering dynamic system – which is true for all real networks – then the way of communication also has a crucial effect on the topology. To make this statement clear, let us define a bit more precisely what we mean by "the way of communication". In every network where there is communication the preferred routes for the information flows emerge eventually. The preferences that drive this process can be described by a set of rules what we collectively call *routing policy*. The routing policy naturally dictates that only those connections are created and kept that are remunerative for efficient communication. To illustrate of this idea consider the simple example in Fig. 1.1.

Based on this mutual correspondence between the network topology and the used routing policy we can say that understanding the topological footprint of a routing policy can give valuable insight into the topological properties of a real network. This insight is somewhat complementary to the existing results as the usual approaches do not really consider routing policies or assume shortest path communication which is a highly simplified view. In this dissertation two fundamental routing policies are discussed: (i) the **Border Gateway Protocol (BGP)** [49], that encompasses the inter-domain routing policies on the Internet autonomous system (AS) level topology and within that the two most important rules called valley-free (VF) and highest local preference (HLP) rule, and (ii) the **greedy routing** policy, that is a promising candidate as the navigation method of natural complex networks [14].

In Chapter 2 I define a game theoretical model that enables to investigate the topology forming effects of the BGP routing policy in an analytical way. As a result I show a subgraph that is present in the topology and can be considered as the natural consequence of the VF rule. As the next step I extend the model by adding the HLP rule and I describe a further refined version of the subgraph obtained from the first model, that I call the *Spiderweb graph*. In Chapter 3 I show the existence of such a subgraph in the Internet's AS level topology through measurements and design a topology generator, based on the game theoretical results. I show that although the model is very simple still the generated topologies reflect the features of the AS level topology along several metrics. I also provide a comprehensive comparison with existing topology generators. In Chapter 4 I broaden the context from the Internet specific routing policies to investigate the greedy navigation that is the most accepted routing policy of general complex networks. For the analysis I create a model that describes the communication more realistically than previous game theoretical models, namely, instead of *shortest path* nodes I use *greedy path* forwarding. I prove that in the Euclidean space the emergence of topologies on which we understand the principles of greedy navigation, and which also inspire many algorithms, cannot be justified through the interaction of rational, selfish players. I also present a brief outlook on how the situation changes if the hyperbolic space is used instead of the Euclidean. Finally, the applicability of the results is discussed in Chapter 5.

1.1 Definitions and Notations

In this chapter I give an overview of the important definitions and notations from the field of graph theory, game theory and hyperbolic geometry that are used throughout the dissertation. It is not my intention to cover any of the following topics in full-depth, the list of definitions and examples are merely the most essential ones.

1.1.1 Graph Theory

Graph theory is a natural framework for the exact mathematical treatment of networks, since networks can be easily represented by grahps. Network science uses graphs extensively for the research of complex networks which led to several new definitions in the last decades. The most important ones necessary for understanding the dissertation are listed in the following table.

Basic Definitions	
graph	A *graph* $G = (\mathcal{V}, \mathcal{E})$ consists of two sets $\mathcal{V}$ and $\mathcal{E}$, where $\mathcal{V} = \{n_1, n_2, ..., n_N\}$ and $\mathcal{E} = \{e_1, e_2, ..., e_K\}$ refer to the nodes (vertices) and links (edges) of the network, respectively. A node is referred to by its order i in the set of $\mathcal{V}$. In a graph each edge is defined by a pair of vertices i and j and denoted as (i, j) or e_{ij} (in a directed graph $e_{ij} \neq e_{ji}$).
subgraph	A *subgraph* $G' = (\mathcal{V}', \mathcal{E}')$ of $G = (\mathcal{V}, \mathcal{E})$ is a graph such that $\mathcal{V}' \subseteq \mathcal{V}$ and $\mathcal{E}' \subseteq \mathcal{E}$. In the following I will refer to a subgraph G' as a proper subgraph of G, so either $\mathcal{V}' \subset \mathcal{V}$ or $\mathcal{E}' \subset \mathcal{E}$.
complete graph	A *complete graph* K_N is a graph with N vertices in which every pair of vertices is connected.
clique	A *clique* $C = (\mathcal{V}', \mathcal{E}')$ of $G = (\mathcal{V}, \mathcal{E})$ is a complete subgraph of G.
path	A *path* from vertex i to vertex j is an alternating sequence of adjacent nodes and edges that begins with i and ends with j and in which no vertex is visited more than once.
shortest path	A *shortest path* is the path of minimal length between two vertices. In the dissertation length is defined as the number of hops.
diameter	The *diameter* of G is equal to the longest shortest path.
degree	The *degree* k_i of a vertex i is the number of edges that end at that vertex. In the case of a directed graph the degree of the vertex has two components: (i) the number of outgoing edges (referred to as the out-degree of the vertex) and (ii) the number of incoming edges (referred to as the in-degree of the vertex).
Network Science Definitions	
small world property	The *small world property* refers to feature that the average distance L between two arbitrarily chosen vertices is proportional to the logarithm of the number of vertices N in G, i.e. $L \sim \log(N)$.
degree distribution	The *degree distribution* $P(k)$ is the most basic topological characterization of a graph G which is defined as the probability that a vertex chosen uniformly at random has degree k, or equivalently, as the fraction of vertices in the graph having degree k.

clustering coefficient	The *clustering coefficient* of a selected vertex is defined as the probability that two randomly selected neighbors are connected to each other, i.e. it measures the degree of how cliquey the network is and it is calculated as the average of $C = \frac{1}{N}\sum_{i=1}^{N} c_i$, where $c_i = \frac{2\cdot\{\textit{number of links among neighbors}\}}{k_i(k_i-1)}$ is the clustering coefficient of vertex i with degree k.
scale-free network	A *scale-free network* is a network whose degree-distribution follows a power-law, at least asymptotically. That is, the fraction $P(k)$ of nodes in the network having k connections to other nodes goes for large values of k as $P(k) \sim k^{-\gamma}$, where $2 < \gamma < 3$.
complex network	A *complex network* is a network which has the following topological features: (i) small world property, (ii) power-law degree distribution and (iii) high clustering coefficient.

1.1.2 Game Theory

Game theory is the study of mathematical models that describe how conflict and cooperation between intelligent, rational decision-makers affect each other's outcomes. Game theory has a very extensive area of use including economics, political science, psychology, computer science and biology that spawned numerous types of games and approaches. To catch the idea of this approach I shortly introduce one of the most well-known and deeply investigated games, the Prisoner's dilemma [77].

Prisoner's Dilemma. This game is about a situation in which two criminals are under interrogation for a crime. The interrogator officer make the same offer for each criminal, which is to confess the crime or to remain silent. If neither of them confesses the charge against them cannot be proved and both will serve a one-year prison term for lesser offenses. If only one of them confesses, his term will be reduced to 1/2 year and in return he will be used as a witness against the other, who will be sentenced to 10 years. If both of them confess, this counts as a mitigating circumstance and they both end up with 8 years. These choices eventuate four possible outcomes, which can be summarized in a 2x2 matrix (Fig. 1.2).

It is easy to see that the common interest of the criminals would be to remain silent because in this case they will be sentenced only to one year each. However, they have to make their choice independently without knowing in advance how the other would decide, so they cannot make sure that the other would be cooperative. In this

C1 \ C2	confess	silent
confess	(8, 8)	(1/2, 10)
silent	(10, 1/2)	(1, 1)

Figure 1.2: Cost matrix of the Prisoner's Dilemma

situation the only stable solution of the game is that both prisoners confess because choosing to confess always ends up with a better outcome from the perspective of both individuals (1/2 and 8 versus 1 and 10 years). This Prisoner's Dilemma type of situation arises often in reality, such as in the case of overfishing, traffic policy of ISPs, nation-states stockpiling nuclear weapons or athletes using performance-enhancing drugs.

Modeling these situations with games can help us to understand which outcomes will likely happen and how the motivations of participants drive these. Another application is to recognize whether we face a well-known, analyzed game-like situation and based on the derived findings act proactively to avoid these. For example, Prisoner's Dilemma-like situations can be resolved by building trust between participants - e.g. with a contract - that ensures cooperative behavior.

Defining Games. Generally a game can be described by three components: *players*, *strategies* and *payoffs*. Formally, such a game consists of a set $\mathcal{P}$ of players (intelligent rational decision-makers) with cardinality N. Each player u has its own set of possible strategies S_u describing the possibilities how u can act. During the game each player u selects a strategy $s_u \in S_u$ and each state of the game is represented by the strategies of players $s = (s_1, ..., s_N)$.

The strategies $s \in S$ selected by the players determine the outcome for each player. To specify the game, we need to give a preference ordering on these outcomes for each player by giving a complete, transitive and reflexive binary relation on the set S of all strategies. Given two elements in S, the relation for player u says which of these two outcomes u weakly prefers; we say that u *weakly prefers* s_1 to s_2 if u either prefers S_1 to S_2 or considers them as equally good outcomes. The simplest way to specify preferences is by assigning for each player a value to each outcome. In some games it is natural to think of the values as the payoffs to players or in others as the cost incurred by players,$I_u : S_u \rightarrow R$ and $C_u : S_u \rightarrow R$, respectively. Actually, costs and payoffs can be used interchangeably, because $I_u(s) = -C_u(s)$ [77].

Since every player has its own incentives and it is rare that one's optimal strategy does not interfere with others' there is often a trade off among players. A desirable game-theoretical solution for this situation when individual players act according to their incentives on maximizing their own payoff. This idea is best captured by the notion of the *Nash equilibrium* (NE), that is the central solution concept in game theory.

Definition 1. *The game is in the state of a NE if no player has anything to gain by unilaterally changing its strategy. In other words, a strategy $s \in S$ constitutes a Nash equilibrium if for all players u and for each alternate strategy $s'_u \in S_u$:*

$$I_u(s_u, s_{-u}) \geq I_u(s'_u, s_{-u}), \tag{1.1}$$

where s_u denotes the strategy played by player u and s_{-u} denotes the strategies played by all other players.

A Nash equilibrium is *unique*, if that is the only Nash equilibrium in a game (e.g. Prisoner's Dilemma). Not all games possess a unique Nash equilibrium and there are many existing games having multiple Nash equlibria [77]. It is important to mention that the NE is not necessarily optimal for players. In games with multiple equilibria, different equilibria can have (widely) different payoffs for the players. In order to be able to evaluate the different equilibria and to get a more precise picture about a game two distinguished metrics are defined, which are the *price of anarchy* (PoA) and the *price of stability* (PoS). The PoA is the most popular measure of the inefficiency of equilibria, resolves the issue of multiple equilibria by adopting a worst-case approach. In order to be able to define the PoA precisely, first we need to define what the optimal outcome of a game is, that is called *social optimum* (SO).

Definition 2. *The SO refers to an equilibrium state that maximizes the social welfare (i.e. minimizes the sum of all cost) even if its emergence requires a central coordination force (i.e. the independent decision-making is taken away from players).*

Formally, a strategy vector $s \in S$ constitutes a SO if:

$$\sum_u I_u(s_u, s_{-u}) \geq \sum_u I_u(s'_u, s_{-u}). \tag{1.2}$$

Definition 3. *The PoA quantifies the loss to selfishness by comparing the performance at the Nash equilibrium to the optimal state of a game. In other words, PoA shows the cost that players may pay for the lack of coordination in a worst-case scenario. It is calculated as the ratio between the worst Nash equilibrium and the optimal outcome (SO). Formally:*

$$PoA = \frac{\min_{s\in\varepsilon} \sum_u I_u(s_u, s_{-u})}{\max_{s\in S} \sum_u I_u(s'_u, s_{-u})}, \tag{1.3}$$

where ε is the set of Nash equlibria.

In the case of the Prisoner's Dilemma discussed with cost matrix shown on Fig. 1.2 with the cost function $C(s_1, s_2) = u_1(s_1, s_2) + u_2(s_1, s_2)$ the PoA is $\frac{16}{2} = 8$.

Definition 4. *The PoS is an optimistic form of the PoA as it shows how far the best-case scenario of the game, that is created by selfish players, lies from the optimum. Formally, the PoS of a game is the ratio between the best Nash equilibrium and the optimal outcome (SO):*

$$PoS = \frac{\max_{s\in\varepsilon} \sum_u I_u(s_u, s_{-u})}{\max_{s\in S} \sum_u I_u(s'_u, s_{-u})}, \tag{1.4}$$

where ε is the set of equlibria.

In the case of the Prisoner's Dilemma the PoS is equal to PoA, since the game has a unique Nash equilibrium.

It is worth noting that a bound on PoS, which ensures that some of the equilibria are close to the optimum, is much weaker than a bound on the PoA, which ensures that every equilibrium is better than or equal to the given result. But despite this, PoS is often worth to be found as (i) in some cases a nontrivial bound is possible only for PoS and (ii) PoS often could serve as a decent solution to envision a concrete design in practice by a central authority based on the game theoretical analysis. If both PoA and PoS close to 1 it indicates that the game is insusceptible to selfish behavior.

1.1.3 Hyperbolic Geometry

Hyperbolic geometry is one type of *non-Euclidean* geometry, it accepts the first four postulates (axioms) of Euclidean geometry but negates the fifth postulate, which is equivalent to "the parallel postulate“:
(1) A straight line may be drawn from any point to any other point.
(2) A finite straight line may be extended continuously in a straight line.
(3) A circle may be drawn with any center and any radius.
(4) All right angles are equal.
(5) The parallel postulate: given any straight line and a point not on it there exists exactly one straight line passing through the point that does not intersect the first line.

Features of the hyperbolic geometry

Even though hyperbolic geometry differs from the Euclidean geometry only in one axiomatic rule, we can list several consequences to which this difference leads:

Intersecting lines: Two intersecting lines have the same properties as in Euclidean geometry. Two lines can intersect in no more than one point, intersecting lines have equal opposite angles and adjacent angles of intersecting lines are supplementary angles. But by adding a third line then the properties of intersecting lines are differ from intersecting lines in Euclidean geometry, e.g. given 2 intersecting lines there are lines that do not intersect either of the given lines.

Non-intersecting lines (parallel lines): Non-intersecting lines in hyperbolic geometry also have properties that differ from non-intersecting lines in Euclidean geometry. For any given line R and point P which does not lie on R there are an infinite number of coplanar lines through P that do not intersect R.

Circles and disks: The circumference of a circle of radius r is greater than $2\pi r$. Let $R = \frac{1}{\sqrt{-K}}$, where K is the (negative) Gaussian curvature of the plane. Then the circumference of the circle of radius r is equal to $2\pi R \sinh \frac{r}{R}$ and the area of the enclosed disk is $4\pi R^2 \sinh^2 \frac{r}{2R} = 2\pi R^2(\cosh \frac{r}{R} - 1)$.

Hypercycles and horocycles: In hyperbolic geometry, there is no line that remains equidistant from another. Instead, the points that have all the same orthogonal distance from a given line are on a curve called a hypercycle. Horocycle is a curve whose normal or perpendicular (the relationship between two lines which meet at a right angle (90°)) geodesics all converge asymptotically in the same direction.

Triangles: Unlike Euclidean triangles in hyperbolic geometry the sum of the angles of a hyperbolic triangle is always strictly less then 180°.

Distances: In hyperbolic geometry distance calculation also differs from the Euclidean case.

Such a geometry is very different from the familiar Euclidean geometry, which cannot be embedded isometrically into the Euclidean space, so different models were created for hyperbolic geometry, the common ones are the Klein model, the Poincaré disk model, the Poincaré half-plane model and the Lorentz or hyperboloid model [86]. In the dissertation the Poincaré disk and hyperboloid models are used.

The Poincaré-disk model is an intuitive representation of the hyperbolic geometry, each of the characteristic shapes (circle, straight line, hypercycle, horocycle) appear as an arc in the model. By using Descartes coordinates things like intersection points and mirrorings can be relatively easily calculated. Furthermore, the tessellations look appealing. In contrast, in the hyperboloid model calculations are easier by using polar coordinates, especially integration, area and distance calculation.

1.2 Routing Policies

1.2.1 Border Gateway Protocol

As the Internet has become increasingly larger over the decades the Exterior Gateway Protocol [85], which was the original routing protocol of the Internet, has become obsolete. EGP was a simple reachability protocol, that was limited only to hierarchical, tree-like topologies. Beyond scalability problems (*distance vector protocols* suffer from count-to-infinity and *link state protocols* must flood information) the support of business policies was an urging issue as well. In order to solve these issues a new exterior gateway protocol, the BGP, was designed and standardized in 1989, which can be classified as a *path vector protocol*. The current version of BGP is version 4 (BGP-4) codified in 2006 [50].

BGP is responsible for exchanging routing and reachability information among the autonomous systems (AS) of the Internet. According to the classic definition of an AS it is a set of routers under a single technical administration, using an interior gateway protocol and common metrics to route packets within the AS, and using an exterior gateway protocol to route packets to other ASs [52]. Consequently, the administration of an AS appears to other ASs as a single coherent interior routing plan and presents a consistent picture about other ASs that are reachable through it. So from the viewpoint of exterior routing an AS is monolithic, namely reachability to networks directly connected to the AS must be equivalent from all border gateways of the AS.

1.2.1.1 External and Internal BGP

ASs exchange reachability information that is based on a set of policies established within an AS, through BGP gateways. The devices that communicate with each other via BGP are known as BGP neighbors. These neighbors can be located either in the same AS or in different ASs. Based on this we can differentiate between two types of BGP communication sessions: (*i*) external BGP (EBGP) and (*ii*) internal BGP (IBGP), which is illustrated on Fig. 1.3. BGP uses the same message types on EBGP and IBGP sessions but the rules for when to send and how to interpret each message is slightly differs. For this reason people usually refer to IBGP and EBGP as two separate protocols.

EBGP is used and implemented at the edge or border router of an AS and during an EBGP communication session the BGP neighbors belong to different ASs but share a common network infrastructure that is used to carry the BGP messages between them. EBGP works in collaboration with the IBGP to transfer the routing information. Usually, if an AS has multiple connections to other ASs then multiple

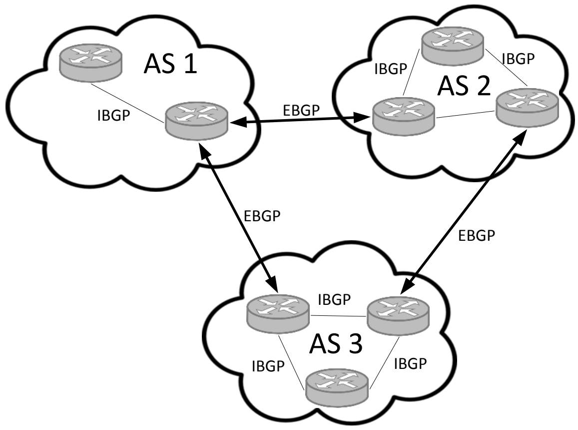

Figure 1.3: BGP communication is carried out either with IBGP or EBGP protocol. If the communicating BGP capable devices are in the same AS (intra-domain scenario), then IBGP protocol is used, otherwise (inter-domain scenario) EBGP.

BGP gateways are needed. In this case all the BGP gateways representing the same AS must give a consistent image of the AS to the outside. This can only be done by assuring that border routers always have consistent routing information and for this purpose internal BGP peering is set up between all of the BGP gateways. During the IBGP session all of the EBGP routes are redistributed among BGP gateways by an interior gateway protocol, such as OSPF [53] or IS-IS [51].

1.2.1.2 AS Relationships and Types of Traffic

A significant volume of traffic is carried within an AS that either originates or terminates at the AS; this kind of traffic can be categorized as *local traffic*. All the other kinds are considered *transit traffic*, which is controlled by BGP. In the Internet, however, the flows of the traffic are not determined only by technical conditions (bandwidth, delay, existence of a connection, etc.) at the first place but are further restricted by business relationships existing between ASs described by service level agreements (SLAs). In the AS ecosystem these business relationships can be quite diverse, still we can classify most AS-AS links into basically two major groups [48]: in a *customer-provider* relationship the customer AS pays the provider for forwarding its traffic, while in a *peering* relationship neighboring ASs voluntarily exchange traffic

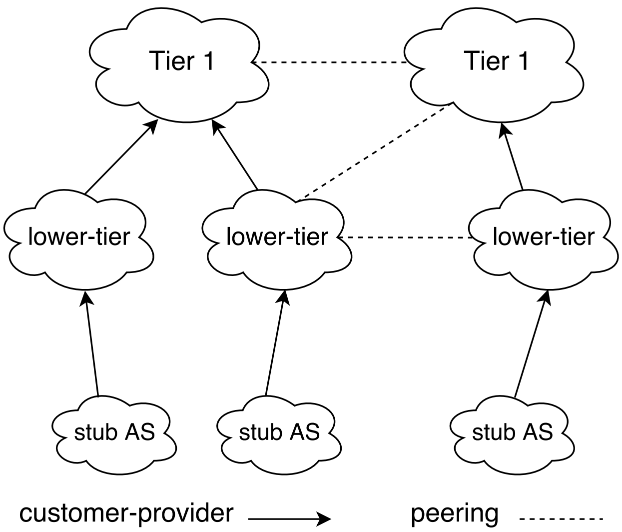

Figure 1.4: Fundamental AS types based on traffic flows. At the bottom of the hierarchy are the stub ASs that do not have any customers, only provider(s). Upper in the hierarchy there are the lower-tier ASs, which have both customer and provider ASs, while Tier 1 ASs have only customers. However, peering is possible between any AS pairs, even if they are not classified to the same tier group.

with each other in a settlement-free manner.

Based on how an AS deals with transit traffic each AS can be placed into one of the following three categories (Fig. 1.4):

- Multihomed AS: an AS that has more than one connection to other ASs, but refuses to carry transit traffic. From the business perspective these ASs are at the top of the hierarchy, the upper-tier ASs, and often referred as Tier 1 networks. By definition a Tier 1 AS is a network that can reach every other network on the Internet without purchasing IP transit or paying settlements and a Tier 1 AS peers with every other Tier 1 ASs.

- Transit AS: an AS that has more than one connection to other ASs and is designed to carry both transit and local traffic. This practically means, that such an AS has its own customer(s), but it is also purchases IP transit (i.e. has provider AS) to reach some portion of the Internet. These ASs are referred as lower-tier networks.

- Stub AS: an AS that solely purchases transit from other ASs to reach the

Internet. Such an AS carries only local traffic.

1.2.1.3 Policy enforcement

BGP enables to enforce policies based on *various constraints and preferences which are tipically non-technical related considerations*. These can be decided by the AS operators and can be set as configuration information. According to the settings the BGP gateways affect the process of best path selection, in the case of multiple alternatives, and take care of redistributing the preferred routing information.

The non-technical constraints are related to political, economic or security considerations that are usually independent from performance-related preferences. For example a multihomed AS is able to decide to avoid the forwarding of transit traffic and can enforce this policy in the form of BGP configuration. Another typical example is when an AS favors or disfavors to carry transit traffic via a certain AS.

The performance-related considerations can be controlled by BGP are (*i*) minimizing the number of transit ASs (i.e. shorter AS paths preferred over longer ones), (*ii*) preferring internal routes over external ones and (*iii*) choosing higher quality ASs for carrying transit traffic. The quality of an AS can be measured by things like diameter, link speed, capacity, tendency to become congested and quality of operation.

1.2.1.4 The Best Path Selection Algorithm

In the case of multiple valid paths exist BGP use the Best Path Selection Algorithm in order to decide the next hop. The algorithm consists of 13 rules in consecutive order and each rule introduces a new filter condition for valid paths. First the algorithm selects valid paths by rule 1, then, if more than one valid path exists, it continues by adding rule 2 and so on until only one valid path left. As the algorithm goes through the rules every step enables a subtler distinction. Among the rules there are Cisco specific rule (e.g. the WEIGHT attribute) and also several lower rules related to the command type used for advertising the path, to timestamps or to router IDs and addresses. Table 1.1 shows a simplified version with the most significant rules of the route selection process [24][40]. The *valley-free* criteria is distinguished as rule No. 0, since BGP path selection works over valley-free paths.

In the following I will concentrate only for the first two steps, i.e. the *valley-free*

Table 1.1: The simplified BGP Best Path Selection Algorithm.

#	Rule
0.	Valley-free route
1.	Highest local preference
2.	Shortest AS path
3.	Lowest origin type
4.	Lowest MED
5.	EBGP-learned over IBGP-learned
6.	lowest IGP metric to the BGP next-hop

and *local preference* policies, as in the dissertation I investigate the effects of BGP based on these rules. The reason behind this is threefold: these rules are (*i*) strictly economically motivated (unlike the other rules), (*ii*) the most general ones affecting the communication [38] and (*iii*) enable to capture non-trivial aspects of inter-domain routing in a clean form that is analytically tractable at the same time.

Valley-Free Routing Policy

The business relationships of different ASs can be diverse, based on exclusive contracts, service-level agreements and other policy issues. Still we can categorize most AS-AS links into basically two major groups [48] which are the *customer-provider* and the *peering* relationship. In the former the customer AS pays another for forwarding its traffic, while in the latter neighboring ASs voluntarily exchange traffic with each other in a settlement-free manner.[1] The valley-free policy manifests the simple economic principle that the flow of traffic must coincide with the flow of cash.

To put it shortly the policy dictates that AS A can use a link to a neighboring AS B to forward the traffic if and only if either the incoming traffic is from a customer or B is a customer of A. In other words, valley-free compliant paths comprise arbitrary (may be zero) number of customer-provider links, zero or one peer link and again arbitrary provider-customer links strictly in this order (Fig. 1.5). The valley-free policy is a typical example on how important an economic, i.e. a non-technical, policy constraint could be.

[1] Sibling and backup relationships are omitted for simplicity.

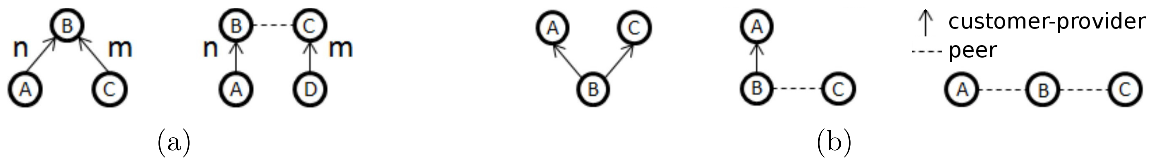

Figure 1.5: Illustration of path types that satisfy (a) and violate (b) the VF policy. A valid path contains n customer-provider, at most 1 peer and m provider-customer link strictly in this order, where $n, m \in \mathbb{N}$. All the other types are invalid paths.

Highest Local Preference Policy

The second rule of the Best Path Selection Algorithm is the other very important economic-based policy, which is the *highest local preference* policy. It is applied on top of valley-free routes meaning that an AS can pick one from the available valley-free routes according to its local interest. Meanwhile these local interests can exhibit high variety the minimalistic rule, that customer and peer paths are favored over provider paths, is contained in basically every local preference setting within the ASs [38] (Fig. 1.6). This is in line with the nature of these routes as customer and peer paths are completely free unlike provider paths in which the provider has to be compensated in some way for the carried transit traffic.

1.2.2 Greedy Navigation

Greedy or geographic routing is a very simple heuristic in which routing decisions (i.e. choosing next hops) are based only on local information. In a networking context local information is generally understood as the coordinates of the nodes but actually it can be any consistently assigned *attribute value* from a metric space (mathematically metric space is a set for which distances between all members of the set are defined). In such a network whenever a node wants to communicate with another one, the only information that has to be considered is the *attribute value* of the destination and the neighboring nodes of the source. The next hop is always the neighbor that, based on the metric space, brings closest to the destination (Fig. 1.7).

The attractiveness of this heuristic lies in its simplicity as it does not require global information yet is able to provide short paths during the communication, especially on scale-free topologies (described in 1.1.1), which are known as the common signature of many large-scale self-evolving complex networks.

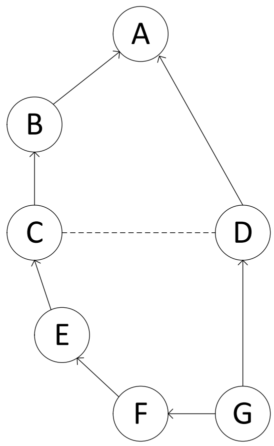

Figure 1.6: Illustration of paths prioritization according to the highest local preference rule. If AS C wants to reach AS G it has multiple options, since C–B–A–D–G, C–D–G and C–E–F–G are equally valid valley-free paths. However, economically, the favorable order is C–E–F–G (as E is a customer of C), C–D–G (as D is a peer partner of C) and C–B–A–D–G (as B is a provider of C). Note that, in this step of the Best Path Selection Algorithm the first hop is more important than the length of the paths, namely it is generally worth to choose a longer but cheaper path than a shorter but more expensive one.

This phenomenon was first confirmed by the noted social psychologist Stanley Milgram in 1967 by an experiment [75] carried out in the USA. He asked participants to deliver letters via their personal acquaintances. For example one in Nebraska had to reach someone in Massachusetts. The only additional information - beyond the address - was the profession of the addressee. It turned out that the average path length of successful forwarding chains were between five and six and in these cases the general approach was to choose someone (as a next hop) based on (*i*) who brings closer to the destination and (*ii*) whose profession is similar (i.e. seems "related") to the addressee. So Milgram's famous experiment was the first that showed empirical evidence that complex networks are small worlds and indeed navigable by distributed greedy routing at the same time [92]. However, this efficiency is not trivial, since being only a heuristic, greedy routing theoretically can stuck during forwarding, as it is described at (Fig. 1.7).

Kleinberg proposed an analytic model and a working algorithm that justifies the

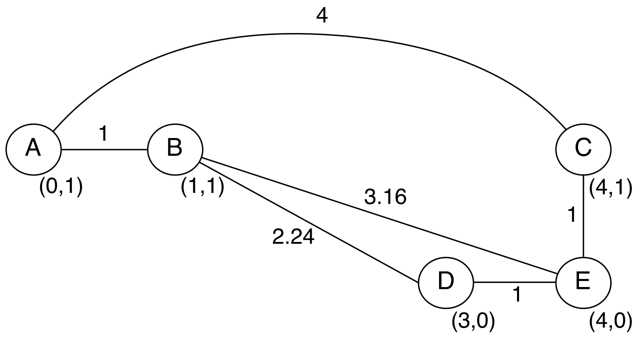

Figure 1.7: Illustration of greedy navigation. Consider this graph comprising 5 nodes and a distance calculation method based on geographical coordinates. If A wants to reach D, then first calculates distances between B – D and C – D, then choose C as next hop, since C brings closer to D than B. C does the same way and from A – D and E – D choosing E (trivially). Finally, E considers B – D, C – D and D – E, then forwards to D. In the case of this specific graph if D – E would not exist greedy routing ended up with E and get stuck there, since there is not neighbor closer to D than E itself.

existence of such a *small-world* message forwarding [58] experienced by Milgram. In this work the world is modeled as a two-dimensional grid (Fig. 1.8), where each vertex is a person and there are local and long-distance edges between vertices. The probability of the existence of a long distance edge is $P_r(u, v) = d^{-r}(u, v)$, where d is the lattice distance and parameter $r \in [0, \infty)$. It is shown that a simple greedy routing algorithm needs $O(\log^2 n)$ time to travel between any pair of nodes if and only if $r = 2$ or more generally $r = D$, where D denotes the dimension of the lattice.

Since metric spaces are *either existent* [94] or *can be efficiently constructed* with regard to social and computer networks [15, 60], greedy routing is a remarkably efficient mechanism. A number of practical routing solutions are based on the greedy routing principle. Perhaps the most successful practical systems using greedy forwarding are the overlay networking solutions based on distributed hash tables, e.g. CAN [83] and Chord [32]. These schemes employ different underlying abstract geometries as a basis for forwarding, torus and circle, respectively. Hamming-distance based greedy routing has been utilized in Microsoft's BCube data center design [45].

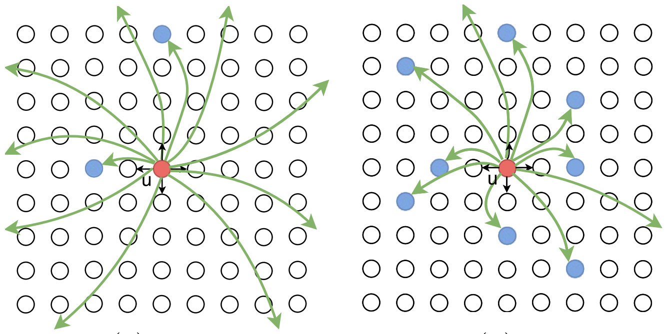

(a) $r << 2$ (b) $r \sim 2$ (c) $r >> 2$

Figure 1.8: The effect of parameter r to the topology. Sub-figure a), b) and c) shows how likely longer connections emerge for different values of r. In the case of $r \sim 2$ most of the connections will lead to nodes relatively close to u, but there are some long connections as well, making possible of bypassing large distances, thereby enabling shortest paths, i.e. small-world property in the network.

1.3 Approaches for Investigating the Topology of the Internet

Knowing more about the topology of the Internet is one of the most popular topics of the network science community. The reason behind the popularity is that (i) compared to other complex networks this topology is highly tractable and both active and passive measurements can be executed on this network, thus "screenshots" can be created easily and (ii) this topology - at least from a macroscopical view - is very similar to other complex networks like the neural and social network, the network of airline routes or the food web. [10, 93]. By similarity we mean three fundamental features in which complex networks are common, these are the *small world property*, the *power-law degree distribution* and the *high clustering coefficient* (described in 1.1.1). Thus knowing more about the topology of the Internet is valuable for researchers from diverse or multidisciplinary areas as well.

Basically, two types of Internet topologies are used during the investigations, one is the *router level* and the other is the *AS level* topology. However, in most of the works on "Internet topology" actually refers to the AS level topology, since it has a more manageable size with its ~50 000 ASs and from a global view the data obtained is more precise and persistent at this level.

1.3.1 Motivation and Benefits

The last decades have supplied us with thousands of stories where topology-related information about the Internet was directly transformed into more efficient architectures and services or more appropriate business decisions. In the following I give an overview about the most significant investigations inspired or connected to the AS level topology.

Content Delivery Networks

The most specific example is Content Delivery Networks (CDNs) that deliver a high volume of commercial content. In 2011 almost 50% of peak-time traffic coming into North American access networks consisted of real-time entertainment, provided by content providers such as Netflix and YouTube via their CDN operators: Akamai, Limelight, Level 3 and Google [79]. This trend holds outside North America as well and CDNs has an increasingly important role and impact on the network traffic.

In order to enable efficient content delivery beyond origin servers CDNs use numerous surrogate servers to support serving the requests. Large providers such as Akamai have more than 200 000 servers in over 100+ countries around the world [91]. The proper placement of these servers is crucial for performance and there are several approaches including Hot Spot [64], Tree-based [80], greedy method based [63] or topology informed [55, 81] replica placement strategies.

Topology informed strategies use existing topological information about the Internet and the actual CDN network for efficient mapping. Basically there are two types: (i) Servers are placed at highly connected hosts in the AS Level topology graph of the Internet according to degrees in descending order. This is based on the assumption that hubs can reach others more easily thereby decreasing the overall latency. (ii) The Internet router-level topology can also be used for mapping servers to places. In this case each LAN associated with a router is a potential site where servers can be placed [79].

Cloud Networks

Similarly to CDNs cloud providers have a constantly increasing number of customers. Large providers such as Microsoft Azure offer 200+ cloud services that are hosted on more than 100 globally distributed data centers, edge computing nodes and service operation centers [28]. Providing performance that satisfies customer needs could be very costly (tens of millions of dollars) in such networks, which are made up of four main components: (*i*) servers (CPU, memory, storage systems), (*ii*) infrastructure (power distribution and cooling), (*iii*) power draw (electrical utility) costs and (*iv*) network (links, transit, equipment) [42].

Geo-diversity of the cloud network influences all these factors as it can lead to lower latency and better reliability. The impact of latency is shown by Google and Amazon through experiments. Google reported that a 500 msecs increase of displaying search results caused 20% revenue loss and Amazon reported about 1% sales decrease due to 100 msecs additional delay. However, geo-diversity also increases the cost of service. So there is a trade-off between the performance and the placing and sizing of data centers (i.e. to design the appropriate network topology), that is a challenging optimization problem closely connected to the emerging topology of the Internet.

Peer-to-Peer Networks

Peer-to-Peer (P2P) overlay networks are another example on how useful the topological knowledge could be. As communication environments became increasingly complex finding new ways to manage distributed systems without central organization and hierarchical control attracted growing interest. The idea of P2P overlay networks is to create a virtual mapping that overlays the physical networks and help locating data as quickly as possible with minimal overhead and maintenance in order to enable efficient, massively scalable, robust and fault-tolerant routing in a self-organized fashion. These features make P2P overlay networks important in data sharing and content distribution applications. Over the years several solutions have been created. Based on how they create the overlay network, they can be categorized into two main groups: (*i*) *structured* and (*ii*) *unstructured* P2P overlays [66].

The fundamental difference between these categories is the way they find the content stored by overlay peers. In the first group the network assigns keys to data items and organizes its peers into a structured graph that enables efficient discovery of these items using the given keys. Popular structured P2P overlay networks are CAN [82], Tapestry [100], Chord [90], Pastry [87], Kademlia [71] and Viceroy [69]. In the second group peers are organized in a random graph and looking up the data item is carried out by flooding, random walks or by expanding-ring Time-To-Live search. Popular networks that belong here include Freenet [27] and Gnutella [41].

Both types of networks have strengths and weaknesses - for example unstructured solutions are less complex, but for a rare data item many peers would have to be involved in the lookup process - hence choosing the best method depends on the application and on its required functionality. However, as data lookup process is based on the virtual mapping of the physical topology it is very unlikely in both types of overlays that the proper peers - which own the searched data item - are reached via an optimal number of intermediate nodes (neighboring peers in the overlay network are not necessarily connected physically).

This matching can be improved by creating topology-aware P2P overlays that aim to exploit network proximity in the underlying Internet. In [22] an improved version of Pastry is presented with a comparison to CAN, Chord and Tapestry. Results supporting the claim that topology-aware routing approaches in P2P overlays can improve application performance and reduce network usage substantially at the same time while incurring only modest additional organization and maintenance overhead.

Epidemic Models

Scale-free networks have aroused interest from the aspect of epidemic spreading as well. In addition to biological networks Internet is again a popular target for researchers due to its technological and economical relevance. The general approach is to use a scale-free topology with a standard epidemic model (e.g. SIS or SIR [6]) in which nodes can be in the state of *susceptible*, *infected* and *recovered* and infection can spread through direct connections (links). Important metrics are the *average lifetime of the virus* spreading and the *epidemic threshold* (a value which shows whether the infection spreads and become persistent or dies out).

In [78] it is shown that (*i*) the average lifetime of viruses is larger if the Internet expands and that (*ii*) scale-free networks surprisingly do not have an *epidemic threshold*. The reason for that is the existence of the hubs and their extremely high connectivity, since the threshold value is inversely proportional on the node's degree, that annuls the threshold.

Navigation of Complex Networks

One of the most fascinating discoveries about natural complex networks is the fact that navigation is efficient, even though nodes communicate with greedy navigation, which relies only on local information without the presence of any global intelligence. This phenomenon is closely connected to the structural properties of these networks. The similarity of the Internet topology to natural complex networks has raised the question *"Is it possible to use greedy navigation on the Internet as well?"*, thereby solving the long-standing scalability problems come from the currently used BGP based routing architecture, that relies on constantly maintaining a coherent view of the global topology. [14]

However, the first step towards this direction should be the clarification on *"How does the topology exactly affect (make possible) the efficiency of greedy navigation?"*. Authors of [14] investigate the question and explain this connection by introducing a general mechanism relies on the presence of a metric space hidden behind the observable network. As described at greedy navigation (Section 1.2.2) this metric space could be anything that enables to calculate distances between two arbitrarily chosen nodes of the network. In Milgram's experiment the *attribute value* of the metric space was the profession of the participants. For example if one needs to send message to a politician via personal acquaintances it is more likely to contact a lawyer or his high school history teacher than a car mechanic as next hop.

Such underlying - yet undiscovered - metric spaces can be found in many complex networks helping the decisions of nodes during the forwarding process and the amount of their usefulness depends strongly on topological properties. In order to justify this the authors of [14] provide a scale-free network model that is similar to the Internet AS topology and by simulations they show evidence that structural properties like clustering coefficient and degree-distribution directly affects the average length of forwarding paths in the network.

Hyperbolic Mapping of the Internet

Recently it was shown that not only natural complex networks but geographically embedded Internet-like synthetic networks also enable efficient communication with greedy navigation [14, 16, 61]. The reason for the efficiency is again due to the metric space. Embedding here actually means to create a synthetic metric space that strongly supports decision making on finding the best next hop during forwarding. The general idea behind this embedding is the fact that routing on the Internet nowadays is somewhat equivalent to forwarding based on a hypothetical road atlas that does not contain the real geographic information, but only lists road network links, which are pairs of connected road intersections, abstractly identified.

Authors in [15] suggest that the routing task could be drastically simplified in this environment by using the real geographic coordinates, since if the coordinates of the starting and destination points are given then it is easy to tell what direction would bring closer to the destination. Furthermore, since geographic coordinates are invariable this information does not need to be exchanged in case of topology changes, which can lead to efficient routing with minimal overhead based on local information. However, constructing the proper map congruent with the network topology is crucial for this kind of routing architecture along with the chosen coordinate system. The authors have shown in recent works [14, 16, 61] that the efficiency of forwarding can be maximized by using hyperbolic space, hence they create the map of the real AS level Internet in a hyperbolic space based on statistical inference techniques and show that embedding this map indeed enables the expected forwarding efficiency.

Robustness and Traffic Handling

The structural consequences of the Internet AS topology is also investigated and showed [84] that there is a trade-off between the *resilience to structural damage* and the *efficient handling of traffic flows*. As a method authors created an AS topology model, based on the data supplied by projects DIMES [88] and ROUTEVIEWS [1], on which they carry out vulnerability tests and traffic dynamics analysis. They also compare their model with a classical random network of similar size.

They summarize the results as follows: whereas the Internet AS topology ensures robustness at a structural level (i.e. random node failures cannot cause significant structural damage) it does not allow to reach the same efficiency for handling traffic flows. In fact, an Internet-like network reaches the congested phase prior to a random network. This comes from the scale-free nature of the network, as large hubs on the one hand ensure excellent connectivity through their connections, however, on the other hand they attract much more traffic than other parts of the network, since most of the short paths lead through them, which can easily lead to congestion.

This effect can be mitigated by scaling out with stronger central routers, but fundamentally it cannot be resolved in this way. That is why authors also discuss strategies the are used to avoid this central-links bottleneck problem by using alternative paths in congested situations based on the topological knowledge.

1.3.2 Approaches

The aforementioned cases represent some examples where topological knowledge is greatly exploited, but this knowledge is partial and researchers still put a lot of effort into further extending it. In the following I give an overview about the different approaches, including projects and models as well, that aim to gather precise and accurate topological information about the Internet. Besides categorizing and shortly introducing them I also place and describe my approach.

1.3.2.1 Measurements

Today we have historical and contemporary measurement data collected continuously and made publicly available according to various approaches, which include (*i*) *Internet mapping* (Rocketfuel ISP Mapping [89], Skitter [72] and Opte Project [67]), (*ii*) *discovery of economic relationships* (CAIDA's Archipelago project [19], IXP anatomy [3]) and (*iii*) *visualizing the infrastructure* of multiple backbone providers simultaneously (MapNet [25]).

These projects mainly rely on the data provided by the University of Oregon Route Views Project [74], looking glass servers of TRACEROUTE.ORG [57] and the Archipelago project. Rout Views originally conceived as a tool for Internet operators to obtain real-time information about the global routing system from the perspectives of several different backbones and locations around the Internet. It uses publicly available BGP tables consists of AS path to destination traceroutes and collects data since 2000. TRACEROUTE.ORG is a collection of looking glass servers that enables running queries on participating ASs, although it provides only a constrained view of the routing system. Archipelago is CAIDA's active globally distributed measurement infrastructure set off in 2007. In order to improve the view about the global Internet

CAIDA is continuously distributing hardware measurement nodes (2nd gen. Raspberry Pi) with as much geographical and topological diversity as possible. Archipelago project is tailored specifically for active network measurement and it is also included in several measurements that are not directly connected to the topology (e.g. The Spoofer Project [12], TCP Behavior Inference [19], IPv4 and IPv6 Stability [19], TCP-HICCUPS [29]).

Meanwhile the data stemming from these measurements is the exclusive source of direct information about the AS topology and thus can be treated as the ground truth we can keep ourselves to, the way these measurement systems work is continuously reported to be imperfect and far from optimal [3]. Additionally the collected data reveals only the current state of the network and cannot give usable predictions and clear characterization of the topology forming processes lying in the background.

1.3.2.2 Internet Models

The other popular approach of getting more knowledge about the Internet topology is the creation of different network models. Nowadays it's commonly accepted [3] that the Internet belongs to the group of scale-free networks, so the relevant models here are those that can produce such topologies. However, there is an interesting historical background that is worth of summarizing shortly the evolutionary milestones:

1. **Random graphs:** the most commonly used models for generating networks algorithmically were the Erdős-Rényi [18] and the Waxman [96] models from 1985 and 1988, respectively. In the Erdős-Rényi model the network is a random graph $G(n, p)$ with n nodes where each possible p edge has a probability p of existing. The number of edges in a $G(n, p)$ is a random variable with the expected value $\binom{n}{2}p$.

   The Waxman model is similar, but here nodes are randomly distributed in a two dimensional grid. Links are added to the graph by considering all possible pairs of nodes and then deciding whether a link should exist according to a probability function involving the distance between nodes and the number of links expected in the network.

   These models are simple but has serious drawbacks when they are used for generating Internet-like topologies [20]: (i) the network aren't resemble to real

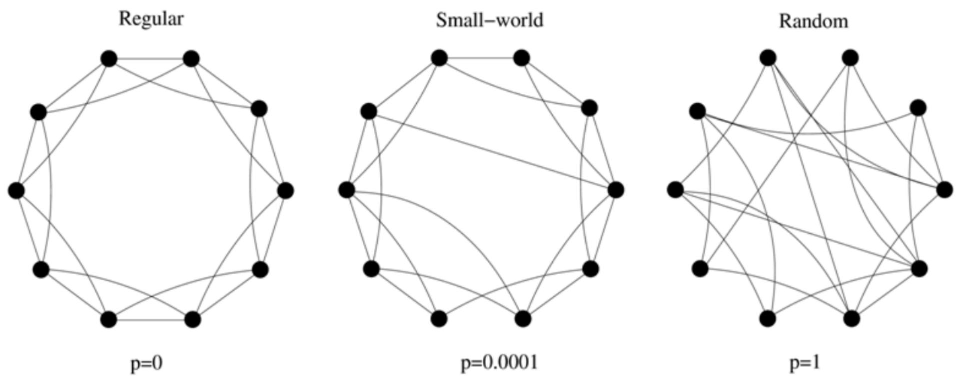

Figure 1.9: Transformation from regular lattice into a small world and random graph.

networks as they lack any sense of backbone or hierarchy, (*ii*) there is no guarantee for a connected network, (*iii*) the number of links is proportional to the number of nodes and (*iv*) although the average path length is small, the clustering coefficient is low and has a Poisson degree distribution instead of power-law.

2. **Watts and Strogatz graphs** [95]: W. and S. considered that real networks are neither entirely regular nor entirely random, so they combined the two extremes, namely the regular lattice and the Erdős-Rényi model. The initial state is a regular lattice, such as a ring, and then some of the edges are rewired in order to introduce a measure of randomness. Each edge is examined and redirected to a random destination by probability p or left in place by probability $1-p$, where $p \in [0, 1]$. If $p = 0$ the lattice remains unchanged and if $p = 1$ it is transformed to a random graph. The interesting cases lie between these values. Besides addressing most of the issues of pure random models this model can generate graphs with high clustering coefficient, thereby creating "small worlds", which is a significant step towards realistic networks. However, the realistic degree distribution is still missing from the generated networks.

3. **Scale-free graphs**: one of the most influential discoveries of the last decades in network science is the identification of scale-free networks, i.e. that the node connectivities in a complex network follows a scale-free power-law distribution [9]. The first clue is connected to the map of the World Wide Web (WWW) finding that the probability of a Web Page has exactly k connections

(i.e. degree k) is $P(k) = k^{-\gamma}$, which is very far from the previously assumed Poisson distribution [11]. This is a consequence of two generic mechanisms that is the networks expand continuously by the addition of new nodes and that the new nodes attach preferentially to already well connected nodes.

The first scale-free model is the Barabási-Albert (BA) model that mimics these behaviors and grows a network from an initial connected network of m_0 nodes. Each time when a new node is added it connects to the existing nodes with a probability proportional to the degree of the nodes already in the network. Formally, the probability that an edge is connected to node i is $p = \frac{k_i}{\sum_j k_j}$. The number of edges a new node connects with can be adjusted by a parameter. This discovery inspired a plethora of new Internet AS-level models over the last decades that can be categorized into three fundamental types that are **causality oblivious**, **causality aware** and **game theoretic** models.

Causality Oblivious Models

These models use various mathematical approaches to generate network topologies but *without any regard about the incentives and individual decision making of nodes* that leads the network to its final state. The topology generation process is carried out based on such premises that ignore realistic decision making considerations and if the result is promising enough the used methods are finally generalized. We can say this is the de facto way of getting information about complex networks. Popular models that belong here are PLRG [4], Inet [97] and the dK-graphs [68].

PLRG produces random graphs with a power-law degree distribution depending on the number of nodes N and an exponent β. It assigns a degree to every node according to the power-law distribution. For creating links between nodes PLRG makes k_i copies of each node, where k_i is the degree picked for node i based on the power-law distribution, then connects copies by randomly picking pairs until no isolated copies remain.

Inet generates topologies in multiple steps: First, it assigns a degree to every node that should be in the graph then grows a spanning tree from nodes that have a degree greater than 1 by connecting each of them according to preferential attachment described at the BA model. Next, it connects all nodes with degree 1 to the graph again using preferential attachment. Finally, it connects the remaining free degrees in G, starting from the node with largest degree first. When making these connections nodes are randomly picked according to preferential attachment. Inet has multiple versions and each contains some improvements based on previous versions.

The dK-graphs are also based on given degree distributions (dK-distributions) that specify node degree correlations within subgraphs of size d. The dK-graphs are the sets of graphs constrained by given values of dK-distributions. Through an incremental generation process a family of dK-graphs, where $d < 4$ can be produced. Each family describe random graphs in a successively finer detail according to the fundamental metrics in the literature. However, the larger the value of d the higher the computation complexity is. According to the authors $d = 2$ is sufficient for most practical purposes, while $d = 3$ reconstructs the AS- and router-level topologies precisely.

While the causality oblivious models are simple and often precise (as far as the power of measurement data can verify) this approach suffers basically from the *inability to capture correlations between node behavior and topology changes and to anticipate evolutionary trends.*

Causality Aware Models

Causality aware models try to *mimic some concrete behavior of the network formation process*, thereby creating relevant topologies. The aforementioned BA model and the Heuristically Optimized Trade-off (HOT) [34] models are the very first topology generators that belong to the causality-aware group. HOT also uses incremental network growing with preferential attachment for producing networks with power-laws, but HOT builds a tree from arriving nodes and the selection process takes into account of minimizing the Euclidean distance between the two connected nodes and to possibly find a centrally located one as well.

BRITE [73] incorporates the findings of power-laws, the skewed node placement and the locality network connection during the topology generation process that can be fine tuned by parameters offering multiple choices. Nodes are first placed in a $HS \times HS$ plane according to either a uniform random or a pareto distribution, then each square is divided further and the assigned nodes are uniformly distributed among them. In the next step a backbone spanning tree is created from the nodes selected one-by-one for each HS square. The remaining nodes are then connected to the backbone with preferential attachment based on locality and/or outdegree.

SIMROT [31] generates hierarchical topologies that include business relationships. For each network it distinguishes four node types: T (tier-1), M (transit), CP (content provider) and C (customer). After calculating the necessary amount of each type of node (and some other constraints) it assigns regions to them and start incremental growing with preferential attachment in a top-down manner, always taking into account that both end-nodes of an edge must be in the same region. This first phase creates only customer-provider edges. In the second phase nodes create peer edges for type M with preferential attachment and for type CP with uniform probability.

Causality aware models not only produce often precise topologies (similarly to causality oblivious models), but also introduce an observable network formation process that can provide even predictions for topological changes. However, this approach is *not applicable to prove that the processes these models are defined upon, are actually present in the real AS network*. For example one cannot really think that preferential attachment in its pure form (where an AS chooses its peers according to their exact nodal degree) happens in the AS ecosystem. This inability makes these models and their predictions somewhat ambiguous. In other words, *the knowledge we can gain through measurements and causality-aware models does not really focus on deeper understanding of the networks, as incentives of nodes are unnoticed and the self-organizing nature of the network formation process remains unrevealed.*

Game Theoretic Models

Game theoretic models concentrate on network-creation from the aspect of individuals. In these models all nodes considered as rational, selfish players whose intention is forming a network along their own interest. The topology formation process can be analyzed by a game defined with a triplet of *players*, *strategies* and *payoffs*. The network creation process is considered as finished if the game reaches a Nash equilibrium state, in which no nodes worth to change its actual strategy provided that others do not change theirs.

These models introduce two important features, (i) by capturing the self-organizing aspect of the AS-level Internet they give us *the possibility to predict certain properties and changes of the network* and (ii) by assuming prudently defined premises such a model enables us to *analytically prove the exact topological consequences* of certain intentions of the players. Nevertheless, several models are not intended to be analytically tractable as they focus only on the first aspect, i.e. capturing network dynamics as realistically as possible. Based on this we can further distinguish between *computational* and *analytical* game-theoretic models. For both types one of the most important design step is to define premises that should capture non-trivial aspects of the system, but for analytical models it is also crucial to balance between complexity and tractability.

Consequently, computational models are able to include more detailed information about the nodes and to play a more realistic network formation game but at the cost of greater complexity, which is usually reflected by the inability of generating large networks (above hundreds of nodes). Analytical models, by contrast, concentrate on capturing some certain behaviors in a simpler yet realistic setting that yields non-trivial and analytically provable conclusions at the same time, even for larger networks. It is also worth noting that the game theoretic approach is somewhat complementary and not supplementary of the other Internet models as those try to tackle the AS level network from completely different angles and can provide totally different insights. By understanding the inherent self-organization of nodes/players we have the ability to predict trends and changes in a complex network that makes us capable to act proactively.

GENESIS [65] is a computational game-theoretic model that puts the focus on capturing inter-domain traffic flows, geographic constraints and economics to model the network formation process. This is carried out by the implementation of several realistic rules referring to IXPs placement, traffic generation and consumption of players, cost structure of business relationships, etc.. Each simulation produces different equilibria (or in about 10% of the runs it gets stuck in oscillations) and aims to answer questions on *"How does the changes of business relationships affect the resulting network in terms of topology, traffic and economics?"*. However, the high complexity of the model enables to experiment only with smaller networks (circa a few hundred nodes). The model of Chang et al. [23] is similar in spirit to GENESIS as it also includes realistic economic, geographic and traffic constraints and aims to model the decision process by which connectivity between ASs in the Internet is established. However, among others, it differs in using hard-wired strategies for creating business relationships.

Analytical game theoretic models have a large body of literature with several results collected in [77]. The range of these models include simple (in the term of realistic design) network formation games. In the *local connection games* nodes face two conflicting desires, namely, to pay as little as it possible, and to have short paths to all other nodes. In *global connection games* players make global decisions, in that they may build edges throughout the network. Unlike the local connection games, here players actually build and maintain large-scale shared networks and are interested in connecting to some specific nodes, called terminals, as cheaply as possible. In *facility location games* there is a more sophisticated cost model as links (connections) still have costs, but players also select prices for users (thereby creating providers and customers) so as to maximize net income, which is the price minus the cost paid.

The above mentioned computational and analytical models provide valuable information about network dynamics and equilibrium states from multiple aspects, but still the *topological footprints inflicted by routing policies, which is one of the most determining factors on the resulting network, is an unacquainted aspect.* Computational models comprise constraints referring to policy rules but - following from their nature - along with so many other influencing factors that makes hard to give a general, yet precise conclusion about the effect of these rules. On the other hand, existing analytical models do not really consider routing policies either, or if they do so then by using unrealistic or oversimplified rules (e.g. shortest path distance for describing business driven forwarding paths between nodes).

In my work I aim to fill this gap by *investigating the topology formation effects of different routing policies connected to complex networks.* To do this I design analytical models in which the policy rules exist in their pure and more realistic form, carefully separated from other factors that could confuse the analytic inference. In the following chapters I concentrate on the rules of BGP and greedy routing policies and discuss their effects in respect of the emerging topologies. In the case of BGP I prove the existence of a special subgraph that is always included to the network due to the incentives of nodes drove by BGP. For greedy navigation I show that currently popular topologies, that enable efficient greedy routing, cannot emerge - at least in an economically verifiable way - under Kleinberg-like, constant dimensional, grid-based models.

Chapter 2

Topological Consequences of the BGP [J1, J2]

Since the AS level topology is formed along the rational business decisions of the individual ASs, game theory is a natural modeling tool of choice. So in the following I consider ASs as rational but selfish players whose incentive is to communicate with each other using the *valley-free* (VF) and *local preference* (HLP) rules (Section 1.2.1.4) for routing policies.

Accordingly, players are defined as the nodes of the network, while the strategies of the players are to create a set of links connecting them to an arbitrary subset of the other players. The goal of the players is to minimize their cost function which consists of two parts: *link cost* and *communication cost*. The cost function for player u is generally defined as [35]:

$$C_u = \underbrace{\sum_{\forall u \neq v} d^{\text{sh}}_{G(s)}(u,v)}_{\text{communication cost}} + \underbrace{\alpha |s_u|}_{\text{link cost}}, \quad u, v \in \mathcal{P}, \tag{2.1}$$

where $d^{\text{sh}}_{G(s)}(u,v)$ denotes the length of the shortest path between players u and v on the graph $G(s)$ characterized by the union of strategies of the players; s_u stands for the strategy of player u (i.e. the links that u creates towards the other players), while α represents the cost of building one edge. Thus such a game effectively analyzes the trade-off between link costs and communication costs for rational, selfish players, as the key incentives of building specific topologies. As defined in (2.1), distance is usually measured as the average length of shortest path from the given node to all the other nodes. Throughout my analysis I always define such network formation games adjusted to the specific policy that is investigated.

2.1 The Valley-Free Game

As described in Section 1.2.1.4 the VF rule is the most fundamental part of the BGP policy routing, since any valid path between ASs has to be a VF path as well. Accordingly my first goal is to define and analyze a game to understand how this rule affects the topology.

Players and routing – Let $\mathcal{P}$ be the set of players (identified as network nodes) with cardinality N. Recall that the rules of the VF policy dictate that a player u can forward traffic originated from player w to a neighboring player v if and only if: (i) the incoming traffic of u is from a customer (in this case the relationship between u and v is indifferent and can be both provider-customer or peering), or (ii) v is customer of u (in this case the type of relationship between w and u becomes indifferent). In other words: *after a provider-customer edge or a peer edge, the path cannot traverse another customer-provider edge or another peer edge, respectively.[38]*

Strategies and topology – A strategy for player $u \in \mathcal{P}$ is to create a set of undirected edges to other players in the network. The created edges can be customer-provider (p) and peer (r) edges in accordance with the relationships in the VF routing. The r-type edges are paid for both ends, however p-type edges are only paid by the customer. Note that in some of the following arguments, according to the direction of traversing customer-provider links, I may write provider-to-customer link but these terms refer to the same kind of edge. Thus the complete strategy space of player u is $S_u = 3^{\mathcal{P}\setminus\{u\}}$, where the number 3 accounts for the third choice of node u, which is to create no edge. Let s be a strategy vector containing the strategies of all players hereby representing the current state of the game: $s = (s_0, s_1, ..., s_{N-1}) \in (S_0, S_1, ..., S_{N-1})$. The graph $G(s) = \bigcup_{i=0}^{N-1}(i \times s_i)$ represents the topology between the players.

Payoff – The goal of the players is to minimize their costs. The cost of player u

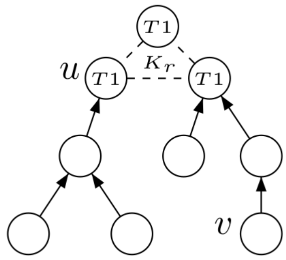

Figure 2.1: Example for a VFF topology. In such a topology there could be two type of nodes, T1 and not T1. T1s are connected by r edges, which are counted on both sides in the cost function, however p edges are paid only by the customer, who requested it. The flow of cash is visualized by arcs. According to this there are two possible cost functions: (i) $C_u = \varphi_r u_r = \varphi_r(|V(K_r)| - 1)$ and (ii) $C_v = \varphi_p$.

is defined as:

$$C_u(s) = \underbrace{\sum_{v \neq u} d_{G(s)}(u,v)}_{\text{communication cost}} + \underbrace{\varphi_p u_p + \varphi_r u_r}_{\text{link cost}}, \quad u,v \in \mathcal{P} \tag{2.2}$$

where φ_x is the cost an edge of type $x \in \{r,p\}$, u_x is the number edges of type x and $d_{G(s)}(u,v)$ is the communication cost between u and v over $G(s)$ given by:

$$d_{G(s)}(u,v) = \begin{cases} 0 & \text{if a VF path exists between } u \text{ and } v, \\ \infty & \text{otherwise.} \end{cases} \tag{2.3}$$

In what follows I identify the Nash equilibrium of the game in different settings of the parameters.

Definition 5 (Valley-free footprint (VFF)). *A graph is a valley-free footprint if it consists of (i) a clique K_r comprising peer (r) links only, and (ii) trees rooted at some subset of $V(K_r)$ having customer-provider links (p) only, such that for all provider-customer connections the provider is always closer to their respective root than the customer (see Fig. 2.1).*

Theorem 2.1. *A VFF is a Nash equilibrium if and only if $\left\lceil \frac{\varphi_p}{\varphi_r} \right\rceil \leq |V(K_r)| \leq \left\lfloor \frac{\varphi_p}{\varphi_r} \right\rfloor + 1$.*

Proof. The proof consists of three parts: first I show that a VFF with $\left\lceil \frac{\varphi_p}{\varphi_r} \right\rceil \leq |V(K_r)| \leq \left\lfloor \frac{\varphi_p}{\varphi_r} \right\rfloor + 1$ is a Nash equilibrium, secondly that a VFF with $|V(K_r)| < \left\lceil \frac{\varphi_p}{\varphi_r} \right\rceil$

or $|V(K_r)| > \left\lfloor \frac{\varphi_p}{\varphi_r} \right\rfloor + 1$ is not and finally that a non-VFF topology cannot be a Nash equilibrium.

1. A VFF topology with $|V(K_r)| = \left\lfloor \frac{\varphi_p}{\varphi_r} \right\rfloor + 1$ is a Nash equilibrium: I show that no player has anything to gain by changing his own strategy if others don't change theirs. Let C_u represent the cost of player u residing in the VFF. It is clear that in a VFF every player can reach others through valley-free paths, hence in their cost functions the communication cost is always zero ($\sum_{v \neq u} d_{G(s)}(u,v) = 0$). Moreover, it is clear that any meaningful state of the game permits valley-free paths between arbitrary pairs of players (otherwise the cost of some player would be infinite), so hereafter writing this term will be omitted.

   (a) if $|V(K_r)| \leq \left\lfloor \frac{\varphi_p}{\varphi_r} \right\rfloor + 1$ and $u \in K_r$:

   $$C_u = \varphi_r u_r = \varphi_r(|V(K_r)| - 1) \leq \varphi_r \left\lceil \frac{\varphi_p}{\varphi_r} \right\rceil \tag{2.4}$$

   If u wants to deviate then it will use some other strategy. The corresponding cost is given by $C_u{}' = \varphi_r u_{r'} + \varphi_p u_{p'}$. If $u_{p'} \geq 1$ then $C_u \leq C_u{}'$, since:

   $$C_u = \varphi_r u_r = \varphi_r(|V(K_r)| - 1|) \leq \varphi_r \left\lfloor \frac{\varphi_p}{\varphi_r} \right\rfloor \leq \varphi_r \frac{\varphi_p}{\varphi_r} = \varphi_p = C_u{}' \tag{2.5}$$

   If $u_{p'} = 0$ then again $C_u \leq C_u{}'$ since $u_r \leq u_{r'}$ must hold to ensure valley-free connectivity to all the other nodes. This holds because if $u_r > u_{r'}$ then there is at least one node $v \in K_r$ to whom u is not connected thus cannot be reached from u via a valley-free path.

   (b) if $|V(K_r)| \geq \left\lceil \frac{\varphi_p}{\varphi_r} \right\rceil$ and $u \notin K_r$:

   $$C_u = \varphi_p \tag{2.6}$$

   $u_{p'} \geq 1$ is trivially not an option for u since that case immediately implies $C_u \leq C_u{}'$. What remains is the case when $u_p' = 0$. In this case to ensure valley-free connectivity u has to create r edges to all other nodes residing in K_r. This would mean ${C_u}' = \varphi_r \left\lceil \frac{\varphi_p}{\varphi_r} \right\rceil$, but even then $C_u < {C_u}'$:

$$C_u = \varphi_p = \varphi_r \frac{\varphi_p}{\varphi_r} \leq \varphi_r \left\lceil \frac{\varphi_p}{\varphi_r} \right\rceil \leq \varphi_r |V(K_r)| = {C_u}' \tag{2.7}$$

The cases (a) and (b) together imply that no node is able to deviate.

2. The VFF topology with $|V(K_r)| < \left\lceil \frac{\varphi_p}{\varphi_r} \right\rceil$ or $|V(K_r)| > \left\lfloor \frac{\varphi_p}{\varphi_r} \right\rfloor + 1$ is NOT a Nash equilibrium: Easily if $|V(K_r)| < \left\lceil \frac{\varphi_p}{\varphi_r} \right\rceil$, then any leaf node u from a rooted tree can lower its cost by joining K_r as:

$$ {C_u}' = \varphi_r |V(K_r)| \leq \varphi_r \left(\left\lceil \frac{\varphi_p}{\varphi_r} \right\rceil - 1 \right) < \varphi_r \frac{\varphi_p}{\varphi_r} = \varphi_p = C_u \tag{2.8}$$

Similarly if $|V(K_r)| > \left\lfloor \frac{\varphi_p}{\varphi_r} \right\rfloor + 1$ then some node $u \in K_r$ can lower its cost by leaving K_r and connect to some other $w \in K_r, w \neq u$ with a p edge.

$$ {C_u}' = \varphi_p < \varphi_r \left(\left\lfloor \frac{\varphi_p}{\varphi_r} \right\rfloor + 1 \right) \leq \varphi_r (|V(K_r)| - 1) = C_u \tag{2.9}$$

3. Finally I show that graphs differing from a VFF cannot constitute Nash equilibria. Let G be an arbitrary graph on which the valley-free connectivity is satisfied, i.e., $\forall u \in G : C_u \neq \infty$. It is obvious that without this property G cannot constitute a Nash equilibrium. In such graph let u be a node whose strategy differs from the nodes in a VFF. The possible cases are:

   (a) u has p edge: The strategy of the nodes in the VFF having p edges is to pay for only a single p edge and nothing more (They can have other p edges attached to them but these are paid by their customers, see Fig. 2.1). Since u differs from this, the corresponding cost function is characterized as:

   $$C_u = \varphi_p u_p + \varphi_r u_r, \quad u_p \geq 1, u_r \geq 0 : (u_p, u_r) \neq (1, 0). \tag{2.10}$$

   Now let the edge (u, w) be a p edge of u. Since w can reach every other player via valley-free path, u also has valley-free paths to all others through w. This means that u can delete its edges except the (u, w) edge, giving $C_u > {C_u}' = \varphi_p$.

(b) u doesn't have a p edge: The cost function of u is:

$$C_u = \varphi_r u_r. \tag{2.11}$$

The valley-free connectivity implies that every path from u to others starts with one r edge which can be followed only by *provider-to-customer* edges. Let t be the number of nodes having no p edges. Such nodes must be the neighbors of u otherwise u cannot reach them. This imposes $u_r \geq t$. Furthermore, for differing from VFF $u_r \neq t$ must also hold. In summary we get $u_r > t$. In this case u has edges (u, w) where w has a provider. Now these edges can be deleted since u can reach w through its provider too. This gives:

$$C_u = \varphi_r u_r > \varphi_r t = C_u{}' \tag{2.12}$$

□

According to Theorem 2.1 aside from two trivial cases – (i) $\varphi_p \geq (n-1)\varphi_r$, when the Nash equilibria is a tree, and (ii) $\varphi_r > \varphi_p$, when the Nash equilibria is a complete graph – the Nash equilibria of this very simple game exhibits some level of structural resemblance with the Internet AS level topology. As described in Sec. 1.2.1.2 on the AS level topology there are only a few nodes in the whole network which can reach every other node without purchasing service (i.e. without having a customer-provider edge), such nodes are called tier-1 (T1). T1s are in a clique at the top of the hierarchy, with having peering agreements set up among them. The rest of the nodes are customers of T1s either in a direct or in an indirect way. These topological features are clearly reflected by the results and now they can be understood as a clear consequence of the VF policy. Theorem 2.1 also gives a rough estimation on the number of T1 nodes as the function of edge costs.

2.2 The Highest Local Preference Game

Moving forward to the next fundamental, economically motivated policy rule I extend the model by adding the HLP rule and define the Highest Local Preference Game.

Players and routing – Let $\mathcal{P}$ be the set of players (identified as the ASs) with cardinality N. Recalling the rule of HLP policy a player always picks from the available VF paths according to its local interest, which is a preference ordering based on the first edge of the path. Customer paths are favored over peer and provider paths and peer paths are favored over provider paths. In this game I use the notations p (or $\overrightarrow{uv}$) and r (or $\overline{uv}$) to denote *customer-provider* and *peer* edges, respectively. This addition is important in order to keep the analysis clear and simple, as in several times referring to edges with their endpoints - instead only their type - is preferable.

Strategies and topology – A strategy for player $u \in \mathcal{P}$ is a vector of the preferred edges to other players in the AS network; i.e. the strategy space is the set $S_u = \{(s_{uv})_{v \in \mathcal{P} \setminus \{u\}} : s_{uv} \in \{0, p, r\}\}$ where $|S_u| = 3^{N-1}$. Easily, player u seeks to contact player v if $s_{uv} \in \{p, r\}$, otherwise $s_{uv} = 0$. Players announce their strategies simultaneously. Any state of the game is represented by an undirected graph $G(s) = (\mathcal{P}, E(s))$ generated by the strategies of the nodes, where $E(s)$ is given by $E(s) = \{\overrightarrow{uv} | s_{uv} = p \wedge s_{vu} = 0\} \cup \{\overline{uv} | s_{uv} \in \{r, p\} \wedge s_{vu} \in \{r, p\}\}$. This settlement of the edges reflects the rational behavior of the ASs as they prefer to create peer edges over customer-provider edges and the instantiation of peer edges requires a *bilateral* agreement between the corresponding players while customer-provider edges can be created *unilaterally*. These can be summarized in Fig. 2.2.

s_{vu} \ s_{uv}	0	p	r
0	0	$\overrightarrow{uv}$	0
p	$\overrightarrow{vu}$	$\overline{uv}$	$\overline{uv}$
r	0	$\overline{uv}$	$\overline{uv}$

Figure 2.2: The created edge according to the strategies of players u and v. The *customer-provider* edges can be created unilaterally, but for *peer* edges both players must bilaterally agree on the establishment. Note that if both players want to create a *customer-provider* edge, then the financial considerations eventuate a *peer* edge.

Payoff – The goal of the players is to minimize their costs, which for a given player u is defined as:

$$C_u(s) = \underbrace{\frac{1}{N}\sum_{v \neq u} d_{G(s)}(u,v)}_{\text{communication cost}} + \underbrace{\varphi_p u_p + \varphi_r u_r}_{\text{link cost}}, \quad v \in \mathcal{P} \tag{2.13}$$

where

$$d_{G(s)}(u,v) = \begin{cases} 0 & \text{if there exists a VF path whose first edge is } \textit{peer} \text{ or } \textit{provider-customer} \\ 1 & \text{if there exists at least one VF path and the first edge of all of them is } \textit{customer-provider} \\ \infty & \text{if a VF path does not exist} \end{cases} \tag{2.14}$$

represents the price of communication between u and v over $G(s)$ in compliance with the VF and HLP policies, φ_p and φ_r are fixed maintenance costs of the provider and peer edges, while u_p and u_r refer to the number of p and r edges of u, respectively. Note that the cost function in Eq. 2.13 is intentionally made as simple as possible for two reasons. First I concentrate purely on the consequences of the two policy rules thus I avoid incorporating cost elements that can mask them. The second reason is simply analytical tractability. Basically the first sum in Eq. 2.13 represents the most simple way of capturing VF and HLP rules and φ_p and φ_r are introduced for setting up a meaningful game (e.g. without attributing costs to the edges the game would end up in producing full graphs) but can be easily justified as inter AS links clearly have maintenance costs. Also note that provider-customer edges are considered to be financed unilaterally by the customer.

The Nash equilibrium of the game is a state such that no player can further reduce her cost by altering her strategy unilaterally. In order to find topologies that are more relevant to a realistic network game I used the following more natural and slightly tailored equilibrium definition for this case:

Definition 6 (Pairwise Stable Nash Equilibrium (PSNE) [54]). *We say $G(s)$ constitutes a pairwise stable Nash equilibrium if (a) it is a Nash equilibrium, (b) $\forall uv \in E(G(s)) : C_u(s) \leq C_u(s') \wedge C_v(s) \leq C_v(s')$, where s' differs from s only in deleting one uv edge from $G(s)$, (c) $\forall uv \notin E(G(s)) : C_u(s) \leq C_u(s') \vee C_v(s) \leq C_v(s')$, where*

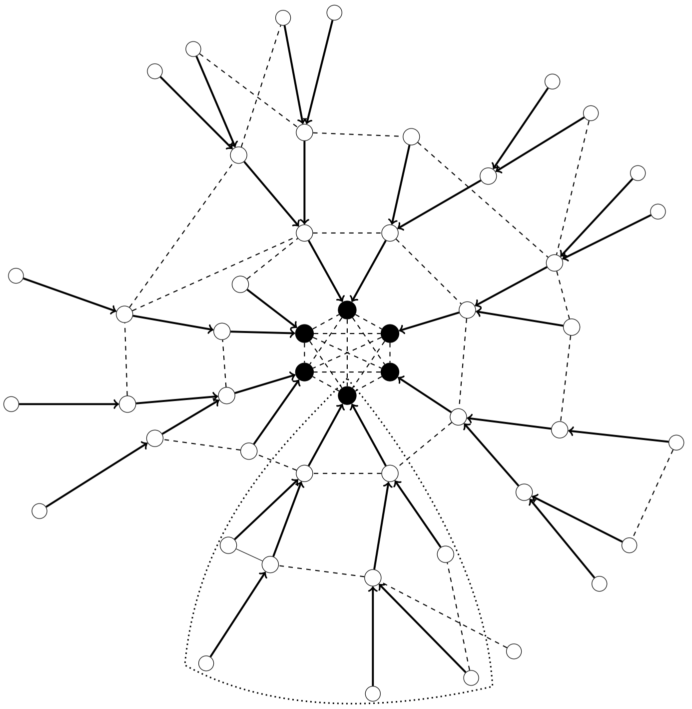

Figure 2.3: An example of the Spiderweb graph, the dashed and directed edges are the peer and customer-provider edges, respectively and the black nodes are the ASs of the clique K, i.e. the tier-1 ASs. The dotted triangle indicates the customer cone of a tier-1 AS.

s' differs from s only in adding uv edge to $G(s)$ and (d) contains no provider loops (cycle of p edges)[1]*.*

Now I am interested in the equilibrium topologies of the game as these topologies will reflect on the consequences of VF and HLP rules. For these claims the following definition is needed.

Definition 7 (Spiderweb graph (Fig. 2.3))**.** *A graph is a Spiderweb graph if it consists of:*

1. *a clique K_r (representing the tier-1 ASs) comprising peer edges only*

2. *trees rooted at some subset of $V(K_r)$ that have customer-provider edges such that the provider in the connection is always closer to the root than the customer*

[1]This requirement is fully in line with the Gao-Rexford conditions [37] ensuring BGP stability.

3. *additional peer edges, such that* $\forall \overline{uv}, \overline{uw} \in G(s) : t(v) \cap t(w) = \emptyset$, *where* $t(x)$ *is the set of nodes in the subtree (i.e. the customer cone) of node* x, *including* x *itself.*

The first claim characterizes all meaningful states (i.e. where all the ASs can communicate with each other) of the above game (and thus the AS topology) by identifying a graph that is omnipresent in the Internet as a subgraph.

Theorem 2.2. *Every meaningful outcome of the game, i.e.,* $\sum C_u \neq \infty$ *contains the Spiderweb graph as a spanning subgraph and every pairwise stable equilibrium (PSNE) of the game is the Spiderweb graph itself.*

Proof. Imagine the customer-provider edges as "directed"[2] edges (from-customer-to-provider). We can say that the subgraph of the customer-provider edges is a spanning DAG, as provider loops are not allowed. By having $\sum C_u \neq \infty$ the sinks of this DAG have to be connected by peer edges in pairs. Hence the set of the sinks correspond to the K_r clique of the Spiderweb graph.

Obviously each AS has a directed customer-provider path to some ASs of K_r. So a spanning forest of the DAG and the K_r clique is a proper spanning Spiderweb graph in the original graph. □

Using Theorem 2.2 I can characterize the pairwise stable equilibria of the game.

Theorem 2.3. *Every pairwise stable equilibrium (PSNE) of the game is the Spiderweb graph.*

Proof. In the following I show that a PSNE topology is necessarily a Spiderweb graph topology. I do this by giving a method which consists of three consecutive steps that refer to the three features described in the Spiderweb graph definition and showing that all of them must be satisfied otherwise there is at least one node that can deviate:

1. *There cannot exist nodes without providers that are not connected by an* r *edge.* I can show this in two steps:

   - *There cannot exist a directed cycle of* p *edges*: Definition 6/d.

[2] This does not influence the traffic as it may still flow in both directions. I use the expression "directed" here only to be able to write the proof in a simpler way.

- *There cannot exist two DAGs with sinks not connected by an r edge*: the cost function of the sinks are infinite because they cannot reach each other via a valley-free route, connecting them with an r edge establishes the valley-free connectivity. As a consequence the sinks must constitute of a clique of r edges which corresponds to the first feature of Spiderweb graph.

2. *There cannot exist nodes with more than one p edge.* Since one p edge is enough for valley-free connectivity, a node can always delete an extra one without increasing its cost function, which corresponds to the second feature of a Spiderweb graph.

3. *There cannot exist an r edge such that $\overline{uv}, \overline{uw} \in G(s) : t(v) \cap t(w) \neq \emptyset$.* This means that either $t(v) \subset t(w)$ or $t(w) \subset t(v)$ and by deleting the edge to the node that is covered by the other the cost function of u can be improved, which corresponds to the third feature of a Spiderweb graph.

□

The following theorem gives a high-level insight into the placement of the peer edges, for which one more definition is necessary.

Definition 8 (Clear-cut Peer Edge (CPE))**.** *An $\overline{uv} \in G(s)$ edge is a clear-cut peer edge if:*

- $\nexists w \in \mathcal{P} : v \in t(w) \wedge \overline{uw} \in G(s)$
- $\varphi_r < \min\{\frac{|t(u)|}{N}, \frac{|t(v)|}{N}\}$.

In other words a CPE is a peer edge that is worth creating for both ASs, since *i*) they have a disjoint set of customers and *ii*) each has a customer cone big enough to make creating the peer edge profitable. A non-CPE is a peer edge that surely won't be created because at least one of the ASs would be better off without that edge.

Theorem 2.4. *If $G(s)$ constitutes a pairwise stable equilibrium (PSNE) of the game then G is a Spiderweb graph with $\max_{u \in K_r} t(u) \leq N(\varphi_p - \varphi_r(|V(K_r)| - 1) + 1)$ and $\forall r \in E_{peer} \setminus E_{K_r}$ is a clear-cut peer edge (CPE).*

Proof. The proof consist of three parts:

1. If G is a PSNE then it is a Spiderweb graph: I have shown this in Theorem 2.3.

2. $\max_{u \in K_r} t(u) \leq N(\varphi_p - \varphi_r(|V(K_r)|-1)+1)$: in other words does not exist node $u \in K_r$ with subtree big enough for leaving the clique. This can be shown easily if we check the actual C_u cost inside K_r and the alternative $C_{u'}$ cost outside K_r. Node u cannot leave if $C_u \leq C_{u'}$:

$$C_u = (|V(K_r)| - 1)\varphi_r \leq C_{u'} = \varphi_p + \frac{N - t(u)}{N} \tag{2.15}$$

   With a simple rearrangement we get:

$$t(u) \leq N(\varphi_p - \varphi_r(|V(K_r)| - 1) + 1). \tag{2.16}$$

3. Every $r \in E_{peer} \setminus E_{K_r}$ is CPE: I prove this indirectly. Assume there exists a PSNE with r which is not a CPE. This means that either (i) $\varphi_{uv} \not< \min\{\frac{|t(u)|}{N}, \frac{|t(v)|}{N}\}$ or (ii) $\exists w \in V(G(s)) : v \in t(w) \wedge \overline{uw} \in G(s)$. For (i) it is easy to see that at least for one node it is worth it to delete the edge, let this node be u and the cost functions before and after deleting $\overline{uv}$ are C_u be $C_{u'}$, respectively:

$$C_u = \varphi_p + u_r\varphi_r + \frac{N - t(u) - t(v)}{N} > C_{u'} = \varphi_p + (u_r - 1)\varphi_r + \frac{N - t(u)}{N} \tag{2.17}$$

   For (ii) it's trivial that for w it is worth it to delete $\overline{uw}$. Let its cost functions before and after the deletion be C_w and $C_{w'}$, respectively:

$$C_w = \varphi_p + w_r\varphi_r + \frac{N - t(w)}{N} > C_{w'} = \varphi_p + (w_r - 1)\varphi_r + \frac{N - t(w)}{N}. \tag{2.18}$$

   Both cases lead to a contradiction.

□

Finally the theorems lead to the following three corollaries.

Corollary 1. *In a PSNE a peer edge appears only if it is in K_r or both its endpoint ASs have sizable customer cones.*

Proof. This is a consequence of Theorem 2.4. □

Corollary 2. *For PSNEs there exists an upper bound for the size of the customer cones of the ASs in* K_r*, or more formally PSNE* $\implies \max_{u\in V(K_r)} t(u) \leq N(\varphi_p - \varphi_r(|V(K_r)|-1)+1)$.

Proof. The cost of a node $u \in V(K_r)$ is $\varphi_r(|V(K_r)|-1)$. However, if u leaves K_r and creates only one customer-provider edge to another node in K_r, its cost would change to $\frac{N-t(u)}{N} + \varphi_p$. Hence for a PSNE

$$\varphi_r(|V(K_r)|-1) \leq \frac{N-t(u)}{N} + \varphi_p, \forall u \in V(K_r), \tag{2.19}$$

and thus

$$\max_{u\in V(K_r)} t(u) \leq N(\varphi_p - \varphi_r(|V(K_r)|-1)+1) \tag{2.20}$$

□

Corollary 3. *In case of a PSNE there exists an upper bound for the size of* K_r *independent from* N*, i.e. PSNE* $\implies |V(K_r)| \leq \frac{\varphi_p+\varphi_r+1+\sqrt{(\varphi_p+\varphi_r+1)^2-4\varphi_r}}{2\varphi_r}$

Proof. According to Corollary 2

$$\max_{u\in V(K_r)} t(u) \leq N(\varphi_p - \varphi_r(|V(K_r)|-1)+1), \tag{2.21}$$

and obviously

$$\frac{N}{|V(K_r)|} = \operatorname*{avg}_{u\in V(K_r)} t(u) \leq \max_{u\in V(K_r)} t(u), \tag{2.22}$$

hence

$$\frac{N}{|V(K_r)|} \leq N(\varphi_p - \varphi_r(|V(K_r)|-1)+1). \tag{2.23}$$

Dividing by N and rearranging the inequality we get:

$$0 \leq -\varphi_r|V(K_r)|^2 + (\varphi_p+\varphi_r+1)|V(K_r)| - 1, \tag{2.24}$$

It's clear from (2.15) that the node with largest $t(u)$ has the best chance for deviation. So to get the tightest upper bound for $|V(K_r)|$ we need to calculate the minimal size of the biggest subtree, i.e. $\min(\max(t(u)))$, and plug it into (2.15). For $\max(t(u)) : \{\text{avg}(t(u)) \leq \max(t(u)) \leq N - |V(K)| + 1\}$, where $\text{avg}(t(u)) = \frac{N}{|V(K)|}$.

By substituting a $t(u) = \text{avg}(t(u))$ in (2.15) we can get the upper bound $|V(K_r)| \leq \frac{\varphi_p+\varphi_r+1\sqrt{(\varphi_p+\varphi_r+1)^2-4\varphi_r}}{2\varphi_r}$.

□

The above theorems deliver the following high-level sketch of the AS topology as a main intuitive message: (*i*) it is a Spiderweb-like graph with a clique (of tier-1 ASs) in the center and trees routed at the nodes of the clique, (*ii*) the peer edges appear more likely between ASs that have sizable customer cones, (*iii*) the size of the clique is constrained by the maintenance cost of peer and customer-provider edges and (*iv*) the largest customer cone size in the nodes of the clique is also driven by these maintenance costs.

2.3 Discussion and double-checking against measurement data

For validating the analytical results I used the AS Relationships dataset of May 2012, provided by CAIDA [2]. Although this dataset received some criticism over the last years, at this moment no other source of data is available that contains more accurate tracing of the peer and customer-provider edges at the AS level.

This dataset contains AS-AS relationships for 41203 ASs with 57158 peer and 83374 customer-provider edges, thus enabling me to build a labeled AS graph. Regarding Theorem 2.2 and 2.3 I investigated the existence of the Spiderweb graph in two steps. First I tracked the customer-provider relationships in a top-down manner proceeding from the top tier-1 clique and kept all the nodes that could be reached, this way I got a 92.5% node coverage which properly validates that the AS graph meets the first two properties (clique inside and trees rooted on the nodes of the clique) of Spiderweb graphs. Second, I examined how typical it is for an AS C with peering neighbors A and B that $t(A) \cap t(B) = \emptyset$. In other words I calculated how typical it is that the customer cones of the peers of an AS are overlapping (this is the direct checking of the third property of Spiderweb graphs see Definition 7). For this I randomly sampled the measured AS graph by choosing 500000 (A, B) node pairs for which $\overline{CA}, \overline{CB}$ exists. In each sample I drew AS C according to a degree-weighted probability function and then I picked the peering neighbors with uniform distribution. The results confirmed that more than 75% of the pairs (Fig. 2.4) have zero overlapping and in other cases the ratio of overlapping vanishes very quickly. These results readily support the claim that the AS level Internet topology is a Spiderweb-like graph.

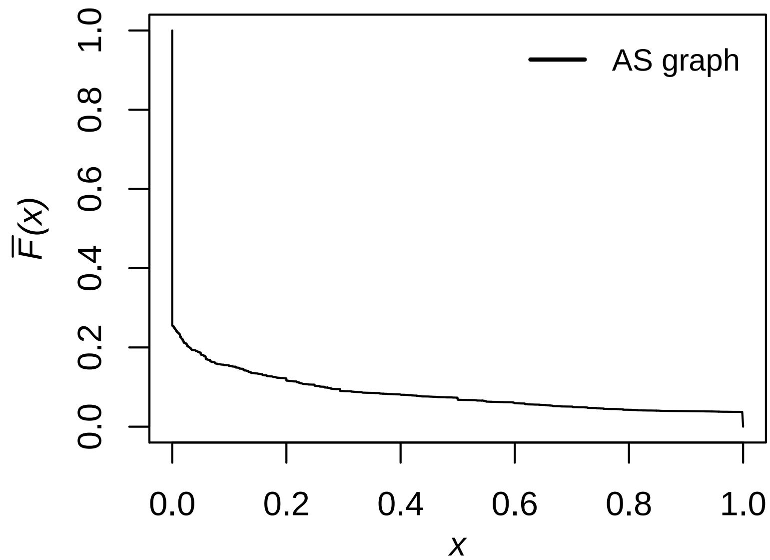

Figure 2.4: CCDF for coverage overlapping of peer edges of an AS defined as $x = \frac{t(A) \cap t(B)}{min\{t(A), t(B)\}}$

After that as a next step I measured the peering likelihood between two ASs as a function of the minimum of their customer cone sizes. The AS graph dataset of Fig. 2.5 shows the empirical probability that two ASs with a given minimum customer cone size ($\min\{t(A), t(B)\}$) are in a peering relationship. The dataset supports that the peering likelihood is highly correlated with the customer cone sizes of the ASs in the peering relationship.

Finally, I present a short argument illustrating predictions on the maximum customer cone size and the max size of the tier-1 clique. To do this I used historical AS datasets from CAIDA. Based on the number of customer-provider and peering relationships I estimated $\varphi_p = \frac{Nc_1}{\text{\#of c-p edges}}$ and $\varphi_r = \frac{Nc_2}{\text{\#of peer edges}}$ with $c_1 = 1.1$ and $c_2 = 0.05$. Using these values I computed the results of the corresponding theorems and measured the max conesize and tier-1 clique size as a function of time in the CAIDA datasets. Fig. 2.6 shows that the rough estimation about the maximal customer cone size in the AS level Internet approximates the measured one based on

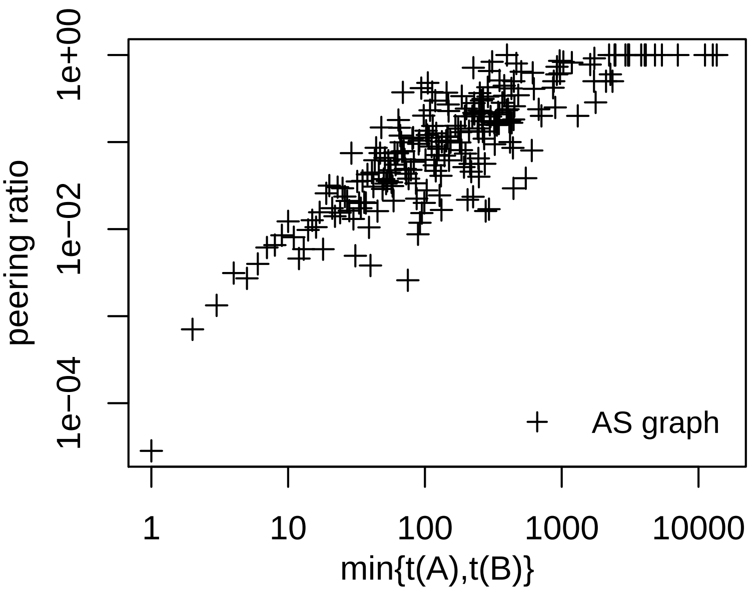

Figure 2.5: Peering likelihood between ASs as the function of their customer cone size.

CAIDA snapshots to a reasonable extent.

Fig. 2.7 shows the prediction of the model regarding the size of the tier-1 clique. Although the simple formulae forecast a more increasing trend, the order of magnitudes are quite the same in both cases.

As a discussion I first kindly call the reader to notice the complementary nature of the game theoretical findings as opposed to the existing causality aware and causality oblivious models. While these models concentrate on degree distribution, clustering, diameter, etc., the game theoretic reasoning gives hints about spanning subgraphs, peering likelihood and constraints on the size of different parts of the network. Also recall that my model is extremely simple and squeezes all maintenance cost related quantities into two constants (φ_r, φ_p). In the light of this simplicity it is remarkable that the model gives practically usable predictions regarding the size of the tier-1 clique and the maximal customer cone of an AS.

One may argue that the results coming out of the analysis are somewhat weak and don't say too much about the AS network. Such criticism may seem to be fair at first, but I find to be important and interesting in itself that the found topological

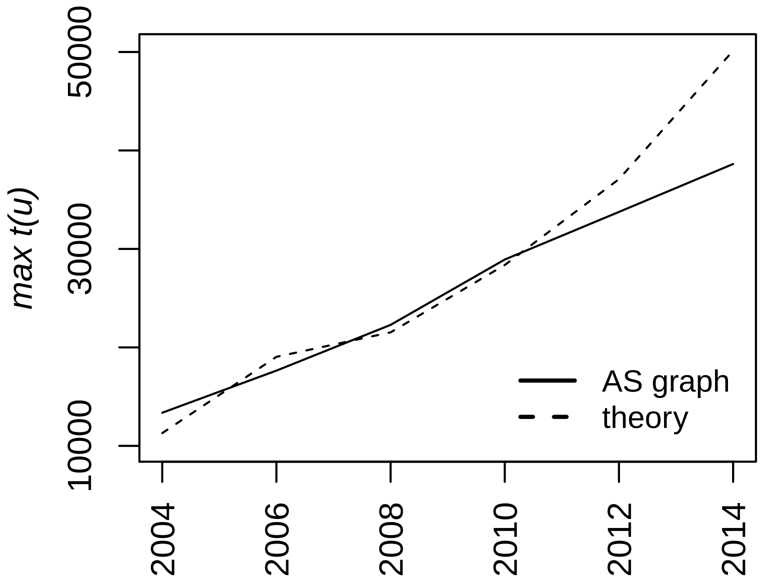

Figure 2.6: Comparing the upper bound for $max\ t(u)$ given by Corollary 2 (theory) with the AS graph over time.

peculiarities (summarized in Theorem 2.2,2.3,2.4 and Corollary 1,2,3) are direct consequences of the used BGP policies and thus will be present on the AS topology as far as these policies are at use. I believe that showing this causality contributes to the very limited amount of information about the Internet AS level topology. Finally, I note that more powerful premises can lead to more precise topology characterization in future works.

2.4 Summary

In this chapter I investigated the consequences of the BGP routing policy through its two fundamental rules, the VF and HLP rules. For this I have designed an analytical game theoretic model in which players are incentivized purely by these rules. In this way I ended up with a model that is still analytically tractable and also allows us to improve our interpretation of the Internet's AS level system as it provides insights complementary to the existing models.

During the analysis I identified a specific subgraph that can be understood as a direct consequence of the VF and HLP policy rules. For future work directions, that

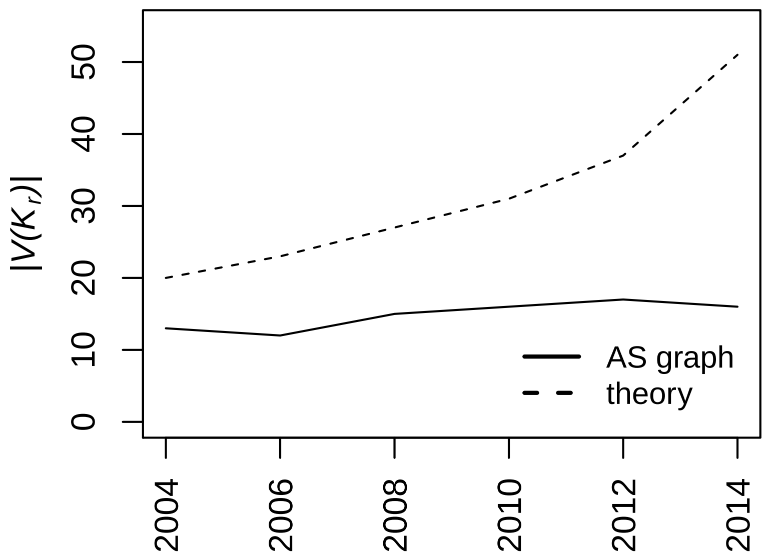

Figure 2.7: Comparing the upper bound for $|V(K_r)|$ given by Corollary 3 (theory) with the AS graph over time.

can shed more light on the AS level topology, I would suggest to dig more deeply into the Best Path Selection Algorithm of BGP and incorporate the AS path length or various sources of traffic engineering in the premises of the models. On the other hand I argue that the basic inter-AS business rules and other technological constraints e.g. the role of IXP-s in the AS-AS connectivity, the multihoming opportunities or security aspects (either in the pure form of supporting secure BGP) can be rich sources of usable premises for future work.

Chapter 3

A Game Theory-Based AS Level Internet Model (YEAS) [J2]

Using the analytical results of Chapter 2, in what follows I define a generative[1] AS topology model called YEAS that is able to create random topologies with similar statistical features. Such a model provides the possibility to furher analyze those statistical features that would be too complex to handle in the game theoretical framework. Besides recovering the usual features of network models (e.g. power-law degree distribution, large clustering, small diameter etc.) I implicitly encode the outcome of the analysis into the node and edge dynamics. Thus finally I require YEAS to produce Spiderweb-like graphs that have correct edge labeling, realistic tier-1 clique size and realistic placement of the peer edges. The framework of YEAS is based on the recently advocated hyperbolic space models presented in [62]. This basically dresses up a very simple hyperbolic model with the findings of Chapter 2.

In the second part of this chapter I introduce the topology generation process of YEAS, which consists of the node placement and edge creation phases. Then I show that YEAS can readily recover *power law degree distribution* and *high clustering coefficient*, which can be observed in a real AS topology. Then I turn to quantities almost never tackled by the existing models. The first one is the *expected customer cone size of the nodes* along with the *cone size distribution of the whole network* and

[1]Generative here means that the created network is the result of a deterministic link creation process, in which the connectivity behavior of nodes is described by an algorithm (the Barabási-Albert model is a good example for generative models).

the second one is the *peering likelihood* of two nodes as the function of their customer cone size. Finally, I compare YEAS generated AS topologies with a potpourri of existing models.

3.1 Topology Generation Process

Node layout

The nodes are distributed (still representing the ASs) quasi-uniformly on the surface of a hyperbolic disk with radius R. This is done by assigning polar coordinates to each node as follows:

$$r = (1/\alpha)\,\mathrm{acosh}\left(1 + [\cosh(\alpha R) - 1]\, U_1\right) \tag{3.1}$$

$$\phi = 2\pi U_2 \tag{3.2}$$

where U_1 and U_2 are independent random variables distributed uniformly over the interval $(0, 1)$ and α is a parameter controlling the heterogeneity of the layout.

Edge creation

To initialize take node u with the lowest radius and initialize a set $\mathcal{K} = \{u\}$[2]. In the first phase take nodes w one by one in an increasing order of their radii r_w and connect them to the others according to the following simple rule. If

$$Q \sum_{v \in \mathcal{K}} l(r_w, \phi_w, r_v, \phi_v) < \min_{v | r_v \leq r_w} l(r_w, \phi_w, r_v, \phi_v), \tag{3.3}$$

then connect w to all nodes in $\mathcal{K}$ with peer edges and add w to $\mathcal{K}$, otherwise connect w to node $\mathrm{argmin}_{v|r_v \leq r_w} l(r_w, \phi_w, r_v, \phi_v)$ with a customer-provider edge. The constant Q is a tunable model parameter controlling the size of $\mathcal{K}$ and

$$l(r_u, \phi_u, r_v, \phi_v) = \mathrm{acosh}(\mathrm{cosh} r_u \mathrm{cosh} r_v - \mathrm{sinh} r_u \mathrm{sinh} r_v \cos(\phi_u - \phi_v)). \tag{3.4}$$

In the second phase every node $u \notin \mathcal{K}$ connects to a node v with a peer edge if

[2] In YEAS this set represents the clique of tier-1 ASs.

$\nexists\ \overrightarrow{uv} \wedge l(r_u, \phi_u, r_v, \phi_v) < \varrho$, where ϱ is a parameter in the interval (0,R) for tuning peering willingness. The pseudocode of the complete process is shown in Figure 3.1 for the sake of reproducibility.

---

initialization:

1. set $r_u = (\text{acosh}(U * (\cosh(R) - 1)) + 1)$ and $\phi_u = 2\pi U$ for each node u;
2. sortedIDs = sort nodes according to r_u;
3. $\mathcal{K} = \{\text{First(sortedIDs)}\}$, $E = \{\}$;

```
Phase 1:
for w → sortedIDs do
    if ∑_{v∈K} l(r_w, φ_w, r_v, φ_v) < min_{v|r_v≤r_w} l(r_w, φ_w, r_v, φ_v) then
        forall v ∈ K do E = E ⋃ wv;
        K = K ⋃ w;
    else
        v = argmin_{v|r_v≤r_w} l(r_w, φ_w, r_v, φ_v);
        E = E ⋃ →wv;
Phase 2:
for ∀(u, v) : u ∉ K ∧ ∄ →uv do
    if l(r_u, φ_u, r_v, φ_v) < ϱ then
        E = E ⋃ uv
Return(E);
```

Figure 3.1: The pseudocode of YEAS

3.2 Features of the Generated Topologies

Due to the fact that YEAS is designed based on the results of the game theoretic analysis and existing hyperbolic models the generated topologies have some implicit features like realistic power law degree distribution, clustering coefficient, customer cone size distribution and peering likelihood. In the following I prove these features one-by-one.

Realistic Power Law Degree Distribution and High Clustering Coefficient

For proving the ability of YEAS to generate a realistic power law degree distribution it is necessary to show that the model generates all edges (u, v) for which $l(r_u, \phi_u, r_v, \phi_v) < \varrho$ but contains edges (u, v) for which $l(r_u, \phi_u, r_v, \phi_v) > \varrho$ with negligible probability. The first statement follows trivially from the edge creation process itself. The second one is the direct consequence of equation (3.7), that is exists if the distance between two points u and v is greater than ϱ, the probability that an edge between these two points is smaller [3] than $\mathrm{e}^{-\delta \mathrm{e}^{\frac{\varrho}{2}}}$. (See the derivation of (3.7) later in this subsection). In the case of some provider edges it is possible that $l > \varrho$ but compared to all the edges the number of these are negligible.

If we simply and arguably ignore (at least from the point of view of degree distribution and clustering) the negligible number of edges of length larger than ϱ, we end up with a model readily analyzed in [62]. Nevertheless I recall the most important claims for making the argument self-contained. The expected degree of a node with coordinates (r, ϕ) is the number of expected points within its $\varrho-$distance vicinity. Equivalently, this coincides with the expected number of points falling inside the intersection of the original $R-$disk and the disk with radius ϱ and center (r, ϕ). In the case when $0.5 < \alpha \leq 1$ the degree of a node with radial coordinate r decays exponentially as the function of r (approximately independently from α), $\bar{k}(r) \sim \mathrm{e}^{-\frac{r}{2}}$, while the node density exponentially increases, $\rho(r) \sim \mathrm{e}^{\alpha r}$. The combination of these two exponentials yield a power law degree distribution $P(k) \sim k^{-2\alpha-1}$, and complement any cumulative degree distribution $\bar{F}(k) \sim k^{-2\alpha}$ [17, 76]. It can be rigorously shown that there exists a constant lower bound on the global clustering coefficient in hyperbolic random graphs, which confirms the high clustering claimed in such networks [43].

Fig. 3.2 shows the cumulative degree distribution of the real AS graph compared to the degree distribution of YEAS with setting $N = 40000$, $Q = 5$, $\alpha = 0.55$, $\varrho = 12.95$ and $R = 18.5$ (I use this setting for all the simulations from now on). The measured AS graph contains 41203 nodes, so I generated a similar-sized topology. The clustering coefficient for the AS graph and for YEAS are both as high as 0.38 and 0.69, respectively. Table 3.1 provides additional metrics for comparison.

[3] For reasonable a value of $N = 40000$ and $\varrho = 12.53$ this probability is 0.00135

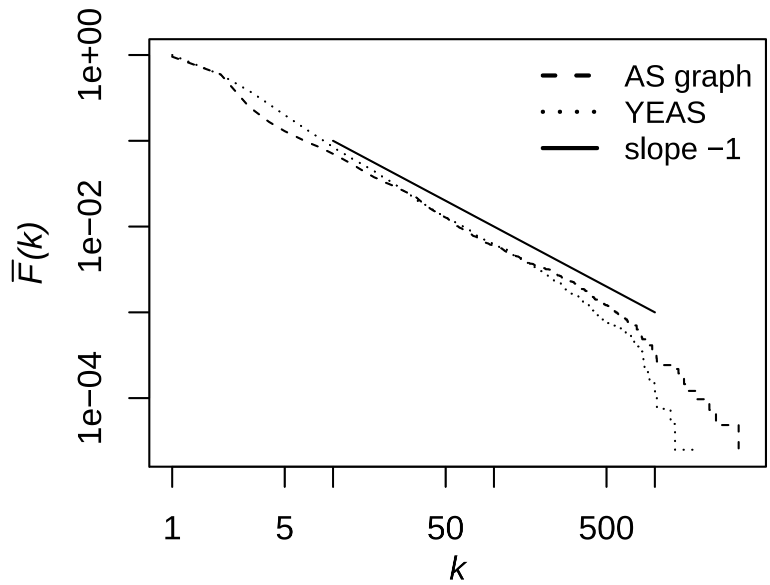

Figure 3.2: CCDF of degrees in the real AS graph and in the YEAS topology.

Realistic Customer Cone Size Distribution

Now turn to the *expected customer cone size of the nodes* and the *cone size distribution of the whole network*. Building upon these results I will be able to quantify the peering probability of the nodes not residing in $\mathcal{K}$. To analyze the average customer cone sizes I temporarily omit the peer edges generated by the model as these do not affect the customer cone sizes.

Let $p(s, \phi, r)$ denote the probability with which node u with radial coordinate s and angle ϕ establishes a provider link to node v with radial coordinate r and angle[4] 0 provided that $s > r$. Recall that node u establishes a customer-provider link to node v if and only if $s > r$ and node v is the closest to node u. The equivalent geometrical meaning of this condition in the generation model is that none of the other $N - 2$ points fall within the intersection of the circle with radius s (with the same center as of the $R-$disk) and the $(s, \phi)-$ centered circle with radius $l(s, \phi, r, 0)$. With using elementary hyperbolic geometrical properties the area of the intersection

[4]The angle coordinate of node v can be assumed to be 0 without loss of generality.

Table 3.1: Comparison of a YEAS generated topology and the CAIDA topology by basic metrics.

Network	Nodes	Edges	C. coef.	Avg. dist.
CAIDA top.	41203	116930	0.38	3.81
YEAS	40000	115309	0.69	4.07

Network	Avg. degree	Diameter	Max. cluster	# Tier-1
CAIDA top.	5.67	14	39327	16
YEAS	5.76	12	40000	16

of these circles can be well approximated (if $l(s, \phi, r, 0)$ is not very close to 0) as

$$A_{intsec} \approx 4\mathrm{e}^{\frac{l(s,\phi,r,0)}{2}} \ . \tag{3.5}$$

Now the probability that none of the other $N - 2$ points fall within this intersection area can be formulated as

$$\left(1 - \frac{A_{intsec}}{A_{R-disk}}\right)^{N-2} \approx \mathrm{e}^{-\delta A_{intsec}} \tag{3.6}$$

where $\delta = \frac{N}{A_{R-disk}}$ is defined as the average node density. The approximation of $p(s, \phi, r)$ is now resulted as

$$p(s, \phi, r) = \mathrm{e}^{-\delta 4 \mathrm{e}^{\frac{l(s,\phi,r,0)}{2}}} \ . \tag{3.7}$$

The expected customer cone size, which is a function of r $\bar{T}(r), r = 0, \ldots, R$, fulfills the following integral equation.

$$\bar{T}(r) = 1 + N \int_{s=r}^{R} \int_{\phi=0}^{2\pi} \bar{T}(s) p(s, \phi, r) \rho(s) \mathrm{d}\phi \mathrm{d}s \ . \tag{3.8}$$

where $\rho(s) = \frac{\sinh(u)}{2\pi(\cosh(R)-1)}$ is the node density function. The intuitive explanation of Eq. 3.8 is the following. The customer cone of a node v with radical coordinate r consists of itself and all other nodes' cones with larger radial coordinate $s > r$ and any angle coordinate ϕ, which are connected to node v by probability $p(s, \phi, r)$. To reformulate this equation the following approximations are used: $\rho(s) \approx \frac{1}{2\pi}\mathrm{e}^{s-R}$,

$\delta = \frac{N}{A_{R-disk}} \approx \frac{N}{\pi e^R}$. By applying these the integral equation becomes

$$\bar{T}(r) = 1 + \frac{\delta}{2} \int_{s=r}^{R} \bar{T}(u) \left(\int_{\phi=0}^{2\pi} e^{-4\delta e^{\frac{l(s,\phi,r,0)}{2}}} d\phi \right) e^s ds \; . \tag{3.9}$$

The inner angle integral can be well approximated as $\frac{1}{\delta} e^{\frac{-s-r}{2}}$. This provides the following (approximate) form of the integral equation:

$$\bar{T}(r) = 1 + \frac{1}{2} \int_{s=r}^{R} \bar{T}(s) e^{\frac{s-r}{2}} ds \; . \tag{3.10}$$

The solution of this integral equation gives the function $\bar{T}(r)$. Unfortunately it can not be solved analytically, however, the solution can be readily characterized as an exponential function. The detailed investigation of the numerical solution confirms this intuition as for a wide range of radial coordinates r the function $\bar{T}(r)$ is approximately proportional to e^{-r}.

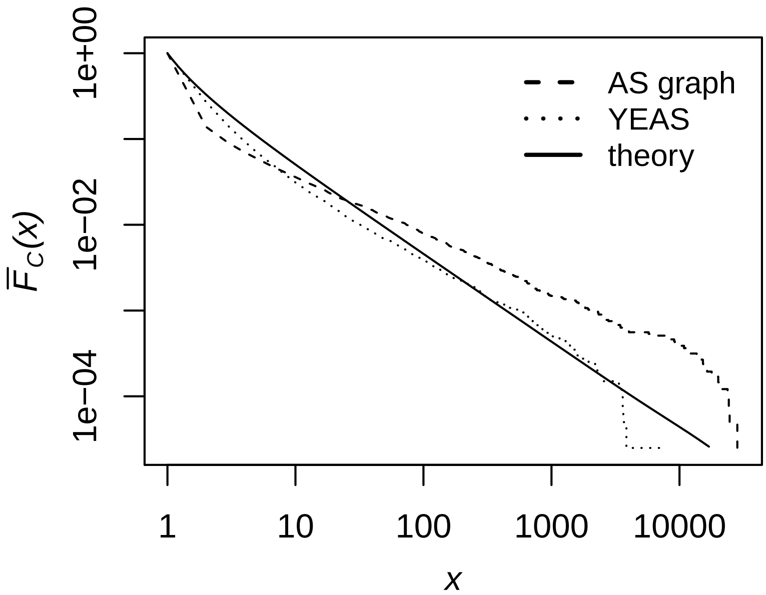

Figure 3.3: CCDF of customer cone sizes in the real AS graph, theory and in the YEAS topology.

Now I can analyze the complement cumulative distribution of cone sizes $\bar{F}_T(x) = P(T > x)$. The CCDF of the cone sizes is approximately[5] power-law with exponent

[5]In case of sparse networks, the conditional distribution of $T(r)$ is Poissonian with mean $\bar{T}(r)$,

-1 provided the expected cone size r is proportional to e^{-r}, that is, $P(T > x) \approx x^{-1}$. Fig. 3.3 readily supports this result as the theoretical result goes hand in hand with the outcome of the simulations. Comparing the real AS topology I detect slightly smaller customer cone sizes produced by YEAS. This is definitely the lack of multihoming in the current version of the model. In the AS graph there are many ASs that have multiple providers in order to increase reliability, thus many AS contributes to the cone size of multiple ASs. Nevertheless the tendency of the cone size distribution is correctly recovered by YEAS, although the exponent is not exactly the same.

Realistic Peering Likelihood

Finally I can turn to analyzing the *peering likelihood* $P_{peering}$ of two nodes having expected customer cone sizes $\bar{T}_1$ and $\bar{T}_2$. More explicitly, I determine the peering probability as the function of $\min(\bar{T}_1, \bar{T}_2)$. For this, first the peering probability of two nodes with radial coordinates r_1 and r_2 as the function of $\max(r_1, r_2)$ is calculated, then the function $r(\bar{T})$ (the inverse function of $\bar{T}(r)$) is applied. Without loss of generality, assume that $r_1 < r_2$. Given r_2, the nodes with smaller radial coordinates $r_1 < r_2$ lie within the circle with radius r_2 and center 0. Clearly, among these nodes those have peer edges to node r_2 which lie in the intersection of this disk and the r_2 centered $\varrho-$radius disk. Therefore, due to the uniform distribution of the nodes the peering probability is the ratio of this intersection area and the area of the 0-centered disk with radius r_2 . Evidently, the peering probability is 1 if $r_2 < \frac{\varrho}{2}$, because in this case the 0-centered disk with radius r_2 is fully contained in the r_2-centered $\varrho-$radius disk. If $r_2 > \frac{\varrho}{2}$ it can be shown by elementary hyperbolic geometry that

$$P_{peering}(r_2, \varrho) \approx \frac{\arccos\left(\frac{\cosh^2(r_2)-\cosh(\varrho)}{\sinh^2(r_2)}\right)}{\pi} + \frac{\exp(\varrho)\arccos\left(\frac{\cosh(r_2)\cosh(\varrho)-\cosh(r_2)}{\sinh(r_2)\sinh(\varrho)}\right)}{\pi\exp(r_2)} . \quad (3.11)$$

$P(T(r) = x) = \frac{\bar{T}(r)^x}{x!}e^{-x}$. Deconditioning this w.r.t. r results in a distribution approximately proportional to x^{-2}, therefore the CCDF of T will be approximately proportional to x^{-1} .

A more detailed analysis of this approximation discloses that it is well approximately proportional to e^{-r_2}. From this a simple approximation can be obtained as

$$P_{peering}(r_2, \varrho) \approx \begin{cases} 1 & \text{if } r_2 < \frac{\varrho}{2} \\ e^{\frac{\varrho}{2}} e^{-r_2} & \text{if } r_2 \geq \frac{\varrho}{2} \end{cases} . \tag{3.12}$$

It follows that

$$P_{peering}(\bar{T}_2, \varrho) \approx \begin{cases} 1 & \text{if } \bar{T}_2 > \bar{T}(\frac{\varrho}{2}) \\ \frac{1}{\bar{T}(\frac{\varrho}{2})}\bar{T}_2 & \text{if } \bar{T}_2 \leq \bar{T}(\frac{\varrho}{2}) \end{cases} . \tag{3.13}$$

This means that the likelihood of peer edges of an AS that have a customer cone size to other ASs which have larger customer cone sizes is proportional to their cone size, and this likelihood tends to be 1, if the cone size is above a certain limit. This characteristic property is also confirmed by the simulations shown in Fig. 3.4 and coincides with results measured on the real AS topology.

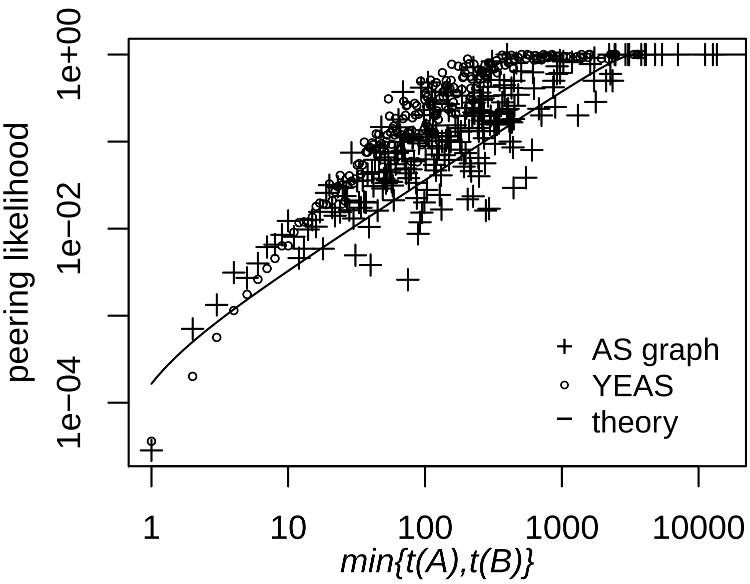

Figure 3.4: Peering likelihood between ASs as the function of their customer cone size (here I extended Fig. 2.5 by adding results about the YEAS generated topology).

The above theoretical results show that YEAS generates realistic complex networks with proper degree distribution, clustering and diameter, yet incorporates the findings of Section 2 as the synthesized topologies are Spiderweb-like (trivially follows from the generation process), with tunable tier-1 clique (through the Q parameter) and realistic peering likelihood.

3.3 Comparison with Existing Work

To compare the AS topologies generated by YEAS I outline its features (Table 3.2) against a potpourri of existing models. The intention behind this by no means comprehensive potpourri was to cover many sides of the spectrum of models used for generating AS topologies. Several of these models were already introduced in Section 1.3.2.2, so here I only shortly recall their most important features. PLRG [4], Inet [97] and dk-Series [68] belong to the group of so called *causality-oblivious [30]* topology generators. All three models introduce various mathematical approaches to generating networks imitating the features of their real-world counterparts in a black-box fashion. The unmissable Barabási-Albert (BA) [9] model is the very first topology generator exhibiting *causality-awareness* thus generating complex networks using processes (preferential attachment and incremental network growing) assumed to take place in networks. BRITE [73] incorporates the findings of power-laws, the skewed node placement and the locality network connection during the topology generation process that can be fine tuned by parameters offering multiple choices. SIMROT [31] generates labeled hierarchical topologies that include BGP relationships. To generate realistic AS topologies it uses a huge number of input parameters which can be determined based on the available AS measurement datasets. GENESIS [65] is a computational game-theoretic model that simulates the AS network formation process and produces different equilibrium topologies. It contains several rules and constraints for realistically mimicking incentives of the ASs like geographical presence, traffic, economic attributes, valley-free routing etc. The model of Holme et al. [47] is similar in spirit to GENESIS but it doesn't contain peering and realistic routing.

The features of these models are summarized in Table 3.2, where in the last row the properties of YEAS are also displayed. One can see that most of the models generate unlabeled graphs thus completely ignore the nature of AS-AS relationships and concentrate only on the complex network face of the AS topology. Closest to my result lie SIMROT and GENESIS which can generate Spiderweb-like labeled AS topologies and realistic peering likelihood, however, these models require a huge number of input parameters adjusted very carefully to produce realistic topologies. Moreover as

GENESIS executes a complex simulation in the background it cannot produce large (> 1000 ASs) topologies within reasonable time limits. The YEAS model can generate large AS topologies with correct labeling and peering statistics while requiring only a handful of input parameters.

Feature	Notation	Feature	Notation
Degree distr.	$\mathbb{P}$	Labeled	$\mathbb{L}$
Clustering	$\mathbb{C}$	Spiderweb-like	$\mathbb{SL}$
Avg. distance	$\mathbb{D}$	Peering likelihood	$\mathbb{PL}$
Large size	$\mathbb{S}$	Few input params	$\mathbb{FP}$

		$\mathbb{P}$	$\mathbb{C}$	$\mathbb{D}$	$\mathbb{S}$	$\mathbb{L}$	$\mathbb{SL}$	$\mathbb{PL}$	$\mathbb{FP}$
C. oblivious	PLRG	✓	-	✓	✓	-	-	-	✓
	Inet	✓	-	✓	✓	-	-	-	✓
	dK-series	✓	✓	✓	✓	-	-	-	✓
Causality aware	BA	✓	-	✓	✓	-	-	-	✓
	BRITE	✓	✓	✓	✓	-	-	-	✓
	SIMROT	✓	✓	✓	✓	✓	✓	✓	-
	H. et al.	✓	✓	✓	✓	-	-	-	-
	GENESIS	✓	✓	✓	-	✓	✓	✓	-
	YEAS	✓	✓	✓	✓	✓	✓	✓	✓

Table 3.2: Comparison of network models.

3.4 Summary

In this chapter I designed a generative AS topology model, based on the previous results about VF and HLP policies called YEAS, that is able to produce networks bearing statistical features similar to the Internet. I also gave a proof for this in the case of power-law distribution, high clustering coefficient, customer cone size distribution and peering likelihood. Finally, I compared the topologies produced by

YEAS with topologies created by CAIDA measurements and several other models along the usual metrics and I got fairly similar results, even though, YEAS was not intended to be a realistic topology generator, but more of a context in which the findings can be verified.

Chapter 4

Topological Consequences of Greedy Navigation [J3, C1]

Greedy navigability is a central issue in the theory of complex networks (Section 1.2.2), as it provides great communication efficiency in small words. A plausible explanation for the favorable navigational properties in such context is the assumed existence of a hidden metric space underneath these networks. Ever since the introduction of Kleinberg's lattice model [58] game theoretical investigation has been focused on explaining how such a network emerges due to the interaction of rational, selfish players. However, existing work assumes shortest path routing when measuring distance between nodes. There are several reasons why this view is limited, but the most important one is that since greedy routing is frequently used in both social and computer networks [13] to great success then it is worth to consider *"Why calculate the shortest path based equilibrium if players know they will route in a greedy manner?"*. Considering greedy routing is also preferable because in the context of the current (and more so the future) Internet, both the need for global topology knowledge hindered by autonomy and policy issues and the linearly scaling router memory requirement [39] raise scalability issues suggesting that shortest path routing has its limitations.

As a consequence, in network games greedy routing looks like a prime candidate to be studied. It is Even-Dar and Kearns [33] who come closest as they present a game played on a Kleinberg-like grid, where nodes create extra edges with a probability decreasing with distance according to a power law. Teh authors determine the equilibrium graphs according to shortest paths then show that greedy routing works reasonably well on these graphs. Therein lies a contradiction: *equilibria are calculated by shortest paths, but players do route in a greedy manner.* In this case, players do not implement their equilibrium paths.

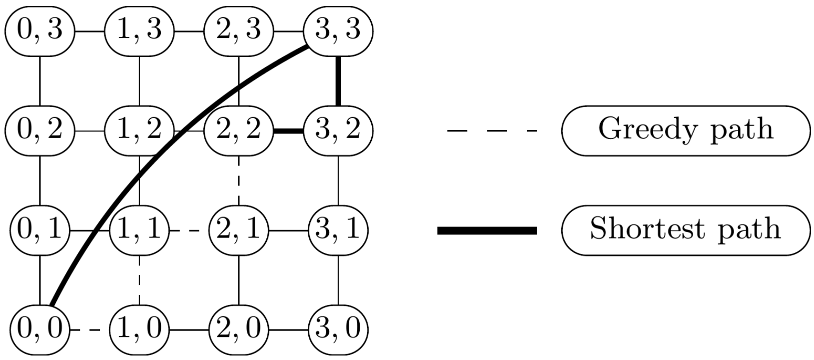

Figure 4.1: Deviation of shortest and greedy paths in the 2D Euclidean grid between nodes $(2, 2)$ and $(0, 0)$.

In the followings I propose the Greedy Network Formation Game (GNFG) to incorporate network creation economics (missing in network models) and navigability (missing in network games) in a single framework. As an extension to network creation games I assume a hidden metric space underneath the network and use the length of greedy paths as the measure of distance between players. Since shortest and greedy paths deviate in essence (see Figure 4.1) this shift substantially change the corresponding equilibria.

Before introducing the game let us recall the pioneering result of Kleinberg [58, 59] on greedy routing on Euclidean lattices, as it is used extensively in the arguments that are based on the analytical results, which are described later in this chapter.

Theorem 4.1. *(Kleinberg) Suppose that network nodes are placed in a 2-dimensional Euclidean lattice. From each node u one shortcut is added to every other node v the topology according to the distribution $P(u, v) \sim l(u, v)^{-r}$, where $l(u, v)$ is the lattice distance between u and v. On this topology the expected delivery time of greedy routing is:*

$$E(t) = \begin{cases} C_1 \log^2(n) & \text{if} \quad r = 2, \\ C_2 n^{(2-r)/3} & \text{if} \quad 0 \le r < 2 \\ C_3 n^{(r-2)/(r-1)} & \text{if} \quad r > 2. \end{cases},$$

As Kleinberg states this result readily generalizes to lattices with higher dimensions.

4.1 The Greedy Network Formation Game

I define the Greedy Network Formation Game (GNFG) using Euclidean lattices, since the question is whether the Kleinberg-like grid network can emerge from the game.

Players, lattice and greedy routing – Let $\mathcal{P}$ be the set of players (identified with network nodes) with cardinality N. Players are placed into the vertices of a D-dimensional $\underbrace{n \times n \times \cdots \times n}_{D \text{ times}}$ lattice (i.e. n is the length of the lattice in each dimensions, so $n^D = N$), which is folded into a torus. The coordinate vector $\mathbf{u} = (u_1, u_2, \ldots, u_D)$ of player u indicates the position of u in the lattice. Distance between two players u and v used in the greedy routing decision is calculated as their lattice distance:

$$l(u,v) \stackrel{\text{def}}{=} l(\mathbf{u},\mathbf{v}) = \sum_{i=1}^{D} \min\{|u_i - v_i|, n - |u_i - v_i|\}. \tag{4.1}$$

A greedy routing step of player u operates over this metric space by choosing the neighbor whose lattice distance is the smallest from target t. If u has no neighbor v such that $l(u,t) > l(v,t)$ then greedy routing is in a local minimum and fails.

Strategies – A strategy for a node $u \in \mathcal{P}$ is to create a set of directed edges (arcs) to other nodes in the network; the strategy space is $S_u = 2^{\mathcal{P}\setminus\{u\}}$. Let s be a strategy vector: $s = (s_0, s_1 \ldots s_{N-1}) \in (S_0, S_1 \ldots S_{N-1})$ and $G(s)$ be the graph defined by the strategy vector s as $G(s) = \bigcup_{i=0}^{N-1}(i \times s_i)$. A mixed strategy is a probability distribution over the above (pure) strategies.

Payoff – The goal of the players is to minimize their cost function which is calculated as follows:

$$C_u(s) = \underbrace{\sum_{u \neq v} d_{G(s)}(u,v)}_{\text{communication cost}} + \underbrace{\varphi|s_u|}_{\text{link cost}}, \quad u,v \in \mathcal{P}, \tag{4.2}$$

where $d_{G(s)}(u,v)$ is the number of nodes involved in the greedy routing process between u and v (including v itself) over $G(s)$ and φ is the constant cost of creating one arc. By definition if greedy routing fails between u and v then $d_{G(s)}(u, v) = \infty$. This setting ensures that we get connected topologies in which there always exists a greedy path between any arbitrary pair of nodes.

Special cases for φ

The following statements characterize the equilibria of the game for special regions of φ.

Theorem 4.2. *If $1 < \varphi = O(N)$, any graph emerging from any NE or social optimum in the GNFG possesses the D-dimensional lattice as a subgraph.*

Proof. I prove this statement by indirection. Suppose that there exists a Nash equilibrium E or social optimum O in which an arc between player u and v is missing, where $l(u, v) = 1$. In this case for all neighbors k of player u, $l(k, v) \geq 1 = l(u, v)$. This means that when u receives a message with destination v, the greedy forwarding process reaches a local minimum and fails causing the cost of player u to be infinite. In this case, it is worth it for u to create arc (u, v) thereby lowering its cost, which means that E cannot be a NE. Since the game has a trivial finite cost solution, namely, the complete graph, O cannot be a social optimum. □

Theorem 4.3. *If $\varphi = \Omega\left(N^{1+1/D}\right)$ then the D dimensional lattice is a unique NE in GNFG.*

Proof. Theorem 4.2 shows that the arcs of the D-dimensional lattice are contained in any possible NE. Now I show that assuming $\varphi = \Omega(N^{1+1/D})$ there exists a sufficiently large lattice for which the NE contains exclusively the arcs of the lattice. It is easy to see that the sum of lattice distances from player u to all other players is given by:

$$\sum_{v \neq u} l(u, v) = Dn^{D-1} \left\lfloor \frac{n^2}{4} \right\rfloor . \tag{4.3}$$

Now consider a state when u (and possibly other players) have additional arcs besides the lattice arcs. It is worth it for u to delete one of its extra arcs if

$$\sum_{v \neq u} d(u, v) + (2D + k_u)\varphi < \sum_{v \neq u} d_e(u, v) + (2D + k_u + 1)\varphi, \tag{4.4}$$

where k_u is the degree of player u, while $d_e(u, v)$ and $d(u, v)$ denote the greedy distance between u and v if there exists an arc between u and e or does not, respectively. From (4.4) we get a satisfactory condition for φ:

$$\sum_{v \neq u} d(u,v) - \sum_{v \neq u} d_e(u,v) < \sum_{v \neq u} l(u,v) - \sum_{v \neq u} d_e(u,v) < \sum_{v \neq u} l(u,v) < \varphi.$$

By using (4.3) we obtain:

$$Dn^{D-1} \left\lfloor \frac{n^2}{4} \right\rfloor < \varphi = \Omega\left(N^{1+1/D}\right)$$

If this condition is met, then u will eventually delete all its extra arcs. □

I illustrate that this bound is tight by calculating the exact threshold for φ if $D = 1$ and N is even and $N \mod 4 = 0$. In this case the average lattice distance is

$$\sum_{v \neq u} l(u,v) = \left(\frac{N}{2}\right)^2, \tag{4.5}$$

and the best possible arc, which minimizes the cost of u is

$$\min_{\forall e} \sum_{v \neq u} d_e(u,v) = \frac{N}{2}\left(\frac{N}{4} + 1\right) - 1. \tag{4.6}$$

This difference of the two gives the condition for φ:

$$O(N^2) = \frac{N^2}{8} - \frac{N}{2} + 1 < \varphi. \tag{4.7}$$

Theorem 4.4. *If $\varphi < 1$ then the full graph is a unique NE in the GNFG.*

Proof. The trivial observation

$$\sum_{v \neq u} d(u,v) - \sum_{v \neq u} d_e(u,v) \geq 1,$$

immediately proves the statement of the theorem. □

4.2 Simplified Greedy Network Formation Game

Deriving results for the GNFG in the region $1 < \varphi = O(N^{1+1/D})$ turns out to be a highly non-trivial problem. For the sake of tractability, in the following I restrict the argument to the one dimensional case and introduce the Simplified Greedy Network Formation Game (SGNFG). I will generalize the results later on. From Theorem 4.2 one can see that any equilibrium or optimum solution of a Greedy Network Formation Game in one dimension always possesses the ring as a subgraph. Therefore I will play the SGNFG on a bi-directional ring, which implies that greedy routing will never fail. On this ring I define the SGNFG as follows: each player can create one directed edge only, which means that the strategy space reduces to a scalar e_u, which indicates the endpoint of the extra edge for player $u \in \mathcal{P}$. This also means that any player u will have a cost of $3\varphi < c_u < \infty$.

When seeking for equilibrium solutions I will use mixed strategies, which means that the strategy of u is a random variable X indicating where to connect its extra edge. As for the distribution $P(X = v) = p_v, p_u = p_{u-1} = p_{u+1} = 0$ and $\sum_{v \in \mathcal{P}} p_v = 1$ hold. Throughout the analysis I - as Kleinberg did - assume that the distribution $P_X \in \mathbb{P}$ is decreasing and monotone, formally, $p_v \leq p_w$ if $l(u,v) > l(u,w)$ and $w \notin \{u-1, u, u+1\}$. This assumptions is fairly realistic, since otherwise the network does not bearing the properties of the underlying space and renders greedy routing meaningless. Let $A(u,v)$ denote the average number of greedy steps required to get from u to v.

Theorem 4.5. *The cost of the optimal solution to the SGNFG is $O(N^2 \log^2(N))$.*

Proof. Theorem 4.1 shows that there exists a distribution X which guarantees that the expected length of greedy paths between any pair of players $A(u,v)$ is $O(\log^2(N))$. Assuming that φ is constant, it follows from this result that the cost for player u is:

$$c_u = \sum_{v \neq u} A(u,v) + 3\varphi = O(N \log^2(N)),$$

which gives $\sum_{u \in \mathcal{P}} c_u = O(N^2 \log^2(N))$ as the optimal solution. □

Price of anarchy in the SGNFG

Theorem 4.6. *The bi-directional Möbius ladder [46], in which the extra edges of each player are directed at exactly the opposite player on the one dimensional ring (see Figure 4.2), is always a Nash equilibrium with total cost* $\frac{N^3}{2}$*. The price of anarchy in the SGNFG is therefore of* $\Omega\left(\frac{N}{\log^2(N)}\right)$.

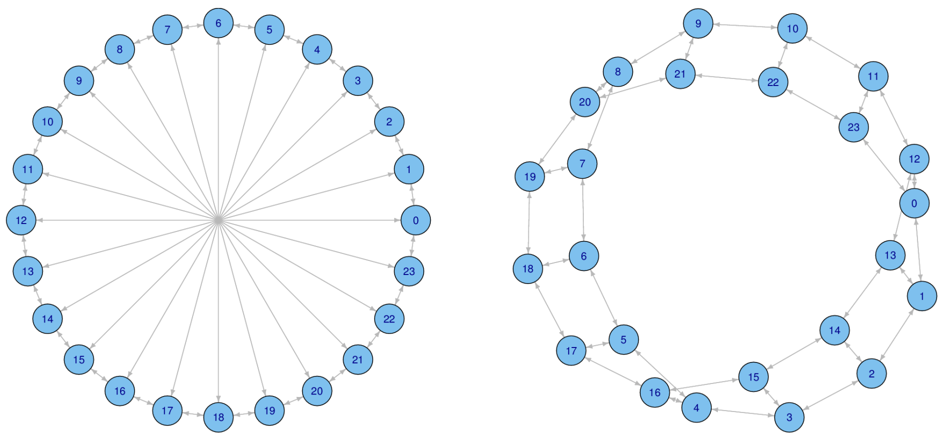

Figure 4.2: Möbius ladder with 24 players.

Proof. If each player is connected to the opposite player in the ring, greedy forwarding starting from player u cannot use the shortcuts of any other player, just its own. Therefore it is enough to show player u's best choice in the empty ring. Without loss of generality consider the player with coordinate 0 and assume that its extra edge is connected to a player whose coordinate $\delta \leq \frac{N}{2}$. To simplify the argument I identify the players with their coordinates and use their name and coordinate interchangeably. For the corresponding greedy step, player 0 chooses:

- player 1 towards players $1, 2, \ldots, \lceil\frac{\delta}{2}\rceil$,
- player δ towards players $\lceil\frac{\delta}{2}\rceil + 1, \ldots, \lfloor\frac{N+\delta}{2}\rfloor - 1$,
- player $N-1$ towards players $\lfloor\frac{N+\delta}{2}\rfloor, \ldots, N-1$.

Let $a = \lceil\frac{\delta}{2}\rceil$ and $b = \lfloor\frac{N+\delta}{2}\rfloor$. The average distance of player 0 from all the other players can be calculated as:

$$\frac{1}{N}\left(\sum_{i=1}^{\lceil\frac{\delta}{2}\rceil} i + \sum_{i=\lceil\frac{\delta}{2}\rceil+1}^{\delta}(1+\delta-i) + \sum_{i=\delta+1}^{\lfloor\frac{N+\delta}{2}\rfloor-1}(1+i-\delta) + \sum_{i=\lfloor\frac{N+\delta}{2}\rfloor}^{N-1}(1+N-1-i)\right)$$
$$= \frac{1}{N}\left(\sum_{i=1}^{a} i + \sum_{i=1}^{\delta-a} i + \sum_{i=2}^{b-\delta} i + \sum_{i=1}^{N-b} i\right) = \frac{N-2}{2N} + \frac{(a^2+(\delta-a)^2+(b-\delta)^2+(N-b)^2)}{2N}. \tag{4.8}$$

With the substitution of a and b:

$$a^2+(\delta-a)^2+(b-\delta)^2+(N-b)^2 = \frac{\delta^2}{2}+\frac{(N-\delta)^2}{2}+\frac{1}{2}\mathbb{I}(N \text{ odd})+\mathbb{I}(N \text{ even and } \delta \text{ odd}).$$

Using the inequality between arithmetic and quadratic means we get:

$$\frac{\delta^2}{2}+\frac{(N-\delta)^2}{2} \geq \left(\frac{\delta+(N-\delta)}{2}\right)^2 = \left(\frac{N}{2}\right)^2.$$

In order for u to optimize its cost it needs to find the δ that minimizes this expression. Since u cannot use arcs of others δ can be chosen as the topology would be an empty ring. The equality holds if and only if $\delta = \frac{N}{2}$ meaning that the best choice is the player at the exact opposite position in the ring, which implies that the Möbius ladder [46] is a Nash equilibrium. In this case each player has a cost $c = \frac{N^2}{8} + O(N)$, and the total cost is $\sum_{u\in\mathcal{P}} c_u = \frac{N^2}{8}N + O(N^2)$. The worst case Nash equilibrium of the game is therefore of $\Omega(N^3)$, which implies that the price of anarchy is $\Omega\left(\frac{N^3}{N^2\log^2(N)}\right) = \Omega\left(\frac{N}{\log^2(N)}\right)$. □

Price of Stability in the SGNFG

To determine the price of stability in the SGNFG I now seek for the best available Nash equilibrium. In the following I give a counting argument on this equilibrium in the space of mixed strategies. Instead of working with particular graph instances and aggregating their properties for getting the expected costs, I conduct the analysis on a “stochastic” graph. When investigating the outcome of a greedy step at a given player I treat the graph such that the extra outgoing edge is generated at the time the greedy routing step is performed by the player. This way I bypass the tedious work with exact deterministic graph structures and determine the expected costs in a direct manner similarly to [59]. To obtain equilibrium solutions the following two lemmas are needed. Let $A_e(u,v)$ denote the average number of greedy steps from u to v if player u has its extra edge connected to player e.

Lemma 1. *The larger the distance between two players, the more number of greedy steps is needed to travel between them on average. Formally: If $l(u,v) \leq l(u,w)$, then $A(u,v) \leq A(u,w)$ for $u,v,w \in \mathcal{P}$.*

Proof. To prove this statement it is enough to show that if $l(v,w) = 1$ and $l(u,w) > l(u,v)$, then $A(u,v) < A(u,w)$. Without loss of generality let $u = 0$. I prove this statement by induction on v. If v is 0, +1 or $N-1$ the statement is true. Assume that for $0, 1, N-1, \ldots, (v-1), N-(v-1)$ the statement holds. I prove the statement for points v and $N-v$, however it's enough to show for $0 \leq v < v+1 \leq \frac{N}{2}$, because $A(0,v) = A(0,N-v)$. This means that the target is to show that:

$$0 \leq A(0,v+1) - A(0,v) = \sum_{j=2}^{N-1} [A_j(0,v+1) - A_j(0,v)]p_j.$$

First, notice that player 0 uses its extra edge as the first hop towards player v if the extra edge is connected to player j and $2 \leq j \leq 2v-2$, otherwise player 0 uses the edge of the lattice which is connected to player 1, since $v < \frac{N}{2}$. Therefore:

$$\begin{aligned} A(0,v) &= \sum_{j=2}^{N-2} A_j(0,v)p_j = \sum_{j=2}^{2v-2} (1 + A(j,v))p_j + (1 + A(1,v)) \sum_{j=2v-1}^{N-2} p_j, \text{ and} \\ A(0,v+1) &= \sum_{j=2}^{N-2} A_j(0,v+1)p_j = \sum_{j=2}^{2v} (1 + A(j,v+1))p_j + (1 + A(1,v+1)) \sum_{j=2v+1}^{N-2} p_j. \end{aligned}$$

If $j > 2v$ then $A_j(0,v+1) - A_j(0,v) = (1 + A(1,v+1)) - (1 + A(1,v)) = A(0,v) - A(0,v-1) \geq 0$ holds because of the induction hypothesis. If $j = 2v$ then $A_{2v}(0,v+$

$1) - A_{2v}(0,v) = A(2v, v+1) - A(1,v) = 0 \geq 0$. If $j = 2$ or $j = 2v-1$, then

$$\begin{aligned}&[A_2(0,v+1) - A_2(0,v)]p_2 + [A_{2v-1}(0,v+1) - A_{2v-1}(0,v)]p_{2v-1}\\ &\quad= [A(2,v+1) - A(2,v)]p_2 + [A(2v-1,v+1) - A(1,v)]p_{2v-1}\\ &\quad= [A(0,v-1) - A(0,v-2)](p_2 - p_{2v-1}) \geq 0\end{aligned}$$

holds because of the induction hypothesis and the third assumption. Finally, if $j \in [3, 2v-2]$ then by investigating them pairwise, $j \in \{v-a, v+1+a\}$ for $a = 0, 1, \ldots, v-3$ we have:

$$\begin{aligned}&[A_{v-a}(0,v+1) - A_{v-a}(0,v)]p_{v-a} + [A_{v+a+1}(0,v+1) - A_{v+a+1}(0,v)]p_{v+a+1}\\ &= [A(v-a,v+1) - A(v-a,v)]p_{v-a} + [A(v+a+1,v+1) - A(v+a+1,v)]p_{v+a+1}\\ &= [A(0,a+1) - A(0,a)](p_{v-a} - p_{v+a+1}) \geq 0.\end{aligned}$$

Hence $A(0,v+1) - A(0,v) = \sum_j [A_j(0,v+1) - A_j(0,v)]p_j \geq 0$ implying that $A(0,v)$ is monotonically increasing. □

Lemma 2. *If player u chooses a more distant player to connect its extra edge then the cost of u reduces. Formally:*

$$\sum_{x\in\mathcal{P}} A_v(u,x) \geq \sum_{x\in\mathcal{P}} A_w(u,x), \text{ if } l(u,v) \leq l(u,w).$$

Proof. To prove this statement it is enough to consider the case when $l(v,w) = 1$. Without loss of generality let $u = 0$ and assume that $w = v+1 \leq \frac{N}{2}$. Similarly to (4.8):

$$\begin{aligned}\sum_{x\in\mathcal{P}} A_v(0,x) &= (N-1) + \sum_{x=1}^{\lceil \frac{v}{2}\rceil} A(1,x) + \sum_{x=\lceil\frac{v}{2}\rceil+1}^{\lfloor\frac{N+v}{2}\rfloor-1} A(v,x) + \sum_{x=\lfloor\frac{N+v}{2}\rfloor}^{N-1} A(N-1,x),\\ \sum_{x\in\mathcal{P}} A_{v+1}(0,x) &= (N-1) + \sum_{x=1}^{\lceil\frac{v+1}{2}\rceil} A(1,x) + \sum_{x=\lceil\frac{v+1}{2}\rceil+1}^{\lfloor\frac{N+v+1}{2}\rfloor-1} A(v+1,x) + \sum_{x=\lfloor\frac{N+v+1}{2}\rfloor}^{N-1} A(N-1,x).\end{aligned}$$

The second sum in the second equation can be rewritten as:

$$\sum_{x=\lceil\frac{v+1}{2}\rceil+1}^{\lfloor\frac{N+v+1}{2}\rfloor-1} A(v+1,x) = \sum_{x=\lceil\frac{v+1}{2}\rceil}^{\lfloor\frac{N+v+1}{2}\rfloor-2} A(v+1,x+1) = \sum_{x=\lceil\frac{v+1}{2}\rceil}^{\lfloor\frac{N+v+1}{2}\rfloor-2} A(v,x).$$

Hence

$$\sum_{x\in\mathcal{P}} A_{v+1}(0,x) - \sum_{x\in\mathcal{P}} A_v(0,x) = \sum_{x=\lceil\frac{v}{2}\rceil+1}^{\lceil\frac{v+1}{2}\rceil} A(1,x) + \sum_{x=\lceil\frac{v+1}{2}\rceil}^{\lceil\frac{v}{2}\rceil} A(v,x) - \sum_{x=\lfloor\frac{N+v+1}{2}\rfloor-1}^{\lfloor\frac{N+v}{2}\rfloor-1} A(v,x) - \sum_{x=\lfloor\frac{N+v}{2}\rfloor}^{\lfloor\frac{N+v+1}{2}\rfloor-1} A(N-1,x).$$

Depending on the parities of v and N the difference is:

	N is even	N is odd
v is even	$A(1,\frac{v}{2}+1) - A(v,\frac{N+v}{2}-1)$	$A(1,\frac{v}{2}+1) - A(N-1,\frac{N+v-1}{2})$
v is odd	$A(v,\frac{v+1}{2}) - A(N-1,\frac{N+v-1}{2})$	$A(v,\frac{v+1}{2}) - A(v,\frac{N+v}{2}-1)$

Using the fact that $A(u,v)$ is a monotonic and increasing function of $l(u,v)$, none of the differences are positive if $v+1 \leq \frac{N}{2}$, which is the maximum distance in the ring. □

From Lemma 1 and 2 one can intuitively conclude that choosing a distant node as the endpoint of the extra edge reduces the cost of a player. Therefore, in case of mixed strategies heavy-tailed distributions are of special interest because of the monotonicity of p_v, since they provide the highest chance for connecting the extra edge to a distant node. Now I show that the best strategy a player can have at any stage of the game is to uniformly choose among other players. The cost of player u is:

$$c_u = \sum_{v\neq u} A(u,v) = \sum_{v\neq u} \sum_{j\in\mathcal{P}\setminus\{u-1,u,u+1\}} A_j(u,v)p_j,$$

which can be transformed to

$$c_u = \sum_{s\in S} p_s f(s). \tag{4.9}$$

Theorem 4.7. *If $f(s)$ is a monotonically decreasing function of $l(u,u+s)$ then in any given situation of the SGNFG, player u's best response to the strategies of the other* *players is to choose the endpoint of its extra link uniformly at random. Formally:* $\operatorname{argmin}_{p\in\mathbb{P}} \sum_{s\in S} p_s f(s) = \text{uniform}$.

Proof. I prove this statement by indirection. Let us suppose that p does not have a uniform distribution, then let be player v and w, that $l(u,v) < l(u,w)$ and $p_v > p_w$, and for every other $x \in \mathcal{P} : l(u,x) \leq l(u,v)$ or $l(u,w) \leq l(u,x)$. Let $T_1 = \{x : l(u,x) \leq l(u,v)\}$ and $T_2 = \{x : l(u,x) \geq l(u,w)\}$, it is easy to see that $S = T_1 \uplus T_2$. The cost c_u can be written as follows:

$$\sum_{s\in S} p_s f(s) = \sum_{x\in T_1} p_x f(x) + \sum_{y\in T_2} p_y f(y) = \mathrm{P}(T_1)\frac{\sum_{x\in T_1} p_x f(x)}{\sum_{x\in T_1} p_x} + \mathrm{P}(T_2)\frac{\sum_{y\in T_2} p_y f(y)}{\sum_{y\in T_2} p_y} \tag{4.10}$$

While f is decreasing then $\min_{x\in T_1} f(x) \geq \max_{y\in T_2} f(y)$ and obviously f isn't constant, so $\max_{x\in T_1} f(x) > \min_{y\in T_2} f(y)$, hence

$$\frac{\sum_{x\in T_1} p_x f(x)}{\sum_{x\in T_1} p_x} > \frac{\sum_{y\in T_2} p_y f(y)}{\sum_{y\in T_2} p_y} \tag{4.11}$$

Now a new distribution can be created p^*, which is also decreasing, so $p^* \in \mathbb{P}$, and

$$\frac{\sum_{x\in T_1} p^*_x f(x)}{\sum_{x\in T_1} p^*_x} = \frac{\sum_{x\in T_1} p_x f(x)}{\sum_{x\in T_1} p_x} \text{ and } \frac{\sum_{y\in T_2} p^*_y f(y)}{\sum_{y\in T_2} p^*_y} = \frac{\sum_{y\in T_2} p_y f(y)}{\sum_{y\in T_2} p_y} \tag{4.12}$$

but $p^*(T_1) < p(T_1)$ and $p^*(T_2) > p(T_2)$, so $\sum_{s\in S} p_s f(s) > \sum_{s\in S} p^*_s f(s)$ p^* can be defined as if $x \in T_1$ then $p^*_x = \alpha p_x$ and $y \in T_2$ then $p^*_y = \beta p_y$, where $\alpha < 1$ and $\beta > 1$. There are two conditions to determine α and β:

$$\alpha \mathrm{P}(T_1) + \beta \mathrm{P}(T_2) = 1 \text{ and } \alpha p_v \geq \beta p_w \tag{4.13}$$

If $\alpha = \frac{p_w}{p_w \mathrm{P}(T_1) + p_v \mathrm{P}(T_2)}$ and $\beta = \frac{p_v}{p_w \mathrm{P}(T_1) + p_v \mathrm{P}(T_2)}$ then p^* has the properties was discussed above. It means if p is not the uniform distribution, it cannot be optimal. □

Corollary 4. *The only Nash equilibrium of the SGNFG with mixed strategies is the case when all players connect their extra edge uniformly at random.*

Now that there is a clue for the structure of the network in equilibrium states, the cost of such equilibria can be calculated by borrowing again from the results of Kleinberg.

Theorem 4.8. *The best Nash equilibrium of the SGNFG is $\Omega(N^{8/3})$, therefore the price of stability is $\Omega\left(\frac{N^{2/3}}{\log^2(N)}\right)$.*

Proof. It is shown in [58] that if every player connects its extra edges by the uniform distribution ($p_v \sim l(u,v)^0$) then the expected length of the greedy paths between two players is of $\Omega(N^{2/3})$. From this, the total cost is $\Omega(N^{8/3})$ and the price of stability is $\Omega\left(\frac{N^{2/3}}{\log^2(N)}\right)$. □

According to [58] from the distributions of the form $p_v \sim l(u,v)^{-r}$, $r = D$ eventuates the only setting that produces a small-world topology, where the length of the greedy paths scales polylogarithmically with N. The conclusion from Corollary 4 is that $r = 0$ is the only possible setting to obtain a Nash equilibrium. This immediately leads to the following observation:

Proposition 1. *Kleinberg's optimal setting is not a Nash equilibrium, therefore small-world equilibrium solution does not exists for the SGNFG.*

This means that the previous investigations that first determine the equilibrium graphs according to shortest path and then show the efficiency of greedy routing actually work with a model that cannot emerge in a self-organizing way.

4.3 Generalization of the Results

In the previous section I presented the in-depth analysis of the SGNFG and drew the negative conclusion that incorporating greedy routing within the network creation game takes the equilibrium topologies very far from the social optimum. Moreover I showed that a small-world network cannot be an equilibrium solution of the game. One might argue that the results may be valid only within the simple framework of the SGNFG. Here I take a quick look at the statements in more general settings of the game.

Multiple edges

In the simplified setting a player could have only one extra edge in addition to its lattice edges, however, in a general case a player can have multiple edges. Now I argue that if each player u can only afford a constant number of edges C_u then the equilibrium solution remains qualitatively the same. In the multiple edge case the cost of player u can be transformed to the form $c_u = \sum_{s \in S} p_s f(s)$ similarly to the single extra edge case (see (4.9)). Theorem 4.7 proves that the uniform distribution minimizes such cost functions. This also means that the best strategy that player u can have is to distribute its C_u edges uniformly in the lattice.

Distance-dependent link costs

In a general setting the cost of an edge may depend on the distance between its endpoints, which gives the more complex cost function

$$c_u = \sum_{v \in \mathcal{P}} p_v \left(\varphi(u,v) + \sum_{x \in \mathcal{P}} A_v(u,x) \right) = \sum_{s \in S} p_s f(s). \tag{4.14}$$

In this case however, $f(s)$ is not necessarily monotonic, which means that the previous argument about the uniform distribution being the only NE does not work in this case. However it can be shown that a distribution which eventuates *strict* Nash equilibrium is uniform until a given lattice distance and zero otherwise.

Theorem 4.9. *If $p \in \mathbb{P}$ then $\exists f()$ for which p is a weak Nash equilibrium. If $p \in \mathbb{P}$ is a strict Nash equilibrium, then $\exists r \in (0,1)$ so that $p_s \in \{0, r\}$.*

Proof. The proof of the first statement is very simple. Let $\varphi(u,v) = C - \sum_{x \in \mathcal{P}} A_v(u,x)$, hence $c_u = C \sum_{v \in \mathcal{P}} p_v = C$ independent of the distribution of the extra link. So p_s is a weak NE.

The second statement can be proved similarly as Theorem 4.7. The only difference is that $\frac{\sum_{x \in T_1} p_x f(x)}{\sum_{x \in T_1} p_x}$ is not necessarily larger than $\frac{\sum_{y \in T_2} p_y f(y)}{\sum_{y \in T_2} p_y}$, in this case α and β should be such that $\alpha > 1$ and $\beta < 1$. The solution of p^* is not worse than the solution of p. If there exist v and w that $p_v > p_w > 0$, than p_s can't be a strict Nash equilibrium. □

Proposition 2. *A small-world topology can't be a strict Nash equilibrium.*

Proof. From Theorem 4.9 we know that if p is a strict NE then $\exists\delta$ that $p_v = r$ for $l(u,v) < \delta$ and $p_v = 0$ for $l(u,v) > \delta$. Now the results can be applied from [58] for the uniform distribution (similarly as in the proof of Theorem 4.8) on the region where $p_v = r$, hereby getting a lower bound $\frac{1}{N}\sum_{v\in\mathcal{P}} A(u,v) = \Omega(\delta^{2/3})$. To achieve a small-world network, δ has to be smaller than $(\log n)^\alpha$, but in this case if $l(u,v)$ is big enough then $A(u,v) \sim \frac{n}{(\log n)^\alpha}$, implying that the average greedy distance can't be polylogarithmic. $\square$

Multiple dimensions

For the sake of simplicity I carried out the proofs for the one dimensional case. In the following I illustrate that the argument can be extended to the finite D-dimensional case. First observe that the simple statement of Lemma 1 (the more distant a player is the more greedy steps are needed to travel between them) is the only result where the one dimensional assumption is exploited. Now I illustrate that Lemma 1 readily generalizes to higher dimensions.

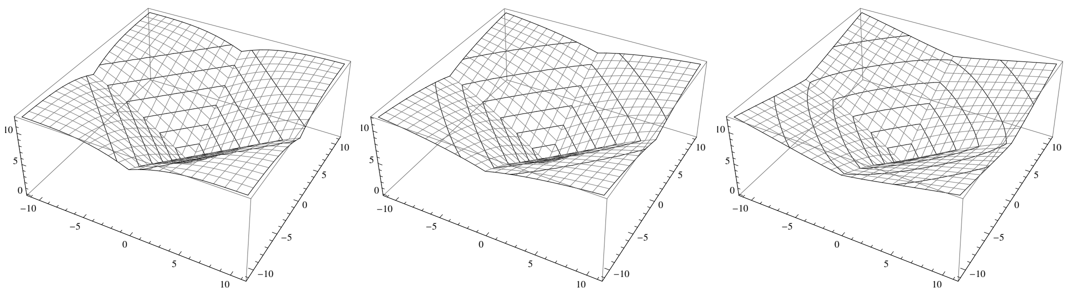

Figure 4.3: The average number of greedy steps ($A(u,v)$) between a reference player $u = (0,0)$ and the other players in the two dimensional lattice, if $p_v \sim l(u,v)^0$ (left), $p_v \sim l(u,v)^{-1}$ (center), $p_v \sim l(u,v)^{-2}$ (right).

Figure 4.3 shows the average number of greedy steps ($A(u,v)$) required to travel between a reference player (at the center of the figure) and the other players in the two dimensional lattice. If w denotes the neighbor of u who is closest to v, then $A(u,v)$ can be calculated by following recursion:

$$A(u,v) = \sum_{x\in\mathcal{P}} p_x A_x(u,v) = \sum_{x:l(x,v)<l(w,v)} p_x(1+A(x,v)) + \left(1 - \sum_{x:l(x,v)<l(w,v)} p_x\right)(1+A(w,v)).$$

Figure 4.3 supports the conjecture that $A(u, v)$ grows with the lattice distance if the game is played in multiple dimensions.

Conjecture 1. *Small-world topologies cannot emerge as equilibria from the SGNFG even if the dimension of the lattice is raised to an arbitrary constant value. This means that the existence of small-worlds cannot be economically justified under the Kleinberg-like constant dimensional grid-based models.*

Note that the situation fundamentally changes when the dimension can depend on the number of players. For example if the number of dimensions can be of $\Omega(\log N)$ the number of steps needed to travel between players trivially drops to $O(\log N)$. Fraigniaud and others show a similar phenomenon for $\log \log N$ [36].

4.4 Hyperbolic Space

The results support the claim that small-world networks cannot be equilibrium solutions of the Greedy Network Formation Game even if the game is played under fairly generalized conditions. So the question arises: *"How can small-world topologies emerge?"* What can then be the incentive of the players to eventuate an asymptotically optimal solution? The recent triumph of hyperbolic space based models in explaining the intricate properties of the internet's topology [62] and the successive application of the hyperbolic space based techniques in problems related to distributed routing techniques like greedy embeddings [60] lead to the idea of investigating the GNFG in the hyperbolic space. In the hyperbolic space I prove that socially optimal solutions can readily emerge from the Greedy Network Formation Game.

For the investigation I use the Poincaré disk model [7] of the two dimensional hyperbolic space as this model makes the calculations easier here. In this space player u has a coordinate vector $\mathbf{u} = u_1, u_2 \in [0, 1)$ and the distance between u and v is calculated according to the Poincaré distance function:

$$d_p(u, v) \stackrel{\text{def}}{=} d_p(\mathbf{u}, \mathbf{v}) = \operatorname{arccosh}\left(1 + 2\frac{||\mathbf{u} - \mathbf{v}||^2}{(1 - ||\mathbf{u}||^2)(1 - ||\mathbf{v}||^2)}\right),$$

where $||x||$ stands for the Euclidean norm of x.

The players are placed at equal distances from each other similarly to the case of the two dimensional Euclidean lattice, thus the players will be located in the vertices of a so-called hyperbolic tessellation (see Figure 4.4). A tessellation [56] can be characterized by a pair (ν, κ) where ν stands for the vertex number of its constituent polygons and κ denotes the number of meeting polygons at a given vertex. For (ν, κ), $\frac{1}{\nu} + \frac{1}{\kappa} < \frac{1}{2}$ must hold. A graph $T(V, E)$ can be constructed from the tessellation if its vertices are considered as the vertices of the graph and the sides of the polygons as edges.

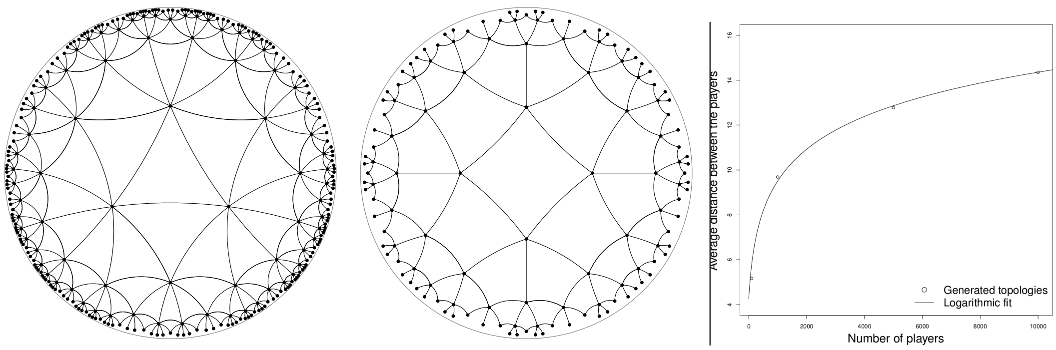

Figure 4.4: (3,8) (left) and (4,5) (middle) hyperbolic tessellations and the average distance between the vertices in the (4,5) hyperbolic tessellation as a function of the number of vertices (right).

In this setting of the GNFG, similarly to the Euclidean game the following lemma holds.

Lemma 3. *Any graph emerging from any Nash equilibrium or social optimum in the two dimensional hyperbolic GNFG possesses the underlying tessellation graph T as a subgraph.*

Proof. The proof of this statement is very similar to the proof of Lemma 4.2 and is thus omitted. □

Of course similarly to Euclidean lattices greedy routing readily routes through the tessellation by finding paths successively between arbitrary pairs of players. In the following I show that the length of these paths are of $O(\log N)$. Let us define the layers of a tessellation as follows: the starting polygon is layer 0, then raising κ polygons over every vertex of the starting polygon yields layer 1, layer $m + 1$ is derived accordingly from layer m. Let us call the newly added vertices in each step as the "perimeter" of layer m.

Lemma 4. *The number of vertices in a (ν, κ) tessellation grows exponentially with the number of layers, thus the diameter of the tessellation graph T is of $O(\log N)$.*

Proof. Let κ ν-gons meet at each vertex. There are three cases according to ν

- If $\nu = 3$ then from $\frac{1}{\nu}+\frac{1}{\kappa} < \frac{1}{2}$ we get $\kappa \geq 7$. At each vertex on the perimeter of the m-th layer at most 3 triangles meet, thus each vertex has at least $(\kappa - 4) \geq 3$ neighbors on the perimeter of $(m + 1)$-st layer. Since two neighbors on the perimeter of the m-th layer have at most one common neighbor on the $(m+1)$-st layer, each vertex on the m-th layer generates at least 2 vertices on the $m + 1$-th layer.

- If $\nu = 4$ then $\kappa \geq 5$. Each vertex on the m-th layer has at least $(\kappa - 4) \geq 1$ tetragon whose other vertices are on the $(m + 1)$-st layer, hence each vertex generates at least three other on the next layer.

- If $\nu \geq 5$, then I calculate the number of edges on the layers' perimeter instead of the number of vertices. Each edge on the perimeter has another polygon on the next layer, these polygons have $\nu - 3 \geq 2$ edges on the perimeter of the next layer, hence the number of vertices is at least twice as large.

□

From Lemma 4 we can see that using only the edges of the tessellation (without any extra shortcut edges) player v can be reached from u in $O(\log(N))$ number of greedy steps (see Figure 4.4). The total cost of the GNFG is therefore of $O(N^2 \log(N))$, which is better than the social optimum for Euclidean case.

4.5 Summary

Numerous theoretical and empirical studies confirm that real-world complex networks lend themselves to be effectively searched by greedy algorithms. However, the existing game theoretical models on network formation, which were summoned to justify the emergence of such networks, rely exclusively on shortest path calculations when estimating the cost of communication among players. I have characterized networks which emerge from the network formation games if players calculate their costs based on the length of the greedy routes between each other. The results support the claim that small-world topologies cannot emerge as equilibria from the interaction of selfish players under the Kleinberg-like grid-based model. I also showed that changing the underlying metric space from Euclidean to hyperbolic simply generated small-world topologies as equilibrium solutions with lower total cost than the optimal solution under the Euclidean version. These findings promote hyperbolic Greedy Network Formation Games for future investigations to explain other network properties like heterogeneous degree distribution and strong clustering.

Chapter 5

Application Opportunities

To conclude the dissertation here I discuss how the obtained results can be used to further extend our knowledge about the topology of complex networks. In this chapter I list some application scenarios for each chapter where results are introduced.

Topological Consequences of the BGP routing policy

Internet specific knowledge can greatly help to improve the performance of the network. The more insight we gather on how BGP drives the topology formation of the Internet's AS level network the easier it is (*i*) to design better routing policies, (*ii*) to understand why and how the traffic emerges and (*iii*) to optimize the current network structure. The most specific example is clearly the area of Content Delivery Networks (CDN) [79], where global topological peculiarities are highly exploited e.g. in surrogate and cache placement strategies or request routing mechanisms. Note that CDN is just a narrow segment of the whole spectrum. To give a few more examples, the placement of data centers [42], peer-to-peer networks [22, 66], traffic engineering [8], business based AS peering strategies [26] can also largely benefit from Internet topology related knowledge. The investigation of the AS topology is also a popular topic in the network science community that consolidates researchers from diverse or multidisciplinary research areas [14, 15, 5, 21, 98, 70, 84].

A Game Theory-Based AS Level Model (YEAS)

Topology generators are often used in diverse testing processes of different applications. Testing of novel routing policies, traffic handling or security algorithms requires realistic topologies that have some randomness but bear similar statistical features to real networks at the same time. Certainly the needs are not exactly the same in all cases, so the topology generators can be categorized along the needs they serve. YEAS can be useful in situations when quickly generated large topologies with the characteristics of the Internet's AS level network are needed, including labeled nodes and connections according to business considerations.

Topological Consequences of the Greedy Navigation

Greedy routing is the most accepted policy in describing the communication process of real-world complex networks. We have empirical evidences (e.g. the Milgram experiment) that in many cases this method enables efficient information distribution among network nodes. However, creating a game that explains the emerging process of such networks had been a non-trivial issue in game theory until recently. My results are cited in two papers, by Yang et. al [99] and Gulyás et. al [44], where the authors manage to explain the emergence of such networks.

Bibliography

[1] *Routeviews project, http://www.routeviews.org/.*

[2] *The CAIDA project http://www.caida.org.*

[3] Bernhard Ager et al. Anatomy of a large European IXP. In *Proceedings of ACM SIGCOMM 2012*, pages 163–174. ACM, 2012.

[4] William Aiello, Fan Chung, and Linyuan Lu. A random graph model for massive graphs. In *ACM STOC*, pages 171–180. Acm, 2000.

[5] Réka Albert, Hawoong Jeong, and Albert-László Barabási. Error and attack tolerance of complex networks. *Nature*, 406(6794):378–382, 2000.

[6] Linda JS Allen. Some discrete-time SI, SIR, and SIS epidemic models. *Mathematical biosciences*, 124(1):83–105, 1994.

[7] J.W. Anderson. *Hyperbolic geometry.* 2005.

[8] Daniel Awduche et al. Overview and principles of Internet traffic engineering. Technical report, RFC 3272, May, 2002.

[9] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.

[10] Albert-László Barabási, Réka Albert, and Hawoong Jeong. Scale-free characteristics of random networks: the topology of the world-wide web. *Physica A: Statistical Mechanics and its Applications*, 281(1):69–77, 2000.

[11] Albert-László Barabási et al. Scale-free networks: a decade and beyond. *Science*, 325(5939):412, 2009.

[12] Robert Beverly and Steven Bauer. The spoofer project: Inferring the extent of source address filtering on the internet. In *Proceedings of USENIX SRUTI workshop*, 2005.

[13] Mm Boguñá and D. Krioukov. Navigating ultrasmall worlds in ultrashort time. *Physical Review Letters*, 102(5):058701, Feb 2009.

[14] M. Boguna, D. Krioukov, and K. C. Claffy. Navigability of complex networks. *Nature Physics*, 5(1):74–80, 2009.

[15] M. Boguna, F. Papadopoulos, and D. Krioukov. Sustaining the Internet with hyperbolic mapping. *Nature Communications*, 1(6):1–8, 2010.

[16] Marián Boguná and Dmitri Krioukov. Navigating ultrasmall worlds in ultrashort time. *Physical review letters*, 102(5):058701, 2009.

[17] Marián Boguná and Romualdo Pastor-Satorras. Class of correlated random networks with hidden variables. *Physical Review E*, 68(3):036112, 2003.

[18] B. Bollobas. Random graphs. *Academic Press, Inc.*, Orlando, Florida, 1985.

[19] CAIDA Monitors CAIDA. The archipelago measurement infrastructure, http://www.caida.org/projects/ark/, 2002.

[20] Kenneth L Calvert, Matthew B Doar, and Ellen W Zegura. Modeling Internet topology. *Communications Magazine, IEEE*, 35(6):160–163, 1997.

[21] Claudio Castellano and Romualdo Pastor-Satorras. Competing activation mechanisms in epidemics on networks. *Scientific Reports*, 2, 2012.

[22] Miguel Castro, Peter Druschel, Y Charlie Hu, and Antony Rowstron. Topology-aware routing in structured peer-to-peer overlay networks. In *Future Directions in Distributed Computing*, pages 103–107. Springer, 2003.

[23] Hyunseok Chang, Sugih Jamin, and Walter Willinger. To peer or not to peer: Modeling the evolution of the Internet's AS-level topology. *Ann Arbor*, 1001:48109–2122.

[24] Cisco. BGP best path selection algorithm: How the best path algorithm works. document id: 13753, May 2012.

[25] K Claffy and B Huffaker. Macroscopic Internet visualization and measurement, CAIDA Mapnet tool.

[26] David Clark, William Lehr, and Steven Bauer. Interconnection in the Internet: the policy challenge. In *Research Conference on Communication, Information and Internet Policy*, 2011.

[27] Ian Clarke, Oskar Sandberg, Brandon Wiley, and Theodore W Hong. Freenet: A distributed anonymous information storage and retrieval system. In *Designing Privacy Enhancing Technologies*, pages 46–66. Springer, 2001.

[28] Microsoft Corporation. www.azure.microsoft.com. 2015.

[29] Ryan Craven, Robert Beverly, and Mark Allman. A middlebox-cooperative TCP for a non end-to-end Internet. In *ACM SIGCOMM Computer Communication Review*, volume 44, pages 151–162. ACM, 2014.

[30] Xenofontas Dimitropoulos et al. Graph annotations in modeling complex network topologies. *ACM TOMACS*, 19(4):17, 2009.

[31] Ahmed Elmokashfi, Amund Kvalbein, and Constantine Dovrolis. Simrot: a scalable inter-domain routing toolbox. *ACM SIGMETRICS Performance Evaluation Review*, 39(2):4–13, 2011.

[32] I. Stoica et. al. Chord: A scalable peer-to-peer lookup service for Internet applications. In *Proceedings of SIGCOMM'01*, San Diego, California, August 2001.

[33] E. Even-Dar and M. Kearns. A small world threshold for economic network formation. In *Advances in Neural Information Processing Systems*, page 385, 2007.

[34] Alex Fabrikant, Elias Koutsoupias, and Christos H Papadimitriou. Heuristically optimized trade-offs: A new paradigm for power laws in the Internet. In *Automata, Languages and Programming*, pages 110–122. Springer, 2002.

[35] A. Fabrikant et. al. On a network creation game. In *Proceedings of PODC'03*, pages 347–351, 2003.

[36] P. Fraigniaud, E. Lebhar, and Z. Lotker. A lower bound for network navigability. *SIAM Journal on Discrete Mathematics*, 24:72, 2010.

[37] L. Gao and J. Rexford. Stable internet routing without global coordination. *IEEE/ACM Transactions on Networking*, 9(6):681–692, 2001.

[38] Lixin Gao. On inferring autonomous system relationships in the internet. *IEEE/ACM Transactions on Networking (ToN)*, 9(6):733–745, 2001.

[39] Cyril Gavoille and Stephane Perennes. Memory requirement for routing in distributed networks. *Proceedings of PODC'96*, pages 125–133, 1996.

[40] Phillipa Gill, Michael Schapira, and Sharon Goldberg. A survey of interdomain routing policies. *ACM SIGCOMM CCR*, 44(1):28–34, 2013.

[41] Gnutella. http://rfc-gnutella.sourceforge.net/.

[42] Albert Greenberg et al. The cost of a cloud: research problems in data center networks. *ACM SIGCOMM CCR*, 39(1):68–73, 2008.

[43] Luca Gugelmann, Konstantinos Panagiotou, and Ueli Peter. Random hyperbolic graphs: degree sequence and clustering. In *Automata, Languages, and Programming*, pages 573–585. Springer, 2012.

[44] András Gulyás, József Bíró, Attila Kőrösi, Gábor Rétvári, and Dmitri Krioukov. Navigable networks as Nash equilibria of navigation games. *Nature Communications*, Vol. 6, 2015.

[45] C. Guo et. al. Bcube: a high performance, server-centric network architecture for modular data centers. *ACM SIGCOMM CCR*, 39(4):63–74, 2009.

[46] R.K. Guy and F. Harary. On the Möbius ladders. *Canadian Mathematical Bulletin*, 10:493–496, 1967.

[47] Petter Holme et al. An integrated model of traffic, geography and economy in the Internet. *ACM SIGCOMM CCR*, 38(3):5–16, 2008.

[48] G. Huston. Interconnection, peering and settlements: Part I. *Internet Protocol Journal*, 2(1), JUN 1999.

[49] RFC-1105 IETF. https://tools.ietf.org/html/rfc1105.

[50] RFC-1105 IETF. https://tools.ietf.org/html/rfc4271.

[51] RFC-1142 IETF. https://tools.ietf.org/html/rfc1142.

[52] RFC-1164 IETF. Application of the border gateway protocol in the internet.

[53] RFC-2328 IETF. https://www.ietf.org/rfc/rfc2328.

[54] Matthew O Jackson. A survey of network formation models: stability and efficiency. *Group Formation in Economics: Networks, Clubs, and Coalitions*, pages 11–49, 2005.

[55] Sugih Jamin, Cheng Jin, Anthony R Kurc, Danny Raz, and Yuval Shavitt. Constrained mirror placement on the internet. In *INFOCOM 2001. Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 1, pages 31–40. IEEE, 2001.

[56] D.E. Joyce. Hyperbolic tessellations. *Clark University. http://aleph0. clarku. edu/˜ djoyce/poincare/poincare. html. Updated December*, 1998.

[57] Thomas Kernen. Traceroute.org, 2008. *UR L http://www.traceroute.org.*

[58] J. Kleinberg. The small-world phenomenon: an algorithm perspective. In *Proceedings of STOC'00*, pages 163–170, 2000.

[59] J. Kleinberg. Complex networks and decentralized search algorithms. In *Proceedings of ICM*, volume 3, pages 1019–1044. Citeseer, 2006.

[60] R. Kleinberg. Geographic routing using hyperbolic space. In *Proceedings of INFOCOM 2007*, pages 1902–1909. Ieee, 2007.

[61] Dmitri Krioukov, Fragkiskos Papadopoulos, Marián Boguñá, and Amin Vahdat. Greedy forwarding in scale-free networks embedded in hyperbolic metric spaces. *ACM SIGMETRICS Performance Evaluation Review*, 37(2):15–17, 2009.

[62] D. Krioukov et. al. Hyperbolic geometry of complex networks. *Physical Review E*, 82(3):036106, 2010.

[63] P Krishnan, Danny Raz, and Yuval Shavitt. The cache location problem. *IEEE/ACM Transactions on Networking (TON)*, 8(5):568–582, 2000.

[64] Bo Li, Mordecai J Golin, Giuseppe F Italiano, Xin Deng, and Kazem Sohraby. On the optimal placement of web proxies in the Internet. In *INFOCOM'99. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 3, pages 1282–1290. IEEE, 1999.

[65] Aemen Lodhi, Amogh Dhamdhere, and Constantine Dovrolis. Genesis: An agent-based model of interdomain network formation, traffic flow and economics. In *INFOCOM, 2012 Proceedings IEEE*, pages 1197–1205. IEEE, 2012.

[66] Eng Keong Lua, Jon Crowcroft, Marcelo Pias, Ravi Sharma, Steven Lim, et al. A survey and comparison of peer-to-peer overlay network schemes. *IEEE Communications Surveys and Tutorials*, 7(1-4):72–93, 2005.

[67] B Lyon. The Opte project, 2003 (www.opte.org).

[68] Priya Mahadevan et al. Systematic topology analysis and generation using degree correlations. In *ACM SIGCOMM CCR*, volume 36, pages 135–146. ACM, 2006.

[69] Dahlia Malkhi, Moni Naor, and David Ratajczak. Viceroy: A scalable and dynamic emulation of the butterfly. In *Proceedings of the twenty-first annual symposium on Principles of distributed computing*, pages 183–192. ACM, 2002.

[70] Sergei Maslov, Kim Sneppen, and Alexei Zaliznyak. Detection of topological patterns in complex networks: correlation profile of the internet. *Physica A: Statistical Mechanics and its Applications*, 333:529–540, 2004.

[71] Petar Maymounkov and David Mazieres. Kademlia: A peer-to-peer information system based on the xor metric. In *Peer-to-Peer Systems*, pages 53–65. Springer, 2002.

[72] D McRobb, K Claffy, and T Monk. Skitter: Caida's macroscopic Internet topology discovery and tracking tool, 1999.

[73] Alberto Medina et al. Brite: An approach to universal topology generation. In *MASCOTS, 2001. Proceedings. Ninth International Symposium on*, pages 346–353. IEEE, 2001.

[74] David Meyer. University of oregon route views archive project, 2001.

[75] Stanley Milgram. The small world problem. *Psychology Today*, 2(1):60–67, 1967.

[76] Mark EJ Newman. Power laws, pareto distributions and Zipf's law. *Contemporary physics*, 46(5):323–351, 2005.

[77] N. Nisan. *Algorithmic game theory*. Cambridge University Press, 2007.

[78] Romualdo Pastor-Satorras and Alessandro Vespignani. Epidemic spreading in scale-free networks. *Physical review letters*, 86(14):3200, 2001.

[79] Mukaddim Pathan and Rajkumar Buyya. *Content Delivery Networks*, chapter A Taxonomy of CDNs, pages 33–77. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[80] Lili Qiu, Venkata N Padmanabhan, and Geoffrey M Voelker. On the placement of web server replicas. In *INFOCOM 2001. Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 3, pages 1587–1596. IEEE, 2001.

[81] Pavlin Radoslavov, Ramesh Govindan, and Deborah Estrin. Topology-informed internet replica placement. *Computer Communications*, 25(4):384–392, 2002.

[82] Sylvia Ratnasamy, Paul Francis, Mark Handley, Richard Karp, and Scott Shenker. *A scalable content-addressable network*, volume 31. ACM, 2001.

[83] S. Ratnasamy et. al. A scalable content-addressable network. In *Proceedings of SIGCOMM '01*, pages 161–172, 2001.

[84] V Rosato, L Issacharoff, S Meloni, D Caligiore, and F Tiriticco. Is the topology of the Internet network really fit to sustain its function? *Physica A: Statistical Mechanics and its Applications*, 387(7):1689–1704, 2008.

[85] Eric C Rosen. Exterior gateway protocol (EGP). 1982.

[86] Boris A Rosenfeld. *A history of non-Euclidean geometry: evolution of the concept of a geometric space*, volume 12. Springer Science & Business Media, 2012.

[87] Antony Rowstron and Peter Druschel. Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. In *Middleware 2001*, pages 329–350. Springer, 2001.

[88] Yuval Shavitt and Eran Shir. Dimes: Let the Internet measure itself. *ACM SIGCOMM CCR*, 35(5):71–74, 2005.

[89] Neil Spring, Ratul Mahajan, David Wetherall, and Thomas Anderson. Measuring ISP topologies with rocketfuel. *Networking, IEEE/ACM Transactions on*, 12(1):2–16, 2004.

[90] Ion Stoica, Robert Morris, David Liben-Nowell, David R Karger, M Frans Kaashoek, Frank Dabek, and Hari Balakrishnan. Chord: a scalable peer-to-peer lookup protocol for Internet applications. *IEEE/ACM Transactions on Networking*, 11(1):17–32, 2003.

[91] Akamai Technologies. www.akamai.com. 2016.

[92] J. Travers and S. Milgram. An experimental study of the small world problem. *Sociometry*, 32:425–443, 1969.

[93] Xiao Fan Wang and Guanrong Chen. Complex networks: small-world, scale-free and beyond. *Circuits and Systems Magazine, IEEE*, 3(1):6–20, 2003.

[94] D.J. Watts, P.S. Dodds, and M.E.J. Newman. Identity and search in social networks. *Science*, 296(5571):1302, 2002.

[95] Duncan J Watts and Steven H Strogatz. Collective dynamics of 'small-world'networks. *nature*, 393(6684):440–442, 1998.

[96] Bernard M Waxman. Routing of multipoint connections. *Selected Areas in Communications, IEEE Journal on*, 6(9):1617–1622, 1988.

[97] Jared Winick and Sugih Jamin. Inet-3.0: Internet topology generator. Technical report, Technical Report CSE-TR-456-02, University of Michigan, 2002.

[98] Bo Xiao, Lian-dong Liu, Xiao-chen Guo, and Ke Xu. Modeling the ipv6 internet as-level topology. *Physica A: Statistical Mechanics and its Applications*, 388(4):529–540, 2009.

[99] Zhi Yang and Wei Chen. A game theoretic model for the formation of navigable small-world networks. In *Proceedings of the 24th International Conference on World Wide Web*, pages 1329–1339. ACM, 2015.

[100] Ben Y Zhao, Ling Huang, Jeremy Stribling, Sean C Rhea, Anthony D Joseph, and John D Kubiatowicz. Tapestry: A resilient global-scale overlay for service deployment. *Selected Areas in Communications, IEEE Journal on*, 22(1):41–53, 2004.

Publications

Journal papers (11.33 p)

[J1] David Szabo, Andras Gulyas. Notes on the Topological Consequences of BGP Policy Routing on the Internet AS Topology. *Lecture Notes in Computer Science Volume 8115, 2013, pp 274-281.* **(6 p)**

[J2] **Dávid Szabó**, Attila Kőrösi, József Bíró, András Gulyás A Deductive Way of Reasoning about the Internet AS Level Topology. *Chinese Physics B.* **(6/3=2 p)**

[J3] Gulyás András, Kőrösi Attila, **Szabó Dávid**, Biczók Gergely. On greedy network formation. *SIGMETRICS PERFORMANCE EVALUATION REVIEW*, 40:pp. 49-52, 2012. **(6/3=2 p)**

[J4] **Szabó Dávid**, Gulyás András, Csernai Márton, Heszberger Zalán. Struktúrafüggetlen címzésen alapuló önszerveződő útvonalválasztási architektúra. *HÍRADÁSTECHNIKA LXVI. évfolyam*, 2011/3: pp. 2-10, 2011. **(4/3=1.33 p)**

Conference papers (4.25 p)

[C1] **Szabó Dávid**, Gulyás András. Jelterjedési utak topológiára gyakorolt hatásának vizsgálata játékelméleti módszerekkel. Mesterpróba konferencia, 1. helyezés, 2012. **(1 p)**

[C2] **David Szabo**, Andras Gulyas, Frank H.P. Fitzek, Daniel E. Lucani. Towards the Tactile Internet: Decreasing Communication Latency with Network Coding and Software Defined Networking. European Wireless 2015. **(3/3=1 p)**

[C3] **David Szabo**, Andras Gulyas, Balazs Sonkoly, Felician Nemeth, Frank H.P. Fitzek. Towards the 5G Revolution: A Software Defined Network Architecture Exploiting Network Coding as a Service. SIGCOMM 2015.

[C4] András Gulyás, Attila Kőrösi, Gábor Rétvári, József Bíró, **Dávid Szabó**. Brief Announcement: Network Formation Games Can Give Rise to Realistic Networks. In: 31st Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing. Madeira, Portugália, 2012.

[C5] **Dávid Szabó**, Csernai Márton. Self organizing structurless routing architecture. In: POSTER 2011: 15th International Student Conference on Electrical Engineering. Prága, Csehország, 2011. **(3/2=1.5 p)**

[C6] Gábor Rétvári, **Dávid Szabó**, András Gulyás, Attila Kőrösi, János Tapolcai. An information-theoretic approach to routing scalability. ACM Hotnets 2014. **(3/4=0.75 p)**

Independent Citations

[J1-1] Lenzner, Pascal. Greedy selfish network creation. Internet and Network Economics. Springer Berlin Heidelberg, 2012. 142-155.

[J1-2] Yang, Zhi, and Wei Chen. A Game Theoretic Model for the Formation of Navigable Small-World Networks. Proceedings of the 24th International Conference on World Wide Web. International World Wide Web Conferences Steering Committee, 2015.

[J1-3] Zolanvari, Maede. SDN for 5G. CSE570S: Recent Advances in Networking, available: http://www.cse.wustl.edu/ jain/cse570-15/ftp/sdnfor5g.pdf.

[J1-4] Do-Duy, Tan, and M. A. Vazquez-Castro. Network Coding function virtualization. Signal Processing Advances in Wireless Communications (SPAWC), 2016 IEEE 17th International Workshop on. IEEE, 2016.

[J1-5] Bouzghiba, Soukaina, et al. Towards an Autonomic Approach for Software Defined Networks: An Overview.

[J1-6] Maier, Martin, et al. The tactile Internet: vision, recent progress, and open challenges. IEEE Communications Magazine 54.5 (2016): 138-145.

[J1-7] Heinonen, Johanna, et al. Mobility management enhancements for 5G low latency services. Communications Workshops (ICC), 2016 IEEE International Conference on. IEEE, 2016.

[J1-8] Vukobratovic, Dejan, et al. CONDENSE: A Reconfigurable Knowledge Acquisition Architecture for Future 5G IoT. IEEE Access 4 (2016): 3360-3378.

[J1-9] Confidential, C. O. OpenLab: Extending FIRE testbeds and tools, FP7 EU project Deliverable D2. 1 Experimental plane–Experiment Controllers.

List of Figures

List of Tables