Comparing Traffic Discrimination Policies in an Agent-Based Next-Generation Network Market

Simon Diedrich and Fernando Beltrán

Abstract Presently, the network neutrality paradigm governs the manner in which most data is transported over the Internet. However, experts often question whether keeping such a policy remains reasonable. In the context of new technologies, such as all-IP Next Generation Networks (NGN), traffic discrimination promises to benefit both network providers and users, but also imposes risks. We develop an agent-based NGN market model, in order to investigate the effects of neutral and non-neutral traffic management policies on the performance of Internet market participants. A simulation-based analysis of different policy and competition scenarios suggests that content providers perform best when network neutrality is imposed, while network providers and consumers may benefit from traffic discrimination, under certain circumstances.

1 Introduction

According to what criteria should network capacity be distributed amongst network users? Should everyone have equal rights of access to this limited resource, or should some users be preferred over others? These are critical questions surrounding the global network neutrality debate. Since the early days of the Internet, the network neutrality paradigm - that all network traffic is treated equally - governs most data transportation practices [2]. Currently, this well-established convention is under "fierce debate" [6], and it is often questioned whether it is still appropriate in telecommunications markets relying on all-IP networks, also known as Next

Simon Diedrich
University of Auckland, 12 Grafton Road, Auckland, New Zealand
e-mail: simon.diedrich@gmail.com

Fernando Beltrán
University of Auckland, 12 Grafton Road, Auckland, New Zealand
e-mail: f.beltran@auckland.ac.nz

Generation Networks (NGN) [1]. Network neutrality opponents argue that allowing for certain forms of discrimination may benefit both network users and providers [17, 20]. Regulators all over the world are reconsidering their positions towards network neutrality and traffic discrimination [9]. In the United States (US), for example, the Federal Communication Commission (FCC) recently affirmed its pro network neutrality position [13], while being backed by the US Senate [16]. Similar measures are currently being evaluated in the European Union [9], and the European Parliament recently called upon the European Commission, to ensure "the preservation of the open and neutral characteristics of the Internet"[12].

Identifying the right regulatory framework is difficult. Until today, network management has generally been neutral, and little experience with non-neutral practices is available. Experts have presented many valid arguments, both in favour, and against, network neutrality. However, their justifications are often hypothetical in nature, and lack sufficient empirical foundation.

This study aims to improve our understanding of the manner in which both neutral and non-neutral traffic management affect different groups of Internet market participants. We focus on welfare implications. Based on an agent-based NGN market model, we simulate and compare policy scenarios.

We believe agent-based modelling to be a most promising approach. The lack of cases of discriminatory traffic management renders an analysis based on real data impossible. Relying on traditional economic modelling would require us to either reduce the model scope, or to omit important elements, in order to produce a market model of manageable complexity. Agent-based modelling permits us to overcome these limitations, and promises to provide more reliable and realistic insights into NGN market dynamics, than do alternative modelling approaches.

The remainder of this paper is structured as follows. Section two sets the study within the broader context of the network neutrality debate, and elaborates on the issues that motivate our research. Section three provides a description of the agent-based model. In section four, we describe how the model is used to simulate a range of market scenarios. In section five, simulation results are discussed. Section six concludes with an evaluation of the agent-based model. It further discusses the implications of simulation results both for regulators and network market participants.

2 Network neutrality and traffic discrimination

The layered, end-to-end architecture of the Internet places network intelligence at the edges rather than at the core [5]. Internet advocates point to such technical features as the sources of the wide range of innovative services as the last 15 years can attest. The network neutrality principle proposes that no operator can discriminate against content or traffic that travels on its network; it also prevents any provider from blocking, degrading or interfering with particular websites and devices used to access the Internet.

This paper seeks to contribute to the literature on the network neutrality debate from a different, yet complementary perspective to what has been offered so far. In Section 3 we will present the details of an agent-based computational simulation model which, replicates selected features of a broadband access market and the content markets that flourish on it. Our motivation is raised by our concern with exploiting the potential power of agent-based technology of being able to account for a large number of consumer-provider interactions in a way that aggregate (emergent) data can be easily collected and analysed.

A divisive issue that has got a lot of recent attention is the capacity that Internet providers have developed to discriminate among different types of traffic. Discrimination can be defined as the unequal treatment of applications and content seeking individual benefit, without the consent of all market participants. Internet service providers (ISPs) can exert discrimination either close to the end-user on an access basis, or far from the end-user at peering points. NGN platforms may in essence exert four different types of discriminatory activities: price discrimination, access tiering, blocking, and service quality discrimination [11]. Even though blocking is the most usual form of discrimination, it seems to be giving way to the more subtle, yet potentially more dangerous practice of giving higher priority to traffic received from or sent to specific content providers.

Network neutrality opponents may argue that allowing discrimination through tiered connections may lead to higher efficiency. On the other hand, price discrimination is claimed to also have potential positive effects. Such unilateral actions by network operators might be justified for two reasons: they may show a path to efficiently internalizing the congestion costs raised by high-volume users; and they may also benefit low-volume users by lowering the access prices they must pay [20].

On a more positive note, as providers respond to different consumers' tastes by offering different sets of service attributes, consumers may benefit from having a wider variety of services that more closely match their preferences. For instance, since ISPs' offers might include different Quality-of-Service (QoS) levels, users must choose one from a QoS menu when subscribing to the network. Quality-sensitive users may benefit from purchasing high QoS, while other users would not mind having lower QoS. Furthermore, discrimination against particular sites, services, or contents by NGN platforms may not only be seen as a network management tool (i.e. a mean to alleviate network congestion), but also as a strategic marketing tool that providers might use to bundle their access services with preferred content generators' offers. An exploratory study on the interconnection of all-IP networks [21], presents a simple model of interconnection settlements with fixed fees and session-based charges. The paper argues that the traditional regulatory goal of promoting competition may be less important than other regulatory concerns. For example, as QoS-enabled networks may find it attractive to degrade the quality or capacity of interconnection between QoS-enabled and best-effort networks, regulatory concerns emerge that must address this issue. Degradation of QoS, even reaching outright traffic blocking, can be regarded as a manifestation of service discrimination, but this depends on the actual drivers that motivate the operators to take such actions; one such driver is congestion.

For the purpose of our analysis here, an alternative view of discrimination is adopted. We compare neutral and non-neutral network management. We further distinguish between two non-neutral network policies: application-based and source-based discrimination [14]. Under network neutrality, all traffic is treated equally, and discrimination and prioritisation are prohibited. Under application-based discrimination, traffic is prioritised, or discriminated, based upon the type of content or application. For example, all VoIP traffic may be prioritised or discriminated. Other criteria, such as the identity of the traffic source or destination, are irrelevant. Since all providers of a certain type of content or application are treated equally, and none of them is prioritised, or discriminated against, individually, application-based discrimination is generally seen as less anticompetitive, or harmful to network user interests [13]

Under source-based discrimination, traffic is prioritised or discriminated based on its source. The identity of a network user serves as prioritisation or discrimination criteria. Network providers are able to assign priority levels to individual network users. They may find it appealing to execute price discrimination, and charge differently for prioritised and non-prioritised network services [19]. Many scholars argue that source-based discrimination may lead to a superior allocation of network capacity and network cost [1]. Others argue that it may result in a strong redistribution of market power in favour of network providers, which these would then use to exploit network users [19]. Regulators, such as the FCC, generally perceive it as anticompetitive [13].

3 An agent-based simulation model of traffic management in all-IP networks

The general NGN market structure is modelled as illustrated in figure 1. Content providers and consumers interact with each other over a network platform, operated and controlled by network providers. Consumers purchase content services from content providers, such as Voice over IP (VoIP) or IPTV. These transactions are limited to upper layers. Network providers act as intermediaries, and transport data from content providers to consumers. They attend to two groups of customers, each connected to opposite ”sides” of the network platform.

This market structure reflects an important characteristic of all-IP networks. Compared to most legacy telecommunications services, NGN-based services break with strong vertical integration, and are characterised by a separation of application-oriented and transportation-oriented functionalities [15]. For example, traditional circuit-switch voice telephony, often involves only a single provider. On the other hand, in VoIP telephony, application-oriented service components are provided by a VoIP provider, and underlying transportation-oriented service components are provided by network operators. Traffic management policies govern the manner in which network providers interact with the two market sides. These rules influence both economic and technical aspects of these transactions, such as pricing or network capacity allocation.

Fig. 1 Structure of all-IP NGN markets

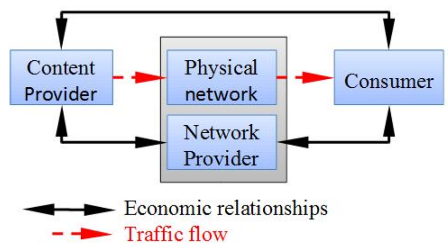

Network capacity is a limited resource, and constrains the amount of network services that can be provided without congesting the network. If demand for network services exceeds capacity, network users will suffer from the negative effects of congestion, e.g. service quality is decreased and consumers gain less utility. Network resources are assigned based on the traffic management policy in place. A distinction is made between high and standard priority, and traffic is either prioritised, or is not.

Under neutral network management, all traffic is assigned standard priority and is transported with equal priority. We distinguish between neutral and non-neutral network management. The analysed forms of non-neutral network management are application-based and source-based discrimination. Under application-based discrimination, all traffic belonging to the same content type is assigned identical priority, e.g. all IPTV traffic may be prioritised. Source-based discrimination allows network providers to prioritise or discriminate based upon the traffic source. Priority levels are assigned to each content provider individually. Under this policy, content providers may choose to subscribe to prioritisation services, and network providers may charge a price premium for granting high priority. Network markets are simulated over a number of time periods. Each period, agents execute a number of consecutive tasks (figure 2). A task is performed by all agents of a certain type, e.g. all content providers undergo the task ”set content service price”. Only when a task is completed by all corresponding agents, the next step starts. Network market participants’ different tasks are described in the following list of consecutive steps. After all tasks have been completed, the period ends and steps are repeated in a new period, until the predefined number of periods is reached.

1. Network providers set the network service prices they charge from content providers and consumers.
2. Content providers select a network provider. Under a source-based discrimination policy, they may also choose between high or standard priority.
3. Content providers set the prices they charge consumers for content services.
4. Consumers evaluate and select their network and content providers.
5. Consumers make their individual consumption decisions, i.e. they select the basket of services that they believe will maximise their utility.

Fig. 2 Consecutive tasks performed each period

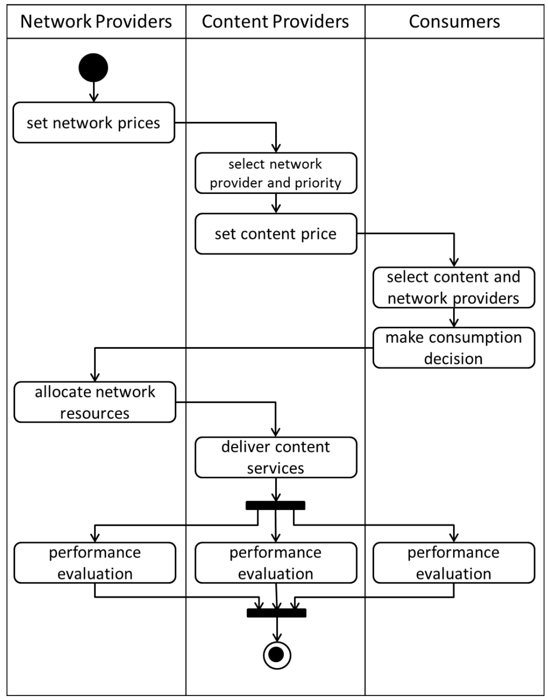

6. Network resources are allocated to network users, i.e. to content providers and consumers, according to the rules mandated by the traffic management policy.
7. Based upon this allocation, content services are sent from content providers to consumers, and payments for network and content services are made. If the demand for network resources exceeds capacity, the network is congested. This has a negative effect on service quality.
8. Agents of all types evaluate their performance, and adopt their expectations and beliefs accordingly. This may affect their future behaviour.

3.1 Consumer behaviour

The consumer side of the market is modelled as a set of independent agents. They act according to their individual interests, and are bound by limited rationality [10]. Consumer agents seek to maximise their utility [7]. They subscribe to the providers they believe to offer the best options. Their decisions are based on their individual preferences, the service quality they expect to receive, and the price charged by a provider.

As illustrated in equation 1, consumer agent i seeks to maximise its aggregated individual utility u_i by consuming the number of service units $x_{i,ct}$ of each content type ct that it believes to be most suitable for its individual needs. The utility gained is determined by the positive effect of service consumption, less the negative effect of service cost. The first half of the equation, $m_{i,ct} * x_{i,ct} - (n_{i,ct} * x_{i,ct}^2)/2$, describes a function of diminishing marginal utility, where $m_{i,ct}$ and $n_{i,ct}$ describe consumer agent i's valuation of content service ct. Similar approaches to model consumers' valuation of services have been employed by [3] and [4]. The values for $m_{i,ct}$ and $n_{i,ct}$ are different for each consumer agent, and preferences are unique. The utility gained from consumption is further mediated by the perceived quality $q_{i,j}$ of the services received from content provider j. Perceived service quality depends on the degree of network congestion and a consumer agent's sensitivity to it. If network congestion and congestion sensitivity are high, utility is reduced substantially, and a consumer only enjoys a fraction of the initial utility. Furthermore, congestion sensitivity varies across content types, e.g. consumers may be less sensitive to congestion for file transferring services, than they are for IPTV or VoIP services.

$$\max_{x_{i,ct}} u_i = \sum_{c=1}^{C} (q_{i,j} * (m_{i,ct} * x_{i,ct} - \frac{n_{i,ct} * x_{i,ct}^2}{2}) - s_{i,ct} * ((p_j + p_k^{con}) * x_{i,ct} + sc)) \quad (1)$$

subject to

$$budget_i \geq \sum_{c=1}^{C} ((p_j + pk) * x_{i,ct} + sc) \quad (2)$$

The utility gained from service consumption is reduced by the negative effect of service costs. The amount of costs depends on the prices charged by content providers p_j and network providers p_k^{con}, along with the amount of services consumed. If a consumer agent decides to switch provider, it also incurs switching costs of (sc). The price sensitivity $s_{i,ct}$ determines how strongly costs affect utility. Consumer agents with high price sensitivity are likely to react more strongly to price changes, than consumer agents with low price sensitivity. The number of service units that consumers can purchase is constrained by their budget (equation 2).

Each period, consumer agents evaluate whether they want to keep a current provider, or if they want to acquire services from a competitor. Principles of bounded rationality are applied, and consumer agents search for an option that is good enough to meet their requirements, rather than making the best possible choice [22]. They do not evaluate all possible options, but stop their search when they find an option that meets their requirements. As an alternative to the concept of a utility maximising, strictly rational homo economicus, bounded rationality offers a more realistic representation of human decision making [10].

At the time consumer agents have to choose their providers or select the basket of services they want to consume, they do not know which service quality they are going to receive, or what level of congestion they are going to experience. They have to predict these values based on their past experiences. Each agent has a unique past, and its predictions are individual. At the end of each period, predicted values are compared to the actually experienced values, and are adjusted if needed.

3.2 Provider strategies

A content provider's goal is to maximise profit [7]. Its profit is the difference between revenue and costs. Revenue is equal to the number of content service units sold x multiplied by the price per content unit p, i.e. the aggregated payments received from consumers in exchange for content services. Content providers incur usage independent fixed costs c^{fix} and variable costs. Variable cost consist of the internal cost of producing one unit of content, c^{var}, and the price charged by the network provider for transporting a unit of content through the network, p_k^{cp}.

$$\max_{p} \Pi = p * x - c^{fix} - (c^{var} + p_k^{cp}) * x \tag{3}$$

As with consumers, content providers subscribe to the network provider they deem most suitable. Each content provider offers a certain type of content service. Content types differ in terms of cost structure, amount of network capacity required for transporting one content unit, and sensitivity to network congestion. For example, the provision of a content service such as IPTV may require more network capacity than a file transfer service, and the quality of the former may be more sensitive to congestion than the quality of the latter. The characteristics of a content type influence content providers' network provider selection and their pricing behaviour.

As with content providers, network providers strive to maximise profit [7]. They allocate network resources according to the traffic discrimination policy in place. Their profit structure is similar to that of content providers (equation 4). However, they generate revenue from two market sides.

$$\max_{p^{cp}, p^{con}} \Pi = p^{cp} * x^{cp} + prem^{cp} * x^{cp,prem} + p^{con} * x^{con} - c^{fix} - c^{var} * x \tag{4}$$

They receive payments from the content provider market side, equal to the price charged from content providers, p^{cp}, multiplied by the amount of service sold to this market side, x^{cp}. Furthermore, they receive payments from the consumer side, equal to the price charged from consumers, p^{con}, multiplied by the amount of service sold to the consumer side, x^{con}. In scenarios where the traffic discrimination policy allows for a distinction between high and low priority traffic, network providers may charge content providers a premium $prem^{cp}$ for high priority network units $x^{cp,prem}$.

Providers compete in price and service quality. Content providers make one pricing decision, while network providers make two independent, yet interrelated pricing decisions, i.e. they set one price for each market side. Providers only have an indirect influence on service quality. The quality of content services depends on the quality of the underlying network services. A content provider may influence the quality of its service by subscribing to a network provider with little or no network congestion, or by paying a premium for the prioritisation of its traffic. Providers directly compete with providers of identical services. They indirectly compete with providers of other services, as consumers have to distribute their limited budget for purchasing different types of services.

Drawing on game theory, we model providers' competitive price setting process as an n-player, extensive-form game. It is reasonable to assume that providers make their decisions sequentially, spread out over a period of time, such as a month or a quarter. Providers make their pricing decision in a random order, and the pricing decision of one provider affects the pay-off of others. They observe the actions taken by others before them, and choose their price accordingly. Furthermore, they evaluate how their own decision affects the actions that others will subsequently rationally take [18]. They chose the pricing strategy which they expect will generate the highest possible profit. The game is solved by using backward induction.

To account for the mediating effect of competition in the relationships between traffic discrimination policies and NGN markets, we distinguish between monopolistic and oligopolistic market structures. A monopolistic market segment is represented by a single provider agent. A monopolistic provider does not have to account for the actions of direct competitors. Its decision problem is reduced to a one player game. However, it may still face indirect competition from other market segments. For instance, the performance of a monopolistic IPTV provider may be affected by the behaviour of VoIP or network providers. An oligopolistic market segment is represented as a number of providers that offer similar services. They have to account for the actions of others. They play an n-player game where n stands for the number of directly competing providers.

Over time, providers learn about the relationships between their pricing decision and their performance. For instance, if a pricing decision results in a profit increase, providers are encouraged to repeat this pricing behaviour. However, a profit increase is not only the result of a provider's individual behaviour, but is influenced by a broad range of factors, such as the prices of other services or the degree of network congestion. These additional factors make it difficult for providers to develop reliable beliefs about the causality between their actions and their performance.

4 Simulation and experimental design

We simulate a range of market scenarios in order to investigate the effects that neutral and non-neutral traffic management have on NGN markets. Scenarios differ in terms of the mandated policy types and the forms of competition. A policy is either network neutrality, application-based discrimination, or source-based discrimination. Furthermore, we investigate scenarios of network market monopolies and oligopolies, and content market monopolies and oligopolies. Agent preferences and behaviour are unique, and repeated simulations of a single scenario may result in different overall model behaviour. In order to account for such variations, each scenario simulation is replicated multiple times, in order to generate a sample that allows for the making of reliable statistic inferences. Scenarios are simulated over a certain period of time, defined as a predefined number of time steps.

Data collected in scenario simulations is analysed based upon an analysis of variance (ANOVA) and corresponding post hoc tests. Agents' performance is measured using aggregated performance measures, such as overall provider profit, the distribution of profits, or overall consumer utility. To account for a warm-up phase that is needed for the agents' initial learning, the first simulation periods are discarded, and only later periods are considered for data analysis.

5 Results

Similar to earlier studies [6, 8], we find that network providers benefit from the ability to execute traffic discrimination in many cases. Their profit is highest in scenarios of application-based discrimination. This supports the suggestion by [7], that network providers may have fewer incentives to engage in controversial forms of discrimination, such as source-based discrimination, than is generally believed. Network providers' share in overall market profit is highest when regulators permit traffic discrimination. Figure 3 illustrates the distribution of profit under different policies and forms of competition. It offers an example of our analysis of simulation data. In network monopoly scenarios, where network competition is already low, non-neutral traffic management has little effect on the distribution of profits. However, in network oligopoly scenarios, where network competition is higher, traffic discrimination enables network providers to significantly increase their relative profit share. This backs the concern, that non-neutral traffic management enables network providers to extract revenues from content markets [19]. It is interesting that this redistribution was strongest under application-based discrimination - a form of non-neutral traffic management that is generally perceived as less controversial and anti-competitive [13].

Simulation results suggest that content providers are unlikely to benefit from traffic discrimination. Non-neutral traffic management results in a decrease of their profits in network monopolies, and has no significant effect in network oligopolies. The effect of discrimination on consumer utility strongly depends on the form of network competition. Source-based discrimination benefits consumers in network oligopolies, but harms them in network monopolies. In markets characterised by low competition, application-based discrimination is more appealing from the consumer's perspective. Under certain circumstances, diverging from network neutrality allows consumers to gain more utility from congestion sensitive content services. This suggests that non-neutral traffic management may have the potential to make such services more appealing to consumers, and thus to favour their profitability.

Fig. 3 Network providers' mean relative share in overall market profit

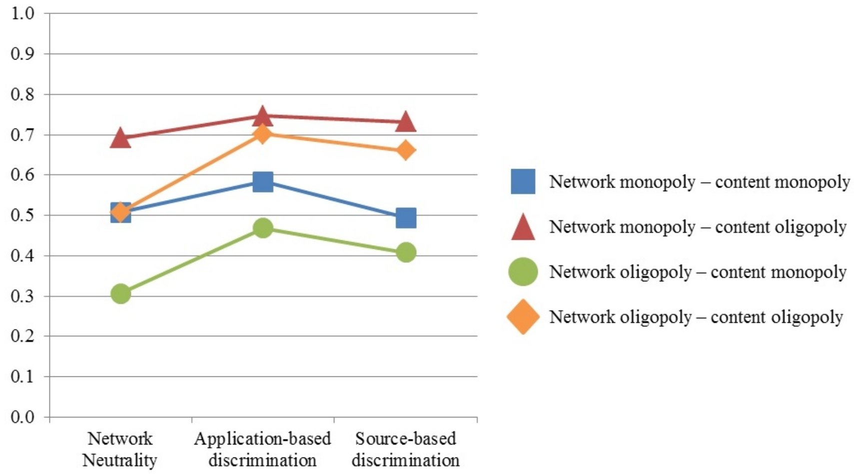

6 Conclusion

The agent-based model developed here allowed us to investigate the manner in which both neutral and non-neutral traffic management policies affect the behaviour and performance of all-IP NGN market participants. In a field where a lack of market experience and data on a sufficiently large scale renders an analysis of real cases impossible, the agent-based model offers a promising alternative to more classical modelling approaches. Most existing network neutrality related models adopt a more traditional, aggregating market perspective, and are tailored towards the analysis of a particular aspect of the network neutrality debate [8, 17]. Our work complements earlier research by offering an alternative perspective. We developed a model that accounts for many of the facets of the network neutrality debate, while maintaining a reasonable complexity. Furthermore, we analyse dynamic market behaviour over time, rather than steady state market equilibria. Even though we aimed to include many facets of the problem, the nature of model development forced us to omit some real phenomena. We understand our model as a flexible basis that future work may build upon, for instance by incorporating further aspects of the network neutrality debate, such as investment in the network.

Based on a simulation of NGN markets scenarios with different traffic management policies, we conclude that none of the three forms of traffic management studied here is strictly welfare superior. None of the policies benefited all NGN stakeholders at the same time. Which policy is most suitable depends on regulatory objectives, and the form of network market competition. For regulators, it is worth evaluating, if allowing network providers to experiment with certain forms of discrimination may in fact be a promising alternative to holding on to the traditional network neutrality paradigm.

Acknowledgements We thank the two anonymous reviewers whose valuable comments allowed us to make significant improvements, and helped us to accentuate the aspects of our work that are most interesting for the reader. Furthermore, we thankfully acknowledge the financial support through the University of Auckland's Thesis and Research Essay Publication Award.

References

1. Becker, G., Carlton, D., Sider, H.: Net neutrality and consumer welfare. J. of Competition Law and Econ. **6**(3), 497–519 (2010)
2. Beltran, F., Gomez, L.: Discrimination in ngn service markets: Opportunity or barrier to digital inclusion? IEEE Communications Magazine **48**(2), 154–159 (2010)
3. Beltran, F., Sharkey, W.: A dynamic model of next-generation networks with two-sided platform competition (2010). Earlier versions of this paper were presented at seminars at the FCC and at the Telecommunications Policy Research Conference in 2008.
4. de Bijl, P., Peitz, M.: Regulation and Entry into telecommunication Markets. Cambridge University Press (2003)
5. Cerf, V.: U.s. senate committee on commerce, science, and transportation hearing on network neutrality (2006)
6. Cheng, H., Bandyopadhyay, S., Guo, H.: The debate on net neutrality: A policy perspective. Inf. Syst. Res. **22**(1), 60–82 (2011)
7. DeMaagd, K., Bauer, J.: Modeling the the dynamic interactions of agents in the provision of network infrastructure. Inf. Syst. Frontiers (2010)
8. Economides, N., Tag, J.: Net neutrality on the internet: A two-sided market analysis. IFN Working Paper (2008)
9. Marcus, J., Nooren, P., Cave, J., Carter, K.: Network neutrality: Challenges and responses in the EU and in the US. European Parliament (Directorate General for Internal Policies) (2011)
10. Marks, R.: Handbook of Computational Econ., vol. 2, chap. Market Design Using Agent-Based Models, pp. 1339–1380. Elsevier (2006)
11. Marsden, C.: Net neutrality: Towards a co-regulatory solution. Bloomsbury Academic (2006)
12. European Commission: The open internet and net neutrality in europe (2011). COM(2011)222
13. Federal Communications Commission: Report and order - in the matter of preserving the open internet - broadband industry practices (2010). FCC 10-201
14. Federal Trade Commision: Broadband connectivity competition policy (2007)
15. Internation Telecommunication Union: Trends in telecommunication reform 2007 - the road to next-generation networks (2007)
16. United States Senate: On the motion to proceed (motion to proceed to s. j. res. 6) (2011)
17. Musacchio, J., Schwartz, G., Walrand, J.: A two-sided market analysis of provider investment incentives with an application to the net-neutrality issue. Rev. of Netw. Econ. **8**(1), 22–39 (2009)
18. Osborne, M.: An Introduction to Game Theory. Oxford University Press (2004)
19. Peha, J.: The benefits and risks of mandating network neutrality, and the quest for a balanced policy. International J. of Communication **1**, 644–668 (2007)
20. Yoo, C.: Network neutrality and the economics of congestion. The Georgetown Law J. **94**, 1847–1909 (2006)
21. Yoon, K.: Interconnection economics of all-ip networks. Rev. of Netw. Econ. **5**(3), 351–365 (2006)
22. Zang, Y., Leezer, J.: Simulating human-like decisions in a memory-based agent model. Computational and Mathematical Organization Theory **16**(4), 373–399 (2004)