Paid Peering and Content Delivery

Research-in-Progress

Sascha Schweitzer
University of Bayreuth
Bayreuth, Germany
sascha.schweitzer@uni-bayreuth.de

Lukas Wiewiorra
University of Frankfurt
Frankfurt, Germany
wiewiorra@wiwi.uni-frankfurt.de

Abstract

Recent conflicts between big content and service providers (CSPs) like Netflix, transit providers (TPs), and Internet service providers (ISPs) have generated considerate media attention and ignited a debate on interconnection agreements, market power of last-mile ISPs and net neutrality. We propose an experimental design to analyze a stylized interconnection market that captures key aspects of actual interconnection markets with a focus on the entry of big CSPs. Participants are invited to assume the roles of ISPs, TPs and CSPs in a computer-aided laboratory experiment. The experiment serves to evaluate potential regulatory tools like transparency and interconnection obligation with respect to the efficiency of the overall interconnection market. Furthermore, we present results of a pre-test of the experimental design and the software implementation. Our preliminary results indicate that operators underinvest into network infrastructure and do not realize the full potential of mutual peering agreements when a CSP participates in the market.

Keywords: Experimental economics, firm performance, future of IS, information economics, Internet, Internet economics, IS economics, IS/IT architecture, IT infrastructure

Introduction

In 1950 the National Science Foundation Network (NSFNET) shut down the existing Internet backbone services and the transition towards a commercially driven Internet backbone began. To guarantee a universal access service, Internet service providers (ISPs) all over the world rely on interconnecting with each other. Those interconnection agreements usually have the form of transit or peering agreements. Transit agreements are typically made with larger operators and include the connectivity to the whole Internet for a certain price per megabit per second. Peering agreements on the other hand cover only the connectivity between the two interconnecting network operators and require the installation of a direct physical connection between them. Peering is a way to reduce transit costs for the amount of traffic that is exchanged between two peering partners. Peering relationships are traditionally so-called “bill-and-keep” agreements, which have the form of settlement free interconnection (Economides, 2006). However, in recent times paid peering emerged as an alternative form of interconnection agreement. In contrast to usual (settlement free) peering agreements, one peering partner is paying the other peering partner for the exchange of traffic between the two networks. Paid peering can emerge if the traffic ratio between two network operators is not balanced, but asymmetric in nature (Besen et al., 2001). Although there is a lot of public information available, the current interconnection agreements and the resulting market performance are by no means transparent. Interconnection between operators is very often established by handshake deals and details of interconnection agreements are in most cases covered by non-disclosure agreements. Therefore it is very difficult for regulators to judge if the market for network interconnection is free of market failures.

With the rise of the net neutrality debate, a heated discussion about data transportation and non-discrimination in the Internet began. Recent events indicate, that the discussion is not only relevant in local access networks, but in the backbone of the Internet as well. Since 2013 Netflix customers in the US complained more often that streaming quality during peak times was distorted with certain ISPs (Brodkin, 2014). Last-mile ISPs exert pressure on content and service providers (CSPs) like Netflix through implicit degradation of their transit business partners. Netflix for instance relied heavily on the transit providers (TPs) Cogent and Level3 to deliver content to its subscribers. The last-mile ISPs Comcast, Verizon and AT&T did not invest sufficiently into the infrastructure facilitating the interconnection with those transit providers. As a result the user experience for Netflix subscribers that are also Comcast, Verizon and AT&T customers was severely degraded. The last-mile ISPs were willing to accept negative side effects on their own installed base to increase the pressure on Netflix and its service partners. Verizon representatives argued that they “are open to negotiation for a commercially reasonable solution that works for both parties.” (Brodkin, 2014). According to a recent traffic analysis (Sandvine, 2013) Netflix was accounting for 31.6% of the North American fixed line downstream traffic at that time.

In July 2013 the European providers Orange, Deutsche Telekom and Telefónica were under suspicion by the European Commission to abuse market power in negotiations with CSPs and their offices were raided to find further evidence (Reuters, 2013). Despite the development in Europe, Netflix and Comcast in the US agreed on a direct paid peering relationship on February 23 2014 and now Netflix pays for high-quality access to the Comcast Internet user installed base. On April 29 2014 Netflix confirmed that it reached an agreement with Verizon as well and soon after also with AT&T (similar to the deal with Comcast). Such procedures are and were never at odds with any Net Neutrality rules in the US. However, Netflix is arguing that those additional payments to last-mile ISPs are nothing else than tolls (c.f. Krämer et al. 2013) and net neutrality regulation should also cover the interconnection agreements between operators in the backbone of the Internet.

Our paper proposes an experimental economics approach to analyze the interconnection market with a special focus on the entry of big CSPs like Netflix. We propose a design for a laboratory experiment in which participants are invited to play the role of ISPs, TPs and CSPs in a computer simulation in an (virtual) Internet interconnection market. The remainder of this paper is structured as follows: First we outline our research questions and the related literature. In the second section we derive our hypotheses and continue with a detailed description of the experimental setup, the timing of events, the relevant parameters and the implementation of the experiment. We conclude with the results of our first pre-test and an outlook on the calibration of a final experimental setup.

Research Questions

With the rise of big CSPs, last-mile ISPs are confronted with new players generating high revenue streams from their installed base. However, those new players have no direct access to their customers and rely on the infrastructure of last-mile ISPs to deliver content to their subscribers. Furthermore, last-mile ISPs (telecommunications providers) are under pressure, because their historic revenue streams (voice, messaging, TV) are eroding due to over the top (OTT) players like Netflix. Last-mile ISPs argue that eroding revenues and imbalanced traffic ratios on the CSP side of the market are forcing them to abandon the established form of interconnection in favor of paid direct relationships. Our research therefore addresses the following questions:

1. Is there a systematic incentive to underinvest into network infrastructure to force CSPs to accept asymmetric interconnection deals?
2. If there is market failure, can regulatory intervention help to improve the market outcome?
3. Which regulatory tools (e.g., transparency- or interconnection obligations), if any, are useful to improve the situation?

Related literature

Cremer et al. (2000) study the incentive of large backbone providers to lower the interconnection quality with smaller competitors under Cournot competition. They find that larger providers have "suboptimal incentives to maintain connectivity" (p.435) and that interconnection degradation is more likely if the difference in market share is high. Since equally sized competitors do not gain a competitive advantage over their rivals through degradation, they do not find those incentives in a model with symmetric providers. Foros et al. (2005), in a closely related paper, look into the incentives of asymmetric ISPs to degrade interconnection quality as well. In contrast to Crémer et al. (2000) they find that a higher interconnection quality can increase the profits of the larger firm. That result is due to the consideration of the positive effect of the quality adjusted network-size. The network provider makes a trade-off between higher profits from locked-in customers against the loss in potential profits of non-acquired new customers. Mendelson and Shneorson (2003) extend the two-sided market approach by Laffont et al (2003) to peering when consumers have costs of delayed data transmission. As a result, operators have to trade off capacity costs and delay costs of consumers. The authors conclude that due to the advantage of larger networks, regulation might be necessary. However, in case of multiple symmetric networks, prices are set similarly to the ones a welfare-maximizing regulator would choose as well. Weiss and Shin (2003) focus on the asymmetry of traffic flows (e.g., between ISPs and CSPs). They find that the industry would adopt the maximum {inbound traffic volume, outbound traffic volume} rule; peering with smaller network operators would be more likely and the market would be more competitive. Jahn and Prüfer (2008) explicitly analyze paid peering as an additional form of interconnection between network operators under price competition. They find that for medium levels of asymmetry between networks paid peering dominates settlement-free peering and transit. However, for large levels of asymmetry providers interconnect via transit. The authors note, that market equilibria in their model not always yield desirable welfare outcomes. This is due to the fact, that the rise of paid peering shifts some consumer surplus to the network operators. Furthermore, if paid peering is facilitated via a variable fee and not only by a lump-sum payment, paid peering is profitable for even larger levels of asymmetry.

Our representation of the following experimental setup relies on graphs. Several papers deal with the problem of finding efficient equilibria in network formation under varying protocols of network formation and preference indication. Most closely related to our problem is a paper by Bloch and Jackson (2007), who look into the case of network formation with transfer payments. They find that the existence "…of positive externalities in payoffs may prevent the formation of efficient networks, because players involved in a link do not internalize the external effects the link has on other players" (p.104). However, the combinatorial problem that arises through the multitude of possible bilateral links between the players can only be solved by imposing additional restrictions on network formation. The existing theoretical literature supports that market power and the popularity of paid peering agreements should raise concerns over the efficiency of those markets.

Experiment

We evaluate the regulatory tools *transparency obligation* and *interconnection obligation* in a stylized interconnection market that captures key aspects of the actual business relations (c.f. Sluijs et al. 2011, Henze et al. 2012). The main focus of our evaluation is on the efficiency of the overall interconnection market. For quantifying this value in the experiment, we rely on the numerical performance measures *network coverage, interconnection capacity* and monetary conditions of *peering agreements*. We state the following hypotheses for the treatment parameters transparency obligation and interconnection obligation:

H 1: ISPs are more likely to postpone investments into interconnection capacity if the negotiation partner is a CSP or a TP serving a CSP.

H 2: ISPs are more likely to demand paid peering if the negotiating partner is a CSP or TP serving a CSP.

H 3: Compared to the benchmark case, a transparency obligation / interconnection obligation:

a) increases the total network coverage

b) increases the number of settlement-free peering agreements

c) decreases the fees for paid peering, if paid peering is established.

Hypothesis 1 states that the likelihood of an ISP degrading the interconnection quality through postponing investments in the interconnection link is higher if it is negotiating with a CSP directly or a TP selling transit to a CSP. Hypothesis 2 states that negotiations with partners without direct access to content consuming customers more often end up with paid peering agreements in favor of the ISP. We want to find out, if the ISPs actually give up potential revenues by degradation and if this strategy is successful in forcing unfavorable deals onto TPs and/or CSPs. Hypotheses 3 a)-c) state that the effect of regulation on network coverage and peering agreements will be positive. This follows from a tendency to equal-split negotiation results known from the bargaining literature, which may be reinforced in a transparent market, because the information available on past negotiations provides a focal point that facilitates future negotiations.

Design of the Peering Game

Our stylized representation of a peering market is inspired by a simulation game introduced by industry advisor William Norton (Norton, 2009). The original game was tested with various industry experts and is considered as a fairly accurate representation of actual peering markets. However, our version of the game abstracts away from some of the particularities of the original design. To reduce noise in our data, we neither include Norton's original chessboard type of network map, nor the geographically located interconnection points. Instead, we replace these elements by a single graph. Further, for the sake of control and replicability, we exogenously define the number of periods and employ formalized Güth et al. (1982) ultimatum game type of negotiations, instead of open oral peering negotiations.

In our CSP setup of the experiment two subjects assume the roles of ISPs, two subjects assume the roles of TPs and one subject assumes the role of a CSP[1]. Each subject is graphically represented by a node in the network graph as depicted in Figure 1b. ISP and TP market coverage is represented by the size of their nodes, i.e. the size of a node is proportional to the ISP's and TP's network coverage. ISPs provide Internet access to end customers, while TPs provide Internet access to enterprise customers and CSPs. In contrast to ISPs and TPs, CSPs do not have unique customers, because all CSP customers are customers of an ISP as well. The CSP recruits its customers from the ISPs' networks and unilaterally delivers services and content. The TPs' customers are not purchasing any CSP's services, because content is not targeted at enterprise customers.

---

[1] That setup can be compared to the situation between Verizon/Comcast/AT&T (ISPs), Cogent/Level3 (TPs) and Netflix (CSP).

In the beginning of the experiment, each ISP and each TP have market coverage of 20 customer fields. Each customer field generates revenue of 2,000 monetary units per period, which the owner of the customer field earns at the beginning of the period. Subjects can increase their coverage by acquiring additional customer fields at a linearly increasing cost for each marginal unit. This represents the firms' infrastructure costs of extending their own coverage. The period-wise cost function for purchasing customer fields is chosen such that a purchase of 10 customer fields per period is optimal in the first period. This number decreases during the course of the game, as the number of periods with future revenues decreases. The total market size, however, is limited to 150 customer fields. After ten periods, given optimal investment decisions, the market of ISP and TP customers would be saturated and no additional customer fields could be acquired. In the original game by Norton, ISPs are not in competition for customers, because all customer fields yield equal profits independent of the number of ISPs covering the field. We introduce a more realistic setting, where ISPs and TPs are in competition for new customers. However, we do not allow firms to acquire existing customers from their competitors. Further, in contrast to end-users and "small" enterprise customers, big CSPs are flexible and can switch interconnection partners easily.

ISPs and TPs, on the one hand, and CSPs, on the other hand, differ fundamentally from each other. Keep in mind that all customers of the CSP are customers of an ISP as well. To indicate that relation the share of ISP customers who are also CSP customers in figure 1b and 1c is shaded red. However, in turn, not all customers of an ISP are customers of the CSP. In contrast to ISPs and TPs the CSP does not have to invest into network coverage. The CSP covers an increasing percentage of the ISP's customer fields, growing endogenously from 20 percent in the first period to 40 percent in the last. Therefore the CSP in our game does not have to invest into new customers because we abstract from advertisement costs in the model. In reality an ISP or TP has to make investments into network coverage and advertisement, whereas the CSP only has advertisement costs. We normalize the advertisement costs for all players to zero. As a result the ISP and TP players have to "buy" customer fields, whereas the CSP receives customer fields "for free".

For the ISPs and TPs customer fields, we assume a symmetric traffic pattern, which means inbound and outbound traffic are equal to each other. In practice inbound/outbound ratios up to 2:1 are considered as balanced in peering relations and the ratio in our model is assumed to be 1:1. When operators grow their networks by acquiring customer fields, network traffic between these operators increases, while the ratio of inbound and outbound traffic remains unchanged. To represent the relationship between network size and traffic in our model, the traffic between ISPs and TPs, is calculated as the sum of the number of customer fields of these operators. Therefore, ISPs and TPs equally contribute to the total units of traffic exchanged. CSPs, however, unilaterally send traffic into the ISPs' networks, either directly or via a TP link. This results in an exceedingly high inbound/outbound ratio and excessive traffic increases caused by the increasing share of (video) content consuming end-users who predominantly request data. To represent the asymmetric relationship, traffic of the CSP increases the overall traffic on the peering link between the TP and the ISP by the factor 5 for each CSP customer field. The same holds for direct traffic between CSPs and ISPs.

The ISPs and TPs in our game represent operators who have already established peering connections among each other in the past. Therefore the ISP and TP networks are completely linked by preexisting peering connections. These connections induce costs of 20,000 monetary units per period, which are split equally among the peering operators by default. However, the preexisting agreements refer only to the status quo extent of network traffic plus a buffer of 50 percent above the current load. Once the network traffic increases, the operators are not obligated to provide any additional capacity. However the operators can choose to provide additional capacity at no cost and may make this capacity extension subject to a renegotiation of the peering agreement.[2]

In the beginning of the game, the CSP has transit access to the Internet through transit, but does not possess any direct connections with the ISPs. Therefore, the CSP has to pay for the transit access for each customer field of the ISPs that it would like to access. The transit prices per period are set by the TPs in

[2] In reality the provision of capacity requires investments into the extension of infrastructure at the interconnection points, such as routers and switches. As this cost is relatively low, we normalize it to zero in the experiment.

the first period of the game and can be adopted during the course of the game. In the experiment as in reality the CSP can purchase transit from multiple transit providers, which is also referred to as multi homing. The CSP player in the experiment decides on the split that defines how much of the generated traffic is purchased from TP 1 and TP 2 at their respective transit prices. The decision on the split of transit traffic can be revised by the CSP in any round, based on the current prices and network conditions.

In the course of the game, ISPs and CSPs can agree on peering connections, which are represented by edges between the interconnected nodes as depicted by the solid lines in Figure 1 a)-c). In contrast to the transit connections between the CSP and the TPs, the establishment of new peering connections requires preceding investments into the underlying network infrastructure. For setting up a peering connection, each party has to make a lump sum investment of 10,000 monetary units. That representation allows us to incorporate the infrastructure investments of network providers in interconnection, without relying on explicit geographical location of customer fields and peering points like in the original peering simulation game by Norton. Once established, peering connections induce operating costs of 20,000 monetary units per period that must be paid by the peering parties. The split of these costs is determined in a negotiation that precedes the initial establishment of a peering connection.

Subjects make their decisions sequentially one after another. At their turn, ISPs and TPs decide on their network investment in the current period and all subjects decide on the provision of additional capacity for the existing peering connections, and on the establishment and (re)negotiation of peering connections. Renegotiations allow the ISPs to account for the increasing share of CSP traffic during the course of the game. In addition, a TP decides on the transit price it offers to the CSP in the current period. ISPs and CSPs can send a peering offer to each other. When sending a new offer or requesting renegotiation, the two subjects involved enter into a Güth et al. (1982) ultimatum peering negotiation about the split of the future operating costs and a potential premium. The costs can be split in steps of 1 monetary unit in the interval from {0:20,000; 1:19,999;...; 20,000:0}. Any asymmetric split of costs results in paid peering relationships. Equal splits result in settlement-free peering. To account for above cost paid peering participants can also suggest a split of costs and add a premium fee on top of the costs {0; 1;..; 100,000} (c.f. Figure 1c).

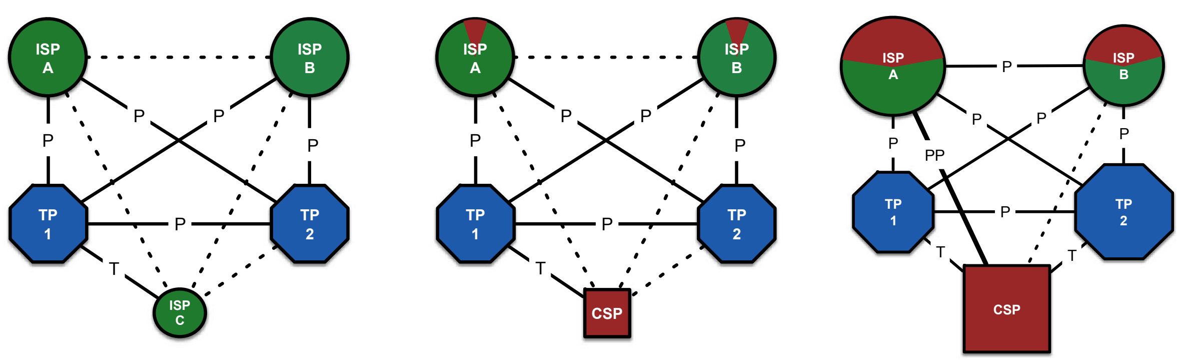

T= transit; P= peering; PP= paid peering above costs
dashed line=possible connection; solid line= active connection

Figure 1a. Benchmark setup

Figure 1b. CSP setup

Figure 1c. Evolved network (CSP setup)

In order to (passively) degrade the quality of interconnection to other players and impact their revenue streams, operators with existing peering connections can decide not to increase the connection capacity to other operators. Missing capacity results in quality degradation. Quality degradation itself requires no action except of not increasing capacity. In that sense, degradation is cost-free. But degradation harms the revenue from the installed base according to the market share of the degraded player, and thus imposes opportunity costs to the degrading as well as to the degraded player.[3] We assume that quality degradation

[3] Imagine there are only two ISPs and customers are connecting to other customers with equal probability. ISP A has five customer fields and ISP B has ten customer fields. If A degrades the transit

reduces the revenues linearly according to the number of customers. However, degradation first affects the CSP customers. Therefore it affects the CSP more strongly than an ISP, as all CSP customers are ISP customers, but not the other way around. In other words, degradation of the CSP harms only some customers of the ISP, but all customers of the CSP at that specific ISP. The revenues from newly acquired customer fields, the transit fees for newly acquired customer fields of the subjects, and the damages from passive degradation take effect with the beginning of the subsequent period. If negotiating parties successfully make a peering agreement, the connection and the terms of the agreement take effect with the beginning of the subsequent period. Once all subjects have taken their turn, a period is completed. The number of periods is limited to ten.

Table 1 depicts the eight treatments of our 3x2 factorial experimental design with regulatory and market treatment parameters. Our regulatory treatment parameters are *control*, *transparency obligation* and *interconnection obligation*. The regulatory tools take on the form of extensions of the control treatment. Interaction effects between the regulatory tools are not investigated. In addition to which regulatory tool we employ, we define a treatment by which market structure we investigate. Our *CSP setup* features traditional ISPs and TPs as well as a CSP as described above, while in our benchmark setup the CSP is replaced by a third ISP who behaves like the existing ISPs except that it enters the market without preexisting peering connections and a smaller network size of 1 (c.f. Figure 1a). That benchmark allows us to compare the interconnection agreements that emerge between CSPs and ISPs with those agreements that may emerge between ISPs and ISPs in a similar situation (later entry to the market).

	Benchmark Setup (c.f. Figure 1a)	CSP Setup (c.f. Figure 1b)
Control	Details of peering agreements are under NDA.	Details of peering agreements are under NDA.
Transparency obligation	Details of all existing peering agreements are visible to the other players before the negotiation phase.	Details of all existing peering agreements are visible to the other players before the negotiation phase.
Interconnection obligation	Negotiation partners are not allowed to reject “reasonable” offers.	Negotiation partners are not allowed to reject “reasonable” offers.

Table 1. Experimental Treatments

When no transparency obligation has been introduced, subjects naturally know their own negotiation history and all firms know each other’s regular price and revenue structure, but they do not know their competitors’ negotiation results and ISPs do not know the individual transit prices that TPs offer to the CSP. A transparency obligation makes this information available to all subjects: they see all opposing peering agreements including the split of the peering costs, the possible above cost premiums, and the transit prices. Thus, the network graph that is shown to the subjects contains more information in the transparency obligation setting. When entering further peering negotiations, subjects can take into account the additional information on what has been agreed on in similar situations and adopt their offers accordingly.

In the interconnection obligation treatment, ISPs are not allowed to reject 0:20,000 offers by the CSPs, i.e. offers by the CSP to bear the full costs of a peering connection. Regulation does not dictate capacity investments and it does not exclude other peering agreements though. Therefore, CSPs may offer paid peering contracts with a premium payment above the peering costs in order to encourage the ISPs to provide future capacity of the peering links when the network traffic grows.

Methodology and Implementation

In order to test the hypotheses stated above, we employ a controlled laboratory experiment. Economic laboratory experiments are used to isolate the effect of economic parameters in actual human decision

connection quality, A’s customers would take 2/3 of the damage, while B’s customers only bear 1/3 of the damage.

situations. Although the laboratory setting is necessarily only a stylized representation of reality, it can provide decision makers with complementary evidence about general mechanisms and potential pitfalls, which may remain relevant even in a more complicated situation. As Plott (1997) argues "If a mechanism does not work acceptably in a simple case created in a laboratory, then there may be no reason to think that it will work in the complex cases found in a field application." (p. 607). Thus, one contribution of economic laboratory experiments is to provide a test bedding environment before a mechanism is rolled out on a larger scale.

To induce an economic decision situation in a controlled way it is crucial to incentivize subjects with sufficient monetary rewards (Smith, 1976) and to ensure that all subjects understand the incentive structure and act accordingly (Schweitzer, 2012). In our experiment, the incentive structure is represented by experimental dollars that correspond to real money at a fixed exchange rate. At the end of the experiment one out of three repeated runs of the game is randomly selected for payoff. Additionally, the subjects receive a lump sum of 100 experimental dollars for their participation in the experiment. If they make a loss in the game round selected for payoff, the loss will be deducted from this lump sum. However, negative total payoffs are not possible.

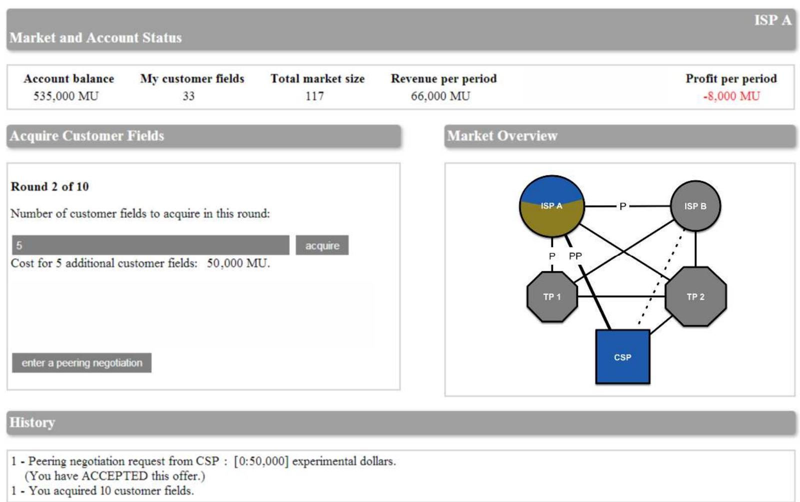

Figure 2. Screenshot of the experiment software in the benchmark with CSPs

The subjects participating in the experiment are students at the participating universities. While critics of economic experiments argue that students are lacking the experience of professionals in their respective fields, Ball and Cech's (1996) extensive review of experimental studies on this issue concludes that there is "little evidence for subject pool effects between students and market professionals acting as subjects in laboratory environments which *strictly* follow experimental economic precepts and test economic questions" (p. 257, emphasis in the source).

The procedure of our experiment is as follows. Each group of subjects participates only in one regulatory treatment, which is repeated in a sequence of three games. In each of these games they are one of altogether five participants in a group. The composition of this group is determined randomly in advance of each game round (random stranger matching). Before the start of each game period the participants are informed about their respective roles. The subjects are instructed before the beginning of the experiment and a standardized test program at individual computer terminals at our computer laboratory is used to test their understanding. After the subjects' comprehension of the decision situation is verified, the peering game is conducted at the same individual computer terminals. Communication between the subjects is not allowed.

The experiment is implemented in a web-based software, which employs the same interface for the computer terminals in the laboratory as for an online version of the game that could be used for gathering data of subject pools that are outside the reach of typical lab experiments. Our general approach allows us to collect data in a controlled lab experiment with students, but also by inviting industry experts to participate in our study via a web-based interface. All actions of the participants and all user interface events are immediately tracked and stored in a server-side database.

Figure 2 provides a screenshot of the prototype. The network graph of the current game round is depicted in the right part of the screen and provides a one-glance graphical overview of the current status of the game. In addition, in the upper part of the screen, the subjects are provided with a summary of their current financial situation and, in the lower part of the screen, with a history of their own past decisions and interactions in the game. In the left part of the screen, subjects make their decisions for the respective game round. These decisions include the number of customer fields to acquire, the optional choice of a quality degradation value and the optional initiation of a peering negotiation.

Results of the Pre-test

The experimental software has been pre-tested with seven independent groups of four students each in June 2014. The pre-test delivered first impressions on how subjects deal with the network and peering situation and on how the game may be configured for the actual experiment. First, as depicted in Figure 3, there was a pronounced tendency to underinvest in the first three periods of the game. In the subsequent Periods 4 to 10, actual investments roughly resembled the theoretically optimal investments. Second, three of the seven groups managed to establish peering connections between all participating subjects. Remarkably, every subject established at least one peering connection. And third, the results of a paper-based survey, conducted after the experiment, suggest that participation was perceived as immersive and the rules as self-explanatory. Note that the pre-test generates only preliminary results from a small sample of participants. The results were used to make important final design decisions and to calibrate the parameters for the actual experiment. The resulting design is still a simplified representation of the actual peering situation, but does capture many of the crucial aspects of the market dynamics in reality.

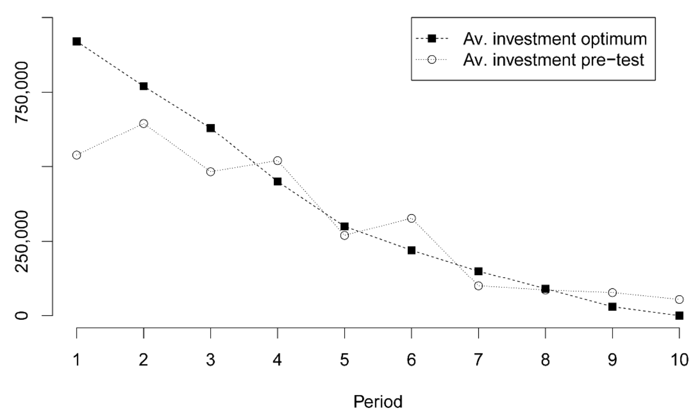

Figure 3. Average investment into customer fields as compared to the optimum investment, (n= 7 groups).

Outlook and Conclusion

In this paper we presented an experimental design to analyze interconnection markets when CSPs negotiate for private peering with ISPs. We introduced the option to strategically withhold capacity expansions, as well as paid peering negotiations. Furthermore, we introduced possible regulatory interventions that will be tested in separate treatments of the experiment and benchmarked against the "old" status quo in interconnection markets. Our research will contribute to a better understanding of ISPs behavior in the changing interconnection environment and shed light on the question if interconnection markets suffer from market failure. The results can help practitioners and policy makers alike to understand the implications of the transition of the interconnection market that started with the rise of streaming video and culminated in the recent events around Netflix's direct peering agreements. Our next steps include to conduct a second pre-test validating the treatment design and to conduct the full experiment afterwards.

References

Ball, S. B., and Cech, P.-A. 1996. "Subject Pool Choice and Treatment Effects in Economic Laboratory Research," *Research in Experimental Economics (6)*, pp. 239–292.

Besen, S., Milgrom, P., Mitchell, B., and Srinagesh, P. 2001. "Advances in Routing Technologies and Internet Peering Agreements," *American Economic Review (91:2)*, pp. 292–296.

Bloch, F., Jackson, M. 2007. "The formation of networks with transfers among players," *Journal of Economic Theory* (133:1), pp. 83–110.

Brodkin, J. 2014. "Netflix packets being dropped every day because Verizon wants more money". *Ars Technica*. Retreived from: http://arstechnica.com/information-technology/2014/02/netflix-packets-being-dropped-every-day-because-verizon-wants-more-money, Last accessed: 22.04.2014.

Crémer, J., Rey, P., Tirole, J. 2000. "Connectivity in the commercial Internet," *The Journal of Industrial Economics* (XLVIII:43), pp.433–472.

Economides, N. 2006. "The economics of the Internet backbone," in S. Majumdar, & I. Vogelsang (Eds.), Handbook of telecommunicationseconomics, Vol. II (pp. 375–412). Amsterdam: Elsevier Publishers.

Foros, Ø., Kind, H.J., Sand, J.Y. 2005. "Do Internet Incumbents Choose Low Interconnection Quality?," *Information Economics and Policy* (17), pp. 149-164.

Güth, W., Schmittberger, R., and Schwarze B. 1982. "An Experimental Analysis of Ultimatum Bargaining," *Journal of Economic Behavior and Organization (3)*, pp. 367–388.

Henze, B., C. Noussair, and B. Willems. 2012. "Regulation of network infrastructure investments: an experimental evaluation," *Journal of Regulatory Economics* (42:1), pp. 1-38.

Laffont, J. J., Marcus, S., Rey, P., Tirole, J. 2003. "Internet interconnection and the off-net-cost pricing principle," *RAND Journal of Economics* (34:2), pp. 370–390.

Jahn, E., Prüfer, J. 2008. "Interconnection and competition among asymmetric networks in the Internet backbone market," *Information Economics and Policy* (20:3), pp. 243-256.

Krämer J., Wiewiorra L., Weinhardt, C. 2013. "Net neutrality: A progress report," *Telecommunications Policy* (37:9), pp. 794-813 .

Mendelson, H., Shneorson S. 2003. "Internet Peering, Capacity and Pricing," *Mimeo*, Stanford University.

Norton, W. 2014. "The Art of Peering: The Peering Playbook (v1.4)," *Peering White Papers,* Retrieved from: http://drpeering.net/white-papers/Art-Of-Peering-The-Peering-Playbook.html, Last Accessed: 28.04.2014.

Norton, W. 2009. "Peering Simulation Game," *Peering White Papers,* Retrieved from: http://drpeering.net/white-papers/Peering-Simulation-Game.html. Last accessed: 21.04.2014.

Plott, C. R. 1997. "Laboratory Experimental Testbeds: Application to the PCS Auction," *Journal of Economics & Management Strategy (6:3),* pp. 605-638.

Reuters, 2013. "Deutsche Telekom, Orange, Telefonica in EU antitrust raids," *Thomson Reuters*, Retrieved from: http://www.reuters.com/article/2013/07/11/us-orange-antitrust-idUSBRE96A0CR20130711, Last accessed: 11.07.2013.

Sandvine. 2013. "Global Internet Phenomena Report 2H 2013," Retrieved from: https://www.sandvine.com/downloads/general/global-internet-phenomena/2013/2h-2013-global-internet-phenomena-report.pdf. Last accessed: 18.03.2014.

Schweitzer, S. M. 2012. Large-scale Multi-item Auctions: Evidence from Multimedia-supported Experiments, Karlsruhe: KIT Scientific Publishing.

Shin, S. J., Weiss, M. 2004. "Internet interconnection economic model and its analysis: Peering and settlement," *Netnomics* (6:1), pp. 43–57.

Sluijs, J. P., Schuett, F., & Henze, B. 2011. "Transparency regulation in broadband markets: Lessons from experimental research," *Telecommunications Policy* (35:7), pp. 592-602.

Smith, V. L. 1976. "Experimental Economics: Induced Experimental Value Theory," *The American Economic Review (66:2)*, pp. 274–279.