University of Pennsylvania
ScholarlyCommons

Departmental Papers (ESE)

Department of Electrical & Systems Engineering

8-14-2011

Exploring User-Provided Connectivity - A Simple Model

M. H. Afrasiabi

University of Pennsylvania, afram@seas.upenn.edu

Roch A. Guérin

University of Pennsylvania, guerin@acm.org

Follow this and additional works at: https://repository.upenn.edu/ease_papers

Part of the [Management Sciences and Quantitative Methods Commons](#), [OS and Networks Commons](#), and the [Systems and Communications Commons](#)

Recommended Citation

M. H. Afrasiabi and Roch A. Guérin, "Exploring User-Provided Connectivity - A Simple Model", . August 2011.

Suggested Citation:

Afrasiabi, M.H. and Guérin, R. "Exploring User-Provided Connectivity - A Simple Model." Proc. ICQT'11 - 7th International Workshop on Advanced Internet Charging and QoS Technology. Paris, France.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/ease_papers/600
For more information, please contact repository@pobox.upenn.edu.

Exploring User-Provided Connectivity - A Simple Model

Abstract

The advent of cheap and ubiquitous wireless access has introduced a number of new connectivity paradigms. This paper investigates one of them, user-provided connectivity or UPC. In contrast to traditional infrastructure-based connectivity, e.g., connectivity through the up-front build-out of expensive base-stations, UPC realizes connectivity organically as users join and expand its coverage. The low(er) deployment cost this affords is one of its main attractions. Conversely, the disadvantages of connectivity sharing and a high barrier-to-entry from low initial penetration create strong disincentives to its adoption. The paper's contributions are in formulating and solving a simple model that captures key aspects of UPC adoption, and in articulating guidelines to make it successful. For analytical tractability, the model is arguably simplistic, but the robustness of its findings is demonstrated numerically across a wide range of more general (and more realistic) configurations

Keywords

connectivity, externalities, sharing

Disciplines

Management Sciences and Quantitative Methods | OS and Networks | Systems and Communications

Comments

Suggested Citation:

Afrasiabi, M.H. and Guérin, R. "Exploring User-Provided Connectivity - A Simple Model." Proc. ICQT'11 - 7th International Workshop on Advanced Internet Charging and QoS Technology. Paris, France.

Exploring User-Provided Connectivity A Simple Model*

M. H. Afrasiabi and R. Guérin

University of Pennsylvania, Philadelphia PA 19104, USA,
(afram@seas,guerin@ee).upenn.edu

Abstract. The advent of cheap and ubiquitous wireless access has introduced a number of new connectivity paradigms. This paper investigates one of them, *user-provided connectivity* or UPC. In contrast to traditional infrastructure-based connectivity, *e.g.*, connectivity through the up-front build-out of expensive base-stations, UPC realizes connectivity organically as users join and expand its coverage. The low(er) deployment cost this affords is one of its main attractions. Conversely, the disadvantages of connectivity sharing and a high barrier-to-entry from low initial penetration create strong disincentives to its adoption. The paper’s contributions are in formulating and solving a simple model that captures key aspects of UPC adoption, and in articulating guidelines to make it successful. For analytical tractability, the model is arguably simplistic, but the robustness of its findings is demonstrated numerically across a wide range of more general (and more realistic) configurations.

1 Introduction

There is no denying that we are a networked society, and the increasing capabilities and versatility of mobile devices has fueled a growing thirst for ubiquitous connectivity, *i.e.*, connectivity everywhere and all the time. This has driven the growth and success of wireless carriers worldwide. These carriers tout comprehensive coverage and connectivity that in some instances approaches that of wired networks. However, their very success has often made it difficult to maintain the connectivity quality their users expect [19]. This is in part because connectivity relies on a costly *infrastructure*, whose deployment calls for careful long-term planning. This together with the relatively high cost of those services has awoken interest in alternative solutions to offering ubiquitous connectivity.

One such promising alternative is that of *user provided connectivity* (UPC), where connectivity grows “organically” as more users join the network and improve its coverage. In UPC, as users gain (local) access to connectivity, *e.g.*, from subscribing to an Internet Service Provider (ISP), they allow others to share that connectivity in exchange for either compensation or reciprocation. More specifically, a UPC user allows *roaming* users to obtain connectivity through its own local access for a small fee or the ability to enjoy the same benefits when itself roaming. This is made possible by the availability of low-cost wireless access solutions

* This work was supported by NSF grant CNS-0915982

(Wi-Fi), as popularized through services such as FON (<http://www.fon.com>). FON users purchase an access router (FONERA) that they use for their own local broadband access, but with the agreement that a (small) fraction of their access bandwidth can be made available to other FON users. In exchange, they receive the same privileges when roaming, *i.e.*, they are able to connect through the access points of other FON users.

The challenge faced by such a service model is that while it has low deployment costs (no expensive infrastructure build-out), it does not offer truly ubiquitous connectivity until it has reached a large enough level of penetration, *i.e.*, there are enough users to offer comprehensive coverage. This high externality in the service's value can, therefore, hamper early adoption and hence eventual success. Consider for example a FON-like service starting with no users. This makes the service unattractive to users that value ubiquitous connectivity highly, *e.g.*, users that roam frequently, because the limited coverage offers little connectivity beyond that of a user's "home base". On the other hand, sedentary users are unaffected by the negative utility associated with low coverage, and if the price is low enough can derive positive utility from the service; hence joining. If enough such (sedentary) users join, coverage may increase past a point where it becomes attractive to roaming users who will start joining. This would then ensure rapid growth of the service, were it not for a negative dimension to that growth.

Specifically, as more roaming users join, they start consuming resources in the home bases of other users. This lessens the utility these users derive from the service. This can cause some (sedentary) users to leave, and the corresponding reduction in coverage makes the service less attractive to roaming users that also start leaving. As a result, the initial period of growth in the service is followed by a decline, and the process repeats. The extent to which such behaviors arise depends on many factors, including the benefits users derive from the system, its cost, the severity of the degradations they experience when other users access the network through their home base, and the possible incentives the service provider offers to compensate for those.

The goal of this paper is to develop a simple model that can help understand how these many factors interact and affect the adoption of UPC services. The paper's contributions include

- Identifying service adoption equilibria (or lack thereof) and how they are influenced by system parameters;
- Characterizing "regions" (ranges of price and user valuation) that result in high or low adoption equilibria;
- Validating the robustness of the findings through the numerical evaluation of more realistic (and more complex) models.

The rest of the paper is structured as follows. Section 2 introduces the model and its parameters. Solution methods and findings are presented in Section 3. Section 4 demonstrates the robustness of the results to generalizations of the model. Finally, Section 5 gives a brief overview of related works, while Section 6 summarizes the paper's results and points to possible extensions.

2 Model Formulation

This section introduces a simple model for the decision process of individual users faced with the question of whether to adopt a UPC service. As commonly done, adoption depends on the *utility* a user derives from the service, with users adopting if their utility is non-negative. Utility depends on several factors, including the number (coverage) and type (roaming or not) of existing adopters. Users are myopic when evaluating the utility they expect to derive from the service, *i.e.*, do not account for the impact of their decision on other users, but as adoption levels change, an individual user's utility varies. In other words, the service value exhibit (positive and negative) externalities that affect adoption decisions. It is those dynamics we seek to capture.

For analytical tractability, the model relies on a number of simplifying assumptions. They are relaxed in Section 4, where we show that the findings remain qualitatively unchanged. The utility $U(\theta)$ of a user considering the adoption of a UPC service is given in Eq. (1), where θ , $0 \leq \theta \leq 1$ represents the roaming characteristics of the user, *i.e.*, a low θ indicates a sedentary user while a high θ corresponds to a user that frequently roams. The exact value of θ is private information, but its distribution (over the user population) is known.

$$U(\theta) = \gamma + \theta f(x) - p - g(m) + g^*(m). \quad (1)$$

The parameter γ denotes the intrinsic utility that all users associate with basic home connectivity, while $\theta f(x)$ represents the utility they derive from being able to connect through the home base of other users while roaming. The function $f(x)$ reflects the coverage that the UPC service offers, which grows with the level of adoption x , $0 \leq x \leq 1$. The factor θ in $\theta f(x)$ accounts for the effect of heterogeneity in the roaming characteristics of users, *i.e.*, low θ or sedentary users derive comparatively little benefits from being able to connect through other users' home base. The impact of heterogeneity could arguably be extended to how users value basic connectivity, *i.e.*, γ , as well as capture the fact that roaming users (high θ values) may in turn put less value than sedentary users on home connectivity. The assumption of a fixed γ value across users is a reflection of our focus on understanding how a UPC service can be attractive to users that require more than just home connectivity, *i.e.*, we are not trying to model the adoption of basic Internet service. Conversely, accounting for the fact that roaming users may value home connectivity less could be accomplished by replacing γ by $(1 - \theta)\gamma$. As shown in Appendix A of [1], this does not affect the overall structure of the model. As a result, we only consider the utility function of Eq. (1) in the rest of the paper.

The parameter p is the price charged for the service, while the factors $g(m)$ and $g^*(m)$ capture how roaming traffic affects users, with m measuring the volume of roaming traffic in the system (a function of how many roaming users have joined). Specifically, $-g(m)$ is the (negative) utility associated with roaming traffic consuming resources in the home base of other users. We note that this penalty depends only on the volume of roaming traffic and not on the availability

of resources at a user's base station. This is reasonable in the context of home based Internet connectivity where access bandwidth is the main resource, and roaming users can connect at any time. Conversely, the quantity $g^*(m)$ represents possible compensation that the UPC service provider may offer to offset the negative impact of roaming traffic, *e.g.*, by logging external accesses to a user's home base and offering payment for each instance.

For analytical tractability, we make several assumptions regarding the form and range of the parameters of Eq. (1) (as mentioned earlier, Section 4 explores the impact of relaxing many of those assumptions).

First, the parameter θ that measures a user's propensity to roam, is taken to be uniformly distributed in $[0, 1]$. This implies that the adoption level, x , of a UPC service is given by

$$x = \int_0^1 I_{[U(\theta)]} d\theta, \quad (2)$$

where $I_{[U(\theta)]}$ is an indicator function that takes value 1 if $U(\theta) \geq 0$ and zero otherwise.

Next, we assume that the distributions of users over the service area and their roaming patterns are uniform. A uniform distribution of users implies that coverage grows in proportion to adoption, x . Similarly, uniform roaming patterns mean that roaming traffic is evenly distributed across users' home bases, *i.e.*, on average all home bases see the same volume of roaming traffic.

The next assumption concerns the shape of the function $f(x)$. Specifically, we expect frequently roaming users, *i.e.*, users with a high θ value, to see little or no value in the service until its penetration is high enough to realize a certain minimum level of coverage. This means that the overall connectivity utility of those users, as measured by $\gamma + \theta f(x)$, should be positive only once x is large enough. For ease of exposition we use the function $f(x)$ below to capture this effect.

$$f(x) = d(2x - 1), \quad d > 0,$$

where the factor d scales the weight of this utility relative to other terms in $U(\theta)$.

The function $f(x)$ is linear in x and negative for small x , *i.e.*, for x below a threshold value of $1/2$ ¹. It should be noted that a similar outcome could be realized while keeping $f(x)$ positive for all x , by assuming instead that the value of home base connectivity decreases for roaming users, *i.e.*, replace γ by $(1 - \theta)\gamma$ in Eq. (1). As discussed in Appendix A of [1], this yields a structurally equivalent model.

With a similar goal of simplicity, both the penalty and the compensation that users receive from providing connectivity to roaming traffic are assumed proportional to the volume of roaming traffic they carry. In other words, the functions $g(m)$ and $g^*(m)$ are taken to be linear functions of m , *i.e.*,

$$\begin{aligned} g(m) &= c m, \quad c > 0 \\ g^*(m) &= b m, \quad b > 0 \end{aligned}$$

¹ Other threshold values obviously quantitatively affect the outcome, but do not qualitatively affect overall *behaviors*.

where

$$m = \int_0^1 \theta I_{[U(\theta)]} d\theta.$$

In practice, the volume of roaming traffic at individual home bases varies. However, users whose home base carries more roaming traffic also receive a proportionally larger compensation (when $b > 0$). This should mitigate the impact of heterogeneity.

Using the above assumptions in Eq. (1), a user's utility becomes

$$U(\theta) = k + l m + \theta (2x - 1), \quad (3)$$

where $k = \gamma - p$ and $l = b - c$, and where for normalization purposes, the maximum roaming utility d was taken to be 1. We also assume that roaming and home base connectivity are of a similar nature, so that the utility γ derived from home base connectivity is no more than the maximum utility from roaming connectivity, *i.e.*, $0 \leq \gamma \leq 1$. From the above expression for k , this then implies

$$k \leq 1 - p \leq 1. \quad (4)$$

Before proceeding with investigating the adoption process that Eq. (3) gives rise to, we note that its parameters $k = \gamma - p$ and $l = b - c$ include both *exogenous* and potentially *endogenous* components. Specifically, γ and c capture external system properties, *i.e.*, users valuation for connectivity and their sensitivity to the impact of roaming traffic, respectively. The values of such exogenous parameters can be estimated, *e.g.*, using techniques from marketing research as discussed in [10], but not controlled. In contrast, the service price, p , and incentives for providing connectivity to roaming users, b , are both under the control of the UPC provider. They can, therefore, arguably be endogenized to optimize some measure of success such as profit. Using the results of this paper to explore such options is a topic of ongoing research.

3 Equilibria and Adoption Dynamics

With Eq. (3) in place, it is possible to investigate the dynamics of user adoption over time. We formulate a discrete-time model that evaluates user adoption decisions at successive epochs. For simplicity², at epoch $(n + 1)$ all users are assumed to know the system state produced by adoption decisions at epoch n . Users with a non-negative utility then proceed to adopt. Specifically, the utility at epoch $(n + 1)$, $U_{n+1}(\theta)$, of a user with roaming value θ is given by

$$U_{n+1}(\theta) = k + l m_n + \theta (2x_n - 1), \quad (5)$$

where x_n and m_n are the adoption level and volume of roaming traffic produced by adoption decisions at epoch n .

² Section 4 gives numerical results for a more realistic, diffusion-based adoption model.

The next proposition (the proof is in Appendix B of [1]) establishes a key result that ensures the analytical tractability of a solution, namely, that as the system evolves adopters remain associated with a continuous set of θ values. In other words, the set of adopters does not fragment.

Proposition 1 *For all choices³ of $k, 0 < k \leq 1$ and l , the set of adopters is characterized by a range of θ values of the form $[0, \hat{\theta}]$ or $[\hat{\theta}, 1], 0 \leq \hat{\theta} \leq 1$.*

Returning to Eq. (5), note that m_n depends not just on the overall adoption level, x_n , but also on *which* users have adopted. This is because the amount of roaming traffic a user contributes depends on its θ value. As a result, characterizing the system state calls for specifying the level of adoption *and* identifying adopters. As shown in Appendix B of [1], adoption at epoch $n + 1$ then depends on adoption levels at *both* epochs n and $n - 1$. Specifically, x_{n+1} depends on x_n , *and* on whether x_{n-1} was in the range $[0, 1/2)$ or $[1/2, 1]$. Although as stated in Proposition 1 adopter sets remain continuous, they can experience abrupt changes when adoption crosses the threshold ($x = 1/2$) of $f(x)$. Abrupt changes are inherent in discrete time models, but as shown in Appendix B of [1], this introduces additional difficulties in characterizing adoption evolution. These are technical in nature, and call for the use of different functional expressions when characterizing adoption after crossing the $x = 1/2$ threshold (in either direction).

Fig. 1. Adoption Evolution as a Function of Initial Adoption.

This is illustrated in Fig. 1, which corresponds to a scenario where depending on initial adoption levels, final adoption can converge to either one of two stable⁴

³ Starting from zero adoption, non-zero adoption is possible only if $k > 0$.

⁴ In this scenario, there is also an unstable equilibrium in each of $[0, 1/2)$ and $[1/2, 1]$.

equilibria in $[0, 1/2)$ or $[1/2, 1]$. The x -axis of the figure is the current adoption level, while the y -axis, $H(x)$, denotes the next adoption level given x . The dash-dot curves of Fig. 1 correspond to expressions that characterize the evolution of adoption just after crossing the $x = 1/2$ threshold, while the solid lines are used to characterize adoption while it progresses inside either $[0, 1/2)$ or $[1/2, 1]$. The dashed arrows illustrate adoption trajectories for different initial adoption levels. For example, when the system starts with no adopters, $x_0 = 0$, adoption increases monotonically until it reaches about 10%, the stable equilibrium in $[0, 1/2)$. If seeding is used, *i.e.*, $x_0 > 0$, the outcome depends on the seeding level. When seeding is “low,” *e.g.*, $x_0 \approx 35\%$, adoption declines back to 10%. If seeding is high enough, *e.g.*, $x_0 \approx 46\%$, adoption enters $[1/2, 1]$ and eventually converges to the higher adoption equilibrium in that interval (around 85%).

Using the approach developed in Appendix C of [1], adoption evolution can be characterized. Possible outcomes are summarized in the table on the left-hand-side of Fig. 2, with the right-hand-side displaying the regions of the (k, l) plane corresponding to each table entry. Region boundaries, *i.e.*, f_1, f_2, f_3 and f_4 , are derived from conditions on the roots of the equation $H(x) = x$ as discussed in Appendix C of [1]. There can be multiple equilibria, both stable ($\bullet$) and unstable ($\circ$), as well as fixed points associated with an “orbit” ($\odot$). Orbits can be convergent, periodic, or chaotic depending on the choice of (k, l) values (in regions 2' and 3'). Finally, some (k, l) values (region 1) altogether result in the absence of any equilibrium (denoted by — in the table).

Fig. 2. Possible Combinations of Equilibria and Associated Regions of the (k, l) Plane.

The scenario shown in Fig. 1 corresponds to Case 8 of Fig. 2, where as discussed earlier, convergence to a stable equilibrium in either $[0, 1/2)$ or $[1/2, 1]$ is possible depending on initial adoption levels. As indicated in the table of

Fig. 2, other behaviors are possible based on which region of the (k, l) plane the system parameters belong to. In the rest of this section, we review the different possible outcomes that can arise, and attempt to provide some intuition in how and why they are associated with different combinations of system parameters.

Behavior (i): Absence of convergence to an equilibrium. This arises in Cases 1, 2', and 3'. Region 1 consists of relatively high values of $k(= \gamma - p)$, *i.e.*, at its offered price the intrinsic value of the service is reasonably high, but rather negative values of $l(= b - c)$, *i.e.*, even accounting for compensation (b), the negative impact of roaming traffic is high. This produces the following dynamics: When the service has few users, coverage is low and frequently-roaming users find the service unattractive in spite of the high k . In contrast, sedentary users are unaffected by the limited coverage, and the high k value entices them to adopt. As they adopt, coverage improves and the service becomes attractive to roaming users, which start adopting. The associated growth in roaming traffic, however, starts to negatively affect sedentary users that derive little benefits from the improved coverage. This leads some of them to disadopt, which reduces coverage so that eventually roaming users start leaving as well. Once roaming traffic has been sufficiently reduced, the service becomes again attractive to sedentary users, and the cycle repeats. A similar, albeit more nuanced process is at work in regions 2' and 3'. Appendix C of [1] offers additional discussions.

Behavior (ii): Convergence to a single stable equilibrium in either $[0, 1/2)$ or $[1/2, 1]$, independent of initial adoption. This arises in Cases 2, 3, 4, and 5. Cases 2 and 4 correspond to low k values and relatively large negative l values. Because of the low k value, few sedentary users adopt and coverage never gets high enough to make the service attractive to frequent roamers. Hence, adoption saturates at a low level in $[0, 1/2)$. Seeding is of no help in this case, as a combination of a low intrinsic value and a high (negative) impact of roaming traffic keeps the service unattractive to frequent roamers even if coverage is artificially increased. A symmetric situation exists in Cases 3 and 5, where adoption converges to a single stable equilibrium in $[1/2, 1]$. The value of k is now relatively high and l boasts only a small negative value. The high intrinsic value of the service initially attracts sedentary users that are not deterred by the limited coverage. Once enough of them have adopted, frequent roamers start joining. Because incentives compensate for the impact of the increasing roaming traffic, few sedentary users leave and adoption stabilizes at a high level.

Behavior (iii): Convergence to one of two stable equilibria in $[0, 1/2)$ or $[1/2, 1]$, as a function of initial adoption. This arises in Cases 6, 7, and 8, which share relatively low k values and marginally negative l values. Under those conditions, while adoption (coverage) is low, frequent roamers are not interested in the service and the small k value limits the number of sedentary users who adopt. Hence, adoption saturates at a low level. However, unlike (ii), this is an

instance where seeding can help. In particular, a high enough level of seeding can lead to a much higher final adoption (in $[1/2, 1]$ as opposed to $[0, 1/2)$). Specifically, if seeding is high enough, frequent roamers will start adopting in spite of the low k value. As their number grows and coverage continues improving more adopt and even some sedentary users might also adopt because of the relatively high level of compensation they receive to allow roaming traffic through their home base. As a result, overall adoption eventually stabilizes at a high level.

Fig. 3. Adoption Outcomes as a Function of k and l .

The above behaviors are illustrated in Fig. 3 that plots the “final” adoption levels for different (k, l) pairs when starting from an initial adoption level of $x_0 = 0$. In scenarios where adoption does not converge, *i.e.*, **Behavior (i)**, the adoption level reported in the figure was sampled at a particular iteration. The figure clearly identifies the regions of the (k, l) plane that correspond to chaotic or at least non-converging adoption (regions 1, 2', and 3'), low adoption (regions 2 and 4, as well as regions 6, 7, and 8 since no seeding was used), and regions of high adoption (regions 3 and 5).

4 Robustness to Model Variations

The model and the analysis behind the paper’s results are predicated on a number of simplifying assumptions that are unlikely to hold in practice. It is, therefore, important to validate that the findings and insight derived from these results remain applicable under more realistic conditions. For that purpose, a number of “perturbations” were introduced to the modeling assumptions, and their impact on the results evaluated. The perturbations that were investigated include

1. Relaxing the synchronized nature of adoption decisions and perfect knowledge of system state, *i.e.*, through a “diffusion-like” process that introduces heterogeneity in how users learn and react to changes in system state.

2. Generalizing the distribution of users' roaming characteristics θ , and therefore sensitivity to coverage, *i.e.*, from uniform to arbitrary distributions;
3. Varying users' sensitivity to roaming traffic and incentives compensating for it, *i.e.*, by considering sub-linear and super-linear utility functions;

Because those perturbations typically imply a loss of analytical tractability, numerical evaluations were used to assess their impact. A representative scenario is shown in Fig. 4 that assumes a diffusion-like adoption process in a configuration where the analysis predicts the existence of both a stable and an unstable equilibrium in $[1/2, 1]$. The paper's analytical model assumes that adoption proceeds by discrete jumps, so that it eventually enters a region where convergence to the stable equilibrium is guaranteed. Adoption progression is different under a diffusion-like model, as there is latency in how changes in adoption affect users' utility, and therefore adoption decisions. As a result, adoption trajectories can "traverse" unstable equilibria, but those traversals can depend on the initial service penetration, *e.g.*, as realized through seeding. This is illustrated in Fig. 4.

Fig. 4. Adoption evolution as a function of initial penetration.

Fig. 4a illustrates the adoption trajectory when the service starts from zero penetration. In this case, adoption starts relatively steeply as the service is attractive to many sedentary (low θ) users. Adoption then slows down as it approaches the unstable equilibrium, but the initial momentum is sufficient to "carry it through" that region. The pace of adoption then picks up again and eventually overshoots the stable equilibrium before it finally converges back to it. Fig. 4b shows a different behavior when initial penetration is high, *e.g.*, because of seeding, but below the unstable equilibrium. In this scenario, the artificially inflated utility of many adopters drops quickly and they disadopt, which trig-

gers a rapid initial drop in service adoption. However, once adoption has dropped sufficiently, a similar process as that followed in Fig. 4a emerges, and adoption proceeds to grow again and finally converge to the same final adoption level.

Details reporting the outcome of investigations of many other perturbation scenarios can be found in Appendix D of [1]. They establish that the main findings of the paper remain valid in those more general and realistic settings.

5 Related Works

The service adoption process that the paper targets exhibits both positive and negative externalities. There is a vast literature investigating the effect of externalities, often called *network effects* [13], but the majority of these works focus on either positive or negative externalities separately. See for example [4, 5, 8, 11] for works exploring the impact of positive externalities on product adoption and competition. Conversely, the impact of negative externalities, *e.g.*, from congestion, has been extensively investigated in the context of pricing for both communication networks [9, 14, 16, 18] and transportation systems [2, 12, 15].

Systems that exhibit both positive and negative externalities have been studied mostly in the context of the theory of clubs [17]. Club-like behaviors also arise in peer-to-peer (p2p) systems where more peers increase the total resources available to store content, but induce a higher load on file-serving peers. This has triggered the investigation of *incentives* to promote resources sharing, *e.g.*, BitTorrent “tit-for-tat” [6] or [7] that also explores a possible application to a wireless access system similar in principle to the one considered in this paper.

This paper differs from these earlier works in a number of ways. It introduces a model for individual adoption decisions of a service, which allows for heterogeneity in the users’ valuation of the service. In particular, certain users (roaming users) have a strong disincentive to adoption when coverage/penetration is low, while others (sedentary users) are mostly insensitive to this factor. Conversely, this heterogeneity is also present in the negative externality associated with an increase in service adoption, which depends not just on the number of adopters, but on their identity as well, *i.e.*, roaming or sedentary users.

6 Conclusion

The paper introduces a simple model that captures the positive and negative externalities of a UPC service. The model’s solution characterizes possible outcomes (equilibria) and when they arise. The robustness of the findings to relaxations in the model’s simplifying assumptions was verified numerically.

There are many extensions of interest to this basic model. The first is to endogenize system parameters associated with prices and incentives, based on a revenue maximization objective. This is the topic of ongoing work. Empirical validation of the approach is also obviously desirable and a target for future work, *e.g.*, by collaboration with a UPC provider such as FON. Other extensions include the introduction of competition between UPC and infrastructure

based services, as well as the investigation of the possible benefits of cooperation between two such offerings *i.e.*, how a UPC service can best complement a traditional 3G or 4G offering. This is a topic that has seen much recent interest [3] because of the rise in bandwidth demand originating from smartphone and other Internet-enabled portable devices, *e.g.*, e-readers.

References

1. M. H. Afrasiabi and R. Guérin. Exploring user-provided connectivity – a simple model. Technical report, U. Pennsylvania, 2011. Available at <http://www.seas.upenn.edu/~afram/tr/upc2011.pdf>.
2. M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation. Cowles Commission Monograph*. Yale University Press, 1956.
3. D. Burstein. 21st Century triple networks: Ubiquitous 4G, WiFi, & Wires. CircleID - Internet Infrastructure. January 28, 2011.
4. L. Cabral. On the adoption of innovation with network externalities. *Mathematical Social Sciences*, 19, 1990.
5. J. P. Choi. The provision of (two-way) converters in the transition process to a new incompatible technology. *The Journal of Industrial Economics*, 45(2), 1997.
6. B. Cohen. Incentives build robustness in BitTorrent. In *Proc. 1st Workshop Econ. Peer-to-Peer Sys.*, Berkeley, CA, June 2003.
7. C. Courcoubetis and R. Weber. Incentives for large peer-to-peer systems. *J. Select. Areas. Commun.*, 24(5), May 2006.
8. J. Farrell and G. Saloner. Installed base and compatibility: Innovation, product preannouncements, and predation. *American Economic Review*, 76, 1986.
9. R. J. Gibbens and F. P. Kelly. Resource pricing and the evolution of congestion control. *Automatica*, 35, 1999.
10. P. E. Green, A. Krieger, and Y. Wind. Thirty years of conjoint analysis: Reflections and prospects. *Interfaces*, 31(3):S56–S73, 2001.
11. M. Katz and C. Shapiro. Technology adoption in the presence of network externalities. *Journal of Political Economy*, 94(3), June 1986.
12. F. P. Kelly. Road pricing. *Ingenia*, 29, 2006.
13. S. Liebowitz and S. Margolis. Network effects. In M. Cave, S. Majumdar, and I. Vogelsang, editors, *Handbook of Telecommun. Econ.*, volume 1. Elsevier, 2002.
14. J. K. MacKie-Mason and H. R. Varian. Pricing the internet. Computational Economics 9401002, EconWPA, January 1994.
15. A. Odlyzko. The evolution of price discrimination in transportation and its implications for the Internet. *Review of Network Economics*, 3, 2004.
16. I. C. Paschalidis and J. N. Tsitsiklis. Congestion-dependent pricing of network services. *IEEE/ACM Transactions on Networking*, 8(2), April 2000.
17. T. Sandler and J. T. Tschirhart. The economic theory of clubs: An evaluative survey. *Journal of Economic Literature*, 18(4), December 1980.
18. R. Srikant. *The Mathematics of Internet Congestion Control*. Birkhauser, 2004.
19. J. Wortham. Customers angered as iPhones overload AT&T. The New York Times. September 2, 2009.