Chapter 9
Comparing and Visualizing the Social Spreading of Products on a Large Social Network

Pål Roe Sundsøy, Johannes Bjelland, Geoffrey Canright, Kenth Engø-Monsen, and Rich Ling

Abstract By combining mobile traffic data and product adoption history from one of the markets of the telecom provider Telenor, we define and measure an adoption network—roughly, the social network among adopters. We study and compare the evolution of this adoption network over time for several products—the iPhone handset, the Doro handset, the iPad 3G and videotelephony. We show how the structure of the adoption network changes over time, and how it can be used to study the social effects of product diffusion. Specifically, we show that the evolution of the Largest Connected Component (LCC) and the size distribution of the other components vary strongly with different products. We also introduce simple tests for quantifying the social spreading effect by comparing actual product diffusion on the network to random based spreading models. As videotelephony is adopted pairwise, we suggest two types of tests: transactional- and node based adoption test. These tests indicate strong social network dependencies in adoption for all products except the Doro handset. People who talk together, are also likely to adopt together. Supporting this, we also find that adoption probability increases with the number of adopting friends for all the products in this study. We believe that the strongest spreading of adoption takes place in the dense core of the underlying network, and gives rise to a dominant LCC in the adoption network, which we call "the social network monster". This is supported by measuring the eigenvector centrality of the adopters. We believe that the size of the monster is a good indicator for whether or not a product is going to "take off".

P.R. Sundsøy (✉) · J. Bjelland · G. Canright · K. Engø-Monsen
Corporate Development, Telenor ASA, Oslo, Norway
e-mail: pal-roe.sundsoy@telenor.com

R. Ling
IT-University, Copenhagen, Denmark
e-mail: rili@itu.dk

T. Özyer et al. (eds.), *The Influence of Technology on Social Network Analysis and Mining*, Lecture Notes in Social Networks 6, DOI 10.1007/978-3-7091-1346-2_9,

9.1 Introduction

This paper is an extended version of [28], where new methodology is developed and applied to existing datasets. A new longitudinal dataset is also added. The new methodology quantifies the social spreading effects for transactional products, and is applied on the video telephony data. In addition we enrich the paper by adding spreading studies of the recent iPad 3G tablet.

The current study is motivated by the question of how people adopt new products and services, and what role the underlying social network structure plays in this process. The effect of the social network on product adoption and diffusion has been well documented in early market research, see e.g. [29] for an overview of this research. Most of the early studies have suffered from limited network data availability, since social networks have traditionally been difficult to measure. Several theoretical network models have been developed. Some are less realistic due to the evolutionary nature and power law degree distributions [5, 14]. Analyses of more realistic models can be found in [6, 18, 31].

In recent years massive social network data have been made available to researchers through electronic phone logs [9, 15, 23, 26] and online social network services [2, 21]. These studies have confirmed that “the network matters” when customers decide to churn [4, 9] and when purchase decisions are made [8, 15]. Most of the existing research on product diffusion on networks has been focused on a single product, with a static snapshot of the social network. For an overview of analyses of large networks, the reader is referred to [1, 10, 11, 18, 24, 27, 30]. A few papers study the evolution of real-world networks [12, 19, 20, 22, 25]. In this paper we will present an empirical study of how the social network among adopters of telecom-products develops over time. In addition we will show how the product diffusion depends on the underlying social network.

We know that, for many products, a person’s adoption probability increases with the number of that person’s friends or contacts that have adopted the same product [9, 15]. This can be interpreted as inter-personal or social influence, and can be measured empirically. These measurements do not typically say anything about the large-scale structure of the social network. In telecommunications it is possible to obtain detailed anonymized mobile traffic data for a large connected network of users. One can then use this telephony network as a proxy for the underlying social network. Studies show that a telephony network is a very good proxy for the real social network [13]. Furthermore, by combining telephone network data over time with the adoption history for a product of interest, it is possible to observe how different products spread over the social network.

Using anonymized datasets from one of Telenor’s markets, we will show how two different handsets have spread over the social network. The cases being used in this study are the highly buzzed iPhone, and the less fancy, but user-friendly, Doro type handset, which is more common among elderly people.

Both handsets are tracked from their early introduction and followed for a period of 2 years. We also present the tracking of a *transactional* product, mobile video telephony, a potentially useful product which allows users to talk to each other, while simultaneously viewing one another (or one another's surroundings)—given certain technological preconditions.

At the end we will show how a computer tablet, specifically iPad 3G, is spreading over the network. We include this as a preview on ongoing research due to the recent introduction in the market.

We start by introducing the *adoption network*, a construction which is readily visualized and which gives insight into the spreading of a product or service. Our figures will include many visualizations, which, we believe, are useful in understanding the product diffusion process on the underlying social network.

9.2 The Adoption Network

We will in this section define what we mean by the adoption network, followed by an empirical example.

9.2.1 Definition of Adoption Network

We define an adoption network as follows. Given a measured telephony network, the node set of the adoption network is the set of subscribers that have adopted a given product, and the links are the communication links belonging to this subset.

Mathematically, an adoption network is thus a subgraph of the whole mobile communication network $\mathrm{C} = (\mathrm{c}_{ij})$, where C represents (most generally) a weighted, directed, and possibly disconnected graph.

The mobile communication matrix C places a link between each pair of communicating subscribers, so that each nonzero element c_{ij} represents communication. The communication can be based for example on a weighted sum of SMS and voice duration (in which case we call these weighted links *W-links*), or other transactional data like video telephony traffic. All the results in this paper depend only on whether the communication link exists or not, without consideration to weight or direction. We consider traffic between Telenor subscribers in one market, which implies that the matrix will be $\mathrm{n} \times \mathrm{n}$ large, where n is the size of the customer base (several million subscribers).

The adoption network is then simply the subgraph of C formed by including only the adopting nodes and their common links. As we will see, for a *transactional* product (video telephony), there are two distinct useful choices for the communication links to be used in defining the adoption graph: (i) the standard (voice + SMS) links, or (ii) the links representing the use of the transactional service.

9.2.2 Introducing the "Social Network Monster" by Example

Figure 9.1 shows the empirical iPhone adoption network from Q4 2007 (This was measured before the iPhone had been introduced into the Telenor net; hence these

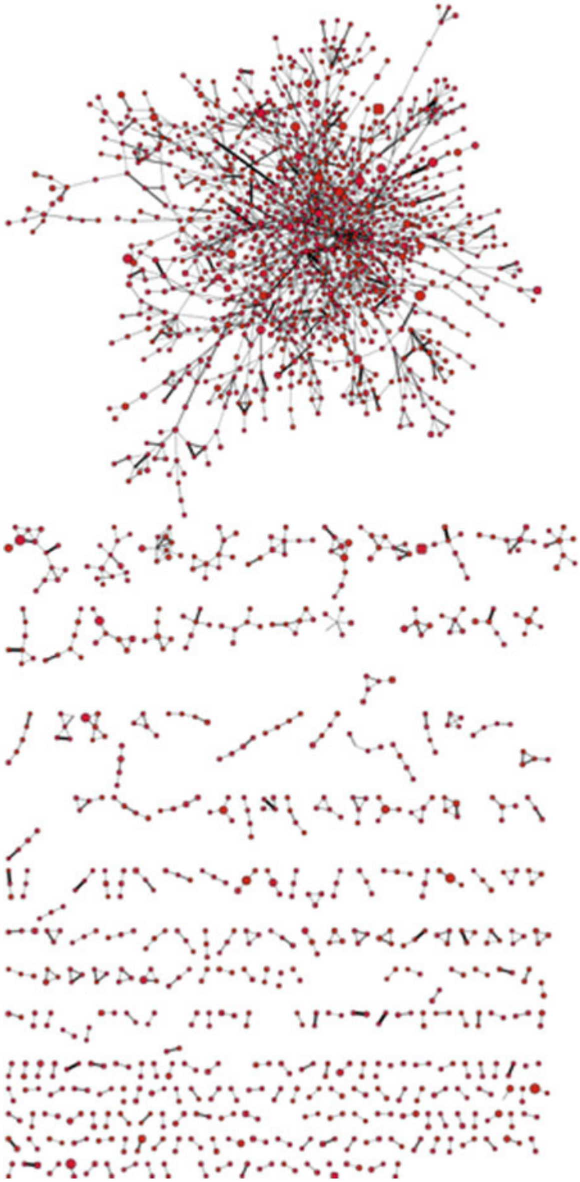

Fig. 9.1 iPhone Q4 2007 adoption network. One node represents one subscriber. Node size represents downloaded internet volume. Link width represents a weighted sum of SMS + voice. Isolates-adopters who are not connected to other adopters-are not shown in the picture

users have presumably bought their iPhones in the US and "cracked" them for use on the Telenor net). The data show that 42 % of the iPhone users communicated with at least one other iPhone-user, which speaks to the social nature of technology consumption, while 58 % did not have any iPhone contacts. We call the latter *isolates*. We do not include isolates in any of our visualizations of adoption networks, but do include them in all results counting number of users. We also study the connected components of the adoption network, where the connected components are subgraphs in which any two nodes are connected to each other by paths. Using this convention, we find for example that the largest connected component (LCC) in the adoption network of Fig. 9.1 includes 24.7 % of the total number of adopters (while representing over half of the nodes visible in Fig. 9.1). When the LCC in the adoption network is much bigger than all other connected components, and also represents a large fraction of all adopters, we will call the LCC a "social network monster". We note that this is not a precise definition; but we find that such monsters are typically found in adoption networks, and hence believe that the concept is useful.

9.3 Time Evolution of Adoption Networks

By studying the time evolution of an adoption network, we can get some insight into how the product which defines the adoption network is diffusing over the underlying social network. In particular we will often focus on the time evolution of the LCC of the adoption network—which may or may not form a social network monster. We recall from Fig. 9.1 that the other components are often rather small compared to the LCC. Hence we argue that studying the evolution of the LCC itself gives useful insight into the strength of the network spreading mechanisms in operation. It also gives insight into the broader context of adoption. As described in [20], two friends adopting together does not necessarily imply social influence—there might also be external factors that control the adoption. In this paper we will not try to separate the 'influence-effects' from external effects such as network homophily. Instead, when we observe a tendency that people who talk together also adopt together, we will use the term 'social spreading'—without making any implicit claim as to the underlying mechanism.

9.3.1 The iPhone Case

The iPhone 2G was officially released in the US in late Q2 2007 followed by 3G in early Q3 2008 and 3GS late Q2 2009. It was released on the Telenor net in 2009. Despite the existence of various models, we have chosen to look at the iPhone as one distinct product, since (as we will see) the older models are naturally substituted in our network. Figure 9.2 shows the development of the iPhone monster in one particular market. We observe how the 2G phone is gradually substituted by 3G (red to green), followed by 3GS in Q3 2009 (yellow nodes). The 2G model falls from 100 to 10 % with respect to all the iPhone subscribers in the adoption network. In Q1 2009 we observe the same amount of 2G as 3G models.

We show only the LCC in Fig. 9.2 because the other components are visually very much like those seen in Fig. 9.1; the main change over time is that the non-LCC components increase greatly in number, but not in size. That is, essentially all

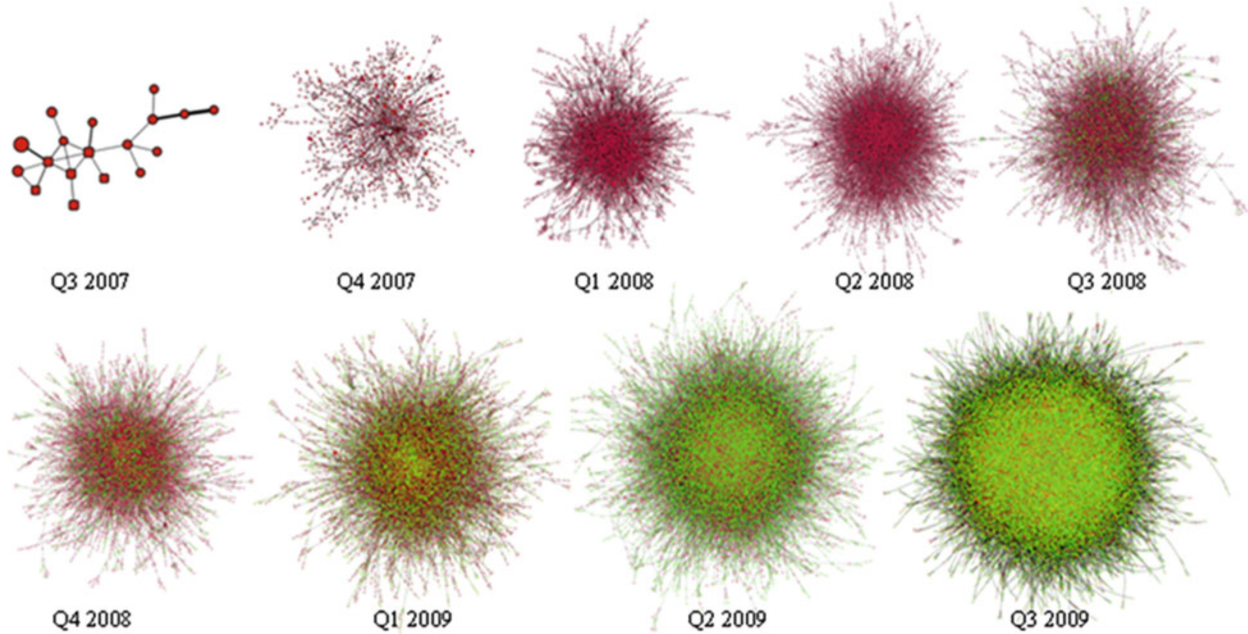

Fig. 9.2 Time evolution of the iPhone adoption network. One node represents one subscriber. Node color: represents iPhone model: *red* = 2G, *green* = iPhone 3G, *yellow* = 3GS. Node size, link width, and node shape (attributes which are visible in Q3 2007) represent, respectively, internet volume, weighted sum of SMS and voice traffic, and subscription type. *Round* node shape represents business users, while *square* represents consumers

significant growth in component size occurs (in the iPhone case) in the LCC. We regard this growth as a sign that the iPhone is spreading strongly ("taking off") over the social network. It is worth noting that there is a significant marketing "buzz" and external social pressure associated with the iPhone that is perhaps unique. We will offer in later Sections other kinds of measurements which support this conclusion.

9.3.2 The DORO Case

The next example is the Doro. As with the iPhone, there are several different models that are considered collectively. It is a handset which is easy to use, and mainly targeted towards elderly people [17]. Since Doro has a relatively low number of non-isolated users in all quarters studied, we present in Fig. 9.3 visualizations of the whole adoption network (minus isolates) over the entire time period, from introduction (in Q4 2007) to Q3 2009. Figure 9.3 shows that most Doro users that are not isolates appear in pairs in the adoption network. The social network monster never appears—the contrast with the iPhone case is striking. We believe that the kind of adoption network evolution seen in Fig. 9.3 is indicative of a product where "buzz" effects—social influence in the spreading of adoption—are weak or absent, whereas what we see in Fig. 9.2 indicates strong buzz effects. It is possible to argue that the adoption of the Doro is more of an individual choice, or perhaps even the choice of the user's children who wish to be in contact with their elderly parents. We note finally that the tiny "monster" (LCC) seen in Q3 2009 of Fig. 9.3 consists entirely of enterprise subscribers. Hence we speculate that these users are not the elderly of the target segment, but rather users with some other interest in the product.

9.3.3 Mobile Video Telephony Case

Compared to iPhone and Doro, video telephony has no value for an isolated user; thus users will always appear in pairs. A similar (pairs-only) constraint may be seen in [3], where the connections are based on romantic relations. In the video telephony case we actually have two distinct link sets which may be used to define an adoption network: W-links (voice + SMS), and video links. Thus for mobile video telephony we create and study two distinct adoption networks:

- The video-link set gives rise to the video adoption network (VAN)
- The W-linkset (voice + SMS) gives rise to the W-Video adoption network (WVAN).

We find that these two networks for video telephony have quite different behavior. We consider first the WVAN. This network connects users who (i) have a communication link (voice and/or SMS) and (ii) both use video telephony—not necessarily with each other. For the WVAN we find consistently a large social monster, much like the one seen in Fig. 9.2 that starts to form. However, differently

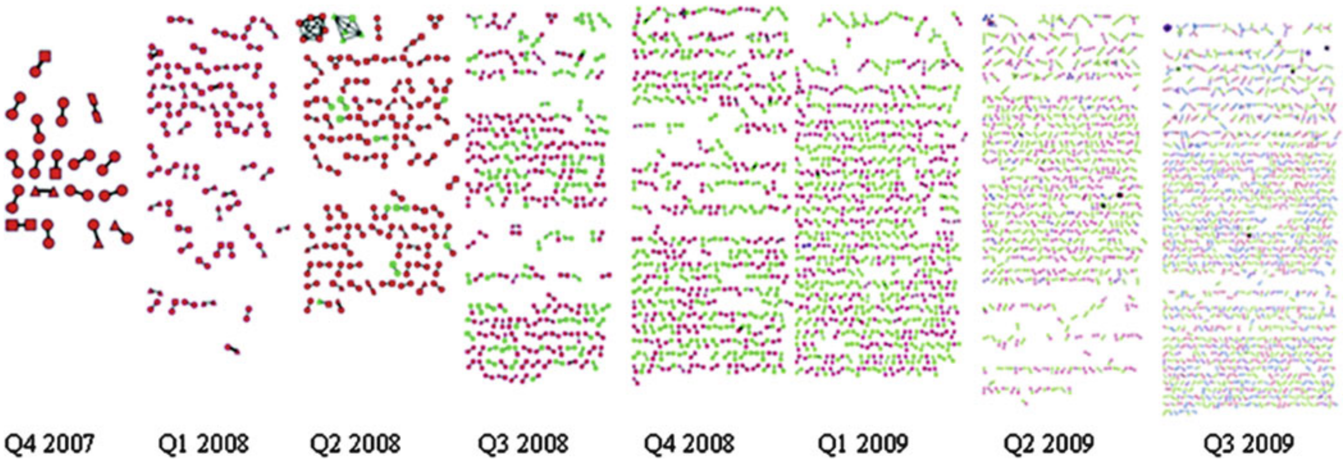

Fig. 9.3 Time evolution of Doro adoption network. One node represent one subscriber. Node color represents Doro Model: *red* = HandleEasy 326,328, *green* = HandleEasy 330, *blue* = PhoneEasy 410, *Purple* = Other Doro models. Node shape represents age of user: A *circle* means that user is older than 70 year. Link width represents weighted sum of SMS and Voice traffic

from Fig. 9.2, the monster in the WVAN actually diminishes in size over time—both in absolute number of users, and in the percentage of users in the LCC. Figure 9.4 illustrates this by showing two WVAN video-monsters which are 2 years apart.

We gain even more insight by looking at the time evolution of the VAN. Figure 9.5 shows the time evolution of the VAN-LCC. We see growth in the monster from Q3 2007 to Q4 2007, followed by a rather dramatic breaking down of the LCC after that time. Hence we see indications that the service itself had the potential to form a real social monster and take off, but some change in the service and user conditions killed that takeoff—in this case, we have found that a new pricing model was introduced.

9.3.4 Comparison of the Social Network Monsters Over Time

Figure 9.6 sums up much of what we have seen in the visualizations of the last subsections. The figure shows the fraction of adopters in various components of the adoption network. Subscribers in the blue area are adopters which have no connection to other adopters. These users (termed isolates here, and referred to as singletons in [21]) have not been visible in our visualizations. The users in the green area correspond to the adopters in the social network monster (there is in every case only one component with >1,000 users).

We first consider Fig. 9.6a, the figure describing iPhones. Here we see that the growth of the monster (green), as a percentage of the total number of users, has not been monotonic. The monster has however grown monotonically in the absolute number of users—see again Fig. 9.2. We conclude from this that the number of isolated subscribers grew more rapidly than did the core. This implies that some change in the offering has induced a large growth in the number of new users in this period (Q2 2008–Q3 2008). One candidate explanation is the appearance of 3G handsets in this time period. Another likely explanation for many new users is the fact that “legitimate” iPhones were first available on the Telenor net at this time.

Figure 9.6b (Doro) simply confirms the picture seen in Fig. 9.3: no monster, essentially no large LCCs. At the same time we see an enormous dominance of isolates. This is consistent with the hypothesis that Doro users are elderly (which we can confirm), and that they speak mostly with other generations, i.e., non-Doro users. Again, it suggests that the adoption of this phone is not based on network influence, but on more ego-based considerations.

Figure 9.6c shows the VAN, while Fig. 9.6d shows the WVAN for the video product. Again we confirm the qualitative picture obtained from Figs. 9.4 and 9.5: the WVAN-monster decays slowly, while the VAN-monster collapses. In the case of the video service, the collapse corresponds to the initiation of payment for the system. Using the lingo of the iPhone example, this would be the same as turning of the “buzz”. We also see in Fig. 9.6c a dominance of two-node components—not surprising for a transactional service—and a complete absence of isolates.

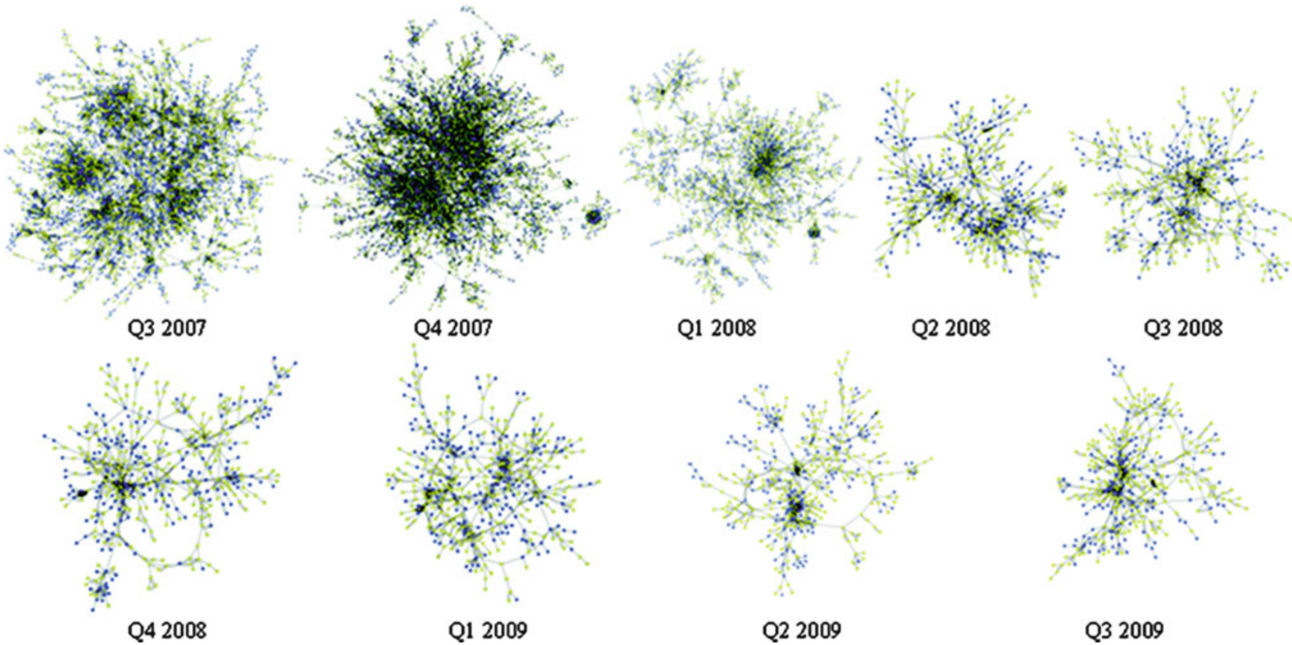

Fig. 9.4 Time evolution of the mobile video telephony adoption network. Links are real video links, where width represents duration of video conversations. Enterprise adopters have *blue* node color, while consumer adopters are tagged yellow

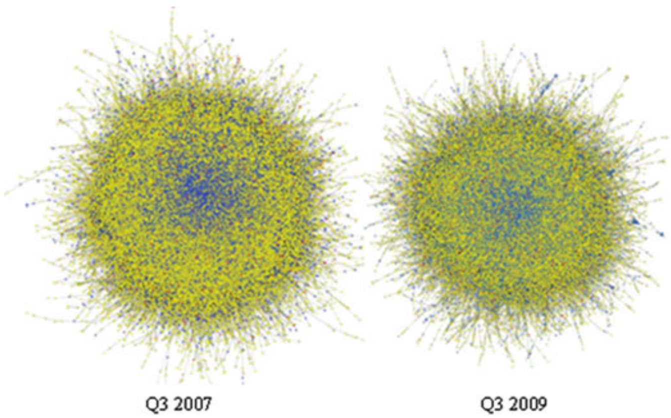

Fig. 9.5 Time evolution of mobile video telephony adoption network (WVAN-where the social links include all communication). Only two quarters are shown due to the fairly stable LCC. *Blue* node color represents enterprise subscribers, while *yellow* represents private subscribers

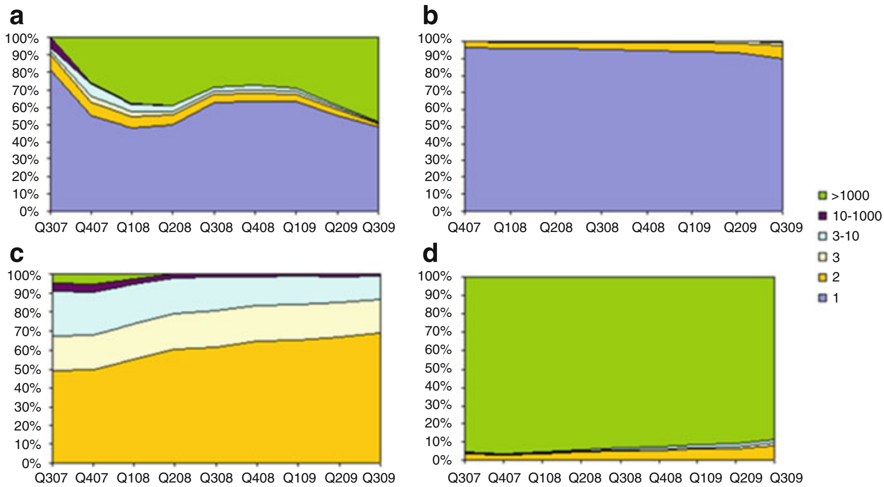

Fig. 9.6 Fraction of subscribers in components of various sizes in the adoption networks for: (**a**) iPhone. (**b**) Doro. (**c**) Video (video-links). (**d**) Video (W-links)

The latter result, while not surprising, is not in fact guaranteed (for WVAN) by our definitions: we will see two isolates in WVAN every time two subscribers use video transactions, but have no other (W) communication, and have no friends using video transactions. We see that this simply does not happen—primarily because every pair that uses video also uses voice, SMS, or both.

We offer some quantitative details illustrating the dynamics seen in Figs. 9.6c, d. We observe that 95.8 % of users are in the core of WVAN in Q3 2007, while only 5.7 % are in the VAN core. Two years later, the corresponding numbers are 88.7 % for WVAN and 0.76 % for VAN.

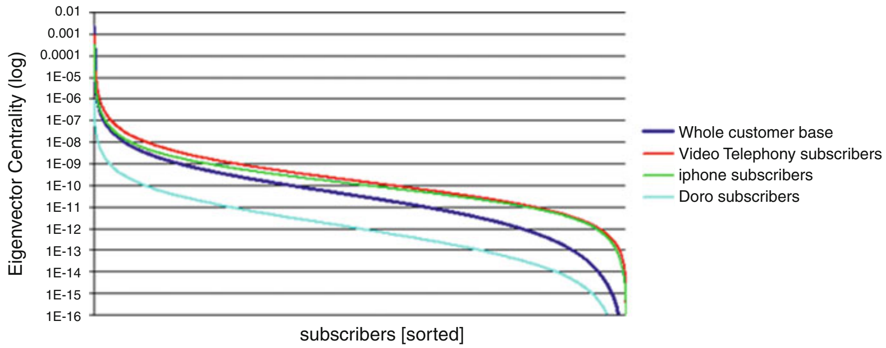

Fig. 9.7 Eigenvector centrality distributions for the whole customer base, Doro users, iPhone users, and video telephony users. Distributions are from Q3 2009. For ease of comparison, the x-scale is normalized so as to run from 0 to 100 % for all displayed distributions

9.4 Centrality of Adopters

We have seen how the iPhone and video adopters form a giant monster with respect to the W-links, while the Doro adopters do not. Motivated by this, we calculate the social centrality for all the adopters in each group, comparing it to the centrality of the whole customer base. Our expectation is that the involvement of highly central users is essential to the development of a large monster. To measure centrality we use the well-known eigenvector centrality (EVC). We believe high EVC will be strongly correlated with presence in the social monster, as it is already known to be correlated with strong spreading [7].

Figure 9.7 shows the EVC distributions for all the adopters. We see that video users are the most central, with iPhone just behind. We also see that the Doro adopters are rather peripheral socially, with their distribution falling well below that for the entire customer base. This supports our expectation that people in the giant monsters tend to be more central than the rest of the customer base. We believe that one may find, among these customers, the influential early adopters—those that adopt new products and services fairly early, and stimulate (or perhaps demand) others to do the same.

9.5 Kappa-Test

In all our results so far we see indirect evidence for social spreading effects (or their apparent absence, in the Doro case). As another test for social spreading, we introduce a simple statistical test, the kappa-test or κ-test.

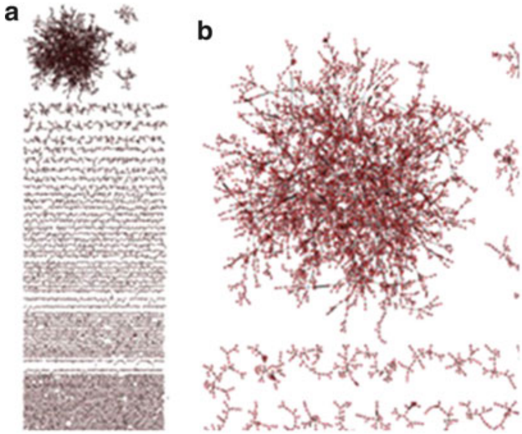

Fig. 9.8 (**a**) iPhone network from random reference model, Q1 2009. (**b**) LCC zoomed in

9.5.1 *Definition of the κ-Test*

We consider again the entire social network (as proxied by our communication graph) and define two types of links:

- A-links: links where neither, or only one, of the two connected nodes have adopted the product.
- B-links: links where the two connected nodes have both adopted the product.

We regard B-links to be the links which can indicate (but not confirm) social influence. We also recognize however that B-links can arise by other mechanisms, and even by chance. In order to evaluate the significance of the B-links that we observe in the empirical adoption data, then, we compare the empirical number of B-links (call it $n_{B,emp}$) with the number found by distributing at random the same number of adopters over the same social network, and then counting the resulting number of B-links $n_{B,rand}$.

We then define$\kappa \equiv n_{B,emp}/n_{B,rand}$. Clearly, if κgs significantly larger than 1, we have strong evidence for social spreading effects. More precisely, $\kappa > 1$ implies that people who communicate with each other tend to adopt together.

Figure 9.8 illustrates what happens when we scatter the iPhone adopters in Q1 2009 randomly over the empirical social network. The monster is still there, but it is smaller (by more than a factor 3) than the empirical monster seen in Fig. 9.2. Comparing the corresponding whole adoption networks of Figs. 9.2 and 9.8 by using the κ-test, we find that there are over twice as many links in the empirical adoption network compared to the random reference model—that is, κgs 2.18. We take this to be evidence that social spreading has occurred—more precisely, that people who talk together adopt together much more often than chance would predict.

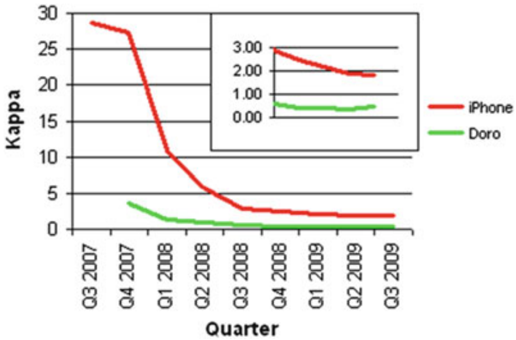

Fig. 9.9 Kappa for Doro and iPhone

Figure 9.8 illustrates an important point which gives insight into both monsters and social network structure. The point is that monsters arise even in the complete absence of social spreading effects. The monster seen in Fig. 9.8 is thus telling us something about the structure of the social network itself—that it has a "dense core" in which a dominating LCC arises even in the case of random adoption. At the same time, this dense core must include the set of users who give rise to the empirical monsters that we observe. The empirical LCC is simply larger than the random one (for the same number of adopters), due to social spreading (arising from mechanisms such as social influence or homophily).

Figure 9.9 shows the evolution of κ over time, both for the iPhone and for Doro. We notice that κ is very large in the early stages of product adoption (e.g., around 28.7 for the iPhone in Q3 2007). We find this to be typical: the first adopters are not randomly distributed, but rather tend to lie in a few small connected social groups. The large value for κ tells us that this observed distribution of the early adopters on a social network is extremely unlikely to have occurred by chance.

We notice also that κgs consistently less than 1 for Doro, after the early phase of adoption. While we argue that $\kappa > 1$ is evidence for social spreading effects, we do not believe that $\kappa < 1$ proves that such effects are not occurring. What $\kappa < 1$ does say is that adopting friends are found less often than a random model would predict. Our explanation for this is that the random model hits the dense core more often than the actual empirical adopters do. In other words, the empirical adopters are socially peripheral. This idea is in agreement with the EVC distribution seen in Fig. 9.6.

Finally we note that our κgest is not performed for the video adoption network in this chapter. The reason is that (as discussed in Sect. 9.3) the transactional nature of the video service constrains both VAN (exactly) and WVAN (empirically) such that there are no isolates. This constraint is not captured by the random reference model of the κgest. The purpose now is to introduce a link-based version of the node-based kappa-test.

9.6 Link Based Kappa-Test

The kappa-test defined in Sect. 9.5 considers products which are adopted nodewise, like e.g. the adoption of handset. In the case of transactional products like videotelephony, product adoption is defined in a pairwise manner: the nature of the product is *to activate links on the social network*. We consider videotelephony as a *link-based* product. The spreading process of a link-based product will be different compared to a node-based product, and hence the random spreading model in the kappa-test has to reflect this. In the node-based kappa-test the *nodes* are spread randomly and then the resulting links are compared to the empirical number of links between adopters. For link-based products, we suggest a random model where links in the social network are randomly activated.

We suggest the following for a link-based kappa-test: We start from the underlying social network, and track the time-evolution (adoption) of a link-based service; e.g. MMS or video telephony. In other words; a link-based service is a service in which the adoption requires both end-points of a link to adopt the service. Both nodes at the end-point of a link are mutually depending on each other in starting using the link-based service—communicating with each over the link. The link adoption of the network-based service, also gives rise to an adoption network.

Again, we are interested in measuring ‘adopting together’, so for links this is interpreted as meaning that two links have adopted together if they have a common node as end-point. In other words; two links are said to adopt together when the two links are adjacent. Figure 9.10 illustrates how we count adjacent links, two links A–C and A–D define an adjacent pair of links, since the node A is a common node for both links. Counting the number of open triangles in the adoption network then measures the occurrence of links adopting together. The random adoption network is generated through randomly choosing the links to adopt the link-based service from the set of possible links in the underlying social network, and the number of open triangles is then counted in the random case. We suggest then to define the link-based kappa test as;

$$\kappa = \frac{\#\, \textit{adjacent links in expirical network}}{\#\, \textit{adjacent links in random network}} \tag{9.1}$$

The above framework can now be used for generating kappa tests for comparing the empirical and random occurrence of links and adjacent links in the true adoption network and a randomly simulated adoption network, respectively. These adoption networks are considered relative to the underlying ‘true’ social network that can be measured. Summarizing, we have the following two kappa tests:

- Node based adoption, counting links: κg empirical number of links/random number of links
- Link based adoption, counting open triangles: κg empirical number of open triangles/random number of open triangles.

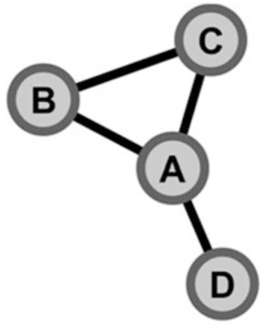

Fig. 9.10 Illustration of open and closed triangles. *A–C* and *A–D* are adjacent links since they share one common node and make an open triangle. *A–B*, *B–C* and *A–C* together make a closed triangle

We count the total number of open links in the network using the binomial coefficient and sum over all nodes i:

$$\begin{aligned}\sum_{i=1}^{n}\binom{k_i}{2} &= \sum_{i=1}^{n}\frac{k_2!}{2!(k_i-2)!}\\ &= \sum_{i=1}^{n}\frac{1}{2}k_i(k_i-1)\\ &= \frac{1}{2}\sum_{i=1}^{n}k_i^2 - \frac{1}{2}\sum_{i=1}^{n}k_i\\ &= \frac{1}{2}\sum_{i=1}^{n}k_i^2 - L\end{aligned} \tag{9.2}$$

where k_i is node degree, n is the total number of nodes, and L is the number of links—all measured relative to the adoption network.

From the above discussion, it is also possible to consider other types of kappa-tests. For example, one can envision kappa-tests that compares;

- The number of closed triangles to the number of random triangles.
- The number of closed loops of some order to the number of random closed loops.

The idea is to consider local network structures, and how these local structures potentially 'adopt' together.

9.6.1 Quarterly Development of the Link Kappa

From Fig. 9.11 we see that the link kappa value defined in (1) is steadily increasing from around 15 to about 20. A kappa value of 15 means that there is 15 times as many adjacent links in the empirical data than in the random reference case! We can then conclude that the distribution of video usage on the social network is far from

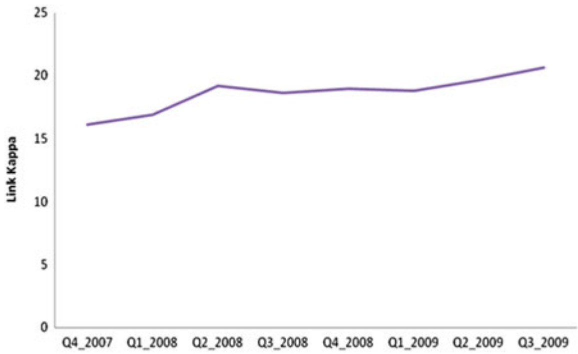

Fig. 9.11 Historical development of link kappa for video telephony. Link kappa is defined as the number of total adjacent links in the empirical data to the number of adjacent links in the random reference case. The random reference network is constructed by choosing an equal number of links randomly from the underlying social network and then count the number of adjacent links

a random process where social relations are randomly 'activated' to become video links.

The number of adjacent links is related to the node degree according to (2). Figure 9.12 shows the degree distribution for Q3 2009—both for the random 'link activation' model and for the real empirical video adoption network (VAN). As expected from the high link kappa value, the degree of the nodes in the empirical adoption graph is higher than in the corresponding random model (higher number of adopting neighbors means higher number of adjacent links). In the random case around 90 % of the nodes have degree 1—the least degree possible since the unit we use in this experiment is *links*. An isolated *link* will give both end nodes degree 1. This is expected, since the random model does not take into account that the videotelephony users, once they have started using the service, will use the service to communicate with more than one network neighbor. A more sophisticated random reference model could also take this into account. The simple random link activation model acts as basic reference and assigns the same activation probabilities to all links on the network. We reserve testing more advanced reference spreading models for future work.

Another way to measure clustering among adopters is to use the global clustering coefficient defined as total number of closed triplets divided by the total number of adjacent links. See e.g. reference [30] for a detailed discussion of the global clustering coefficient. Using the same 'kappa' logic as before, a completely random adoption process will also show some clustering. To compensate for this effect, we use the kappa test logic on the global clustering coefficient and define a new link kappa measure as C_{emp}/C_{rnd}, where C_{emp} is the clustering coefficient calculated from the empirical network and C_{rnd} is calculated from the random video adoption network. Figure 9.13 shows the development of the ratio C_{emp}/C_{rnd}.

In summary—measuring the social network dependence on video telephony is more difficult than measuring effect on node attributes like choice of handset. Video telephony is a link attribute involving choices of both end nodes. A simple

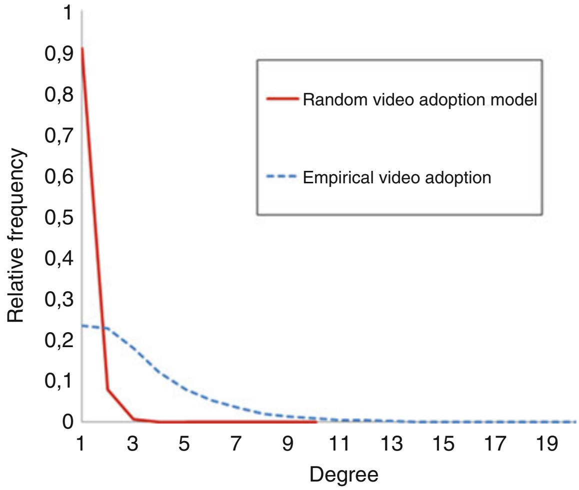

Fig. 9.12 Node degree distribution for the 'random link activation' reference model and the empirical video adoption network (VAN), Q3 2009

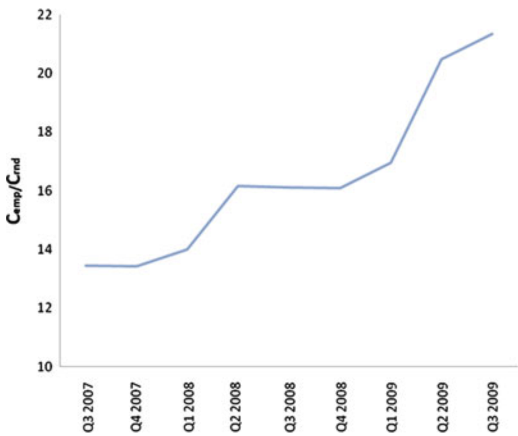

Fig. 9.13 Clustering coefficient kappa test defined as the empirical global clustering coefficient divided by the global clustering coefficient for the random network. The global clustering coefficient is defined as total number of closed triangles divided by the total number of adjacent links in the network

reshuffling of individual adopter status will not take this into account. Our simple link kappa test can be used for quantifying network dependent adoption for services of a transactional nature. In the case of video telephony we find that the empirical relations are clustered on the social network compared to a random link activation model.

9.7 Correlated Adoption Probability

Our final test for social spreading effects is to measure the probability p_k that a subscriber has adopted a product, given that k of the subscriber's friends have adopted the product. This conditional probability does not indicate causation, because it makes no reference to time order—it simply measures (again) how

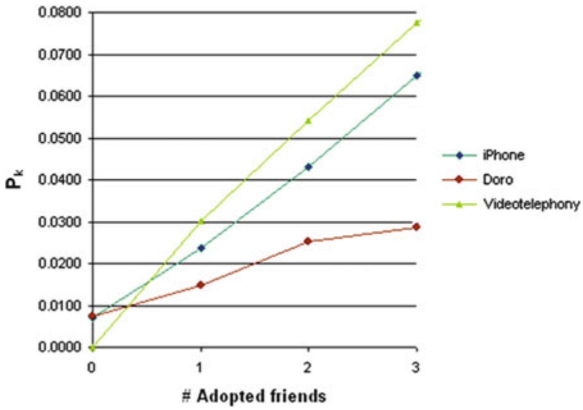

Fig. 9.14 Adoption probability pk vs. the number k of adopting friends, for three products. In each case we see a monotonic growth of pk with k-indicating that some kind of social spreading is occurring

strongly those that communicate together tend to adopt together. We measure p_k simply by first finding all subscribers with k adopting friends, and then finding the fraction of these that have themselves adopted.

Figure 9.14 shows p_k vs. k for the three products, for $0 \leq k \leq 3$. For higher k, the Doro data are too dominated by noise (a low n) to be useful. (The results for iPhone and video have better adoption profiles, and so are meaningful at least out to $k = 10$; but their qualitative behavior—monotonic increase, at roughly constant slope, with increasing k—is like that seen in Fig. 9.14.) Figure 9.14 supports our claim that there are some social spreading effects operating on Doro adoption—since we see a steady increase of p_k with k. Such effects are not visible from our κgest, for reasons given above; yet we see for example that, if we know that a subscriber has one friend using a Doro phone, that subscriber's probability of using one him- or herself is roughly *twice* the adoption probability for a subscriber with no adopting friends.

The iPhone and video p_k results lie, not surprisingly, considerably higher than the Doro curve. This is consistent with our claim that social spreading effects are much stronger for these products. For example, knowing that a person has one friend using an iPhone roughly triples the probability (compared to someone with no adopting friends) that this person also uses an iPhone. This comparison is however not possible to make for the video service, because the probability of using video telephony and having no (W-)friends who use video telephony is (as discussed above) empirically zero. Otherwise the video results are qualitatively the same as the iPhone results.

We also note that the p_k for video-telephony is consistent with the degree distribution for video adopters in Fig. 9.12, where we see that the nodes in the empirical VAN generally have a higher degree than expected from the simulated random link activation model. Since the adoption probability, p_k, increases with number of adopting neighbors, one will expect a higher number of activated neighbors in the adoption graph compared to a random model. A higher number of neighbors will again result in a higher number of adjacent links in the empirical data and thus the high link kappa value we see in Fig. 9.11.

9.8 Time Evolution of Computer Tablet Adoption Networks

Until now we have looked into the social interaction between users of different handsets, as well as video telephony. Recently, so called computer tablets, such as Apple iPad 3G and Samsung Galaxy Tab, have received increased attention. These tablet computers are particularly marketed as a platform for audio and visual media such as books, periodicals, movies, music and games, as well as web content. As a preview of current research we will show the social network of iPad 3G-adopters, and its time-evolution from its release date, as well as some early statistical results. We are not able to study uptake of WiFi-only tablets, since we depend on a SIM to map tablets to the social network. Computer tablet networks address some fundamental differences from the adoption networks already mentioned. An important difference from traditional handsets is that the SIM card placed in a tablet is not necessary the same SIM card which is being used for social interaction, It is typically a "twin SIM" solution, where a customer gets two SIMs on the same subscription. We use the SIM in the traditional handset as the node and map the social network. This will give us a picture of the social interaction among iPad 3G users. We will show that the 'non-pad' SIM cards are most often placed in an iPhone.

9.9 iPad 3G Adoption Network

As with the iPhone, the release date of iPad 3G varies with region—it was officially released in the US Apr 3rd 2010. In Telenor net it was released in November 2010. As the visualizations will show, thousands of users bought their iPad in the US and used it on Telenor net before they actually were officially released by Telenor in November 2010. Since such tablets are quite new in the given market, we have chosen to study the time evolution with a more fine-grained granularity—month by month.

Figure 9.15 shows the time evolution of the iPad 3G adoption network for the first 7 months. The structure reminds us of the iPhone evolution. We have here also included smaller social components, but for visualization purposes we have not included the isolates (iPad 3G-adopters with no iPad 3G friends), but we remind the reader that these are also part of the adoption network. From the visualization we see how the largest connected component gradually turns into a "monster" within a couple of months. When a specific iPad-user also uses an iPhone, the node is colored blue. Visually it is seen that the connected iPad-users have a high percentage of iPhone's. In general we find that 53.4 % of the users also are using iPhone for social interaction. We observe that in the cases where the iPad-user has at least one iPad-friend (adopters shown in Fig. 9.15) there are as many as 72.1 % also having an iPhone. The remaining users, which represent the disconnected users in the adoption network (isolates), there are only 38.6 % which have adopted an iPhone. This shows that iPhone is much more common phenomena in socially connected groups of iPad users. This also supports the idea of an Apple "tribe" of users.

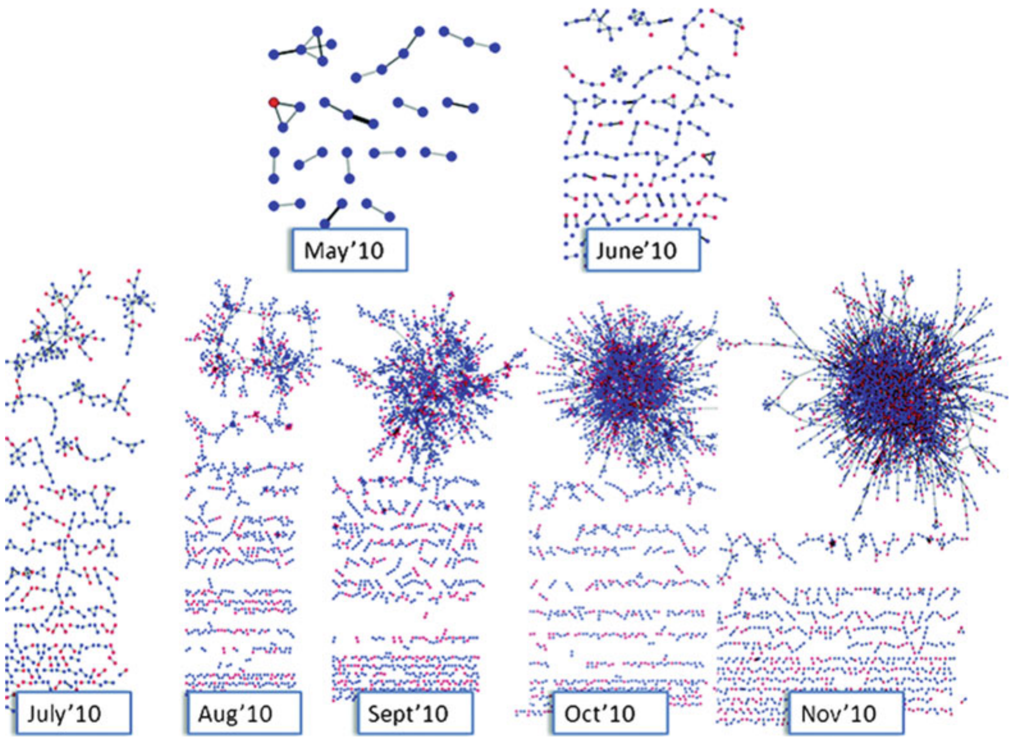

Fig. 9.15 Time evolution of iPad 3G adoption network. One node represent one subscriber. *Blue* node color represent an iPad 3G user which also uses an iPhone, while *red* color represent an iPad 3G user with another handset. Link width represents weighted sum of SMS and Voice traffic. Isolates are not included in the visualization

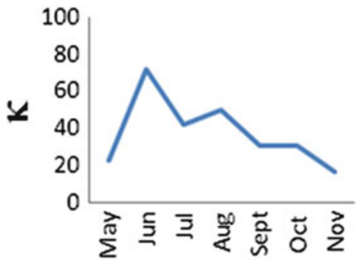

Fig. 9.16 Kappa for iPad 3G

9.10 Kappa-Test for iPad 3G

As stated in Sect. 9.5, $\kappa >$ xgmplies that people who communicate together also tend to adopt together. Figure 9.16 shows the time evolution of κgor iPad 3G. We notice that it is significantly above 1, which indicate strong evidence for social network effects. κgaries between 16.9 and 72, which means that at most there are 72 times more empirical iPad-relations compared to the relations we find when spreading the adopters randomly on the whole communication network. We notice that the numbers are significantly higher than for iPhone—and from the perspective of κsghis trend supports that the spreading is stronger compared to iPhone, which is also reflected through the number of sales [16].

The rapid increase of kappa from May to June can be explained by the time it takes from the first adopters got their pads until friends are influenced and are able to get one. It is a product of how the individuals got their iPads since it involved travel to the U.S. or knowing a friend who was traveling there.

The decrease from june can be due to increased exposure of the product in the media which can lead to less dependence on social spreading, in combination with the increased number of sales. Due to the nature of κsgt will approach 1 if we have complete saturation of iPads in the market.

9.11 iPad 3G Adoption Probability

Figure 9.17 corresponds to the correlation adoption probability which we measured in Sect. 9.6. We measure the probability p_k that a subscriber has adopted iPad 3G, given that k of the subscriber's friends have adopted iPad 3G. The results shows that if you have one iPad friend the probability to adopt iPad is 14 times higher than with zero friends. If you have two friends the probability is 41 times higher, three friends 96 times. We observe a steep monotonic growth which should indicate strong social spreading effects. The adoption is highly dependent on the number of friends.

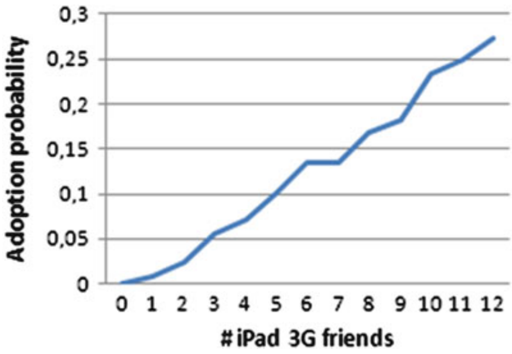

Fig. 9.17 Adoption probability pk vs. the number k of adopting friends, for iPad 3G. We observe a monotonic growth of pk with k-indicating that some kind of social spreading is occurring

9.12 Summary and Future Work

All of our results support a simple and fairly consistent interpretation:

- The iPhone has very strong social spreading, and has truly taken off
- The Doro handsets have only very weak social spreading. This device will probably never take off in the same sense as the iPhone.
- Video telephony use also has strong social effects, and started spreading very strongly; however its early takeoff was stopped by an external factor—here, a new price model.
- Preliminary results from studying the newly released iPad 3G reveals even stronger network-dependent adoption patterns than iPhone.

Standard whole-network measures, such as total number of users, or total traffic over time, can also give useful information on these same questions. We believe however that our measurement methods give new and useful insight into how and why these services have performed so differently.

We have also constructed two different link-based κ-tests for the VAN, because the VAN has a fundamental constraint—that adoption occurs only in pairs—that our simple κ-test does not capture. We conclude from these κ-tests that the remaining subscribers are well clustered due to the increase of κ over time—both for the test based on the clustering coefficient and κgased on adjacent links. These tests should be useful for any transactional graph with the pair constraint.

Also, we suspect that social spreading effects for the Doro handsets may be more visible at the two-hop level. A typical scenario might be an adult child would buy this type of phone for one or more of their elderly parents when, for example, use of a more traditional mobile phone becomes difficult because of the size of the display or keypad, or the complexity of the device. In such a scenario, it may well be there is little or no direct communication among the adopters, but a strong two-hop connection via the younger generation. We plan to test this idea in the near future.

Acknowledgements All of our visualizations were produced using the open-source visualization platform from cytoscape.org. We would also like to thank Dr. Ellen Altenborg at Telenor ASA, and professor Øystein D. Fjeldstad of the Norwegian School of Management, for fruitful discussions and inspiration related to network economy and network analysis.

References

1. Albert, R., Barabasi, A.-L.: Statistical mechanics of complex networks. Rev. Mod. Phys. **74**, 47 (2002)
2. Aral, S., Muchnik, L., Sundararajan, A.: Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. Proc. Natl. Acad. Sci. **106**(51), 21544–21549 (2009)
3. Bearman, P., Moody, J., Stovel, K.: Chains of affection: the structure of adolescent romantic and sexual networks. Am. J. Sociol. **110**(1), 44–91 (2004)
4. Birke, D., Swann, G.M.P.: Network effects and the choice of mobile phone operator. J Evol Econ **16**, 65–84 (2006)
5. Bollobas, B.: Random Graphs. Cambridge University Press, Cambridge/New York (2001)
6. Bollobas, B., Riordan, O.: Mathematical Results on Scale-Free Random Graphs, pp. 1–37. Wiley, Weinheim (2002)
7. Canright, G.S., Engø-Monsen, K.: Spreading on networks: a topographic view. Complexus **3**, 131–146 (2006). doi:10.1159/000094195
8. Choi, H., Kim, S.-H., Lee, J.: Role of network structure and network effects in diffusion of innovations. Ind. Mark. Manag. **39**, 170–177 (2010)
9. Dasgupta, K., Singh, R., Viswanathan, B., Chakraborty, D., Mukherjea, S., Nanavati, A.A., Joshi, A.: Social ties and their relevance to churn in mobile telecom networks. In: Proceedings of the 11th International Conference on Extending Database Technology: Advances in Database Technology, EDBT '08, Nantes, vol. 261, pp. 668–677. ACM, New York (2008)
10. Dorogovtsev, S., Mendes, J.: Evolution of networks: from biological nets to the internet and WWW. Oxford University Press, Oxford (2000)
11. Dorogovtsev, S., Mendes, J.: Evolution of networks. Adv. Phys. **51**, 1079–1187 (2002)
12. Fetterly, D., Manasse, M., Najork, M., Wiener, J.: A large-scale study of the evolution of web pages. Softw. Pract. Exp. **34**(2), 213–237 (2004)
13. Eagle, N., Pentland, A., Lazer, D., Alex, P.: Inferring friendshop network structure by using mobile phone data. Natl. Acad. Sci. **106**(36), 15274–15278 (2009)
14. Erdös, P., Rényi, A.: On random graphs I. Publ. Math. Debr. **6**, 290–297 (1959)
15. Hill, S., Provost, F., Volinsky, C.: Network-based marketing: identifying likely adopters via consumer networks. Stat. Sci. **21**(2), 256–276 (2006)
16. http://www.apple.com/pr/library/2010/06/22ipad.html
17. http://www.doro.com/
18. Kleinberg, J.: Complex networks and decentralized search algorithms. In: International Congress of Mathematicians. European Mathematical Society, Zürich (2006)
19. Kumar, R., Novak, J., Raghavan, P., Tomkins, A.: Structure and evolution of blogspace. Commun. ACM **47**(12), 35–39 (2004)
20. Kumar, R., Novak, J., Raghavan, P., Tomkins, A.: On the bursty evolution of blogspace. Worldw. Web J. **8**(2), 159–178 (2005)
21. Kumar, R., Novak, J., Tomkins, A.: Structure and evolution of online social networks. In: KDD, Philadelphia. ACM, New York (2006)
22. Leskovec, J., Faloutsos, J.K.C.: Graphs over time: densification laws, shrinking diameters, and possible explanations. In: 11th KDD, Chicago, pp. 177–187. ACM, New York (2005)
23. Nanavati, A.A., Gurumurthy, S., Das, G., Chakraborty, D., Dasgupta, K., Mukherjea, S., Joshi, A.: On the structural properties of massive telecom call graphs: findings and impli-

cations. In: Proceedings of the 15th ACM International Conference on Information and Knowledge Management, CIKM '06, Arlington, pp. 435–444. ACM, New York (2006)
24. Newman, M.: The structure and function of complex networks. SIAM Rev. **45**(2), 167–256, 2003
25. Ntoulas, A., Cho, J., Olston, C.: What's new on the web? The evolution of the web from a search engine perspective. In: 13th WWW, pp. 1–12. ACM, New York (2004)
26. Onnela, J.-P., Saramäki, J., Hyvönen, J., Szabó, G., Lazer, D., Kaski, K., Kertész, J., Barabási, A.-L.: Structure and tie strengths in mobile communication networks. Proc. Natl. Acad. Sci. U.S.A. **104**(18), 7332–7336, 2007
27. Strogatz, S.: Exploring complex networks. Nature **410**, 268–276 (2001)
28. Sundsøy, P., Bjelland, J., Canright, G., Engø-Monsen, K., Ling, R.: Product adoption networks and their growth in a large mobile phone network. IEEE Advanced in Social Network Analysis and Mining (ASONAM 2010), Odense. IEEE, Los Alamitos (2010)
29. Van den Bulte, C., Wuyts, S.: Social networks and marketing. Marketing Science Institue, Cambridge (2007)
30. Wasserman, S., Faust, K.: Social network analysis: methods and applications. Cambridge University Press, Cambridge/New York (1994)
31. Watts, D.J., Strogatz, S.H.: Collective dynamics of small-world networks. Nature **393**, 440–442 (1998)