Modeling User Behavior in P2P Live Video Streaming Systems through a Bayesian Network

Ihsan Ullah, Grégory Bonnet, Guillaume Doyen, and Dominique Gaïti

ERA/Institut Charles Delaunay – UMR 6279

Université de Technologie de Troyes

12 rue Marie Curie – 10000 TROYES – France

{`ihsan.ullah,bonnet,doyen,gaiti`}@utt.fr

Abstract. Live video streaming over a Peer-to-Peer (P2P) architecture is promising due to its scalability and ease of deployment. Nevertheless, P2P-based video streaming systems still face some challenges regarding their performance. These systems are in fact overlays of users who control peers. As peers depend upon each other for receiving the video stream, the user behavior has an impact over the performance of the system. We collect the user behavior studies over live video streaming systems and identify the impact of different user activities on the performance. Based on this information, we propose a Bayesian network that models a generic user behavior initially and then adapts itself to individuals through learning from observations. We validate our model through simulations.

Keywords: P2P, IPTV, user behavior, Bayesian networks.

1 Introduction

P2P-based live video streaming has become popular in recent years. Unlike IP multicast, it does not require any major change in the current network infrastructure. Moreover, it reduces the need of deploying streaming servers as compared to the Client/Server (C/S) architecture, which requires an increase in the number of servers as the number of users grows. P2P approach allows cooperating end-hosts (called peers or nodes) to self-organize into an overlay network. Peers in these networks share their computing and upload resources by caching and relaying the video content to each other. Currently, several P2P video streaming systems have been deployed on the Internet and they have attracted a large number of users. Nevertheless, these systems still suffer from some performance problems such as startup and playback delays [1,2].

P2P systems are in fact networks of users who control peers. Since peers depend upon each other, activities of users have a direct impact over the performance of these systems. For example, the departure of a peer disrupts the stream availability to dependant peers. Similarly, the low quality stream received by one peer is forwarded to its descendants. Therefore, due to the importance of the user behavior in P2P streaming systems, several intensive measurement studies

have been performed over them. Based on the measurements, some approaches have been proposed for performance improvements. But on one hand, these approaches consider some aspects of user behavior while ignoring others. On the other hand, measurement studies provide insights into an average user behavior that can be used for deducing a generalized model. However, a global model is not suitable for users having different preferences and interests. Therefore, an individually adaptable model is required.

In this paper, we collect the information provided in measurement studies. We identify metrics of a user behavior and their impact on one another. With the help of this information, we model the user behavior through a Bayesian network. Our network encodes the available knowledge about the user behavior and can learn an individual user behavior, as new observations become available to it. This network can be potentially used by video streaming service providers, network operators and peers in a P2P IPTV systems. It will help there in carrying out decisions that improve the performance.

The remainder of the paper is organized as follows. Section 2 summarizes the related work. In section 3, we analyze the measurement studies over video streaming systems for identifying the user behavior metrics, their impacting factors and relationships among them. In section 4, we propose a Bayesian network for modeling a user behavior and describe results of our preliminary experiments. Finally, section 5 draws conclusions and gives directions for the future work.

2 Related Work

Work on the user behavior in P2P live video streaming systems is mostly in the form of measurement studies. Moreover, some measurements have been extended with propositions for performance improvement.

Tang *et al.* [3,4] first study a C/S and a P2P video streaming system and observe that a user who has spent more time watching a channel is statistically willing to stay longer. Based on this observation, they propose an approach that enables a joining peer to choose a provider peer that is present since a longer time. Wang *et al.* [5], after justifying the importance of stable peers, propose a mechanism for their identification, based on their elapsed time in the current session. The identified stable peers are put in the backbone for minimizing the impact of churn. Nevertheless, both the above mentioned approaches consider the stability of a peer that is one metric of user behavior. Moreover, for the estimation of stability they take into account only its one impacting factor (session's elapsed time) while ignoring others such as streaming quality and channel popularity.

Liu *et al.* [6] first measure a P2P IPTV system and identify the impacting factors of users' session durations and bandwidth contribution ratios. They propose a mechanism for predicting the longevity of peers and their bandwidth contribution ratios. Since their measurement analyzes all users together, their model expects all users to behave in a similar way which is not the case.

Horovitz *et al.* [7,8] propose a machine learning approach based on Support Vector Machines (SVM) for detecting actively the load in the uplink of provider

peers. This approach considers only one metric of the user behavior (i.e. upload bandwidth).

A shortcoming of the above approaches is that they either do not consider all metrics of the user behavior and/or use static models which expect all users to have a similar behavior.

Concerning our adopted approach, Bayesian networks have been also used in other domains for anticipating user behavior. For example, Wenzhi *et al.* [9] use Bayesian networks for predicting the future behavior of a mobile user in an intelligent Location-Based Services (LBS) publish/subscribe system. Laskey *et al.* [10] present an approach based on Bayesian networks to detect threatening user behavior in an information system.

3 Analysis of User Activities in Live Video Streaming

Since a user is not interested in the underlying mechanism of video streaming, we collect the information provided in studies over P2P, C/S and telco-managed IPTV systems to understand a user behavior. These studies measure different metrics of user behavior and provide valuable insights. In this section, we present the metrics of a user behavior and their impact on one another according to the observations given in the measurement studies.

3.1 User Behavior Metrics and Their Impact on Each Other

After a synthesis of measurement studies, we collect components and impacting factors of the user behavior that are observed in these studies. They are time-of-day, user arrival/departure rates, session durations, user population, surfing (activity of channel browsing) probability, bandwidth contribution ratios, channel type, channel popularity, elapsed time, streaming quality (in terms of the buffer level), failure rate (the departure of a peer before starting the playback) and playback delay (the lag between the time of a packet generation at the source and the playback time at peers [11]). Based on the observations given in measurements, we discuss them altogether and identify causal relationships among them.

Time-of-day. High arrival/departures rates are critical in determining the scalability of the system. In case of high arrival rates, the newly joined peers must find the provider peers that can immediately supply the video stream to them. Similarly, higher departure rates can disrupt the stream availability to dependent peers. Jia *et al.* [12] observe that peers join and leave in groups indicating the start and end of a TV program. Hei *et al.* [1] observe high joining and leaving rates at peak times. Moreover, they find higher departure rates at the end of a program. Similarly, a study over telco-managed IPTV system [13] find that STBs (Set-Top-Boxes) turning on and off events occur in larger number during certain times of a day. All these observations reveal an impact of time-of-day on user arrival departure rates.

The session duration of a peer determines its stability which is important for the continuity of stream to consumer peers. Liu *et al.* [14] explore the time-of-day effect on session durations and they observe that peers joining in peak hours stay longer watching the same channel. Similarly, peers have longer session durations within evening leisure times [6]. On the other side, Veloso *et al.* [15] argue that time-of-day has no impact on the session duration while day-of-week has an impact which is in contrast to the results given by Liu *et al.* [6]. One reason may be the difference in the types of applications they study since Liu *et al.* [6] analyze P2P live video streaming system while Veloso *et al.* [15] study a C/S system which stream both live audio and video content.

Concerning user population, measurement studies [1,13,6,15,16,17,18,19] consistently show two peaks during a day one around noon time and another during early night. This clearly indicates an impact of time-of-day on the user population. Finally, according to Cha *et al.* [19] and Qiu *et al.* [13] the surfing probability is also impacted by the time-of-day. They observe a sharp increase in surfing probability after specific time periods because of commercials or end of certain programs.

Channel type. Channel type impacts surfing probability and departure rate. Cha *et al.* [19] observe that surfing probability changes with the type of the channel. For sports and news channels it is higher than other types. Moreover, Hei *et al.* [1] observe that users watching a movie channel possess the behavior of batch departures. However, they do not observe the same behavior for the users watching other TV programs.

Channel popularity. Channel popularity impacts session duration, user population, availability of neighbor peers and surfing probability. We discuss them one by one. Analysis studies [1,14,6] report longer session durations for popular channels and shorter ones for unpopular channels. Concerning user population Hei *et al.* [1] find that popular channels get more users than the unpopular ones. Moreover, while monitoring a peer, they observe that it faced difficulty in finding partner peers during watching a less popular TV channel. Since unpopular channels get less users, finding partner peers becomes difficult. Regarding surfing probability, Cha *et al.* [19] observe that users tend to remain connected when they join a popular channel hence reducing the surfing probability.

Streaming quality. Streaming quality impacts session duration and bandwidth contribution ratios. Liu *et al.* [6,14] find that a peer receiving a good buffer level initially tends to stay longer in that session. Similarly, a peer initially receiving a good quality (buffer level) produces better bandwidth ratios [6] through contributing more upload bandwidth.

Elapsed time. Elapsed time impacts the session duration. Studies [3,19] show that statistically, the time spent on watching a channel is positively correlated to the remaining online time on that channel.

Arrival/departure rates. Arrival/departure rates impact streaming quality and failure rate (the departure of a peer before it starts playback). According to

Liu *et al.*, the streaming quality is impacted by the flash crowd that downgrades at peak hours of the day [6]. Since streaming quality is measured in terms of the buffer level, its degradation means that peers will face longer startup delay leading to a high failure rate as observed by Li *et al.* [17]. The stated reasons by authors are the random partnership making algorithms and a high percentage of peers behind NAT and/or firewalls. Since random partnerships do not prefer one potential parent over another, during high arrival rates a peer can choose a parent peer that itself has joined the system recently and unable to relay the content to its child peers. Similarly, high departure rates impact the stream continuity to child peers.

User population. User population impacts playback delay and bandwidth contribution ratios. Jia *et al.* [12] observe that the average delay is correlated with the number of online peers. It remains lower with the smaller number of online peers and vice versa. Li *et al.* [20] observe a strong correlation between the average bandwidth contribution ratios and the size of the system. It may be due to the fact that in a larger community, peers get more chance to contribute than in a smaller one.

Fig. 1. An abstract causal graph of user activities in live video streaming systems

To provide an abstract view of the findings given in the measurement studies, we combine the related metrics into a group and depict the resulting abstract graph in Figure 1. The arrows from one group to another shows the impact of one group on another.

4 Modeling User Activities with a Bayesian Network

To model the user behavior in a P2P live video streaming system, we use Bayesian networks also called belief networks.

A Bayesian network [21] is a pair (G, P) , where $G = (V, E)$ is a directed acyclic graph of vertices V and edges E . Vertices/nodes represent random variables and directed edges show direct dependencies between variables. Let $U = \{X_1, \dots, X_n\}$ be a set of random variables, then P is a set of conditional probability functions $p(X_i | X_j, X_j \in \text{parents}(X_i))$. A directed link in G from X_i to X_j means X_i is the parent of X_j and X_j is the child of X_i . Child nodes are directly dependent on parent nodes. Variables having no parent are independent. In a Bayesian network each variable is conditionally independent from all its descendants given its parents. A variable can have several states, each of them with a probability value. Each node has a Conditional Probability Table (CPT) that relates it to its parent nodes. A CPT contains probability values for all combinations of the states of a variable and states of its parents. Initial probabilities are called prior probabilities, which are deduced from data or provided by an expert. With the availability of new observations, this prior knowledge can be updated to posterior knowledge. The set P defines a set of joint probability distribution over U as shown in (1).

$$P(U) = P(X_1, \dots, X_n) = \prod_{i=1}^n p(X_i | X_j, X_j \in \text{parents}(X_i)) \quad (1)$$

A Bayesian network can compute the probability of a variable X_i to be in state S_k given that X_j is in state S_l , where $S_k \in \text{states of } X_i$, $S_l \in \text{states of } X_j$ and $X_j \subseteq U \setminus X_i$.

Since measurement studies analyze user behavior generally for all users, an individual can behave differently. Using this information in a global model will expect all users to behave in a similar way. Therefore, we make use of Bayesian networks that enable to combine expert's knowledge (knowledge deduced from the measurements) with new observations (observations received from the system). Indeed, in the absence of any observation, our network is general. It updates itself for each user with new observations through learning.

Based on the synthesis of measurement studies, the Bayesian network we propose is depicted in Figure 2. It involves 12 nodes each of them representing a user behavior metric or an impacting factor. To be able to carry out a preliminary evaluation through a simplified model, we discretize each variable into two states. In Figure 2, the name of the node is given on top and its states are described below it with prior percent probability values. In case of nodes having parents, an average of all the given probabilities for each state is shown. Since, for the moment we do not have access to any real trace of a live video streaming system, we reported the prior probabilities from figures given in measurement studies. We explain the discretization and approximation of prior probabilities in the following.

For instance, channel popularity has two states namely popular and unpopular. To assign them the prior probabilities, we consider the 80–20 rule given in [19] that states that 80% of user requests come for the 20% most popular videos. Hence, a user connecting to a channel has 80% chance to connect to a popular channel. Similarly, user activities are higher from 12 noon to 00 midnight, therefore we divide

Fig. 2. Our proposed bayesian network for user behavior¹

the whole day into two parts namely $T1 = 12 - 00$ and $T2 = 1 - 11$ hours. We divide channel types into reality and fiction, where the former represents channels like news, music and sports while the later represents movies and serials. User behavior in reality type is different from fiction type where the surfing probability decreases [19]. Since no clear information is available about their distributions in measurement studies, we assign them a uniform distribution. Arrival rate has two discrete states (high and low) that is dependant of time-of-day. Since the probability of $T1$ is slightly higher than $T2$, high arrival rate has the same.

Session durations are discretized into two simple states namely stable and unstable. To decide the stability of a peer, we use the concept given by Wang *et al.* [5], where they term a peer as stable if it stays up to 40% of the observed period. In our Bayesian network, session duration has five parent nodes and we assign them prior probabilities approximated from the user behavior studies. The overall average is biased to unstable session durations which is consistent with the measurement studies. We set the states of departure rate and user population in the same way. For elapsed time, we consider the study [3], which states that about 50% of user sessions last for less than 200 seconds. So we put a threshold ' $Th = 200$ seconds' to decide the cut off between newly arrived peers and peers who have elapsed more time in the system. Since, we cannot get any information about the probabilities for the states of streaming quality, delay, neighbor peers and bandwidth contribution ratios, we choose approximated prior

¹ The screenshots are taken from Netica (<http://www.norsys.com/>)

probabilities for them. However, we do not consider the impact of NAT/Firewall on bandwidth contribution ratio.

The process of estimating probabilities for a variable is called inference. If the parents of a variable are given, probabilities of its values can be estimated. As an example, if the probability of user population ($UP = Large$) is required to be estimated, given its parents time-of-day ($ToD = T_1$) and channel popularity ($CP = Popular$), the following process is carried out.

$$\begin{aligned}
 &P(UP = Large | ToD = T_1, CP = Popular) \\
 &= \frac{P(UP = Large, ToD = T_1, CP = Popular)}{P(ToD = T_1, CP = Popular)} \\
 &= \frac{\text{Number of samples with } (UP = Large, ToD = T_1, CP = Popular)}{\text{Number of samples with } (ToD = T_1, CP = Popular)}
 \end{aligned}$$

4.1 Experimental Evaluation

To evaluate our proposed network, we use the Bayes Net Toolbox (BNT) [22]. We carry out the evaluation process for two cases. In the first case, we test the global model that we construct with prior conditional probabilities deduced from the studies and perform queries to observe certain states of some dependant variables. In the second case, we let the network learn its parameters from a data set representing an individual behavior and perform similar queries as in the first case.

In both cases we provide evidences to independent nodes and observe results returned by the network for dependant ones. Provided evidences for time-of-day consist in a representative day divided into two parts that are high usage time and low usage time. For channel popularity and type, we select a sample TV channel² where the programs schedule is given. Also a list of popular programs is given. We assign popularities to each hour according to this information. Moreover for channel type, we find only two hours as reality type while the rest as a fiction type. Concerning the elapsed time, we provide uniform inputs for its two states. The evidences provided to each of these four nodes are shown in Figure 3 (a).

We query our network for computing the states of dependant variables each time an evidence is provided to independent ones as discussed above. Some of the obtained results are shown in Figure 3 (b). Starting from the top-left plot, we describe them one by one. $P(AR = High)$ gives a two valued result, showing a low probability of high arrival rate in low usage hours that clearly indicates its dependency on time-of-day. Similarly, $P(DR = High)$ changes according to the values of its parent variables (time-of-day and type of program). The dependance of user

² <http://www.mtv.com/>

(a)

(b)

(c)

AR: Arrival rate; DR: Departure rate; UP: User population; SQ: Streaming quality; SD: Session duration; BCR: Bandwidth contribution ratio.

Fig. 3. Evaluation of our network (a) Evidences given to the network ; (b) Results returned by the global network ; (c) Results returned by the network after learning

population on time-of-day and channel popularity is clear from $P(UP = Large)$, which changes accordingly. Results regarding streaming quality are interesting. Since it is negatively correlated with the arrival and departure rates which are high during peak usage hours, results suggest that users will get low quality during these hours. It shows the importance of handling churn to minimize its effect over the performance. Concerning stability of peers, we consider $P(SD = Stable)$ equal or greater than 0.5 as a stable peer. Results computed by our proposed network are consistent with the study of Wang *et al.* [5], where they found 5.5% to 15.5% stable peers during different analyses. Bandwidth contribution ratios are relatively smoother because they depend upon two contrasting variables (i.e. user population and streaming quality).

To evaluate the network for learning an individual user behavior, we create a training data set from the information given in the measurement studies. It contains 10 samples while there is a possibility of 4096 samples. We let the network learn its parameters from the given data set. After learning we perform the same process we did previously i.e. providing the evidence variables and observing the investigated ones. We provide the same set of evidence variables we did in the previous example. The results we get are depicted in Figure 3 (c). A comparison with Figure 3 (b) indicates that the departure rates are less impacted by time-of-day. Similarly, user population is less impacted by channel popularity. We can notice similar differences in streaming quality, peer stability and bandwidth contribution ratios. It shows that the network can learn a user behavior from limited observations and can adapt to individual behaviors.

4.2 Applications

Our Bayesian network integrates all identified user behavior metrics in such a way that through inference an unobserved metric can be estimated from other fully or partially observed metrics. For example, user population can be estimated from time-of-day and/or channel popularity. Conversely, channel popularity can be inferred from user population and/or session duration. This enables our Bayesian network to be used by video streaming providers, Content Delivery Networks (CDNs), telco-managed IPTV and P2P IPTV systems.

A video streaming provider can use this network for anticipating high arrival rates and user population that can help to manage their resources efficiently. Similarly, a CDN providing live video streams can anticipate the increase in number of users in a particular region and can proactively plan the assignment of servers accordingly. Some proposition exist that state that peering Set-Top-Boxes (STBs) for providing the rewinding facility will help to reduce the load on the network [23]. In that scenario, our approach can help in predicting the on-time of an STB.

P2P IPTV systems can use it both on the source side and users side. On one hand, the source can decide to accept a child peer which is stable and contributes good bandwidth ratios. On the other hand, its integration into peers can help them react proactively and carry out decisions that improve the overall performance. For example, a peer while choosing a stream provider can consider the

estimated bandwidth contribution ratios and streaming quality of the potential providers. Moreover, churn can be managed in an improved way.

5 Conclusion and Future Work

User behavior in P2P live video streaming systems plays a critical role in the performance of these systems. That is why it is extensively studied in recent years. Nevertheless, the proposed approaches are static and model all users in a similar way. In fact, individuals behave differently according to their interests and habits. Therefore, in this paper, we analyzed the measurement studies on video streaming systems and modeled the relationships among different metrics of the user behavior. After getting a generalized view of causal relationships of various user activities, we depicted them in an abstract causal graph. Based on the information given in the measurement studies, we proposed a Bayesian network that models the user behavior generally and then learns from observations to adapt itself to individual users. Our Bayesian network allows to infer an unobserved metric from other metrics, even if some of the other metrics are unobserved. This model can be used by video streaming service providers (centralized, CDN-based or telco-managed-IPTV) for better managing their resources and can also be integrated into P2P IPTV systems. It will enable peers to carry out useful decisions for performance improvement. We also validate our model through simulations.

Currently, we are working on extending our model to multinomial variables for better estimations and predictions of user behavior metrics. We will also integrate it into a P2P network for enabling peers to make decisions based on the estimations of this model. Possibilities we are looking into are, whether to have one Bayesian network that is managed in a centralized way or to integrate one in each peer. Similarly, we are investigating types of decisions that will be carried out by the stream provider peers and stream consumer ones.

References

1. Hei, X., Liang, C., Liang, J., Liu, Y., Ross, K.W.: A measurement study of a large-scale P2P IPTV system. *IEEE Transactions on Multimedia* 9(8), 1672–1687 (2007)
2. Hei, X., Liu, Y., Ross, K.W.: Inferring network-wide quality in P2P live streaming systems. *IEEE Journal on Selected Areas in Communications (JSAC)* 25(9), 1640–1654 (2007)
3. Tang, Y., Sun, L., Luo, J.G., Zhong, Y.: Characterizing user behavior to improve quality of streaming service over P2P networks. In: Zhuang, Y.-t., Yang, S.-Q., Rui, Y., He, Q. (eds.) *PCM 2006. LNCS*, vol. 4261, pp. 175–184. Springer, Heidelberg (2006)
4. Tang, Y., Sun, L., Luo, J.G., Yang, S.Q., Zhong, Y.: Improving quality of live streaming service over P2P networks with user behavior model. In: Cham, T.-J., Cai, J., Dorai, C., Rajan, D., Chua, T.-S., Chia, L.-T. (eds.) *MMM 2007. LNCS*, vol. 4352, pp. 333–342. Springer, Heidelberg (2006)

5. Wang, F., Liu, J., Xiong, Y.: Stable peers: Existence, importance, and application in Peer-to-Peer live video streaming. In: IEEE INFOCOM, pp. 1364–1372 (2008)
6. Liu, Z., Wu, C., Li, B., Zhao, S.: Distilling superior peers in large-scale P2P streaming systems. In: IEEE INFOCOM, pp. 82–90 (2009)
7. Horovitz, S., Dolev, D.: Maxstream: Stabilizing P2P streaming by active prediction of behavior patterns (2009)
8. Horovitz, S., Dolev, D.: Collabrium: Active traffic pattern prediction for boosting P2P collaboration. In: IEEE International Workshops on Enabling Technologies, pp. 116–121 (2009)
9. Wenzhi, C., Liubai, Zhenzhu, F.: Bayesian network based behavior prediction model for intelligent location based services. In: 2nd IEEE/ASME International Conference on Mechatronic and Embedded Systems and Applications, pp. 1–6 (2006)
10. Laskey, K., Alghamdi, G., Wang, X., Barbara, D., Shackelford, T., Wright, E., Fitzgerald, J.: Detecting threatening behavior using Bayesian networks. In: Conference on Behavioral Representation in Modeling and Simulation (2004)
11. Dán, G., Fodor, V.: An analytical study of low delay multi-tree-based overlay multicast. In: Workshop on Peer-to-Peer streaming and IP-TV, pp. 352–357 (2007)
12. Jia, J., Li, C., Chen, C.: Characterizing PPStream across internet. In: IFIP International Conference on Network and Parallel Computing Workshops, pp. 413–418 (2007)
13. Qiu, T., Ge, Z., Lee, S., Wang, J., Xu, J., Zhao, Q.: Modeling user activities in a large IPTV system. In: 9th ACM SIGCOMM conference on Internet measurement, pp. 430–441 (2009)
14. Liu, Z., Wu, C., Li, B., Zhao, S.: Why are peers less stable in unpopular P2P streaming channels? In: Fratta, L., Schulzrinne, H., Takahashi, Y., Spaniol, O. (eds.) IFIP-TC 6. LNCS, vol. 5550, pp. 274–286. Springer, Heidelberg (2009)
15. Veloso, E., Almeida, V., Meira, W.J., Bestavros, A., Jin, S.: A hierarchical characterization of a live streaming media workload. IEEE/ACM Transactions on Networking 14(1), 133–146 (2006)
16. Vu, L., Gupta, I., Liang, J., Nahrstedt, K.: Measurement and modeling of a large-scale overlay for multimedia streaming. In: International Conference on Heterogeneous Networking for Quality, Reliability, Security and Robustness, pp. 1–7 (2007)
17. Li, B., Xie, S., Keung, G.Y., Liu, J., Stoica, I., Zhang, H., Zhang, X.: An empirical study of the coolstreaming+ system. IEEE Journal on Selected Areas in Communications 25(9), 1627–1639 (2007)
18. Xie, S., Keung, G.Y., Li, B.: A measurement of a large-scale Peer-to-Peer live video streaming system. In: International Conference on Parallel Processing Workshops, pp. 57–62 (2007)
19. Cha, M., Rodriguez, P., Crowcroft, J., Moon, S., Amatriain, X.: Watching television over an IP network. In: 8th ACM SIGCOMM conference on Internet measurement, pp. 71–84 (2008)
20. Li, B., Xie, S., Qu, Y., Keung, G.Y., Lin, C., Liu, J., Zhang, X.: Inside the new Coolstreaming: Principles, measurements and performance implications. In: IEEE INFOCOM, pp. 1031–1039 (2008)
21. Jensen, F.V.: Bayesian Networks and Decision Graphs (2001)
22. Murphy, K.P.: The Bayes Net Toolbox for MATLAB. In: Computing Science and Statistics, vol. 33 (2001)
23. Cha, M., Rodriguez, P., Moon, S., Crowcroft, J.: On next-generation telco-managed P2P TV architectures. In: IPTPS (2008)