COVER PAGE

How Helpful Can Social Network Friends Be in Peer-to-Peer Video Distribution?

Maria Luisa Merani
Department of Information Engineering
University of Modena and Reggio Emilia
Via Vignolese, 905 – 41125 Modena, Italy

Abstract

Today two keywords more and more frequently recur over the Internet: Peer-to-Peer (P2P) and social networks. P2P, in all of its different declinations, represents a widely-adopted approach for content distribution, particularly for video diffusion. In parallel, the proliferation of social networks is an analogously stunning phenomenon, of unprecedented popularity and scope.

In the present work we examine a mesh-based P2P overlay, specifically designed for video streaming, and put forth some modifications to the neighborhood creation and chunk scheduling algorithm the platform adopts, with the goal of favoring peers belonging to a social network and granting them better performance. The improvements that such modifications attain are measured in terms of delivery ratio (throughput) and playback delay. We find that it is possible to guarantee a clear service differentiation, so that social network peers experience an improved viewing experience at the expense of ordinary overlay members, and that the scheduling mechanism modifications warrant the more consistent gains; we also show the role that different percentages of peers belonging to the social network have on the considered metrics. We finally suggest that the attained differentiated service level can be leveraged as an incentive to convince peers of the video overlay to join the social network.

Keywords - Peer-to-peer, Video Streaming, Online Social Networks, Differentiated Service Levels

How Helpful Can Social Network Friends Be in Peer-to-Peer Video Distribution?

Abstract— **Today two keywords more and more frequently recur over the Internet: Peer-to-Peer (P2P) and social networks. P2P, in all of its different declinations, represents a widely-adopted approach for content distribution, particularly for video diffusion. In parallel, the proliferation of social networks is an analogously stunning phenomenon, of unprecedented popularity and scope.**

In the present work we examine a mesh-based P2P overlay, specifically designed for video streaming, and put forth some modifications to the neighborhood creation and chunk scheduling algorithm the platform adopts, with the goal of favoring peers belonging to a social network and granting them better performance. The improvements that such modifications attain are measured in terms of delivery ratio (throughput) and playback delay. We find that it is possible to guarantee a clear service differentiation, so that social network peers experience an improved viewing experience at the expense of ordinary overlay members, and that the scheduling mechanism modifications warrant the more consistent gains; we also show the role that different percentages of peers belonging to the social network have on the considered metrics. We finally suggest that the attained differentiated service level can be leveraged as an incentive to convince peers of the video overlay to join the social network.

Keywords – Peer-to-peer, Video Streaming, Online Social Networks, Differentiated Service Levels

I. Introduction

Social networks are a pervasive phenomenon of the last decade, along with the increased adoption of P2P architectures to share and distribute heterogeneous contents, ranging from posts to television channels.

In this work we imagine the following setting: a P2P streaming platform that distributes within its overlay a video, originating from a server with limited upload capacity. As in [1] and in similar architectures, video diffusion relies upon a pull-based approach, where peers periodically advertise the video chunks they own and conversely request the ones they are missing. Among the peers, some are also members of a social network, a parallel world where long-term relationships exist, and strong bonds tie direct friends.

We wonder how the P2P system can leverage the social links, and accordingly propose some structural modifications to the overlay construction: the aim is to favor social network users, boosting their throughput and reducing the delay they experience in downloading the video. To attain these goals, we begin increasing the neighborhood of peers belonging to the social community, placing a significant number of their friends in it. Next, we alter the scheduling algorithm that rules the distribution of the missing video chunks: to speed up the download of social nodes, every social peer stores the chunks that it has been requested in two separate queues, one for the chunks solicited by social mates, one for those being asked by ordinary peers. Social network requests are served first, and only when the corresponding queue is empty, can the needs of non-social peers be satisfied, therefore implementing a priority queueing type of service. The rationale behind both choices is that peers are expected to be more willing to help their friends, rather than perfect strangers.

We find that both proposals are successful in achieving the desired outcomes: indeed, social network users experience a higher delivery ratio and a lower playback delay than ordinary peers; the gap in performance becomes wider when both modifications are introduced within the overlay and the more the overlay operating conditions depart from the safe region where there is excess upload capacity to flawlessly serve the download requests of all peers.

We finally investigate how system behavior varies when the percentage of peers also belonging to the social network is modified, concluding that this parameter is not particularly relevant when the extended neighborhood only is introduced, whereas its role is significant once the priority scheduling algorithm is implemented.

The remainder of the paper is organized as follows. Section II frames the current contribution among the existing literature; Section III illustrates the most salient features of the examined P2P platform and the modifications we introduce to it; Section IV describes the model adopted to replicate the social network; Section V critically presents the achieved results and Section VI reports our concluding remarks.

II. Related Work

Social networks and P2P platforms have evolved independently, and only recently a few proposals have arisen that aim at merging some of their features, taking advantage of the strengths of both. One of the first, milestone steps in this direction is represented by TRIBLER, the social-based P2P system proposed by Pouwelse et al. in [2], where long-term social relationships act as the base layer of a P2P system, not only for content discovery or recommendation, but also to improve download performance. The improvement is achieved thanks to the cooperative downloading implemented by the users that join the same social groups, where members who trust each other cooperate. Along a parallel path, [3] suggests that social networks will be the next main application field for the P2P paradigm: through their prototype, the authors show that a P2P-based social network is feasible and testify that the distributed approach is indeed profitable. More recently,

[4] examines the feasibility of social network based content delivery in BitTorrent, relying upon Twitter initialized/shared torrents and demonstrating that the improvement are meaningful even with when the set of socially active peers is small.

However, to the authors' knowledge, there is still a limited number of contributions that merge the social network friendship concept with P2P streaming. The concise work of [5], indicating how social communities can be used to build new incentive mechanisms, represents a first attempt in this direction; also our work in [6] proposes a delivery mechanism that privileges social network peers whenever critical conditions start building up within the P2P overlay, i.e., when the overall bandwidth is scarce. In detail, when the system operates in such conditions, a peer belonging to the social network requesting video contents and not finding any, asks for the help of friends (and possibly friends of friends), that discard non-social peers currently served to make room for their mate.

In the current work we take a more systematic approach, modifying both the neighborhood creation and the video chunk scheduling algorithm of social network members, with the intent of guaranteeing them better performance regardless of the overlay condition: we demonstrate that a distinct service differentiation is achieved, and that social network peers become truly privileged video overlay members.

III. The examined P2P protocol and the proposed modifications

The system we take under consideration is a mesh overlay, whose characteristics closely recall GridMedia, a P2P TV system that experienced a wide popularity in China around year 2006 [7], [8]. It is a pull-based streaming protocol, where the video stream is divided into chunks of appropriate size (1250 bytes), sequentially numbered and spread within the overlay with the help of all peers. New peers entering the system initially rely upon a tracker to obtain a list of potential parent nodes, that they start contacting to build their neighborhood and to receive video chunks. Peers within the neighborhood periodically exchange buffer maps, i.e., specific control packets where the sending node indicates the sequence numbers of the chunks it currently owns within its buffer; likewise, each peer periodically asks its neighbors the missing packets within its current request window. Typically, the request window size is lower than the size of the buffer, and the sending rate for the missing packets requests is higher than the buffer map sending rate. As regards the scheduling of the requests for the missing video packets, we implemented the "rarest first" strategy: a peer first requests the rarest chunks, i.e., those that less frequently appear among the buffer maps it has collected from its neighbors.

In accordance with existing literature, we define the resource index σ of the overlay as the sum of the capacities of all peers and of the streaming server divided by the number of peers and further normalized by the streaming rate. When $\sigma > 1$ the system is said to be underloaded, meaning that there is – in principle – enough overall bandwidth to support all peers' requests; on the contrary, $\sigma < 1$ indicates a condition where not all peers can receive the video at the maximum streaming rate.

In the original platform, all nodes exhibit the same neighborhood maximum size; instead, in our first proposal nodes that are also members of the social network enjoy a privileged condition, as their neighborhood is extended by an additional set of nodes that exclusively lie within their community friends. In our second proposal, we grant peers belonging to the social network a further advantage, modifying the scheduling algorithm of the missing packets at the receiver side: in detail, whenever a social network peer receives packet requests from other nodes within its neighborhood, it piles them up in two separate queues, and first serves packets requests coming from its friends; only when the corresponding queue is empty, it will start serving packet requests of ordinary peers.

The metrics we take into consideration to demonstrate how such proposals influence the achievable performance of both peers belonging to the social network and ordinary peers are the delivery ratio and the playback delay: the first is defined as the single peer received streaming rate normalized to the original video streaming rate, the second indicates the interval elapsed from the time a video chunk is sent out of the video server until it reaches the peer, both averaged over the streaming session duration. Such metrics are further averaged with respect to the entire population of social and non-social peers.

IV. The Social Network Model

In order to recreate the graph of the social network, where nodes represent the community members and edges the relationships between them, we resorted to the network evolution model proposed in [9]. Other approaches could have been possible (see e.g. [10] and [11]), but in this work we are interested in scouting if there is a performance gain, rather than assessing to what extent the amount of advantage differs in several communities.

Briefly, the model in [9] that we adopted relies on three distinct processes: (i) the process ruling the arrival of new users joining the social network; (ii) the process of initiating new edges within the social network; (iii) the edge destination selection process, indicating what node will be the destination of a newly created edge. As our interest lies in videos, we exclusively concentrated on the data reported in the above work that refer to Flickr, a social network born to share pictures, later extended to allow the sharing of videos [12].

Given $N(t)$ indicates the number of nodes within the social network at a generic time t of its evolution, its behavior is exponential, with $N(t) = e^{0.25t}$. The lifetime of a node is exponentially distributed with parameter $\lambda = 0.0092\ s^{-1}$. During its lifetime, a node adds a new edge every δ time steps, where δ is a temporal gap described by a truncated exponential probability density function [9]. Finally, the way adopted to choose the destination edge is the *random-random triangle closing model*: during the active lifetime of a node, when δ expires, the originating node s randomly picks one of its neighbors, say u, then it randomly picks one of u's neighboring nodes, say v, and a new edge is initiated between the origin node s and v, therefore closing – at random – the (s, u, v) triangle.

We actually had to introduce a few amendments to the model, to have it satisfyingly fit the experimental data.

The first modification dealt with $N(t)$ behavior: data show that Flickr would count up to approximately 580.000 nodes in a relatively short time span of observation, 25 months. However, it is impossible to reach such significant size exclusively relying upon the proposed $N(t)$ behavior. We therefore rapidly and forcedly increased the number of social network members at the beginning of the simulation, to later converge to the steady trend described by the $N(t)$ function. Only in this manner the artificial data would fit the experiments, as demonstrated in [6].

The exponentially distributed lifetime of the nodes has also been modified: once more, experimental data highlight that there is a relatively high probability that a node, once invited to join Flickr, creates very few edges and then never returns, a condition that the exponential distribution cannot capture. We believe this difference is non negligible and according to the data set reported in [9], we put forth a possible solution to this problem, modifying the exponential density function into a mixed one: the latter displays a a significantly non-null probability that a newly invited node has a null lifetime ($p(0) = 0.55$), a probability $p(1) = 0.05$ that the node lifetime is 1 day and a probability $p(2) = 0.05$ that the node lifetime is 2 days. [6] indicated that with this choice the fitting of the recreated data to the real ones is very satisfying and much closer than that obtained imposing the choices of [9].

Last, the truncated exponential density employed to determine the rate at which a node updates its contacts has also been amended. When employing the setting suggested by the authors of [9], we found that the values of the generated time gaps δ are much larger than in reality: the longer the node time gap, the lower the number of edges the node creates, hence the lower its degree, i.e., the number of edges incident to the vertex representing the node in the graph. After several tests, β was set equal to 0.02, a choice that guarantees a slope for the network degree distribution of -1.76, very close to the real measured value, -1.74 [6]. Thorough details about all the amendments we introduced and the improved fitting we achieved are reported in [6].

V. Numerical results

The performance gains guaranteed by the techniques we have devised are assessed over a simulative replica of GridMedia that exhibits the following main features: the P2P overlay members are 530; the size of the conventional neighborhood is 15, whereas the maximum number of additional neighbors that social network peers can select among their friends is 40. We investigate the behavior of a static overlay, where the nodes, once created, remain within the overlay for the entire duration of the simulation. Moreover, their upload and download capacities are distributed in accordance with the values and the percentages that Table I summarizes; the video server upload capacity is 600 kbit/s. Peers send out their buffer maps every 1 s, and request missing video chunks every 400 ms; their buffer is 1 minute wide and their request window size is 20 s. Overall, we simulate 8 streaming sessions, with a duration of 7 minutes each.

Network Access Type	Upload	Download	Percentage
High speed peers	1.5 Mbit/s	3 Mbit/s	5%
Medium speed peers	768 Mbit/s	1.5 Mbit/s	10%
Low speed peers	384 kbit/s	768 kbit/s	85%

TABLE I

Distribution of the upload and download bandwidths among P2P nodes

As regards the social community, in accordance with the model that we briefly recalled in Section II, we generate a social network graph with a size equal to 10^6 nodes (intentionally, a far greater value than the social network size). In the first set of simulations presented below, the probability that a peer of the streaming platform is also a member of the social network is taken equal to $p = 0.4$.

A. Average Packet Delivery Ratio and Delay with an Extended Neighborhood

The first results we discuss are obtained when the only difference between peers that are social network members and those that are not resides in the extended neighborhood of the former ones. Figs.1 (a) and (b) report the average packet delivery ratio and the average playback delay both categories of peers attain, as a function of the streaming rate of the video that has to be delivered, ranging from 400 to 600 kbit/s; in parallel, Fig.1 (c) reports the value of the resource index σ, indicating that the system is overloaded ($\sigma < 1$) when the streaming rate is higher than approximately 480 kbit/s.

These figures indicate that the first proposed modification already attains a noticeable performance improvement: simply relying upon a larger neighborhood, social network peers experience a higher delivery ratio and a lower playback delay than ordinary peers, both in the underloaded ($\sigma > 1$) and overloaded ($\sigma < 1$) conditions: on the contrary, non-social peers are penalized with respect to the conventional operating conditions where no social relationships are exploited. As a limiting example, when the streaming rate is equal to 600 kbit/s, their average playback delay raises to 54 s, to be compared against the 39 s of social network nodes. The 95% t-Student confidence interval are also computed and shown in these figures, indicating that the average values are definitely valuable indicators.

B. Priority Scheduling Effects

The same set of evaluations are performed when, in addition to the extended neighborhood, the priority scheduling approach is introduced in the overlay management. The corresponding results, reported in Figs.2 (a) and (b), highlight that this further modification guarantees a remarkable throughput

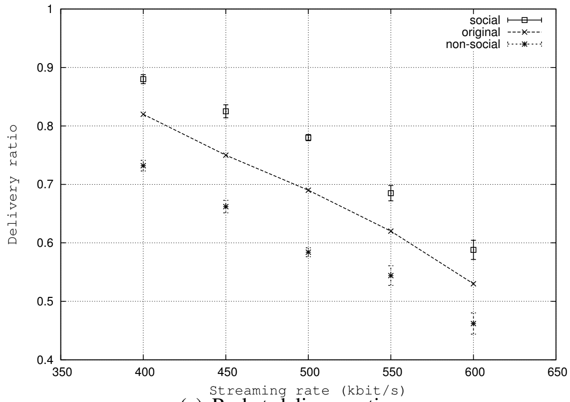

(a) Packet delivery ratio

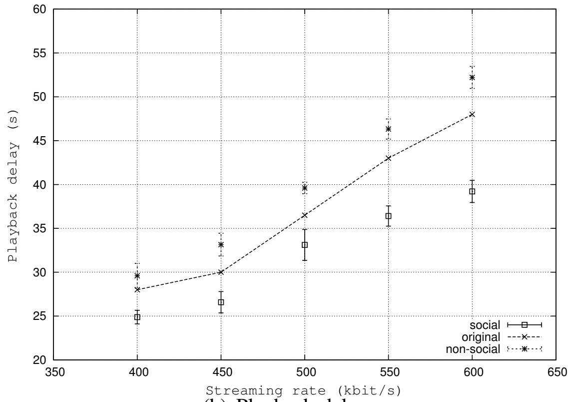

(b) Playback delay

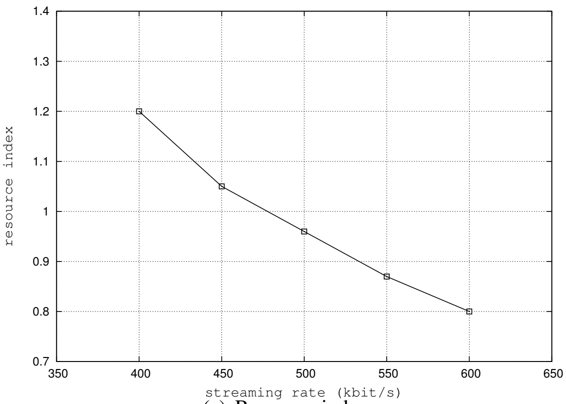

(c) Resource index

Fig. 1. Extended neighborhood case

increase and delay decrease to social network nodes, at the expense of a penalization for ordinary peers: in particular, their average delay swiftly increases the more the system departs from the "safe" condition where $\sigma > 1$ (in the graph, the higher the streaming rate considered). This is a consequence of the fact that the chunk requests of ordinary peers will be served only after social network peers requests have been satisfied, and when bandwidth is scarce, this will take more and more time to occur.

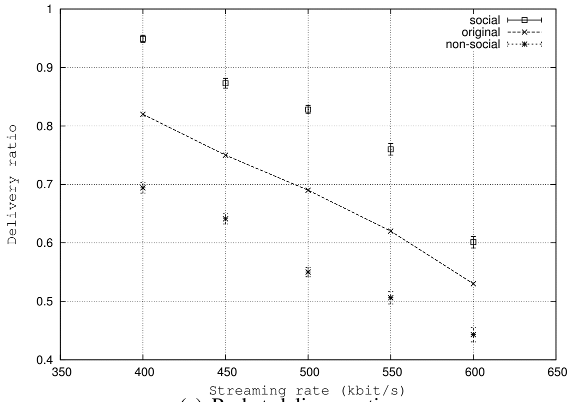

(a) Packet delivery ratio

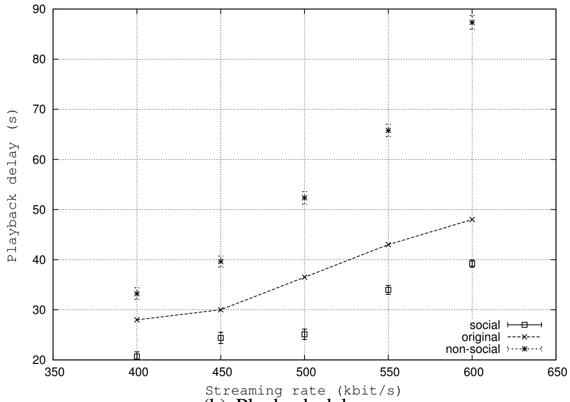

(b) Playback delay

Fig. 2. Extended neighborhood combined with the modified priority scheduling

C. Modifying the Percentage of Social Network Members

Our last investigation regards the effects that different values of the probability p of a P2P node to be also member of the social network have on the examined metrics. The evaluations refer to a fixed value of the resource index, i.e., the limiting $\sigma = 1$ case, achieved for the same parameters set employed before, setting the value of the video streaming rate equal to 480 kbit/s. Figs.3 (a) and (b) refer to the choice of extending the neighborhood of social network peers: varying p from 0.2 to 0.8 causes no relevant variations in either the average packet delivery ratio or the playback delay of both classes of peers: as a matter of fact, the additional neighborhood that social nodes leverage is almost fully utilized for $p = 0.2$ and once it is filled, no further improvements can be attained.

On the contrary, the effects of different p values are noticeable when priority scheduling is also introduced, as Figs.4 (a) and (b) indicate: the more numerous the social network peers within the overlay, the better performance they attain; in parallel, the higher the penalization ordinary peers suffer, mainly in terms of increased delay.

VI. Conclusions

This paper has proposed some modifications to the neighborhood creation and chunk scheduling algorithm of a mesh-based P2P overlay, in order to favor peers belonging to a social

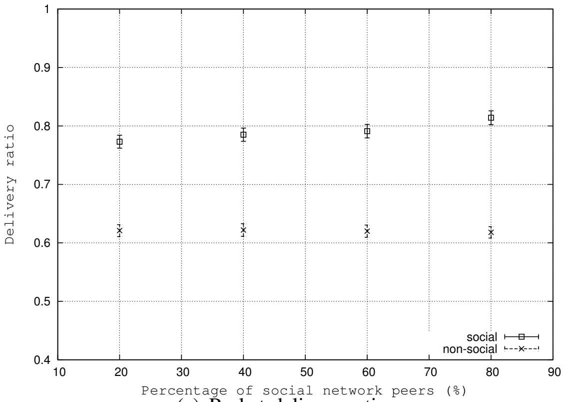

(a) Packet delivery ratio

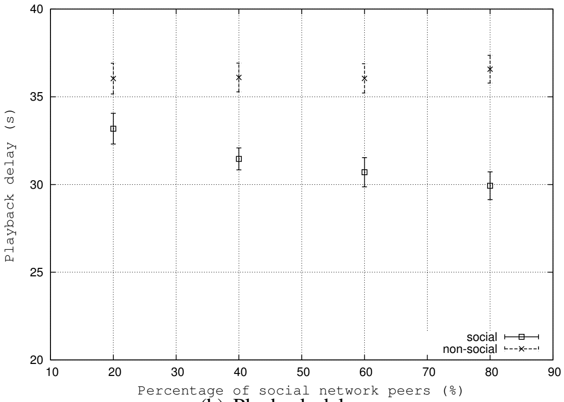

(b) Playback delay

Fig. 3. Effects of different p percentages, for the extended neighborhood case

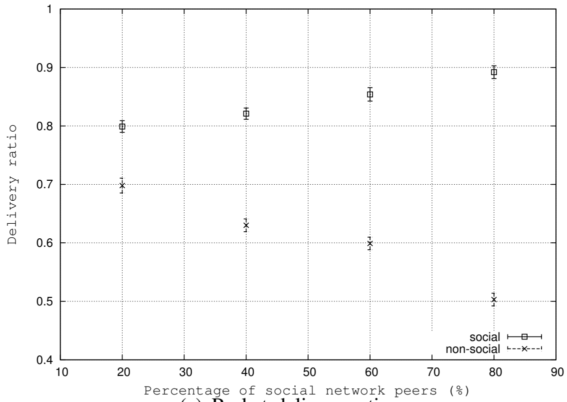

(a) Packet delivery ratio

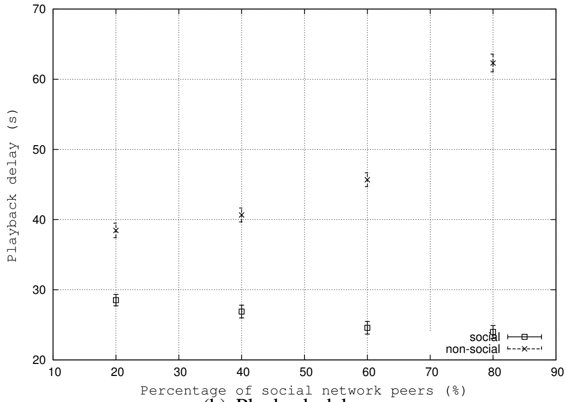

(b) Playback delay

Fig. 4. Effects of different p percentages, for the extended neighborhood combined with the modified priority scheduling

network and to grant them better performance. It has been shown that the introduction of an extended neighborhood for social peers attains such goal, both in terms of delivery ratio and playback delay, and that more consistent improvements are achieved when also priority is introduced in the video chunk scheduling algorithm of social nodes, serving their friends requests before those of ordinary members of the overlay. Hence, a clear service differentiation can be guaranteed, acting as an incentive to join the social network and experience a better viewing experience, at the expense of those pees that do not belong to the social community.

ACKNOWLEDGMENTS

The author would like to thank his former student Alessandro Valli for having performed some of the numerical assessments presented in this paper.

REFERENCES

[1] B. Li, H. Yin, "Peer-to-peer live video streaming on the internet: issues, existing approaches, and challenges", IEEE Communications Magazine, Vol. XX, No. YY, pp. 94-99, June 2007.
[2] J.A. Pouwelse, P. Garbacki, J. Wang, A. Bakker, J. Yang, A. Iosup, D.H.J. Epema, M. Reinders, M.R. van Steen and H.J. Sips, "Tribler: A social-based Peer-to-Peer System", in Recent Advances in Peer-to-Peer Systems and Security (IPTPS 2006), pp.127-138.
[3] K. Graffi, C. Gross, D. Stingl, D. Hartung, A. Kovacevic, R. Steinmetz, "LifeSocial.KOM: A secure and P2P-based solution for online social networks", Proceedings of IEEE Consumer Communications and Networking Conference (CCNC), January 2011, pp.554-558.
[4] H. Wang, F. Wang and J. Liu, "On long-term social relationships in peer-to-peer systems", Proceedings of 2011 IEEE 19th International Workshop on Quality of Service (IWQoS), June 2011, pp.1-8.
[5] O. Abboud, T. Zinner, E. Lidanski, K. Pussep, R. Steinmetz, "StreamSocial: a P2P Streaming System with Social Incentives", Proceedings of IEEE International Symposium on a World of Wireless Mobile and MultiMedia Networks, Montreal, Canada, June 2010.
[6] M.L. Merani, D. Saladino, M. Capetta, "Taking Advantage of Social Network Relationships in P2P Streaming Overlays", to be presented at IEEE Globecom 2011, Houston, Texas, December 5-8.
[7] Y. Tang, J.G. Luo, M. Zhang, S.Q. Yang; Q. Zhang, "Deploying P2P Networks for Large-Scale Live Video-Streaming Service", IEEE Communications Magazine, Vol. 45, No.6, pp.100-106, June 2007.
[8] M. Zhang, Q. Zhang, L. Sun, S. Yang, "Understanding the Power of Pull-based Streaming Protocol: Can We Do Better?", IEEE Journal on Selected Areas in Communications, December 2007, Vol.25, No.9, pp.1678 - 1694.
[9] J. Leskovec, L. Backstrom, R. Kumar, A. Tomkins, "Microscopic Evolution of Social Networks", Proc. of ACM SIGKDD, 2008 International Conference on Knowledge Discovery and Data Mining, August 2008, Las Vegas, Nevada, USA, pp.462-470.
[10] J. Guillaume, M. Latapy, "'Bipartite Graphs as Models of Complex Networks', Physica A: Statistical and Theoretical Physics, vol. 371, No.2, pp.795-813, 2006.
[11]) L. Botha, S. Kroon, "A Community-Based Model of Online Social Networks", SNA-KDD, 2010.
[12] Flickr, "Flickr Homepage", available at http://www.flickr.com/