

# BUILDING A COOPERATIVE P2P-TV APPLICATION OVER A WISE NETWORK: THE APPROACH OF THE EUROPEAN FP-7 STREP NAPA-WINE

E. LEONARDI, M. MELLIA, A. HORVATH, L. MUSCARIELLO, S. NICCOLINI, AND D. ROSSI

## INTRODUCTION

TV services over the Internet can be provided by either exploiting IP multicast functionalities or relying on a pure peer-to-peer (P2P) approach. The first technique will only work on a network infrastructure controlled by a single broadband operator due to limitations of IP multicast facilities. Fortunately, the P2P approach has been successfully exploited to overcome these limits and can potentially offer a scalable planetary infrastructure. Recently, several P2P-TV systems have emerged, with the last generation offering high-quality TV (P2P-HQTV) systems, providing ubiquitous access to the service. The emergence of P2P-TV systems constitutes a potential problem for network carriers since the traffic they generate may grow without control, causing degradation of the quality of service perceived by Internet users or network collapse (and the consequent failure of the P2P-HQTV service itself!).


Starting from these considerations, the Network-Aware Peer-to-Peer Application over Wise Network (NAPA-WINE) project, a Specific Targeted Research Project (STREP) funded by the European Commission within the Seventh Framework Programme, aims at: (i) providing a careful analysis of the impact a large deployment of both general P2P-TV and P2P-HQTV services may have on the Internet, through a detailed characterization of the traffic they generate; (ii) guidelines for P2P-TV developers regarding the design of systems that minimize the impact on the underlying transport network while optimizing the user-perceived quality; and (iii) a road map for Internet service providers (ISPs) to better exploit the network bandwidth by identifying simple and minimum-cost actions that can be taken in the presence of P2P-TV traffic.

The NAPA-WINE Consortium comprises three nationwide operators: France Telecom, Magyar Telecom, and Polish Telecom; a multinational company, NEC; a small Italian enterprise active in the market of videoconferencing over Internet, Lightcomm srl; a medium-size Hungarian enterprise active in the telecommunication market, Netvisor Ltd; and five top European technical universities: Politecnico di Torino, Università di Trento, ENST Télécom Paris — INFRES, Budapest University of Economics and Technology, and Technical University of Warsaw. The project is three years in duration and has a total budget of approximately €6 million.

## THE NAPA-WINE PROJECT

The main goal of the project is the study of a future system suitable for HQTV live streaming over the Internet based on P2P technology, or a P2P-HQTV system. The major focus is on overcoming today's pure layered approach through a cooperative paradigm in which the application and network layers cooperate to optimize the quality of service offered to end users. In more detail, the specific project goals are:

1. Characterization of P2P-HQTV traffic. Both theoretical and experimental activities will contribute to the accomplishment of this objective.
2. Definition and implementation of a distributed monitoring tool. Its aim is integration within P2P-PHQTV clients.
3. Definition and implementation of algorithms for the control of cooperative P2P-HQTV systems. The algorithms will define a middle-layer control overlay and distribution management.
4. Definition of a set of public functional interfaces between the network and P2P-HQTV systems. These inter-



■ **Figure 1.** Schematic representation of the NAPA-WINE cooperative approach.

faces define the rules according to which the network and the application will cooperate.

5. Definition of mechanisms to better exploit the network bandwidth in the presence of P2P-HQTV traffic.

6. Integration of some of the above algorithms within an existing P2P HQTV client.

7. Experimental evaluation of the effectiveness of our approach.

In particular, NAPA-WINE makes substantial and significant steps toward a deep understanding of the interactions between the transport network and P2P-HQTV systems, providing answers to fundamental questions raised by operators, such as:

- Can the current Internet infrastructure support a massive deployment of large-scale high definition video streaming distribution systems?
- How can an operator predict/control what happens in its transport infrastructure?
- What are the costs in terms of the backbone bandwidth needed to successfully support such applications?
- How can the above costs be minimized?
- Will P2P video traffic be harmful to other applications? How can this be avoided?

Moreover, our research will provide answers to important questions raised by P2P-HQTV application developers and content producers, such as:

- Can users' satisfaction be improved by providing a better service?
- Can cooperation be fairly defined, and what kinds or forms of cooperation are at our disposal?

In the following sections a more detailed description of the NAPA-WINE project is given, identifying actions at the application and network layers.

## THE APPLICATION LAYER

Focusing on an unstructured P2P-HQTV system, a schematic representation of the NAPA-WINE cooperative and network-aware application is represented in Fig. 1. An important element of the cooperative P2P-HQTV system is represented by the built-in distributed monitoring tool that allows the application to continuously gather real-time information on both network conditions and user-perceived performance. Information collected by the monitor tool will potentially trigger reconfiguration algorithms acting at the levels of both the chunk distribution mecha-

(Continued on page 22)

(Continued from page 20)

nism (scheduling) and the overlay topology reconfiguration. In addition, some potentially useful information on the system state might also be exported toward the network so that the network is aware of the status of the P2P-HQTV system.

We stress that a key factor for beneficial cooperation is the mutual exchange of information between the application and network layers. Indeed, in the context of P2P file sharing, a recent study [1] suggests that ISPs should offer an “oracle” service to P2P application to bias the overlay neighbor selection. Basically, when a peer supplies the oracle with a list of potential neighbors, the oracle ranks them according to certain performance criteria (proximity to the user, available bandwidth, etc.), thus allowing the ISP a simple handle to cooperatively manage P2P traffic. At the same time, since the oracle interface is unidirectional, the ISP misses a large amount of very relevant information on the system status: the end-to-end quality of service perceived by P2P-HQTV users. This clearly represents a waste of potentially useful information that, by reason of its end-to-end semantic, constitutes a natural complement to the ISP point of view, making the network aware of the status of the P2P-HQTV system.

Within the NAPA-WINE project, first a list of both layer 7 and layer 3 performance metrics to be continuously monitored will be defined. These metrics will include quality of service, overlay topology, and network-related aspects. Then a study of the most suitable and convenient measuring technique for each metric will be conducted; typically, passive measurements will be preferred, since they do not require the injection of additional traffic.

We will also define what exactly awareness and optimization means in the different contexts, starting from the measurements provided by the monitoring tool and interacting with the normal “functions” of the P2P-HQTV application: streaming, distribution management, and overlay management. Several reconfiguration algorithms will be devised, with particular attention to how to exploit the information available at peers, at the swarm, and at global levels gathered from the monitoring blocks.

## THE NETWORK LAYER

The introduction of high-capacity network-aware peers in the P2P system can play a role that benefits the transport network. The application has information on the overlay status, while the owner of the transport infrastructure has detailed information on the layer 3 network status. When a subset of peers is operated by the owner of the network, these two information sets can be merged and exploited. Based on this integrated view, the peer operated by the owner of the network infrastructure, or *network-peer*, can perform actions to optimize the P2P overlay taking into account the status of the transport network. It can cooperate in the stream distribution by downloading and redistributing chunks. It can help in positioning network-peers in the overlay network based on its knowledge of the actual peer position in the transport network. Finally, it can eventually signal the transport network control plane of possible congestion due to P2P traffic, to trigger reconfiguration of the transport network itself. An extension of this solution is the introduction of an entire network of network-peers to achieve the same goal. In this case much better knowledge of global P2P traffic can be collected, constituting a distributed control plane.

Even without the explicit presence of network peers, network nodes will likely monitor several performance indices (link load, losses, per-flow performance, etc.). Due to the huge potential benefit of the collaborative cross-layer paradigm promoted by

NAPA-WINE, some of the information gathered by the network equipment can be shared with the P2P application. However, in this case particular attention must be devoted to the evaluation of two aspects: first, the costs of network monitoring; second, and more important, the potential risks of sharing sensitive information with peers. Similarly, we remark that part of the performance metrics collected by the peers at the application layer can be shared with the transport network. The list of items that could be of some interest for the network control mechanisms that make the network aware of the P2P-TV overlay will therefore be investigated in NAPA-WINE. One particularly important aspect is the possible interaction between overlay dynamics and network dynamics. In the presence of traffic-dependent automatic management algorithms (e.g., dynamic routing protocols), this interaction requires special attention. Indeed, when two dynamic systems make independent decisions, there is a risk of instabilities due to global positive feedback. Given the complexity of these interactions, the project goal is mainly devoted to the avoidance of the potentially unstable situations pointed out in [2, 3] by analytically studying the problem.

## CONCLUSIONS

In this column we presented NAPA-WINE, a three-year-long STREP supported by the 7th Framework Programme of the European Commission. Within NAPA-WINE we envision a cooperative paradigm in which the P2P application will collaborate with the network layer to support high-quality TV streaming over the Internet. We believe that cooperation between the application and network layers must be achieved to offer high-quality service to a very large population of users. Following these goals, NAPA-WINE aims to define, study, and implement a new P2P-HQTV system exploiting monitoring and control techniques. Please visit <http://www.comsoc.org/VeryLargeProjects> for questions and answers on the Very Large Projects column.

## REFERENCES

- [1] V. Aggarwal, A. Feldmann, C. Scheideler, “Can ISPs and P2P Users Cooperate for Improved Performance?” *ACM SIGCOMM Comp. Commun. Rev.*, vol. 37, no. 3, July 2007.
- [2] R. Kerelapura, N. Taft, C.N. Chuah, and G. Iannaccone, “Can ISPs Take the Heat from Overlay Networks ?” *ACM HotNets-III*, San Diego, CA, Nov. 2004.
- [3] Y. Liu, H. Zhang, W. Gong, D. Towsley, “On the Interaction between Overlay Routing and Underlay Routing,” *IEEE INFOCOM '05*, Miami, FL, Mar. 2005.

## BIOGRAPHIES

EMILIO LEONARDI is currently an associate professor at the Dipartimento di Elettronica di Politecnico di Torino.

MARCO MELLIA graduated from the Politecnico di Torino with a Ph.D. in electronic and telecommunication engineering in 2001. He is now an assistant professor at Politecnico di Torino, working on traffic measurement, classification and monitoring, and peer-to-peer networks.

AKOS HORVATH He received an M.Sc. degree in computer science from Budapest University of Technology and Economics in 2007. He is currently working as a research fellow at the Department of Telecommunications of the same university.

LUCA MUSCARIELLO [M] works at France Telecom R&D on performance evaluation of communication networks.

DR. SAVERIO NICCOLINI received his M.S. and Ph.D. degrees in telecommunication engineering from the University of Pisa, Italy, in 2000 and 2004, respectively. In 2004 he joined NEC Network Laboratories, Heidelberg, Germany, where he is currently a senior researcher and team leader of the VoIP Security team.

DARIO ROSSI has been an associate professor of the Computer Science and Networks Department (INFRES) at ENST, Paris, France since 2006. He received his M.Sc. and Ph.D. in electrical engineering from Politecnico di Torino in 2001 and 2005, respectively, while he held a visiting researcher position in the Computer Science Department of the University of California at Berkeley between 2003 and 2004.