# Overlay/Underlay Routing Issues in Wireless Mesh Networks

Stefano AVALLONE[1], Roberto CANONICO[1], Francesco Paolo D'ELIA[1]
[1]*University of Naples "Federico II", Via Claudio 21, Naples, 80125, Italy*
*Tel: +39 081 7683821, Email: (stavallo, roberto.canonico, francescopaolo.delia)@unina.it*

**Abstract:** Recent analyses of the network traffic show that an ever growing portion of it is generated by user communities that share their own resources in a peer-to-peer fashion. Internet Service Providers' attitude towards peer-to-peer applications, so far, has been of considering them more as potential threats rather than as an opportunity and a customer demand to second. This is mostly due to the fact that peer-to-peer applications still represent a significant traffic engineering challenge for ISPs, since they tend to build their own overlays in a way that is largely independent of the Internet routing and topology. Wireless Internet Service Provider (WISPs) are making large use of the Wireless Mesh Network model to easily and quickly offer their services in urban and rural areas with reduced investments. In the WMN context, the problem of controlling P2P traffic is even more stringent, due to the limitedness of resources. In this paper, we discuss these issues and invoke the adoption of an architectural model that allows the exchange of information among peer-to-peer applications and mesh routers, according to a cross-layer approach, where not only lower layers (data link) but also upper layers (application) influence the network layer.

## 1. Introduction

Overlay networks are virtual communications infrastructures implemented 'on top of' an underlying physical network such as the Internet. Routing of packets in an overlay is usually performed at the application level. Overlay networks are usually created either to force some special routing in the network (e.g. multicast routing, or QoS routing) or to organize application-level entities in order to efficiently provide some form of service to large scale communities (e.g. peer-to-peer networks for file sharing or video streaming).

Systems based on overlay of peers are inherently distributed and, if properly designed, more robust and scalable, due to their decentralized nature and to the absence of single points of failure. Nonetheless, overlay routing decisions collide with those made by underlay routing, i.e. ISP routing decisions. While overlay routing objectives are ultimately applications objectives, i.e. compliance with a given QoS specification or latency optimization, underlay routing decisions are guided by objectives such as load balancing, minimization of cross-ISP traffic and minimization of links utilization. As a consequence of such a dichotomy, several un-efficiencies may result. For example, one problem that can arise is traffic routing oscillations: if the overlay routing move a great amount of traffic from one path to the other, ISP as a result may notice a bandwidth decrease and change the weight link balance in order to redistribute traffic. This, in turn, could induce a subsequent overlay routing change that can trigger repeatedly underlay routing change, bringing the overall system to an endless traffic imbalance: this is an example of negative consequences of no information exchange between the two routing layers. Moreover, due to the fact that overlay routing does not know physical nodes positions, it is not uncommon that adjacent nodes of an overlay network are in different ASes. Such a topology arrangement leads to traffic crossing network boundaries multiple times, thus overloading links which are frequently subject to congestion, while an equivalent overlay topology with nodes located inside the same AS could have had same performance.

From what we described above, it emerges that overlay routing may benefit from some form of underlay information recovery, or in general from cross-layer information exchange. When creating an overlay network, the choice of the nodes (i.e. the network topology) can be done by taking advantage of information from the underlay network. The problem of creating network informed overlays has been studied both for general overlays and for specific peer-to-peer applications. Since peer-to-peer traffic is nowadays predominant in the Internet, this problem is of great interest for Internet Service Providers.

Cross-layer optimization techniques are also widely adopted in the context of Wireless Mesh Networks, where they are often applied to make optimal routing decisions in combination with the spectrum resource allocation. In such a context, negative interactions among overlay and underlay routing decisions may jeopardize the whole network ability to keep operating. Since a couple of nodes can communicate only if a common communication channel exists, in WMN. When nodes share a common channel, though, the total link capacity is shared as well: the number of nodes in the same interference range, using the same channel implicitly limits the link bandwidth. It is then clear, then, that channel assignment implicitly limits the link bandwidth. The channel assignment main objective, then, is to realize link bandwidths so that is possible to route traffic inside the network. The more accurate the description of the traffic incoming into the network, the more efficient the bandwidth usage benefits. To this purpose, one important parameter is the traffic demand of a given P2P application. Providing such information to the underlay network could result in better resource utilization in the underlay. In the case of a WMN, overlay applications could give information to mesh routers, such as peer nodes sharing a certain software content and the bandwidth requirements for the file transfer. The Wireless Mesh Network, in turn, could use this information to readapt channel assignment in order to satisfy new traffic requests.

The overlay building procedure often employ application-level measurements to most suitably organize the overlay topology. Such an approach has been proved to be suboptimal and inefficient. A possible alternative would be to build some kind of ”oracle” service, realized and maintained by the underlay network, i.e. the ISP, that could instruct users about the best way to organize their overlay. The adoption of such an approach is seen as mutually beneficial for both users and Internet Service Providers. In the rest of this paper we review some of most recent approaches for underlay/overlay cooperation that have been proposed in the literature. Subsequently we further describe the benefits of supporting cross-layer interactions in a WMN and we illustrate our proposal of adapting an oracle service in the specific context of Wireless Mesh Networks.

## 2. Application Layer Traffic Optimization

Several recent works have analyzed the problem of improving network performance through Overlay/Native networks collaboration.

In [1] the authors study the routing performance of a particular layer when very few information is known from the other layer, i.e. non-cooperative interaction occurs. In particular, in the work one layer is considered as Leader and it tries to predict and react to the other layer’s actions (the Follower), trying to avoid performance deterioration. The leader layer tries to optimize its own performance, trying to preserve the optimization objective of the other layer (friendly approach) or with no regard for the other layer performances (hostile approach). In particular, for the overlay layer the strategies consist in paying attention to not trigger underlay optimization decisions, trying to keep the load balance of the underlay network intact (that is, the friendly approach), or, for the hostile approach, to completely defeat the underlay optimization objectives. For native layer, the friendly strategy consists in modifying the optimization step, with the extra consideration of keeping the native route of the same length as in the last optimization step: in this way, the overlay layer will have no need in changing routing paths, thus maintaining good load balance in the network. The hostile strategy, contrarily, consists in strongly limiting the overlay layer free path choice: in this way, underlay optimization do not need to take into account possible traffic matrix variations due to overlay traffic switching. To manipulate overlay routing, link latency is modified in such a way to push the overlay nodes to keep the same routing table. At the same time, the stability of the network is studied in terms of route flaps numbers, which is the result of the conflicts between layers and it is determined by the number of multi-hop overlay paths: it is shown that for both strategies stability is reached in few steps. The optimization objectives are end-to-end delay minimization for overlay users and network cost minimization for underlay network. In the paper, no attention is dedicated to information exchange between the two layers. Despite the great performance obtained by the hostile strategies, the damage brought to the other layer is unsustainable; on the other hand, the friendly performance gives stability to the network, at cost of reduced performance. In the end, performance improvement can be obtained by just one layer, and in general the performance is strictly dependent of the other layers strategy, with poor information knowledge. This convinced us that, through cross-layer collaboration it is possible to achieve a more balanced use of network resources, leading to better performance for both overlay and native routing.

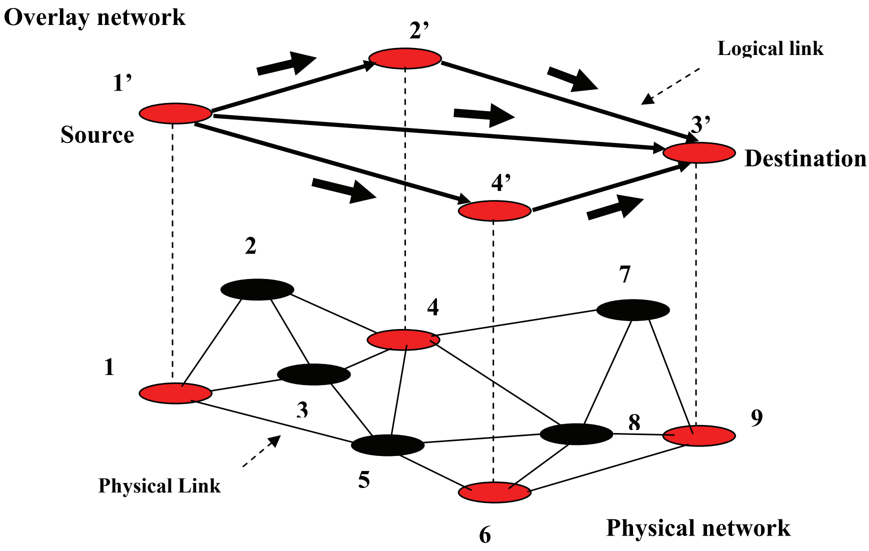

**Figure 1 – An example of overlay network**

In [2], the interaction between overlay and underlay routing is studied for a single AS (Autonomous System). Again, overlay routing objective is to minimize end-to-end delay, while underlay routing tries to minimize a given cost function. Game-theoretic models are used to study the interaction between the layers, where each layer, for each step, tries to optimize the given objective. Nash Equilibrium is found for simple-case topologies and it is shown that underlay performance will never be improved in the Nash routing game when overlay routing is present. Moreover, routes oscillations are found for both overlay and underlay, due to differences in optimization objectives. It is shown that the selfish behavior of the overlay layer can rise the value of the global cost for the underlay layer and, moreover, the optimization steps taken by the overlay routing in each round, i.e. when it plays a best-reply Nash routing game with underlay, can worsen even the performances of the overlay user instead of improve them. It is then revealed that when no explicit collaboration is present, overlay optimization criteria is not always the one that really minimize the objective function, i.e. best-reply strategy is not the best strategy in this case. Besides, when overlay traffic percentages are not small, underlay's cost is shown to be dramatically increased by the interaction with overlay routing: a certain coordination between the two layers in regards of the optimization moments would lessen the misalignment degree. Even the authors feel the importance of knowledge exchange and mutual awareness in cross-layer routing optimization, and an example would be information given to the overlay about the optimization algorithm of the underlay routing: in this way the upper layer would be able to predict underlay decisions, thus optimizing the overlay routing decisions. Still, information estimation and cross-layer strategies are strongly called.

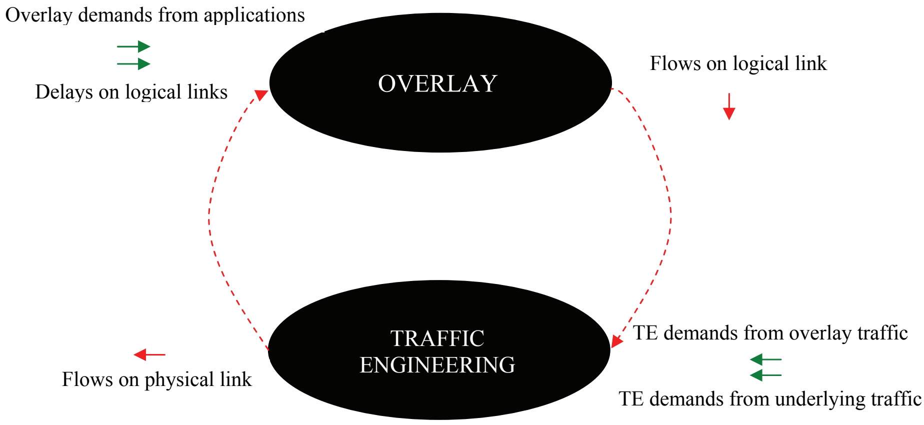

**Figure 2 – Interaction between overlay optimizer and TE's optimizer**

In [3] P2P traffic is directly addressed. Once again P2P user perceived performance is the main issue and it is solved thanks to the overlay network: differences in overlay and underlay routing are analyzed, like topology creation, routing criteria and different dynamics. In particular, the overlay topology creation problem is addressed. Each node should pick the neighbor with smallest delay possibly highest throughput, and in the same AS: in the paper the cross-layer information exchange is welcomed and an ”oracle” service is studied. The ”oracle” is a service supplied by the ISP (Internet Service Provider, i.e. the underlay network manager) to the overlay P2P users: very different information is provided, like link delay, bandwidth estimations, etc. The ”oracle” service takes as input a list of P2P nodes sharing a known content (even from different ASs) and gives back the nodes list ranked according to different performance metrics. In this case, the P2P application won't have to perform such measurements by itself (and these measurements would be available to all P2P applications) and thus, measurements overhead traffic is avoided and the topology creation process is optimized. Moreover, the ”oracle” gives ISPs a way to control overlay topology creation decisions, and, ultimately, overlay routing. The paper shows how explicit cross-layer information exchange (in this case information flow goes from underlay to overlay) can really improve overlay operations (topology creation in this case) and at the same time improve underlay network usage (AS hops are minimized and congested link could be unloaded). Despite the difficulties in oracle management, the results suggest that cross-layer information exchange is a valid starting point in solving the overlay-underlay frictions.

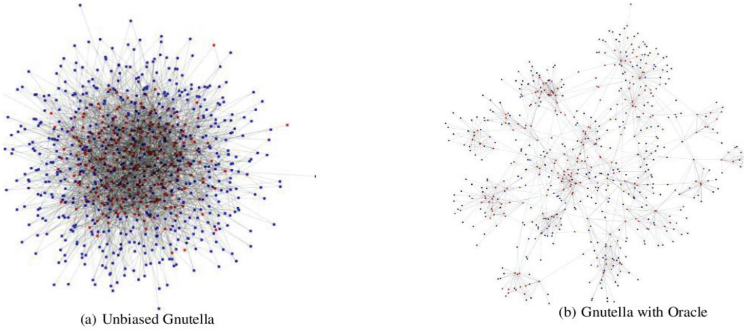
(a) Unbiased Gnutella

(b) Gnutella with Oracle

**Figure 3 – Topologies obtained without and with Oracle**

In [4] the problem of the interaction between routing layers is addressed when route failures take place at the physical layer. For both underlay and overlay layers a dynamic routing protocol (Open Shortest Path First) is employed and average route flaps number is considered while evaluating performances. In particular, two scenarios are considered: one where no interaction is present between layers, with sub-optimal results with regards of route flaps and network cost (due to unawareness between layers and overlap of functionalities between levels); another scenario, where awareness of underlay routing at the overlay is considered, and better results are found in terms of route flaps and cost function optimization. Once again, independent operations of the two layers returns insufficient performances, high number of route flaps and general system instability. Moreover, two approaches are followed to improve interaction between layers: overlay awareness of the underlay routing through information access at the native layer, and underlay routing awareness of the overlay layer and consequent parameters tuning in order to favor overlay routing performances. In particular, better results are obtained with the second approach, where the native layer tries to suit the overlay traffic, helping the overlay layer to reach the best performances according to different metrics.

In [5] collaboration between layers is invoked, especially in terms of control traffic. The results of introducing collaboration between layers is to obtain optimality in terms of underlay network costs, while protecting P2P users performance. The paradigm introduced by the authors, P4P architecture, consists in explicit information exchange between providers and application users, i.e. a portal. The main feature of this architecture is the *p*-distance information interface, which collects and gives information about physical network's status and indications about possible application traffic choices. The authors call these portals iTrackers, and each network is responsible for each own iTracker, since each portal collects information for one particular domain. Each provider can set the p-distance value based on link cost or link congestion, while applications can combine the provider information with delay or bandwidth requirements, through application level optimization algorithms. Moreover, thanks to the iTracker, it is possible to divide management responsibilities between provider and application. In particular, in the paper it is shown how the MLU minimization problem is decomposed in many different and independent problems, each one linked to a given application session: in this way a distributed solution to the optimization problem is possible, thanks to the information exchange between layers. Once again it is shown that through cross-layer information exchange it is possible to reduce completion time and bottleneck traffic, both in terms of intra-domain and inter-domain transmissions.

In [6] the authors work on an adaptation for Wireless Mesh Networks of the distributed resource index protocol introduced in [7]. In particular, the exploited peculiarity of Wireless Mesh Networks is the 1-hop broadcast transmission, while at the same time particular care is directed towards mobility and signalling overhead. Like in [7], MeshChord maps each node ID and each shared resource ID into a hash table entry key. In [6], though, the authors try to exploit node location inside the network, when assigning IDs to nodes: observing packets exchange due to the Chord protocol, the traffic is seen to be more intense when peers are close-by in terms of ID, and then nodes that are physically close-by are assigned to close-by IDs. This approach, though, requires nodes positions. Moreover, to speed up the search operations, each node captures a *lookup* packet even if it is not the destination of the *lookup* packet. In this context, the *lookup* packet is used to find the IP address of a node with a known ID, and it is sent to . The concept is that the node capturing the packet could be able to answer to that request, and in that case it could answer to the requesting node directly, thus speeding up the requesting phase. The paper, finally, shows how a protocol designed for wired-connected networks can be adapted to the features of the wireless networks, thanks to cross-layer techniques, and that it needs to be adapted in order to improve performances.

The problem of providing “standard” ways of interoperation between overlay and underlay networks has been recently considered in the IETF, where a Working Group on “Application Layer Traffic Optimization” (ALTO) has been recently chartered [9].The objective of this working group is “to design and specify a service that will provide applications with information to perform better-than-random initial peer selection”.

## 3. P2P in Wireless Mesh Networks

In the previous section, we have discussed the issues stemming from application layer routing decisions being unaware of the underlying network. In particular, the work appeared in the literature has shown the negative consequences of selfish routing strategies at both the overlay and the underlay level. The aim of this section is to revisit such issues in the context of wireless mesh networks. WMNs are expected to provide Internet connectivity to end-users and, consequently, to carry a large amount of P2P traffic. Such volume of traffic, in the absence of cooperation between overlay and underlay layers, may have even more negative consequences in WMNs than in wired networks.

Firstly, P2P applications generate a large amount of overhead traffic for maintenance operations. For instance, each peer may autonomously probe the underlying network in the attempt to guess some network performance. In wireless multi-hop networks, where the available bandwidth is cut down by interference, overhead traffic has much more harmful consequences than in wired networks. Hence, one of our goals is to define an architecture where peers are aided by the ISP in the selection of a group of peers to be contacted. Thus, peers no longer need to probe the network by themselves.

Secondly, the wireless nature of links in WMNs requires the tuning of parameters such as the frequency channel, the transmission rate and the transmission power for a link that can determine the existence of the link and the bandwidth available on the link itself and on the links in the neighbourhood. As shown in the literature (e.g., [8]), the configuration of such parameters is strictly inter-dependent with the traffic expected to flow on every wireless link. Indeed, each router in a WMN is often endowed with a certain number of radios to alleviate the interference problem. In order to optimize the network performance, flows distribution needs to be taken into account. Typically, the amount of traffic flow that crosses a network link is found so that a given objective function is optimized (e.g., minimize the costs inside a network, maximize the throughput, etc.). In wireless mesh networks, a new challenge is represented by the channel assignment problem, as flows optimization and channel assignment are not independent problems. Indeed, when determining the optimal traffic flowing across a link, link capacity has to be known. In wireless networks, however, the channel assignment settles the number of links sharing the medium and, hence, their link bandwidth. Channel assignment, thus, should be solved prior to the flows assignment. But, one of preconditions of a channel assignment algorithm is usually the knowledge of the traffic flow expected on every network link, in order to allow for enough bandwidth on every link. Thus, flows assignment should be solved prior to the channel assignment. It turns out that channel assignment and flows optimization should be jointly solved. The joint problem, however, is known to belong to the NP-complete class of problems. Thus, the typical approach is to find an approximate solution by solving the two problems separately. The flows optimization problem is solved first. Then the channel assignment problem turns out to be the problem to assign a channel, a transmission power and a transmission rate to every link so that the flows returned by the first step are schedulable. Such a channel assignment problem is also found to be NP-complete and hence effective heuristics are to be designed.

It is thus clear that a given channel assignment is tailored on a certain distribution of flows inside the network. A different distribution of flows instead may result in poor network performance. Since P2P applications can move large amount of traffic in the attempt to improve the performance they experience, it is crucial that their decisions match the underlying network configuration. This last point motivates us to study the overlay/underlay routing interaction problem from a somewhat different perspective than previous research. Indeed, the underlay routing has always been looked at as a sort of black box, giving a little chance, if any, to modify its behaviour. Our goal, instead, is to design an architectural model for the exchange of information between the overlay network and the underlying wireless mesh network. We envision that some (or even all of the) mesh routers are equipped with some enhanced functionalities that allow them to interact with overlay applications. In a typical scenario, overlay applications may provide a mesh router with a list of neighbors sharing the desired resource and an estimate of the bandwidth required by such a communication. The mesh router will then rank each peer based on how the corresponding traffic matches the actual network configuration. Alternatively, the mesh router may re-compute a new network configurations (channels, transmission powers and rates) if one exists that is better suited to accommodate the new and the previous traffic requests.

Another important piece of such an architecture is represented by the underlying routing protocol, which should take into account the bandwidth available on every link as a result of the computed channel assignment. If routers were unaware of the flows distribution for which the current network configuration has been computed, then the resulting network performance would be poor. The routing protocol should be robust to minor variations in the traffic offered to the network, in order to maintain the network performance at an acceptable level. Also, the routing protocol should be able to rapidly react to temporary link unavailability or quality degradation, which may often occur in wireless environments.

The architecture we envision thus largely adopts cross-layer principles, given that parameters at lower layers heavily impacts the behaviour at the application layer and vice versa. The need of mixing information from different layers to take proper decisions stems from the strict inter-dependence that exists in wireless mesh networks between the distribution of flows inside the network and the configuration of data link parameters such as the frequency channel, the transmission rate and the transmission power.

## 4. Conclusions

From the recent works, inefficiency and instability following independent interaction between P2P/overlay and underlay routing emerges. When dealing with wireless mesh networks, even more problems arise, suggesting the need for a cooperative cross-layer routing approach, with particular attention to the radio channel assignment. The cross-layer routing protocol would take routing decision based on heavy information exchange between all the layers, from the MAC level up to the application level, in the attempt to optimize both network and application performance. The first step will be to define and evaluate underlay metrics that could be useful in helping applications. Then, a new architectural model is to be designed in order to exchange information between layers. Finally, applications should be designed in order to interact with the proposed architecture. The model we envisage is characterized by overlay nodes giving information about traffic bandwidth request to the underlay network. The underlay routing protocol then analyzes the requests, allocates the needed resources and notifies the upper level about the new changes.

## References

[1] S. Seetharaman, V. Hilt, M. Hofmann, M. Ammar, *Preemptive Strategies to Improve Routing Performance of Native and Overlay Layer*s – IEEE INFOCOM 2007, pages 463-471

[2] Y. Liu, H. Zhang, W. Gong, D. Towsley, *On the Interaction Between Overlay Routing and Underlay Routing* – IEEE INFOCOM 2005, pages 2543-2553

[3] V. Aggarwal, A. Feldmann, C. Scheideler, *Can ISPs and P2P systems co-operate for improved performance?*. In ACM SIGCOMM Computer Communications Review (CCR), 37:3, pp. 29-40, July 2007.

[4] S. Seetharaman, M. Ammar, *On the Interaction between Dynamic Routing in the Native and Overlay Layers* – IEEE INFOCOM 2006, pages 1-12

[5] H. Xie, R. Yang, A. Krishnamurthy, Y. Liu, A. Silberschatz, *P4P: Provider Portal for Applications.* In ACM SIGCOMM Computer Communications Review (CCR), 38:4, pp. 351-362, October 2008.

[6] S. Burresi, C. Canali, M. Renda, P. Santi, *MESH CHORD: A Location-Aware, Cross-Layer Specialization of Chord for Wireless Mesh Networks* – IEEE Percom 2008 pages 206-212

[7] R. Morris, D. Karger, M. Kaashoek, H. Balakrishnan, *Chord: A Scalable Peer-to-Peer Lookup Service for Internet Applications,* – SIGCOMM 2001, pages 149-160

[8] S. Avallone, I. Akyildiz, G. Ventre, *A Channel and Rate Assignment Algorithm and a Layer-2.5 Forwarding Paradigm for Multi-Radio Wireless Mesh Networks,* – IEEE/ACM TRANSACTIONS ON NETWORKING, to appear in the February 2009 issue

[9] Application-Layer Traffic Optimization (alto) IETF Working Group, Applications Area. http://www.ietf.org/html.charters/alto-charter.html