Routing Schemes in Heterogeneous Wireless Networks Based on Access Advertisement and Backward Utilities for QoS Support

Leonardo Badia, IMT Lucca Institute for Advanced Studies
Marco Miozzo, Michele Rossi, and Michele Zorzi, University of Padova

ABSTRACT

This article reviews routing algorithms in heterogeneous wireless networks with the goal of including QoS awareness. After a general overview of the issues and challenges of QoS provisioning over heterogeneous networks, classic routing strategies are revisited, and how they can be exploited to achieve QoS efficiency is discussed. In particular, for the considered scenarios our proposal is to account for some proactivity in the routing algorithms, as well as a QoS-driven control selection, which are shown to improve throughput, delay, and energy consumption. Finally, we introduce a general backward utility formulation for user satisfaction as a tool to capture complex and dynamic QoS variations.

INTRODUCTION

Current advances in hardware and microelectronic design have led to a new generation of very small radio communication devices. This, together with the decrease in the cost of network interface cards, makes it possible to supply the same device with multiple access techniques, often with different characteristics. In the near future, the coexistence of multiple and heterogeneous access techniques is expected to transition toward high-level integration [1]. This evolution will probably be carried out by means of common network management strategies by exploiting, in a cross-layer fashion, ambient and context information about location, mobility, surrounding access techniques, as well as user profiles.

In a fully integrated scenario, we might think of a network where a number of nodes, each of them being either a fixed access point (AP) or a mobile terminal (MT), coexist, and all own several wireless interfaces. We may further assume that the APs (or at least some of them) are interconnected through wires to the Internet, such that both wired and wireless communications take place in the same network. For the wireless part, we may exploit multihop communication for improved performance in terms of, say, coverage extension and reduced transmission power. Such a scenario includes both purely ad hoc networks, connected to one or more sinks; mesh networks, where the backhaul interconnecting the APs may be entirely wireless; and hybrid solutions [2].

In recent research the availability of multiple channels/access techniques is increasingly regarded as a viable solution to improve network capacity, by means of reduced interference and increased transmission parallelism. However, the harmonization and coordination of multiple radio technologies (RTs) is a challenging and still unsolved issue. Different RTs, due to the specific features of their access and physical layers, usually provide diverse quality of service (QoS), which impacts the performance of upper layers (e.g., transport and application) and protocols (e.g., routing).

For example, wireless LANs (WLANs) based on IEEE 802.11 technology offer relatively high bit rates, but the service is inherently best effort. This is due to the medium access being based on carrier sensing and therefore affected by collisions [3]. Hence, their actual performance very much depends on the state of the network, and it is not so trivial to devise schemes for achieving proportional QoS [4]. On the other hand, cellular networks, such as Universal Mobile Telecommunications System (UMTS), offer better coverage and stability, but generally at lower data rates. While similar differences can be discussed for other RTs, in the remainder of this article and without loss of generality, we focus on the two aforementioned wireless systems, whose integration represents a case of practical interest.

 0163-6804/07/$20.00

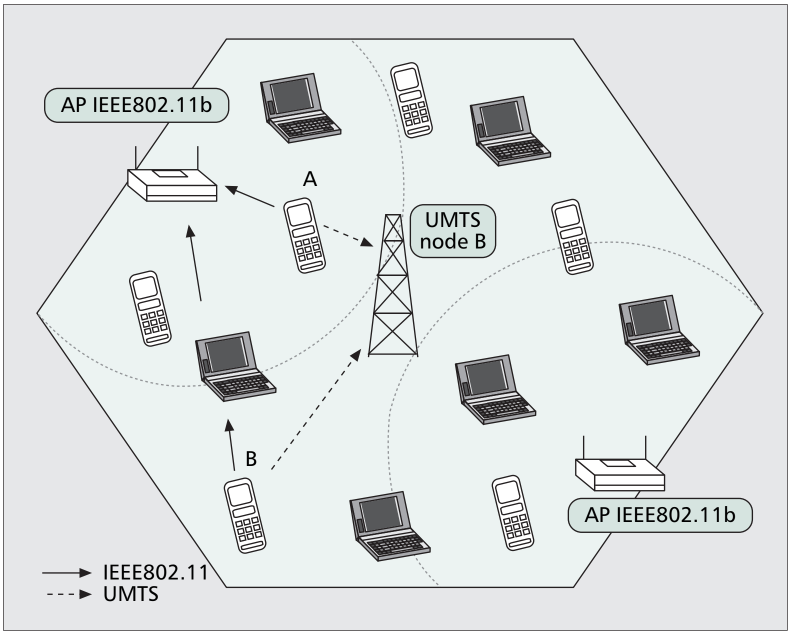

Figure 1. *A heterogeneous network scenario.*

Routing Flows through a Heterogeneous Network

Multihop routing is one of the main problems of the complex scenario described above. Typically, current cellular networks and WLANs implement single-hop transmissions. However, it is well known that the introduction of multihop communication is beneficial in many cases (e.g., to reduce interference and increase transmission parallelism). Multihop communication is also good for providing network coverage to those users who do not own certain network capabilities, by means of relaying through other nodes such as APs or more powerful MTs [5]. For these reasons, it is realistic to assume that multihop capability is enabled.

Figure 1 shows an example scenario where users are interested in connecting to the Internet using their MTs. A routing algorithm must select the access technique for users under direct coverage of multiple APs (e.g., user A) and decide how to exploit multihop relaying, as might be the case for user B. The routing strategy may first select the most suitable RT to use (e.g., IEEE 802.11b or UMTS in the figure) and then route the flow according to a given policy. In addition, this should be done in a QoS-aware manner (i.e., by picking a feasible path that satisfies rate/delay constraints and meets users' preferences).

In past literature, many papers have dealt with the issues of extending coverage and providing connectivity to APs through multihop communication. For example, the problems of routing data flows with different requirements in a heterogeneous scenario are considered in [6], where the main focus is on defining metrics for the routing scheme and evaluating their effectiveness in terms of load balancing. The authors of [7] propose and evaluate a two-tier architecture, where MTs exploit cellular base stations as gateways to gain access to the fixed Internet and use an ad hoc (IEEE 802.11-like) wireless technology to route data via multihop to the gateways. This solution deals with probabilistic routing, where routing probabilities are calculated based on link costs. The whole line of research dealing with QoS routing, a survey on which can be found in [8], aims at modeling routing strategies as solutions of a constrained optimization problem, which can be solved through either exact or approximate algorithms.

As to the QoS modeling issue, [9] shows that the efficiency of a given routing strategy in terms of QoS very much depends on factors such as mobility, propagation parameters, application requirements, and so on. Furthermore, it is sometimes useful to capture QoS for wireless applications with a more general but still accurate formulation than just referring to throughput or delay requirements. To this end, we may consider utility functions, which are a concept borrowed from economics but also recently used for analyzing wireless communication systems. In particular, in [10] we presented an original model of *backward utilities* to describe quality degradations that occur dynamically because of terminal mobility or changes in the network operating conditions (e.g., a given RT becomes unavailable).

In this article we review and merge several ideas coming from the above contributions and discuss how QoS issues can be practically accounted for in heterogeneous wireless networks. We argue that, with respect to completely distributed and purely ad hoc networks, in our hybrid scenario it is sometimes worth introducing some proactivity in the routing scheme to enhance performance. We show, through example results, that proactivity may be beneficial to improve latency, throughput, and overhead (and thus also energy consumption). In addition, we discuss a utility-based framework to account for variability in the perceived QoS.

Routing Strategies and QoS Issues

Suitable routing strategies to transmit flows over heterogeneous networks should not only find any possible route (provided that multiple solutions are available), but also check whether the user's QoS requirements are met. In a network with multiple RTs and multihop capability, in general, the following situations are possible: either users transmit their data flows through multiple hops to the APs (which have a wired connection to the Internet), by relying on the support of other users who forward their traffic, or users transmit data directly to an AP within transmission range. This can, however, be not directly connected to the wired network; in this case multihop routing may be required to deliver the flow to a wired AP. These two situations might even coexist.

Concerning the general philosophy of the routing strategy, two basic and different approaches exist: proactive and reactive algorithms [11]. Proactive protocols try to keep every node informed with up-to-date information about the routes within the network so that these are known in advance and ready for use. In contrast, reactive protocols invoke the route search on demand (i.e., only when needed). In the present treatment, and according to the classification in [11], we consider so-called *flat routing strategies*, where all nodes have the same role. Our description can easily be extended to account for *hierarchical* and *geographic* routing, where different roles are either assigned to particular nodes or determined exploiting location information, respectively.

For general ad hoc networks, it is commonly believed that proactive schemes offer lower delay but are impractical, because of their usually large signaling overhead. Due to network mobility, it is also expected that route information often becomes obsolete. Thus, one should use reactive protocols in order not to waste network capacity. However, we claim that this is true for distributed and uniform networks, where any two nodes are potentially interested in communicating with each other. In our service provisioning scenario, however, this seems less of a problem. In fact, if the main task of the routing algorithm is to forward packets to/from any of the available APs, as we show shortly, the overhead introduced by proactive solutions may be acceptable. Also, in these settings reactive solutions may become less effective because of their long evaluation time. Consider the case where, due to the sudden absence of compatible RTs, a given MT is unable to communicate with any of its neighbors. Even though neither proactive nor reactive strategies can solve this problem, the former schemes can at least detect it immediately, whereas the latter need more time to even become aware that such a situation has arisen.

As for QoS awareness, in [12] it was shown that the hop count (HC) metric is a simple but effective criterion for routing over distributed networks. In this article we discuss and analyze a proactive routing scheme where HC is used together with QoS-based utility functions to accounti for both link quality and node congestion. A key point of heterogeneous networks is that they are extremely variable not only in access, but also in the services users enjoy and their requirements. For this reason, to accurately represent a user's perceived QoS we might think of a general approach, where utility functions are introduced to model user satisfaction. In particular, a generic value q may be exploited to this end, normalized from 0 to 1, where 1 means a good link with perfect QoS and 0 represents a completely unusable link. For instance, utility values may describe link qualities as well as the congestion status of the route, whose estimation can be initiated at the AP from, say, power measurements or number of detected collisions. These issues are treated in greater detail later. If these techniques are used, a proactive solution is also indicated, as AP-initiated congestion detection makes it possible to solve the bottleneck problem at the destination (APs): severe congestion may be promptly propagated to the nodes, thus suggesting the selection of alternative paths.

QoS and Multi-Technology-Aware Routing Algorithms

According to the above discussion, we let the APs advertise their presence in a proactive fashion, periodically sending routing information (e.g., about the link quality to get to them).

A practical routing algorithm for heterogeneous networks might therefore work in the following way, which more or less follows the framework of Destination-Sequenced Distance Vector (DSDV) [13]. Each AP periodically sends a message, called a *beacon*, including its own identifier, a sequence number (increased sequentially to detect stale beacons), and an HC number (initially set to one). In addition, to account for QoS issues, each AP estimates the capacity still available (*AP congestion level*) and includes it in its beacons. For a UMTS AP, it is reasonable to estimate this residual capacity with reference to the maximum bit rate it can allocate to a new data connection, whereas for the IEEE 802.11b RT such an estimate may be derived by calculating the fraction of idle time multiplied by the effective transmission rate (and scaling it down by a factor to account for protocol overhead).

Upon receiving a beacon from a given RT, a node looks for the identifier of the AP that generated the message and for its HC number. It subsequently checks in its local routing table whether a beacon from the same AP with either a higher sequence number or an HC number equal to or smaller than this exists. If so, the beacon is discarded. Otherwise, the beacon is referred to as *fresh*, and the MT stores the information carried in the beacon in its local routing table (including the AP identifier, the AP congestion level, and the sequence and HC numbers). Fresh beacons are rebroadcast by MTs (or non-gateway APs acting as relays) after increasing their HC field by one and including in the messages the *local congestion level* they currently experience. Also, a refresh mechanism is introduced to remove old entries from the routing tables.

Long-range RTs (e.g., UMTS) may use single hop communication only, but they distribute their beacons to all MTs so as to provide feedback about the APs' congestion status. Short-range RTs (e.g., IEEE 802.11b) still propagate their beacons network-wide. In addition, they may also exploit their multihop routing capabilities. In our specific scenario (Fig. 1), only the MTs that own IEEE 802.11b rebroadcast the beacons to build a multihop tree, which is used to route data packets back to the originating AP. In addition, in order to reduce the collision probability, beacons are retransmitted after a random time interval, accounting for the average number of interfering neighbors within range.

Next, we present a refined beacon propagation policy referred to here as *QoS control*. Upon receiving a beacon, the MTs get an indication of its signal strength from the physical layer and use it to derive a quality parameter q, corresponding to the reliability of the link over which the beacon was transmitted (e.g., its packet error rate). These measurements are used within a

A key point of heterogeneous networks is that they are extremely variable not only in the access, but also in the services the users enjoy and their requirements. For this reason, to accurately represent the user's perceived QoS we might think of a general approach, where utility functions are introduced to model the user satisfaction.

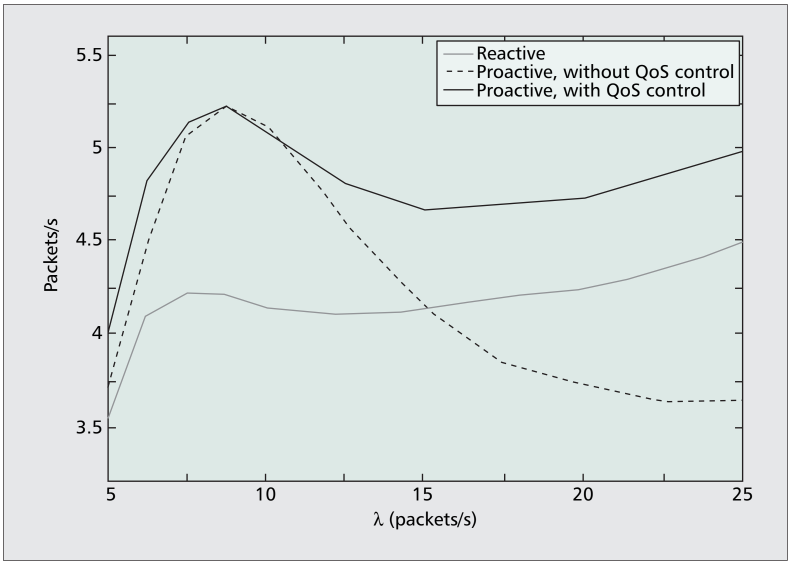

Figure 2. *Average throughput as a function of traffic load.*

simple threshold mechanism to prune unstable (weak) links. Specifically, links are no longer used for forwarding if their quality parameter q is worse than a target quality q_T. For simplicity, we focus here on a single fixed value of q_T, as we simply want to show that pruning weak links considerably improves performance. A deeper discussion could involve the selection of more values of q_T, as well as the investigation of different strategies (threshold policies may lead to instability due to time variability of the links, caused by, e.g., channel fading, interference, mobility). These issues can be interesting directions for further research.

For routing, we adopt the following strategy: each MT checks its own routing table by excluding all APs whose AP congestion level does not allow them to accommodate the user's traffic. The MT then selects the AP with the lowest HC among the remaining APs in the table. Once the destination AP is selected, the node picks the next hop for data forwarding by choosing the least congested neighbor who recently received a beacon from this AP. This last decision is made thanks to the *local congestion levels* included in the beacons.

To numerically evaluate the effectiveness of the above QoS-aware routing algorithm, we discuss next some results obtained with a multi-RT event-driven simulator. The transmission channel is modeled by accounting for path loss, shadowing and multi-path fading. Signal-to-interference-plus-noise ratios (SINRs) are tracked for every received bit according to users' transmission activities and link gain matrices. We consider a scenario with two RTs, IEEE 802.11b and UMTS, both carefully modeled. For IEEE 802.11b we adopt a rate adaptation strategy similar to RBAR [14]. The choice of these technologies in our case study is not restrictive as it allows us to understand the impact of their different characteristics on both performance and protocol design. MTs are randomly scattered at the beginning of the simulation within an area of 240×240 m^2. We consider a single UMTS AP, placed in the center of the area so as to give coverage to all users, and four IEEE 802.11b APs, randomly placed with the constraint of being at least 70 m apart from each other. We account for 35 MTs: 10 of them are equipped with both RTs, while the 25 remaining users are equipped with IEEE 802.11b only. This gives an average of 4.6 users in the IEEE 802.11b coverage area, which is enough to perform relaying when needed. Only 18 nodes (active users) generate data traffic, whereas the remaining devices only perform relaying. For source nodes, traffic is generated according to a Poisson generator whose intensity is λ packets/s/node. Packets have a constant size of 512 bytes. Finally, beacons are propagated from each AP to the MTs every 5 s.

We compare the performance of three routing algorithms. The first two exploit the AP-initiated and proactive strategy outlined above. In the first scheme, *with QoS control*, beacons are propagated to all nodes, and links are pruned if their error rate q is higher than $q_T = 0.01$. The second proactive scheme, *without QoS control*, does not prune weak links. For comparison, we also simulate a purely reactive strategy following the principles of dynamic source routing (DSR) [15].

Figure 2 shows the average throughput experienced by active users. At low traffic loads, proactive algorithms give similar performance and outperform DSR. At high loads, instead, the proactive scheme without QoS control performs poorly, since not all paths found by the proactive procedure are stable as they may include low-quality links. Such a problem is more evident under the high traffic regime, where multi-user interference is higher. Note that unstable links are likely to be pruned by DSR as well, thanks to the two way procedure it uses to set up routes [15]. For this reason, the performance of DSR is not reduced much as the traffic load increases. Also, from Fig. 2 we observe that after an initial decrease, the throughput increases for increasing λ for both DSR and the proactive scheme with QoS control. In fact, for high traffic loads the interference is high, and in turn, only a few users, usually placed close to the APs, can still successfully send their data. In addition, these users usually experience high data rates as they do not have to share the channel with data coming from MTs located further away, who now experience heavy interference. In practice, fairness is substantially decreased in these cases [4].

Figure 3 shows that proactive routing also provides substantial improvements in end-to-end delay, the main reason being that reactive schemes employ route discovery (RD) procedures to find a route to a given AP. This introduces additional delay and also increases system interference, thereby reducing the bandwidth available for data transmission. For the reactive case, the instantaneous data rate frequently drops to zero due to RD activation. This clearly reduces both throughput and delay performance. However, at very low λ (not shown in the figures), reactive solutions are the best option. Hybrid schemes may therefore be designed, where the APs estimate the traffic load and switch between proactive and reactive mechanisms to achieve the best performance under any setting.

In Fig. 4 we plot the overall number of bits per second sent by a generic MT to set up and maintain routing paths (overhead) as a function of the traffic load λ. Note that the overhead is tightly related to the energy consumption and that, for the considered values of λ, proactive solutions achieve better trade-offs. Also, the overhead is a decreasing function of λ. In fact, higher λ leads to more network congestion which, in turn, allows for the propagation (network wide) of a lower number of routing messages.

To sum up, QoS control effectively improves the performance. In addition, especially when the heterogeneous network comprises terminals with very low capabilities, pruning links with poorer QoS reduces interference and increases the reliability (and stability) of multihop routes.

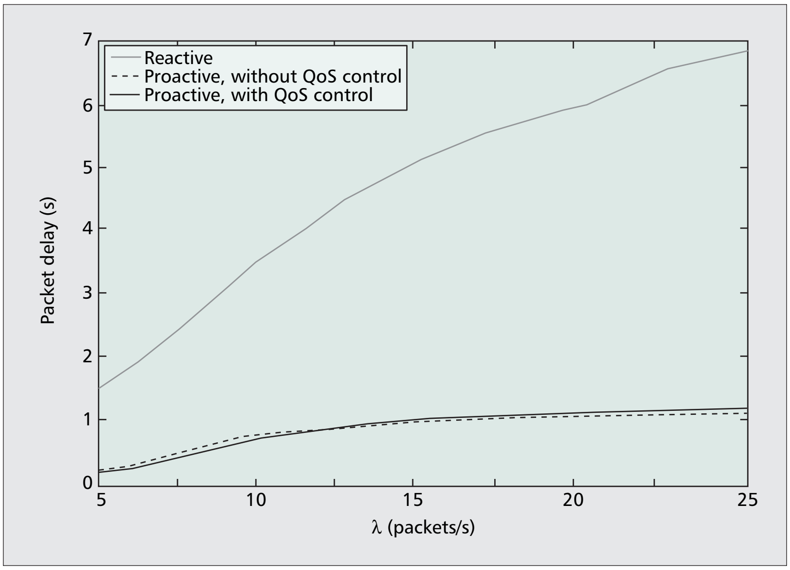

■ **Figure 3.** *Average end-to-end (MT → AP) packet delay.*

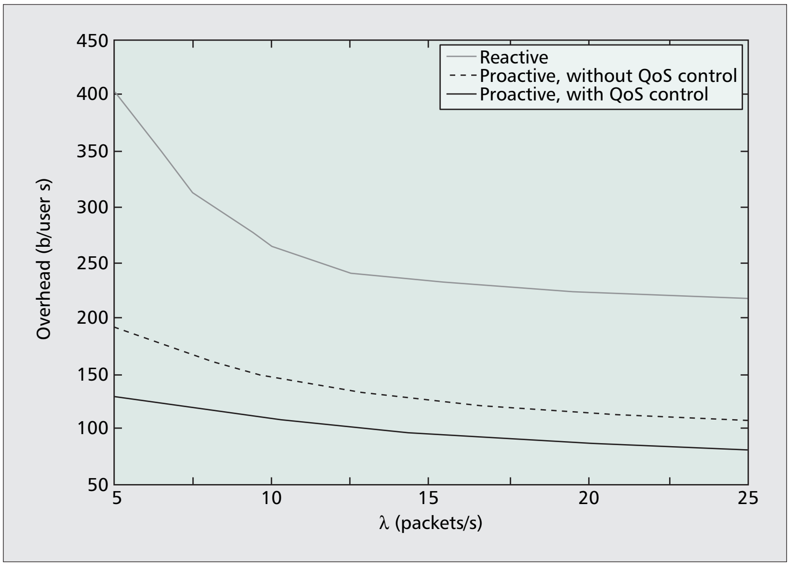

■ **Figure 4.** *Overhead (energy consumption) as a function of traffic load.*

Dynamic QoS and Backward Utilities

The mechanism outlined in the previous section relies on quantities that affect QoS. However, the quality issue may need to be addressed in more detail. In fact, applications may also require QoS stability, and it is sometimes possible that degradations occur while a connection is ongoing, which results in an unsatisfactory service. This phenomenon is very likely to happen in a heterogeneous network (e.g., due to mobility or unavailability of compatible RTs within range).

Such a degradation is not visible at the routing level, since it is application-dependent. However, we might think of operating in a cross-layer fashion so that a smart routing strategy does not forward packets if they do not provide satisfactory QoS to the user. This can be included in the routing framework described above, by adding another QoS control in the beacon selection at the sources, where we try, in addition to the threshold mechanism that prunes weak links, to guarantee a flow-based *stable* QoS. It is worth noting that in this way the actual throughput is decreased, but this would happen anyway, since users experiencing unstable and hence unsatisfactory QoS are likely to terminate their connections before completion. Next, we outline a possible way to effectively model this. This QoS characterization can be validated more deeply in future research.

In general, if service quality is improved, we can think of the users' perception as more or less constant, as it is not easy for the service to exploit the increased capacity when the connection is already established. This is especially true if the QoS perception at the application layer is verified at the beginning of the connection to either accept or refuse the service. On the other hand, if a quality parameter of the route decreases during service, this will make the service less valuable. Hence, it makes sense to consider a lower QoS value when degradations occur after the service has already started. For certain services, we could even consider degradation to be worse when it happens toward the end of the connection, whereas in other cases the opposite may be true. For the sake of simplicity, we limit ourselves to considering the *quantitative amount* of degradation by neglecting its temporal occurrence.

We speak of *backward utility*; that is, we assign a lower score to decreasing quality [10]. The quality evaluation is determined at first as a *forward utility* value, which describes the QoS according to a priori criteria. A sudden degradation of service quality such as a forced handover from IEEE 802.11b to UMTS (thus experiencing lower transmission rates) is accounted for with a backward utility value that is lower than the forward one. We model the added annoyance of QoS degradation through quantities called *quality loss parameters*. A simple characterization accounts for a linear scaling of the quality according to a loss parameter L. Its value can, for example, be seen as the relative weight of the two different events of being served at first with a given quality level or experiencing degradation to that same quality during an ongoing connec-

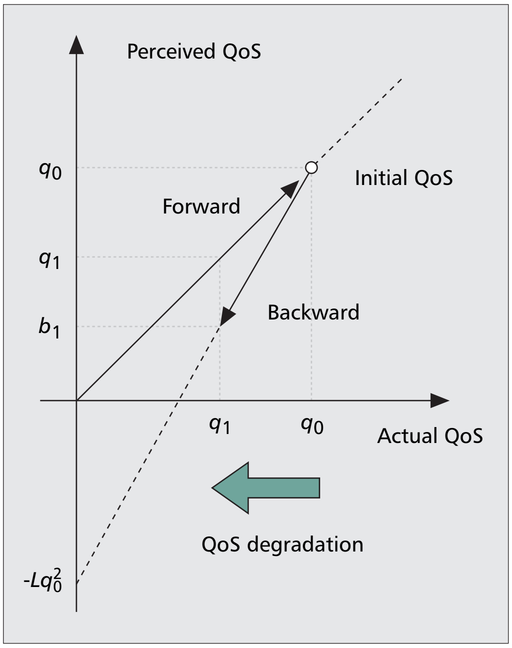

■ **Figure 5.** *Forward and backward utilities.*

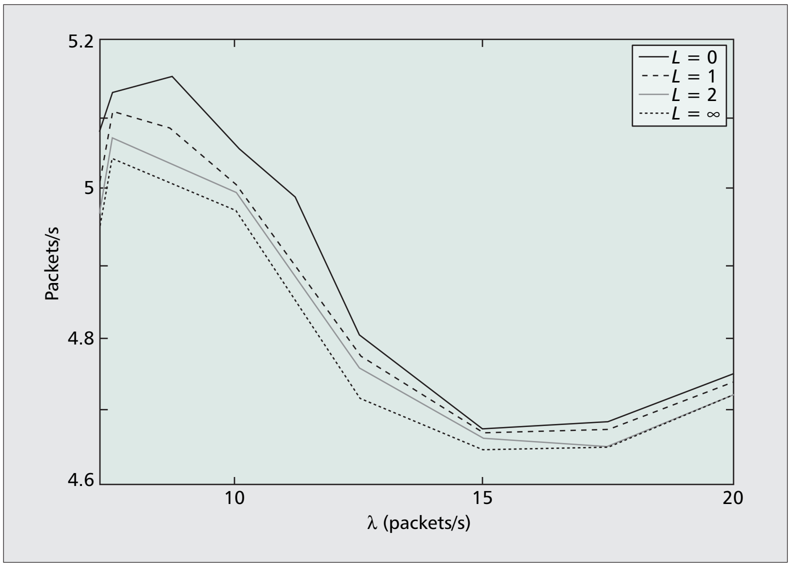

■ **Figure 6.** *Throughput by proactive strategy with admission control based on users' perception.*

tion. This generalizes the well-known trade-off between blocking and dropping probability in admission control, or between performance guarantee vs. best effort service [10]. For the sake of general evaluation, we proceed as follows. We consider that a degradation of the actual quality from q_0 to $q_1 < q_0$ results in lower perceived quality (backward utility) b_1, where $b_1 = q_1 - Lq_0(q_0 - q_1)$ (Fig. 5). This means that in the threshold comparison of the beacon quality parameters, we actually account for a lower perceived quality in case of degradation.

The exact amount of decreased service perception depends on the previous quality q_0 (the higher q_0, the higher the disappointment) and the parameter L. When $L = 0$, the users are insensitive to QoS degradations, so we can speak of *robust* QoS. Instead, if $L > 0$ we have *fragile* QoS (i.e., the perceived quality is lower if the current assignment results from degradation). An infinitely fragile QoS ($L = \infty$) means that the perceived quality suddenly drops to zero as any degradation occurs, no matter how small. A correct evaluation of L would require detailed investigation of the service perception.

In order to quantify the effect of introducing QoS dynamic awareness into the routing strategy, we can repeat the analysis of an earlier section under the backward utility framework. Note that we do not adapt the quality threshold q_T, but rather we evaluate the quality of the connection accounting for the backward utility instead of the forward one. The important result, emphasized in Fig. 6, is that the selection of the routes based on a users' perceived QoS, which also takes into account QoS degradation, does not imply a dramatic decrease of throughput, which is only slightly affected. On the other hand, from the application standpoint, the backward utility concept is beneficial in that it avoids strongly dissatisfactory QoS. That is, Fig. 6 does not indicate different strategies to choose from, but rather shows that proper QoS control should apply link selections also according to users' perception of degradation; from the point of view of raw throughput, the loss is only marginal, but a strong impact on the application layer can be expected. This is not trivial to quantify but can be assumed to be more significant, as it involves the perceived QoS of the whole connection.

Conclusions

In this article we review routing solutions to provide coverage extension in wireless heterogeneous networks. A strong integration of different access techniques is key to guarantee portability, scalability, and fairness, which result in overall user satisfaction. However, this can be achieved only with the support of QoS-aware efficient protocols. We argue that periodic advertisements from APs are beneficial in making network management aware of the quality that can be supported. Also, we discuss selection procedures for route advertisement and propose a utility-based approach, easily tunable to different objectives, to model the QoS of routing paths. Furthermore, to account for dynamic changes in the offered QoS, we outline a *backward utility* framework that could make the model even more realistic.

We observe that existing theoretical work does not entirely capture all the details of the considered scenario. Here, we partially address this problem by discussing practical schemes to effectively provide a satisfactory QoS in heterogeneous wireless networks. We stress that further research work is required to achieve complete and scalable solutions, also addressing the problems of service/access technology discovery and user behavior (e.g., altruistic vs. selfish) in a possibly competitive scenario.

References

[1] A. Furuskär and J. Zander, "Multiservice Allocation for Multiaccess Wireless Systems," *IEEE Trans. Wireless Commun.*, vol. 4, no. 1, Jan. 2005, pp. 174–84.

[2] R. Bruno, M. Conti, and E. Gregori, "Mesh Networks: Commodity Multihop Ad Hoc Networks," *IEEE Commun. Mag.*, vol. 43, no. 3, Mar. 2005, pp. 123–31.
[3] H. Zhu *et al.*, "A Survey of Quality of Services in IEEE 802.11 Networks," *IEEE Wireless Commun.*, vol. 11, no. 4, pp. 6–14, Aug. 2004.
[4] I. Aad and C. Castelluccia, "Differentiation Mechanisms for IEEE 802.11," *Proc. IEEE INFOCOM*, Anchorage, AK, Apr. 2001, pp. 209–18.
[5] M. Rossi *et al.*, "On the Effectiveness of Logical Device Aggregation in Multi-Radio Multi-Hop Networks," *Proc. 3rd IEEE Int'l. Wksp. Mobility Management and Wireless Access*, vol. 1, Maui, HI, June 2005, pp. 354–61.
[6] S. Nelakuditi *et al.*, "Adaptive Proportional Routing: a Localized QoS Routing Approach," *IEEE/ACM Trans. Net.*, vol. 10, no. 6, Dec. 2002, pp. 790–804.
[7] C.-F. Huang, H.-W. Lee, and Y.-C. Tseng, "A Two-Tier Heterogeneous Mobile Ad Hoc Network Architecture and Its Load-Balance Routing Problem," *ACM Mobile Networks and Apps.*, vol. 9, no. 4, Aug. 2004, pp. 379–91.
[8] F. Kuipers *et al.*, "An Overview of Constraint-Based Path Selection Algorithms for QoS Routing," *IEEE Commun. Mag.*, vol. 40, no. 12, Dec. 2002, pp. 50–55.
[9] P. Johansson *et al.*, "Scenario-Based Performance Analysis of Routing Protocols for Mobile Ad Hoc Networks," *Proc. ACM MobiCom*, Seattle, WA, Aug. 1999, pp. 195–206.
[10] L. Badia and M. Zorzi, "On Utility-Based Radio Resource Management With and Without Service Guarantees," *Proc. ACM MSWiM*, Venice, Italy, Oct. 2004, pp. 244–51.
[11] X. Hong, K. Xu, and M. Gerla, "Scalable Routing Protocols for Mobile Ad Hoc Networks," *IEEE Network*, vol. 16, no. 4, July/Aug. 2002, pp. 11–21.
[12] R. Draves, J. Padhye, and B. Zill, "Comparison of Routing Metrics for Static Multi-Hop Wireless Networks," *Proc. ACM SIGCOMM*, Portland, OR, Aug./Sept. 2004, pp. 133–44.
[13] C. E. Perkins and P. Bhagwat, "Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for Mobile Computers," *Proc. ACM SIGCOMM*, London, UK, Aug./Sept. 1994, pp. 234–44.
[14] G. Holland, N. Vaidya, and P. Bahl, "A Rate-Adaptive MAC Protocol for Multi-Hop Wireless Networks," *Proc. ACM MobiCom*, Rome, Italy, July 2001.
[15] D. B. Johnson, D. A. Maltz, and J. Broch, "DSR: The Dynamic Source Routing Protocol for Multi-Hop Wireless Ad Hoc Networks," Ch. 5, *Ad Hoc Networking*, C. E. Perkins, Ed. Addison-Wesley, 2001, pp. 139–72.

BIOGRAPHIES

LEONARDO BADIA (leonardo.badia@gmail.com) received a Laurea degree (with honors) and a Ph.D. in information engineering from the University of Ferrara, Italy. He had collaborations with the University of Padova, Italy, and Wireless@KTH, Royal Institute of Technology, Stockholm, Sweden. Currently, he is with the Institutions Market Technologies (IMT) Institute for Advanced Studies, Lucca, Italy, where he is a research fellow. His research interests include ad hoc and mesh networks, transmission protocol modeling, and radio resource management.

MARCO MIOZZO received his Laurea degree in telecommunication engineering from the University of Ferrara in 2005. He is currently working for the Consorzio Ferrara Ricerche (CFR) and the Department of Information Engineering of the University of Padova. He is part of the European Ambient Networks project. His main research interests are in performance analysis and simulation of multitechnology-multihop wireless networks and new-generation wireless systems.

MICHELE ROSSI (mirossi@gmail.com)received both a Laurea degree and a Ph.D. in information eengineering from the University of Ferrara in 2000 and 2004, respectively. During 2003 he was on leave at the Center for Wireless Communications at the University of California at San Diego (UCSD). Since November 2005 he is an assistant professor at the Department of Information Engineering of the University of Padova, Italy. His research interests include channel access, routing, and data distribution algorithms for wireless networks.

MICHELE ZORZI received a Laurea degree and a Ph.D. in electrical engineering from the University of Padova in 1990 and 1994, respectively. During academic year 1992–1993, he was on leave at UCSD, attending graduate courses and doing research on multiple access in mobile radio networks. In 1993 he joined the faculty of the Dipartimento di Elettronica e Informazione, Politecnico di Milano, Italy. After spending three years with the Center for Wireless Communications at UCSD, in 1998 he joined the School of Engineering of the University of Ferrara, Italy, where he became a professor in 2000. Since November 2003 he has been on the faculty at the Information Engineering Department of the University of Padova. His present research interests include performance evaluation in mobile communications systems, random access in mobile radio networks, ad hoc and sensor networks, energy constrained communications protocols, and broadband wireless access. He was Editor-In-Chief of *IEEE Wireless Communications*, 2003–2005, and currently serves on the Editorial Boards of *IEEE Transactions on Communications*, *IEEE Transactions on Wireless Communications*, Wiley's Journal of Wireless Communications and Mobile Computing, and ACM/URSI/Kluwer *Journal of Wireless Networks*. He has also been a Guest Editor of special issues in *IEEE Personal Communications* (Energy Management in Personal Communications Systems) and *IEEE Journal on Selected Areas in Communications* (Multimedia Network Radios).