A Framework for the Admission Control of QoS Multicast Traffic in Mobile Ad Hoc Networks*

Elena Pagani
Computer Science Dept.
Università degli Studi di Milano
via Comelico 39
I-20135 Milano, Italy
pagani@dsi.unimi.it

Gian Paolo Rossi
Computer Science Dept.
Università degli Studi di Milano
via Comelico 39
I-20135 Milano, Italy
rossi@dsi.unimi.it

ABSTRACT
Recently, QoS issues have initiated to be studied in both wired and wireless networks, to support multimedia and real-time applications. In this paper, we propose a framework for the admission control of multimedia multicast traffic and for the system configuration in MANETs. We present a mechanism to ensure bandwidth guarantees to multicast sessions (*Call-Admission Multicast Protocol for* MANETs, M-CAMP). M-CAMP is scalable, operates on a per-call basis and supports the group membership dynamics. It adopts a measurement-based approach to evaluate the end-to-end bandwidth availability between the traffic source and the group of destinations. M-CAMP is independent of the underlying wireless technology and protocols, as far as a multicast routing service is available. It does not require any maintenance of status information in the mobile hosts.

Categories and Subject Descriptors
C.2 [**Computer Systems Organization**]: Computer - Communication Networks

General Terms
Algorithms

Keywords
Quality-of-Service, admission control, multicast

1. INTRODUCTION
The recent advances in the cellular telephony technology make feasible the connection of mobile terminals to the Internet, allowing the mobile users to access distributed applications. The natural evolution is towards the support of those applications in mobile ad hoc networks (MANETs). A main research field in the Internet engineering concerns the characterization of policies and mechanisms to provide a proper level of *Quality-of-Service* (QoS), to support multimedia and real-time applications. According to the trend of porting those applications from wired to wireless networks, QoS issues must be investigated in MANETs as well. Some works have been recently proposed in the literature concerning these problems. Many of the proposed approaches are tailored to the peculiar characteristics of a particular wireless technology, or focus on unicast communications. Yet, many multimedia applications are characterized by a multicast communication pattern. Moreover, in the future many wireless technologies will be contemporarily available in the same MANET, with mobile hosts (MHs) dynamically switching amongst them to adapt to the current environment and achieve the best performance [19].

In this paper, we propose a framework for the admission control of multimedia multicast traffic and for the system configuration in MANETs. We present a mechanism to ensure bandwidth guarantees to multicast sessions (*Call-Admission Multicast Protocol for* MANETs, M-CAMP). M-CAMP is scalable, operates on a per-call basis and supports the group membership dynamics. It adopts a measurement-based approach to evaluate the end-to-end bandwidth availability between the traffic source and the group of destinations. M-CAMP is independent of the underlying wireless technology and protocols, as far as a multicast routing service is available. It does not require the maintenance of status information in the MHs.

The paper is structured as follows: in the next section, we introduce the MANET environment. In section 3, we present the framework we consider to provide QoS to multimedia applications, while in section 4 we describe M-CAMP. In section 5, we give a survey of the proposals so far appeared in the literature for the QoS support in MANETs, and we discuss the behaviours of M-CAMP in comparison with INSIGNIA [16, 1].

2. THE SYSTEM MODEL
The system we consider is composed of N mobile hosts, or MHs for short, with unique identifiers p_1, ..., p_N, that communicate via a packet radio network. We assume that all the MHs have similar computing and storage resources;

*This work was supported by the MURST under Contract no. MM09265173 "Techniques for end-to-end Quality-of-Service control in multi-domain IP networks".

they act as routers.

The system model we propose to provide QoS in MANETS derives from the wireless network architecture described in [19]; in figure 1 we show the MANET architecture at different layers. At the MAC layer (figure 1(a)), the system is composed of mobile nodes connected via multiple wireless technologies (represented by the different dashed lines). The MHs may be equipped with either omni-directional or highly-directional antennas. The wireless links may be unidirectional. The data-link layer protocols manage the code/frequency assignment and they prevent/solve possible frame collisions. At the IP layer (figure 1(b)) the system is composed by the mobile routers connected via the paths characterized by the routing service. The network and data-link protocols have the task of appropriately forwarding the packets along those paths. At the end-to-end layer (figure 1(c)) the applications exchange data, exploiting the services of the QoS modules. Figure 1 represents our architectural choices. In the next section, we detail the functional components providing QoS guarantees at layer 1(c).

Actually, in the literature, two approaches have been proposed to provide QoS: either exploiting QoS-aware routing protocols (layer 1(b)) [18, 7, 25], or add QoS functionalities on top of the existing routing service (layer 1(c)) [16, 1]. The latter approach allows to have lower overhead and better interoperability amongst heterogeneous routing domains. Moreover, different routing protocols behave differently in response to topology changes: with the latter approach the QoS modules can guarantee greater homogeneity in the service provisioning, although they have to properly behave in spite of heterogeneous underlying multicast services. Multicast routing protocols can be classified into two families, depending on how they deal with the tree disconnections:

1. receivers are unaware of the reconfiguration: with these protocols, the root of a disconnected subtree tries to reconnect to the tree (or, viceversa, the upstream node of a broken link tries to reconnect to its downstream node via an alternative branch). Protocols of this family are for instance AODV [22] and DDM [13];

2. receivers drive the reconfiguration: with these protocols, the disconnected subtree is flushed, and the receivers, informed of the tree break, must re-join the group. A protocol of this family is for instance AMRIS [27].

The mechanisms we propose in this paper work correctly with both types of routing service.

Multimedia applications require particular guarantees from the network with respect to the data transmission and delivery service. A set of functionalities must be available in the MHs, to characterize suitable paths, estimate the available resources along those paths, reserve the needed resources and configure the network elements traversed by the multimedia traffic, so that the applications are supplied with the requested service level. For the sake of simplicity, we assume that all the MHs belong to the same *domain*, that is, they operate with a common set of service differentiation models, provisioning policies and service definition. QoS is provided according to the *Diff-Serv* model [3].

MHs may move around the system, or they may turn off their network interface. Movements and disconnections dynamically change the network topology. Yet, if topology changes are frequent, QoS could become impossible to provide: the paths between two MHs change before they can be characterized or configured to support the QoS traffic. For this reason, in this work we require that the system satisfies a *mobility assumption*, according to which a path eventually exists between two nodes, and it does not change for the time needed to exchange $k+1$ packets, with k a given (small) constant[1]. We consider a system where hosts do not suffer from permanent failures: all types of failures are temporary, although it is not specified how long the faulty condition will last. Both host disconnections and network partitions are tolerated, under the assumption that they are soon or late repaired. Group membership changes occur when hosts voluntarily join or leave the group. These assumptions do not specify any time constraint to mobility, nor they give some upper bound to the disconnection time of a given host. The path stability is not required to be fulfilled contemporarily for all the pairs of MHs. The assumptions simply aim at reducing the complexity of the framework description; their role is explained in section 4.

3. ARCHITECTURE FOR QOS IN MANETS

In figure 2, we show the QoS framework architecture considered in this work. We hypothesize that all the MHs use the priority packet scheduling policy. We assume that unicast and multicast routing services are available that find a path if any, or inform of the disconnection otherwise. We do not make any assumption about the quality-of-path metrics considered by those protocols; the adopted metrics could affect the performance of the system, but not its correctness.

The sources generate Constant-Bit-Rate (CBR) traffic (e.g., CBR Mpeg-encoded video, audio without silence suppression). We assume that a session announcement protocol (e.g., `sdr` [11]) is used to announce the needed session information, such as the transmission start time and the multicast address of the destination group $\mathcal{G}$. Receivers perform an explicit join procedure.

The QoS modules are located between the application layer and the routing services. The sources of multimedia traffic require to those modules the set-up of a QoS session. The source specifies the traffic profile T_s it wants to generate, described for instance as a `TSpec` [24], and the list of receivers. The multimedia traffic may be formed by multiple flows (e.g., an audio flow and a video flow). All the recipients receive the same set of flows, that is all the traffic generated by the source for the group. M-CAMP supports *homogeneous* recipients, having the same QoS requirements. The QoS modules use the transmission specification to evaluate the resources needed to manage the QoS traffic so as to guarantee that it is received at the destinations in a suitable form.

According to the previous considerations and the QoS frame-

[1]The meaning of k will be explained in the sequel.

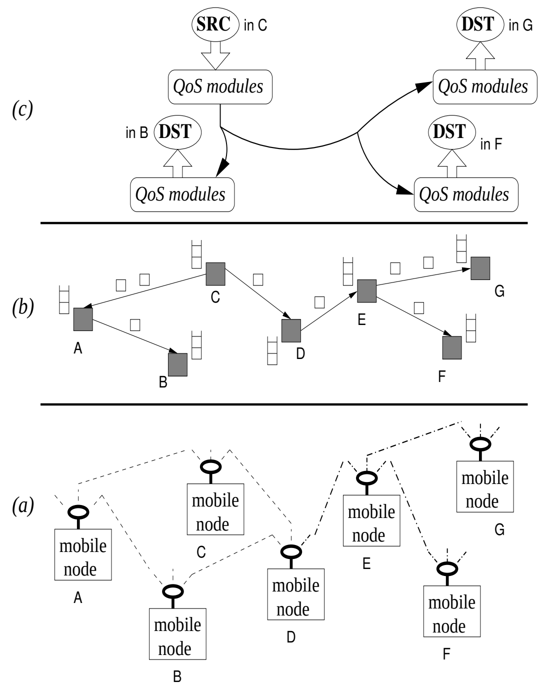

Figure 1: MANET architecture.

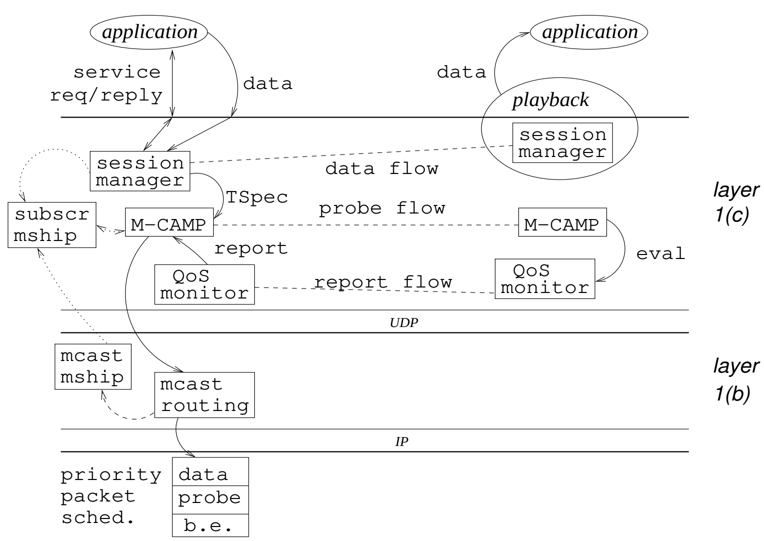

Figure 2: Layout of the considered QoS architecture.

work architecture proposed in [3], our architecture involves the following modules at the end-to-end layer:

session manager: it provides an abstraction of the underlying QoS implementation. It may configure the source and destination hosts. At the source, it ensures that the messages are appropriately marked, so that the the transmitted information may be re-built at the destinations (for instance, it labels the messages with timestamps that allow the re-synchronization of the different medias). It also performs the traffic shaping, so that the traffic produced by the source conforms with the claimed profile. At the destination, buffer space may be reserved and the system may be configured so that the messages are delivered to the destination application so as to *play back* the transmission generated by the source;

QoS monitor: at the destination, this module measures the received QoS, and reports it to the source. At the source, the reports may be used to dynamically re-configure the system. They may be forwarded to the source application, that may adapt to the current receiver status, for instance by dynamically changing the data encoding;

admission control and reservation: this module has in charge the estimation of the available resources and their reservation along the multicast routing tree.

The functionalities of the first and second modules are partially performed by the RTP/RTCP protocols [23]. In the next section, we propose a mechanism to provide the functionalities of admission control and system configuration. This mechanism operates on an on-demand basis. It is independent of the lower layer protocols and works on top of the existing multicast routing service. It configures the in-tree MHs without requiring any maintenance of status information at the nodes.

4. CALL ADMISSION MULTICAST PROTOCOL FOR MANETS

In this section we describe the end-to-end *Call Admission Multicast Protocol* (M-CAMP), that can be used to ensure bandwidth guarantees to multicast sessions in MANETS. M-CAMP is scalable, operates on a per-call basis and supports the group membership dynamics. It performs the call set-up of a multimedia session. It performs a test to verify whether enough bandwidth is available along the multicast tree to support the traffic profile specified by the source. It represents a possible implementation of a distributed *bandwidth broker* [20].

In the following, we assume that M-CAMP is implemented as a module of the RTP library, which is linked to the applications. RTP activates M-CAMP by forwarding it the information received by the application. RTCP is used to monitor the QoS provided at the recipients. In this section we firstly describe M-CAMP in the case of both static group membership and static network topology. Then, we outline the mechanisms used to deal with dynamic membership and topology changes due to the host mobility.

4.1 Static Membership and Topology

In this section, we assume that the recipients join the destination group, and the corresponding multicast routing tree is formed, before the start of the multicast session. A multicast session is divided into two phases: the *session set-up* phase, performed by M-CAMP, and the *data transfer* phase. In figure 3 we show the finite state automata for the source and destination MHS; we use the roman type style to represent the events, and the italic type style for the actions performed in response.

In the set-up phase, M-CAMP uses the information supplied by the application to test whether the network can guarantee the requested bandwidth. The test is carried out by sending `probe` packets to the destinations, so that the generated probing traffic has the same profile of the specified data traffic. The `probe`s are forwarded along the multicast routing tree. The basic idea consists of differentiating `probe`, QoS data and best effort packets by means of three priority levels. The `probe` packets are marked with a higher priority than the best effort packets and a lower priority than the QoS data packets. This priority assignment ensures that the probing traffic does not affect the existing QoS flows. On the other hand, `probe` packets can drain the available bandwidth for new QoS flows at the expenses of the best effort traffic.

At the receiver side, the M-CAMP entity evaluates the quality of the received probing traffic (`probing` state). The source diffuses its own traffic profile by using the periodic RTCP reports[2]. After the reception of k samples, each receiver d matches that profile against the received probing traffic profile. If the difference between the two profiles is within an acceptable threshold, whose value depends on the application, the receiver accepts the whole transmission; otherwise, it refuses. d informs the source about its decision with its own RTCP reports. A receiver that refuses the service unsubscribes from the group. As a consequence, the corresponding router may be pruned from the multicast routing tree (transition from the `probing` to the `idle` state).

If the service is accepted by at least one recipient, then the source switches from the transmission of `probe` packets to the transmission of the data packets generated by the source application, *without discontinuity*. Data packets are forwarded along the pruned tree (transition from the `probing` to the `data transfer` state).

The probing phase ensures resource reservation: the bandwidth that has been preempted by a given flow during its probing phase cannot be later assigned to other `probe` or best effort packets. Moreover, priority naturally provides flow aggregation.

The knowledge of the group membership is needed because the source *must* wait for the reports from *all* the receivers, before switching to the data transmission. The example of figure 4(a) explaines this point. The source generates a QoS flow $flow_1$. If it receives positive reports from a subset of receivers, let us say those in the subtree A, it knows that some destination exists for the data flow. If the source switches

[2]The generation of the reports is not represented in the automata.

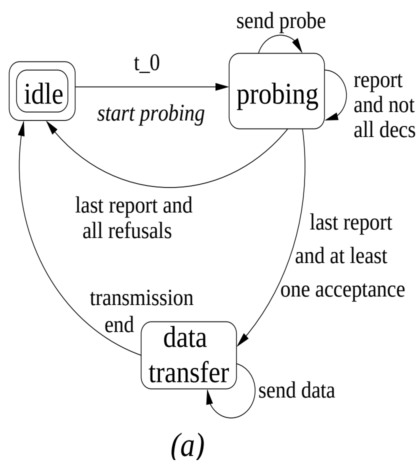

Figure 3: Finite state automata for (a) the source MH, and (b) the destination MH.

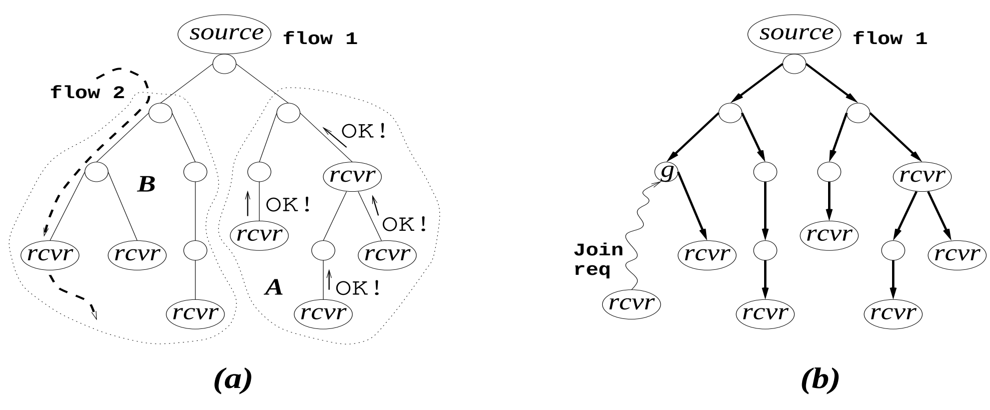

Figure 4: (a) Example of concurrent flows. (b) Example of dynamic group membership.

to the data transmission before receiving the reports from the receivers in the subtree B, other established QoS flows could be negatively affected. In fact, let us suppose that a flow $flow_2$ exists that traverses the B's branches, and does not leave enough bandwidth for $flow_1$. As a consequence, the receivers in B would eventually send negative reports and prune to refuse $flow_1$. If the switch is performed before receiving all the reports (i.e., before the needed prune operations), the data packets of the aggregate flowing on the subtree B will be dropped, independently of the flow they belong to.

The need of obtaining all the decisions requires that report loss is prevented, for instance by adopting TCP as the transport protocol used by RTCP. The source can notice when all the reports have been received, because, by hypothesis, it knows the (static) destination group membership beforehand. The case of dynamic group membership is discussed in the sequel.

4.2 Dynamic Group Membership and Static Topology

In several applications, such as news services and video broadcasting, the source does not need, or desire, to be aware of the receiver group membership, while receivers are free to join and leave the group dynamically. In this new model, the source announces the multicast session via `sdr` and starts transmitting at the scheduled time if at least one receiver is listening[3]. In this case, the M-Camp as described above does not work, because of arguments similar to those discussed above: the multimedia traffic cannot be forwarded along new branches before configuring the MHs (figure 4(b)).

The probing phase cannot be started at the root, to avoid the duplication of the resource reservation already installed for the QoS traffic on the path from the root to g. Rather, it must be started at the graft point g. This seems in contrast with the Diff-Serv model, according to which the routers interior to a domain are QoS-unaware. In fact, our solution does not require that MHs are QoS-sensitive. We install a M-Camp *proxy* in the in-tree nodes involved in membership changes. When a new node d_{new} wants to join a QoS multicast session, it explicitly joins the proper group of receivers and the corresponding routing tree. Let us indicate with $\mathcal{G}$ the group of receivers. As soon as an intree router g, belonging to the tree T for $\mathcal{G}$, creates a new downstream interface for T, it instantiates the M-Camp proxy, say p_g, whose lifetime lasts until the set-up phase for the new host terminates. Once activated, the proxy receives the identifier of the new interface and the list of downstream destinations. It generates, by gathering the information from the local routing table, the data structure of figure 5, where it temporarily maintains the status of the active probing phases. If a join request from a new MH is received while the proxy is active, a new entry of the data structure is created for it.

During its activity the proxy re-marks as `probe` packets all the incoming QoS data packets that should be routed to the probing interfaces. Moreover, it adds its own IP address to the packet. The destination M-Camp entity reliably sends point-to-point its accepting/refusing report to the proxy indicated in the `probes`. When p_g receives one report from a probing downstream destination, if it is a positive report, then the proxy changes the corresponding entry from the state *probing* to *data*. Otherwise, it removes the entry. The proxy turns off when all the entries in the data structure have been deleted (service refused) or switched to data transfer. This simple mechanism solves the problem due to a dynamically changing group membership and achieves the following goals:

1. the "`probe`" packets allow to evaluate the resource availability only along the new branch;
2. the new destinations immediately start receiving the real data flow, although possibly with a lower quality than required;
3. the data sent to the new destinations do not affect previously established flows traversing the new branch, as the data packets are marked and treated as `probe` packets;
4. the membership is hidden to the source. Anyway, it must be known by the proxy because of the same arguments discussed in section 4.1.

Clearly, this mechanism can be exploited at the session initialization as well. In this case, the source is allowed to switch to data transmission as soon as the first acceptance report is received.

4.3 Dynamic Topology

MHs movements and disconnections may partition the multicast routing tree in every moment during the session. A moving MH may later reconnect to the tree at a different graft point. We deal with those topology changes by exploiting the proxy mechanism, slightly modified. In the case the multicast routing protocol belongs to the class *(2)* described in section 2, the reconnecting recipients are conscious of the tree reconfiguration and of the need of re-starting the set-up phase. With protocols belonging to the class *(1)*, we need additional mechanisms to inform the reconnected routers that the probing must be restarted to perform admission control along the new branch. The key point here is to effectively evaluate the bandwidth availability by guaranteeing that a recipient decides about the service acceptance/refusal only after having received a sufficient number k of samples (i.e., `probes`).

If the tree must be reconfigured when the data transmission is ongoing, then the reconnected destinations, that were receiving data packets, start receiving (data packets remarked as) `probes`. This indicates them that a tree reconfiguration has occurred and the set-up must be performed.

If the reconnected destinations were receiving `probes`, to allow them to notice that the `probes` are now covering a different branch we use a *probing sequence number* $ps\#$. In figure 6 we report the finite state automata for the proxy and for the destination MHs in this case.

Initially, the source generates `probes` marked with $ps\# = 1$. Each destination records the $ps\#$ of the `probes` it is currently receiving (transition from the `idle` to the `probing`

[3] That is, if the recipient group exists.

probing interfaces	*state*	*destinations*
output interface 1	`data`	{downstream rcvs}_1
...	...	...
...	...	...
...	...	...
output interface n	`probing`	{downstream rcvs}_n

Figure 5: Data structure maintained by a proxy.

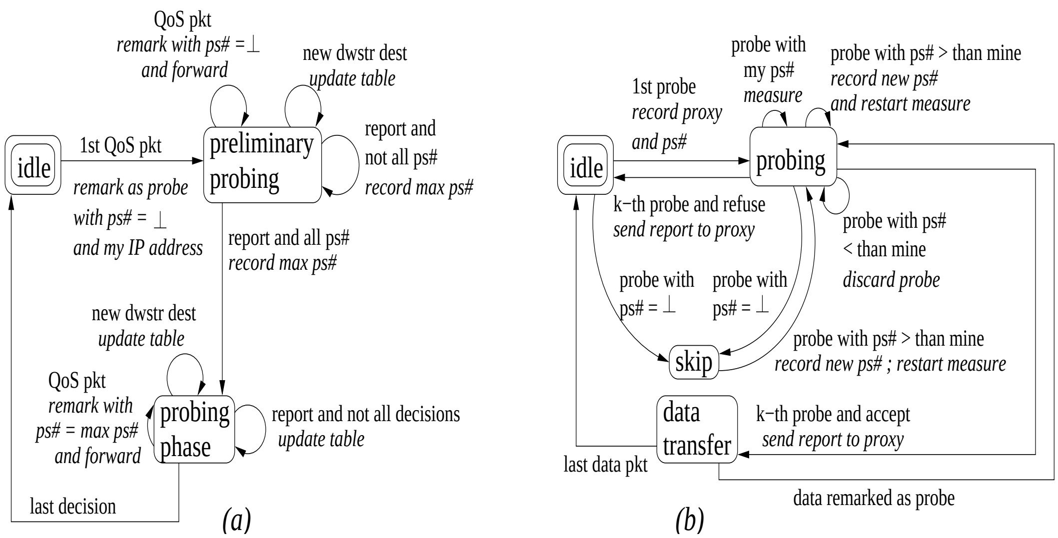

Figure 6: Finite state automata for (a) the M-CAMP proxy, and (b) the destination MH.

state); $ps\#$ is also sent within the RTCP reports to the current proxy as an application-dependent field [23]. The proxy in the graft point initially sends remarked `probe`s carrying $ps\# = \perp$ on the probing output interfaces (`preliminary probing` state). Once the proxy has received a report from each downstream probing destination[4], it computes the maximum among the $ps\#$'s, $max\_ps\#$, and labels the forwarded `probe`s with $ps\# = max\_ps\#{+}1$ (transition from the `preliminary probing` to the `probing phase` state). The destinations do not consider `probe`s with either $ps\# = \perp$ or $ps\#$ lower than the currently recorded value to evaluate the received QoS. If `probe`s with $ps\#$ greater than the currently recorded value are received, the destination records the new $ps\#$ and restarts the probing. It is worth to notice that only the recipients have to maintain status information, represented by the highest received $ps\#$.

In [21], we prove that the described mechanism guarantees that each destination receives at least k `probe`s *along a given path* before deciding. We prove that M-CAMP effectively and efficiently solves the admission control problem; we show that, thanks to the mobility assumption, the eventual termination of the set-up phase is guaranteed, in spite of mobility and dynamic membership. We discuss some implementation issues.

To guarantee the termination in the case of continuous topology changes, a source can stop the probing after αT, with T the data transmission length and $\alpha \in (0, 1]$. If within αT the source does not succeed in collecting all the recipient replies, then there is a highly changing topology, and neither the data transmission could be carried out so as to guarantee the required QoS. The value of α should be chosen according to the trade-off between the probability of setting up the service and an efficient usage of the network resources.

[4]Recall that it must know all the downstream destinations.

5. RELATED WORKS

Recently, some works appeared in the literature studying the QoS problem in MANETs. A survey about QoS support in MANETs is provided by [28].

Some frameworks have been proposed. [6] presents the analysis of the problems that must be dealt with to support QoS in MANETs. The authors outline how to change the proposed routing protocols to provide resource reservation, exploiting a RSVP-like mechanism. In [29], the authors investigate the possibility of providing service differentiation with TCP in MANETs, by means of simulation techniques. In [30], a QoS model is proposed, which combines Int-Serv and Diff-Serv: the differentiation amongst traffic classes is provided by assigning them the currently available bandwidth according to a proportional policy. All these models do not examine multicast issues.

The IETF MANET WG has proposed several routing protocols. Anyway, some of them support multicast routing but not QoS (e.g., [12, 13, 27, 17]). Others support QoS but not multicast: FSR [8] may use functions that compute routes according to different metrics and each node maintains information concerning the path quality in its immediate neighbourhood. IARP [10] can be configured to consider different link quality metrics. In [15, 26, 18], QoS unicast routing protocols are proposed, that heavily depend on the underlying MAC layer (respectively, DS-CDMA, CSMA and TDMA), thus being not easily portable. In [7], a QoS unicast routing protocol is described, which is independent of the underlying services and is multi-path. CEDAR [25] is a QoS unicast routing that characterizes a virtual backbone in the MANET, used for the packet routing. It exploits limited-scope link-state, concerning high-bandwidth stable links, for the route characterization.

By contrast, TBRPF [2] and AODV [22] support both multicast and QoS. TBRPF is a link-state based routing protocol, that in its full topology mode may use the link-state information to compute paths that satisfy QoS constraints. Anyway, its cost is of the order of magnitude $O(N^2)$ with N the number of nodes in the MANET, as all the MHs are involved in the link-state communication. AODV [22] allows to specify delay and bandwidth requirements in the process of characterizing suitable paths. The path requests are flooded through the network.

5.1 Comparison between Insignia and M-CAMP

Among the several protocols proposed in the literature for the MANETS [22, 9, 14], M-CAMP is only comparable with INSIGNIA [16, 1], as both protocols do not perform routing functions but rather add QoS capabilities to the underlying protocols.

Both protocols guarantee the application bandwidth requirements. INSIGNIA uses in-band signaling to perform resource reservation: the source generates RES packets carrying its bandwidth requirements, until it receives the destination reply about the service acceptance. RES packets are checked for admission control at each intermediate node, against the local resource availability. The resource reservation information is maintained by a soft-state mechanism: the state is refreshed by every packet, or deleted upon the timer expiration. The timer value should be chosen according to the flow rate. QoS is supported according to the *Integrated Services* model [5]: a node maintains state information and a timer for each QoS flow that traverses it. Multicast issues are not addressed.

In table 1, we report the main characteristics of INSIGNIA compared with M-CAMP. To exchange control information, INSIGNIA requires the modification of the IP header to accommodate the bandwidth requirements. As control information is carried over every packet, a node has to process all the packets it receives to possibly update either its local reservation state, or the information about resource availability recorded in the header. Moreover, INSIGNIA could not to scale well, as its memory overhead grows with the number of ongoing QoS flows. On the other hand, M-CAMP has a longer set-up phase. This is the cost to be paid for being stateless: the available bandwidth must be estimated anew every time a new service must be established. However: traffic with high rate (such as video) produces a high number of samples (i.e., `probe` packets) in a few seconds. For instance, a Mpeg flow with 1.5 Mbps rate and packet size of 1 KB produces 100 samples in 0.534 secs. With VoIP traffic, having 32 Kbps rate, the probing phase would last 3.125 sec. As a consequence, there is a low probability that topology changes may disrupt the tree structure during those few seconds.

6. CONCLUDING REMARKS AND OPEN ISSUES

In this paper we propose a framework to provide QoS support for multimedia applications in MANETS. We describe a mechanism to perform both the admission control for the multimedia traffic and the system configuration. M-CAMP operates on-demand. It has a distributed control and is independent of the wireless technology, thus supplying QoS support in spite of the system heterogeneity. It exploits the service of an underlying multicast routing protocol and involves mechanisms to cope with mobility and disconnections. It is stateless and adopts a measurement-based approach to estimate the resource availability. It involves only the MHs that must forward the QoS traffic, thus guaranteeing scalability and energy saving.

We have implemented a version of M-CAMP for wired networks in the NS-2 simulation environment [4]; the simulation results are promising. We are currently implementing M-CAMP for the MANET environment in the NS-2 framework, to evaluate the performance of the proposed approach and compare it with INSIGNIA. Some issues remain to be further investigated. One of them concerns how to estimate the length of the set-up phase, that could possibly depend on the lowest rate of the ongoing QoS flows. We are also extending the protocol with mechanisms to deal with different types of traffic (e.g., Variable Bit Rate), with heterogeneous recipients and with different service levels.

7. REFERENCES

[1] Ahn G.-S., Campbell A., Lee S.-B., Zhang X. "INSIGNIA". *Internet Draft, draft-ietf-manet-insignia-01.txt*, Oct. 1999. Work in progress.

[2] Bellur B., Ogier R. G., Templin F. L. "Topology Broadcast Based on Reverse-Path Forwarding (TBRPF)". *Internet Draft, draft-ietf-manet-tbrpf-01.txt*, Mar. 2001. Work in progress.

[3] Blake S., Black D., Carlson M., Davies E., Wang Z., Weiss W. "An Architecture for Differentiated Services". *RFC 2475*, Dec. 1998. Work in progress.

[4] Borella A., Cancellieri G., Pagani E., Rossi G.P. "Implementation Schemes for Multicast Bandwidth Brokers in Multidomain Networks". *To appear in International Journal of Computer and Telecommunications Networking*, 2001. Elsevier Science, Amsterdam.

[5] Braden R., Clark D., Shenker S. "Integrated Services in the Internet Architecture: an Overview". *RFC 1633*, Jun. 1994. Work in progress.

[6] Cansever D. H., Levesque A. H. "A Framework for QoS Support in Mobile Ad-hoc Networks". *Internet*

Table 1: **Insignia and M-Camp characteristics.**

	Insignia	M-Camp
supported QoS	bandwidth guarantee (2 levels)	bandwidth guarantee
signaling	in-band	out-of-band
underlying protocol changes	fields added to the IP header	no
control overhead	∀ pkt: IP header fields periodic QoS reports	`probe` packets periodic RTCP reports
QoS state	per-flow reserved bandwidth	stateless
state maintenance	soft-state (timer ∀ flow)	-
multicast support	?	yes
unidirectional links	yes	yes
computation overhead	∀ received pkts: header processed to possibly update local state or header info	pkts remarking during the proxy lifetime (set-up for downstream dests)
resource release	at the timer expiration	immediate

Draft, draft-ietf-manet-qos-framework-00.txt, Feb. 1999. Work in progress.

[7] Chen S., Nahrstedt K. "Distributed Quality-of-Service Routing in Ad Hoc Networks". *IEEE Journal on Selected Areas in Communications, Vol. 17, No. 8*, pages 1488–1505, Aug. 1999.

[8] Gerla M., Pei G., Hong X., Chen T.-W. "Fisheye State Routing Protocol (FSR) for Ad Hoc Networks". *Internet Draft, draft-ietf-manet-fsr-00.txt*, Nov. 2000. Work in progress.

[9] Haas Z. J., Pearlman M. "The Zone Routing Protocol (ZRP) for Ad Hoc Networks". *Internet Draft, draft-ietf-manet-zone-zrp-02.txt*, Jun. 1999. Work in progress.

[10] Haas Z. J., Pearlman M. R., Samar P. "The Intrazone Routing Protocol (IARP) for Ad Hoc Networks". *Internet Draft, draft-ietf-manet-zone-iarp-00.txt*, Jan. 2001. Work in progress.

[11] Handley M. "The sdr Session Directory: An Mbone Conference Scheduling and Booking System". `http://mice.ed.ac.uk/mice/archive/sdr.html`, Apr. 1996.

[12] Jetcheva J. G., Hu Y.-C., Maltz D. A., Johnson D. B. "A Simple Protocol for Multicast and Broadcast in Mobile Ad Hoc Networks". *Internet Draft, draft-ietf-manet-simple-mbcast-00.txt*, Nov. 2000. Work in progress.

[13] Ji L., Corson M. S. "Differential Destination Multicast (DDM) Specification". *Internet Draft, draft-ietf-manet-ddm-00*, Jul. 2000. Work in progress.

[14] Johnson D. B., Maltz D. A., Hu Y.-C., Jetcheva J. G. "The Dynamic Source Routing Protocol for Mobile Ad Hoc Networks". *Internet Draft, draft-ietf-manet-dsr-04.txt*, Nov. 2000. Work in progress.

[15] Lal S., Sousa E. S. "Distributed Resource Allocation for DS-CDMA-Based Multimedia Ad Hoc Wireless LANs". *IEEE Journal on Selected Areas in Communications, Vol. 17, No. 5*, pages 947–967, May 1999.

[16] Lee S.-B., Ahn G.-S., Zhang X., Campbell A. "INSIGNIA: An IP-Based Quality of Service Framework for Mobile ad Hoc Networks". *Journal of Parallel and Distributed Computing, Vol. 60*, pages 374–406, 2000.

[17] Lee S.-J., Su W., Gerla M. "On-Demand Multicast Routing Protocol (ODMRP) for Ad Hoc Networks". *Internet Draft, draft-ietf-manet-odmrp-01.txt*, Jun. 1999. Work in progress.

[18] Lin C. R., Liu J.-S. "QoS Routing in Ad Hoc Wireless Networks". *IEEE Journal on Selected Areas in Communications, Vol. 17, No. 8*, pages 1426–1438, Aug. 1999.

[19] Macker J. P., Corson M. S. "Mobile Ad Hoc Networking and the IETF". *ACM Mobile Computing and Communications Review, Vol. 2, No. 4*, pages 9–13, Oct. 1998.

[20] Nichols K., Jacobson V., Zhang L. "A Two-bit Differentiated Services Architecture for the Internet". *Internet Draft, draft-nichols-diff-svc-arch-00.txt*, Nov. 1997. Work in progress.

[21] Pagani E., Rossi G.P. "An Admission Control Policy for QoS Multicast Traffic in Mobile Ad Hoc Networks". *Technical Report, Dept. of Computer Science, Università degli Studi di Milano*, 2001.

[22] Perkins C., Royer E. M., Das S. R. "Ad Hoc On Demand Distance Vector (AODV) Routing". *Internet Draft, draft-ietf-manet-aodv-07.txt*, Nov. 2000. Work in progress.

[23] Schulzrinne H., Casner S., Frederick R., Jacobson V. "RTP: A Transport Protocol for Real-Time Applications". *RFC 1889*, Jan. 1996. Work in progress.

[24] Shenker S., Wroclawski J. "General Characterization Parameters for Integrated Service Network Element". *RFC 2215*, Sep. 1997. Work in progress.

[25] Sivakumar R., Sinha P., Bharghavan V. "CEDAR: a Core-Extraction Distributed Ad hoc Routing algorithm". *Proc. IEEE Infocom'99*, Mar. 1999.

[26] Sobrinho J. L., Krishnakumar A. S. "Quality-of-Service in Ad Hoc Carrier Sense Multiple Access Wireless Networks". *IEEE Journal on Selected Areas in Communications, Vol. 17, No. 8*, pages 1353–1368, Aug. 1999.

[27] Wu C. W., Tay Y. C., Toh C.-K. "Ad hoc Multicast Routing protocol utilizing Increasing id-numberS (AMRIS) - Functional Specification". *Internet Draft, draft-ietf-manet-amris-spec-00.txt*, Nov. 1998. Work in progress.

[28] Wu K., Harms J. "QoS Support in Mobile Ad Hoc Networks". *Unpublished paper*, 2000. `http://www.cs.ualberta.ca/~wkui/research-/QoSReview.ps`.

[29] Xiao H., Seah W., Lo A., Chua K. C. "On Service Differentiation in Multihop Wireless Networks". *Proc. ITC Specialist Seminar on Mobile Systems and Mobility*, Mar. 2000.

[30] Xiao H., Seah W., Lo A., Chua K. C. "A Flexible Quality of Service Model for Mobile Ad Hoc Networks". *Proc. 51st IEEE Vehicular Technology Conference*, May 2000.