Internet Connectivity for Mobile Ad Hoc Networks: Solutions and Challenges

Pedro M. Ruiz, DIIC, University of Murcia
Francisco J. Ros, DIIC, University of Murcia
Antonio Gomez-Skarmeta, DIIC, University of Murcia

ABSTRACT

The interconnection of mobile ad hoc networks to fixed IP networks is one of the topics receiving more attention within the MANET working group of the IETF as well as in many research projects funded by the European Union. Several solutions have recently been proposed, but at this time it is unclear which ones offer the best performance compared to the others. In addition to introducing the main challenges and design options that need to be considered, we perform a simulation-based evaluation aiming at providing some insight on the performance of these approaches. These simulation results have proven themselves valuable by showing that some of the most eye-catching features of the proposed approaches have practical performance issues which need to be enhanced.

INTRODUCTION AND MOTIVATION

Ad hoc networks consist of a number of mobile nodes that are free to move and communicate one with each other wirelessly. These mobile nodes have routing capabilities that allow them to create multihop paths connecting nodes which cannot directly communicate. These networks are extremely flexible, self-configurable, and do not require the deployment of any infrastructure for their operation. However, the idea of facilitating the integration of mobile ad hoc networks (MANETs) and fixed IP networks has gained a lot of momentum within the research community. In such integrated scenarios, commonly known as hybrid ad hoc networks, mobile nodes are witnessed as an easily deployable extension to the existing infrastructure. Some ad hoc nodes act as "gateways" that can be used by other nodes to seamlessly communicate with hosts in the fixed network.

This view of hybrid ad hoc networks as a cost-effective extension of existing access networks is receiving a lot of attention from researchers working on future beyond third-generation (3G) wireless and mobile networks. For instance, most of the European manufacturers and mobile operators are conducting research on this topic in the framework of many EU-funded research projects such us MIND[1] and DAIDALOS,[2] in which our group participates. In these beyond 3G scenarios ad hoc nodes are expected to join the ad hoc network and roam across different heterogeneous access technologies at any time.

This topic has also created a lot of interest within the Internet research community, as witnessed within the MANET working group at the 61st and 62nd Internet Engineering Task Force (IETF) meetings. This working group is currently working on standardizing proactive and reactive routing protocols for MANETs. Proactive protocols are those in which nodes periodically interchange routes. Reactive routing protocols discover routes to destinations only when an active source does not know how to reach to a destination. Optimized Link State Routing (OLSR) and Ad Hoc On-Demand Vector (AODV) routing are the baseline protocols for the proactive and reactive variants, respectively. The MANET working group has agreed that both the proactive and reactive routing solutions being designed will incorporate features to interwork with fixed IP networks.

During the last year a number of schemes have been proposed to tackle the interworking of MANETs with fixed networks. Unfortunately, there is not a clear understanding of the performance trade-offs and limitations of these solutions because they have not been properly evaluated and compared yet. Within the IETF MANET working group, there is a need to evaluate and compare existing approaches to avoid the current situation where it is unclear which solutions perform better. In this article we present some simulation work to shed some light on the performance of existing approaches under different mobility and traffic load scenarios.

The intended contributions of this article are twofold. First of all, we compare and describe the operation of the most well-known approaches proposed for hybrid MANETs within the IETF. In addition, we describe the general technical challenges and open issues. Second, and most important, we discuss our simulation results

[1] *www.ist-mind.org*

[2] *www.ist-daidalos.org*

0163-6804/05/$20.00 © 2005 IEEE

on the performance of existing interworking mechanisms, and provide valuable information about the previously unknown performance trade-offs of these solutions.

The remainder of the article is organized as follows. We describe the most relevant technical challenges in hybrid MANETs. We describe and compare existing solutions, pointing out their strengths and drawbacks. We show the performance of existing solutions through simulation. Finally, we give some conclusions and draw some future directions.

Design Alternatives and Technical Challenges in Hybrid MANETs

It is generally accepted that ad hoc nodes use IP addresses. So it may seem that interworking with IP networks is straightforward. Unlike the fully hierarchical addressing scheme used in the Internet, MANETs have a completely flat addressing model. In fact, ad hoc routing protocols like OLSR and AODV do not employ the concept of IP subnet. They assume that a node in a MANET may use any IP address provided that it is not duplicated. In fact, ad hoc routing protocols use host-based routes rather than network prefixes. Unlike in traditional IP networks, two neighboring nodes are not required to have addresses belonging to the same IP subnet to be able to directly communicate one each other. So, all these differences with traditional IP networks create some interworking issues as indicated below [1].

Discovering Internet gateways: The gateway discovery function can be done proactively, reactively, or in a hybrid way. Proactive gateway discovery corresponds to the case in which the gateways use periodic flooding to advertise themselves. In reactive gateway discovery, when a node has to communicate with a destination in the fixed network, it floods a control message throughout the entire network asking for a gateway. There are also hybrid approaches in which part of the nodes discover the gateways proactively and the rest do it reactively. As we show in our simulations, the gateway discovery mechanism has a strong impact in the overall performance, and it is a key component to provide interworking with fixed networks.

Address auto-configuration: This refers to the process by which a mobile ad hoc node obtains and configures an IP address. There are different approaches that can be considered. In stateful auto-configuration the IP address of a node is assigned by a central entity (e.g., a Dynamic Host Configuration Protocol [DHCP] server). However, centralized approaches are not suitable for MANETs due to possible network partitions. Another option is to use stateless auto-configuration. For instance, the gateway can advertise within its control messages a network prefix from which the nodes can derive an IP address. By integrating the auto-configuration information into gateway discovery messages, the overall overhead is reduced. However, the network might still be periodically flooded with prefix information, which may be a waste of resources. In general, it is advisable to configure the nodes with an IP address belonging to the subnet of its default gateway. This guarantees that the access router does not need to perform network address translation (NAT) to get messages routed back from the Internet. In addition, it avoids any conflict with traditional anti-spoofing packet filters that are generally configured for security reasons in access routers.

Reaching a Destination: When a mobile node within the MANET wants to contact a destination, it needs to find out whether the destination is in the fixed network or the MANET. If we do not impose any condition on the addresses that can be used by MANET nodes, the sender cannot check if the destination is in the MANET simply by looking at the network prefix of the destination IP address. In proactive protocols the sender can check if there is an entry in its routing table for the destination. For reactive protocols, a node may need to flood the network and assume that the destination is in the fixed network if no answer is received. A better approach might be to assign IP addresses belonging to the same subnet to all nodes within the MANET. Thus, a source can know that the destination is outside the MANET just by doing a prefix match with its own netmask without requiring control messages. Another issue is whether to use routing headers to reach the gateways or just simple headers. The former avoids reconfiguring IPv6 addresses to change from one access router to another. However, it also requires some additional overhead due to the additional header in each data packet.

Duplicate address detection (DAD): The goal is to avoid two nodes ending up with the same IP address. Existing IPv4 and IPv6 mechanisms were not designed for multihop networks, and they would require extensions. However, if we just allow existing schemes to send their control messages over multiple hops, the performance can be poor because the network might be flooded every now and then. Other ideas like passively listening to data and control messages can help reduce the overhead of DAD. In addition, the DAD algorithm should also guarantee the uniqueness of the addresses and resolve conflicts when network partitions and merges occur.

Name resolution: A mobile node might also need some kind of name resolution service in addition to the address auto-configuration mechanism. The simplest approach consists in integrating the information about available DNS servers within the gateway advertisement or auto-configuration messages. So, when a node gets an IP address and a default gateway it also knows the IP address of the DNS. However, it would be desirable to have a distributed MANET-internal DNS service being able to work even if the MANET gets disconnected from the fixed IP network.

Most of the Internet connectivity mechanisms proposed so far are mainly focused on the problems of address auto-configuration and gateway discovery. These are the elements most relevant to guarantee smooth interworking. In the rest of the article we focus on these particular issues.

Unlike in traditional IP networks, two neighboring nodes are not required to have addresses belonging to the same IP subnet to be able to directly communicate with one each other. So all these differences with traditional IP networks create some interworking issues.

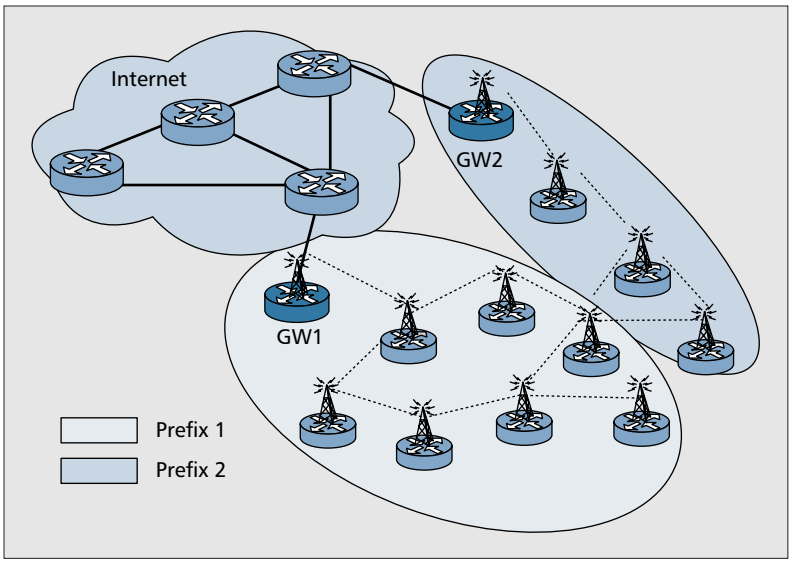

Figure 1. *Prefix continuity in Jelger's proposal.*

Evaluation of Existing Proposals

There have been a number of proposals in the literature to provide Internet connectivity for MANETs. One of the first proposals by Broch *et al.* [2] is based on integration of Mobile IP and MANETs employing a source routing protocol. MIPMANET [3] followed a similar approach based on AODV, but it only works with Mobile IPv4, because it requires foreign agents (FAs). In addition, Ammari *et al.* [4] analyzed the performance of mobile gateways in a MANET based on the DSDV routing protocol. In general, existing proposals are not general enough to work with different routing protocols. Thus, within the IETF the work focused on seeking more general solutions.

In general, auto-configuration proposals are tightly coupled with gateway discovery mechanisms. The same control messages used to discover gateways are usually challenged with network prefix information allowing ad hoc nodes to derive an IP address. Thus, the gateway discovery mechanism strongly influences the overall performance.

Most of the existing proposals use either a reactive or proactive gateway discovery mechanism. However, there are also a few research papers [5, 6] that propose the idea of using hybrid gateway discovery. In [5] the authors propose a scheme in which advertisements are only propagated up to a certain number of hops. Nodes located out of that scope will reactively find the gateways when needed. However, as the authors show, the optimal time to live (TTL) depends very much on the particular scenario and network conditions under consideration, as does the performance of this approach. In [6] the authors propose a more sophisticated approach in which advertisements are sent out only when changes in topology are detected. However, they rely on a source-based routing protocol, which limits the applicability of their approach to this particular type of routing protocol. In general, it is a desirable property for an Internet connectivity solution to be able to work with different routing protocols. Recently, we also proposed in [7] an adaptive gateway discovery mechanism that outperforms existing hybrid approaches. The key is that the TTL for proactive gateway advertisements is adjusted dynamically to network conditions. Those ad hoc nodes out of that scope reactively find the gateways.

Unfortunately, these hybrid approaches are only focused on gateway discovery and are not a complete solution for interworking with fixed IP networks. Thus, in the rest of the section we focus on the proposals that deal both with auto-configuration and gateway discovery.

The proposal from Singh *et al.* [8] defines the gateways as those nodes that are one hop away from the access router. Access routers are expected to be equipped with a wireless interface. The first node becoming a gateway is called the *default gateway*, and it is responsible for sending out periodic gateway advertisements. The other gateways for a given router are called *candidate gateways*. A candidate gateway becomes a default gateway when it stops receiving gateway advertisements from the default gateway for some time. It defines gateway selection based on bandwidth balancing, but unfortunately bandwidth parameters and their use are not described in the current version of the specification.

The proposals receiving more attention within the IETF and the research community in general are those from Wakikawa *et al.* [9] and Jelger *et al.* [10]. We refer to these solutions using the surname of their first author from now on, and they will be the focus of the rest of the article.

Wakikawa

This proposal entitled *Global Connectivity for IPv6 Mobile Ad Hoc Networks* [9] defines how a mobile node can derive an IPv6 global address based on a prefix from a gateway and use it to communicate with nodes in the Internet. This proposal defines two different mechanisms to discover the gateway:

- Periodic flooding of gateway advertisement (GWADV) messages from the gateways
- Reactive flooding of a gateway solicitation (GWSOL) message from the node and subsequent unicast GWADV from the gateway

The first approach is completely proactive, whereas the latter is completely reactive. According to the Internet draft, these messages can be implemented by simply adding an I flag to existing RREQ and RREP messages, indicating that this particular route refers to external connectivity. However, the specification allows for the use of other special control messages such as ICMP packets. In our simulations we used the I flag.

A GWADV message contains the global IPv6 address of the gateway, the network prefix advertised by the gateway, the prefix length and the lifetime associated with this information. When a node receives a GWADV message proactively or reactively, the mobile node is able to derive an IPv6 address which belongs to the advertised network prefix. After that, the mobile node creates a default route in its rout-

ing table, pointing to the selected gateway. In addition, it also adds a host-based route toward the IPv6 address of the gateway. These route entries are updated periodically to guarantee a seamless interworking.

Data packets are sent toward the destinations in the fixed network using a routing header which includes the address of the gateway. It is like a tunnel from the mobile node to the gateway, allowing for the selection of a new gateway without reconfiguring the IPv6 address. The only change is to simply reconfigure the tunnel endpoint to the new gateway. The cost is a little bit of additional overhead in each data packet due to the additional header.

Unfortunately, the authors do not recommend any metric nor give any concrete algorithm for selecting a gateway. In our simulations we have used the number of hops because it is the metric supported by every Internet connectivity approach. Thus, it serves as a common ground for a fair comparison among different approaches.

JELGER

C. Jelger *et al.* proposed an Internet draft entitled "Gateway and Address Auto-Configuration for IPv6 Networks" [10]. This proposal is a proactive approach in which the gateways advertise themselves by periodically flooding gateway information (GW_INFO) messages. Unlike other approaches, the flooding of these messages is based on the idea of prefix continuity. We can see its operation in Fig. 1. The idea is very simple. A node selects, according to some metric, only one of the GW_INFO messages received as its best path toward the gateway. Then it configures an IPv6 address based on the advertised prefix, and forwards only the GW_INFO that corresponds to the selected prefix. All other GW_INFO messages containing different prefixes are not forwarded. This guarantees that given a node, A, every node in the path from A to the gateway has the same prefix. In addition, it also reduces the overhead associated with the proactive flooding of the network. The neighbor from which node B receives its best GW_INFO is called the *upstream neighbor* of B.

As we can see in Fig. 1, the result is basically a set of nodes sharing the same prefix that form a tree rooted at the gateway.

This proposal defines three different metrics to select an upstream neighbor based on distance, stability, and delay, respectively. According to the F_DISTANCE algorithm, node A selects the closest neighbor to any of the gateways as the upstream neighbor. It might require changing the IPv6 address and prefix used by A (and consequently many children of A). This distance is based on the number of hops. The F_STABILITY algorithm is like the F_DISTANCE. but only selects among those upstream candidates sharing the same prefix. With the F_STABILITY algorithm the current prefix of the node is always preserved. Finally, the F_DELAY algorithm selects the upstream neighbor that has the same prefix and from which the GW_INFO has been received first.

In this case data packets just travel back from the source to the gateway following the intermediate upstream neighbors. This approach does not require the use of routing headers. In addition, the prefix continuity guarantees that each gateway only receives IPv6 data packets whose source IPv6 address belongs to its prefix. Thus, problems with anti-spoofing filters are avoided.

	Wakikawa	Jelger	Singh	Broch	MIPMANET
Proactive/reactive	P/R	P	P/R	R	P/R
Multiple gateways	Yes	Yes	No	No	Yes
DAD	Yes	No	n/a	No	No
Stateless/stateful	Stateless	Stateless	n/a	n/a	Stateful
Routing header	Yes	No	Yes	Yes	Yes
Complete spec.	Yes	Yes	No	Yes	Yes

Table 1. *Summary of features of well-known existing proposals.*

Table 1 summarizes the different features provided by each solution. As we can see in the table, each approach has its benefits and its drawbacks. For instance, we consider the prefix continuity offered by Jelger a good advantage in terms of interworking. In addition, its no dependency on routing headers is also a good point. The idea of balancing traffic across gateways is a strong advantage of Singh's proposal. But unfortunately, this proposal does not specify some key aspects of its operation such as selection of gateways based on traffic load or modification of IPv6 packets performed by candidate gateways. Thus, given that both Wakikawa and Jelger are more mature and well defined we use those two in our simulations as representative of reactive and proactive Internet connectivity approaches.

PERFORMANCE EVALUATION

To assess the performance of the proposed approaches, we have implemented them within version 2.27 of the ns2 network simulator. In addition, we have also implemented the OLSR protocol according to the latest IETF specification.[3] We have set up a scenario consisting of 25 mobile nodes using 802.11b at 2 Mb/s with a radio range of 250 m, two gateways, and two nodes in the fixed network. These nodes are placed in a rectangular area of 1200×500 m^2. For an increasing number of active UDP sources (5, 10, 15, 20) within the MANET, we evaluate the performance of the different approaches. These UDP sources send out a constant bit rate of 20 kb/s using 512 bytes/packet. The gateways are located in the upper right and lower left corners, so we can have long enough paths to convey useful information. In addition, we use the two different routing schemes that are being considered for standardization within the IETF: OLSR as a proactive scheme, and AODV as a reactive one. This will help us to determine not only the performance of the proposals, but the type of routing protocols for which they are most suitable. The case of OLSR with a reactive gate-

[3] *Code available at http://ants.dif.um.es/masimum/um-olsr_ns-2.27_v0.8.7.tar.gz*

Scheme	Parameter	Value
OLSR		
	HELLO_INTERVAL	2 s
	TC_INTERVAL	5 s
	NEIGHBOR_HOLDTIME	6 s
AODV		
	RREQ_RETRIES	2 s
	ROUTE_TIMEOUT	10 s
	REVERSE_ROUTE_LIFETIME	6 s
Jelger		
	GW_INFO_INTERVAL	2 s
	GW_INFO_LIFETIME	10 s
Wakikawa pro		
	GWADV_INTERVAL	2 s
	GWADV_LIFETIME	10 s

Table 2. *Detailed simulation parameters.*

way discovery has not been simulated because in OLSR all the routes to every node in the MANET (including the gateways) are already computed proactively. Thus, there is no need to reactively discover the gateway, because it is already available at every node. The simulation parameters for the different gateway discovery approaches and the routing protocols are summarized in Table 2. In both AODV and OLSR we activated link layer feedback.

We also evaluate the effect of mobility by using the random waypoint model in which nodes pick a destination point in the simulation area, wait for a pause time, and then start moving toward that point at a random speed. Once the destination is reached, the process is repeated. In our case speeds are uniformly distributed between 0 and 20 m/s, and pause times range between 0 and 900 s. A pause time of 0 represents a network in which the nodes are always moving, whereas a pause time of 900 s represents a completely static network. In addition, to avoid accumulation of nodes in the center of the simulation area, we let the nodes move for 4500 s of which we only consider the last 900 s for our simulations. Moreover, we conducted additional simulations using the Gauss-Markov model to validate that the results presented in this article are congruent across mobility models.

The main performance metrics under consideration are those recommended in RFC 2501:

- Packet delivery ratio (PDR): Computed as the percentage of successfully delivered data packets over all the data packets sent out by the sources.
- Normalized control packet overhead: Defined as the total number of control packets sent and forwarded over the number of successfully delivered data packets. This metric represents the cost in terms of overhead of the different approaches.

In addition, we also evaluate the gateway discovery overhead, defined as the total number of control messages associated with the discovery of gateways. This metric provides information regarding how much of the control packet overhead is due to the Internet connectivity scheme.

Packet Delivery Ratio

The PDR is mainly influenced by the routing protocol under consideration, although Internet connectivity mechanisms also have an impact. Our analysis of the end-to-end throughput shows a strong correlation with the PDR; so in this section we only show graphs regarding the PDR. Similar to previous simulations of OLSR in the literature, we can see in Fig. 2a and 2b that as mobility increases, AODV simulations exhibit much better PDR than OLSR ones. The reason is that an increment in the rate at which links break makes the protocol send more control packets, and the topology also takes more time to converge than AODV. In addition, according to RFC 3626, when link layer feedback informs OLSR about a broken link to a neighbor, the link is marked "lost" for 6 s. During this time packets are dropped in OLSR. This behavior also affects the routes toward Internet gateways, which is the reason the PDR is so low in OLSR simulations. This effect can be clearly observed in Fig. 3, which shows that most of the packet losses in OLSR occur during these periods of time (i.e., when there is no route available).

In the case of OLSR, Jelger performs surprisingly worse than the proactive version of Wakikawa. Given that Jelger has lower control overhead, we expected the results to be the other way around. The reason is that Jelger is strongly affected by the mobility of the network. After carefully analyzing the simulations (Fig. 3), we found out that the selection of next hops and gateways makes the topology created by Jelger very fragile to mobility. The problem is that the restrictions imposed by prefix continuity in Jelger concentrates the traffic on a specific set of nodes. In AODV, this problem is not so dramatic because AODV, rather than marking a neighbor as lost, starts finding a new route immediately. Thus, we can conclude that although prefix continuity has very interesting advantages, it has to be carefully designed to avoid data concentration and provide quick reactions to topological changes.

As the number of sources increases, the PDR is reduced due to the additional congestion. If we look at the different Internet connectivity schemes in AODV, we can see that they offer a very similar performance. However, it is interesting to note that proactive Wakikawa produces a lower PDR than reactive Wakikawa and Jelger when the number of sources is high. The reason is that it has very high control overhead, which, added to the high volume of data traffic, ends up saturating the ad hoc network. This is not so heavily evident in Jelger, thanks to its lower

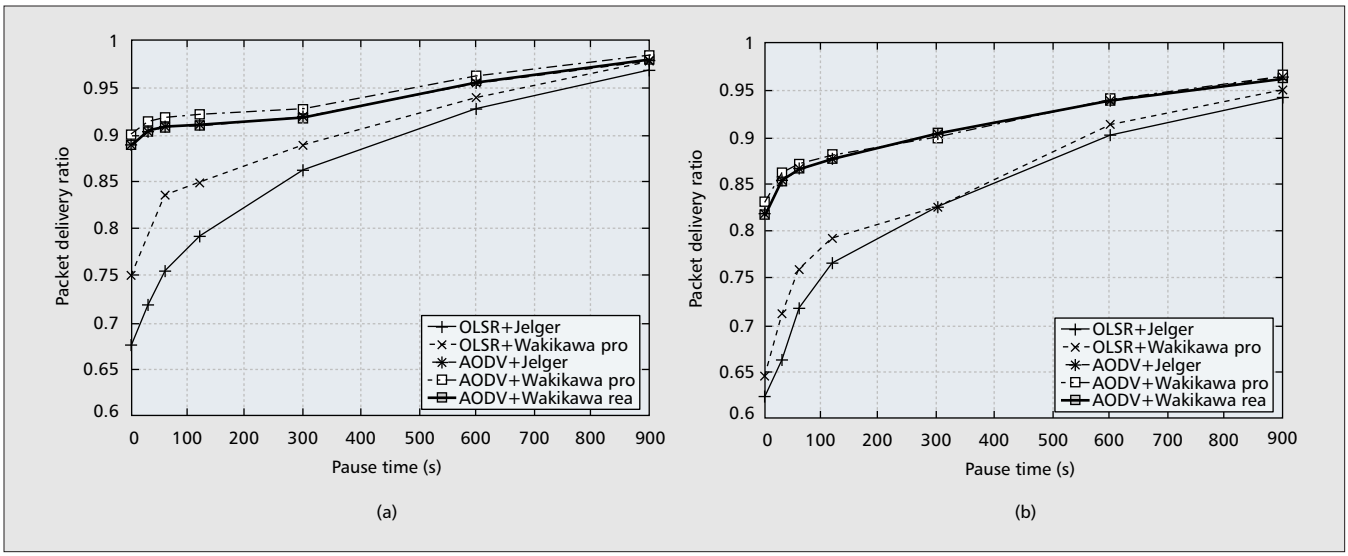

Figure 2. *Packet delivery ratio using different pause times: a) 10 sources; b) 15 sources.*

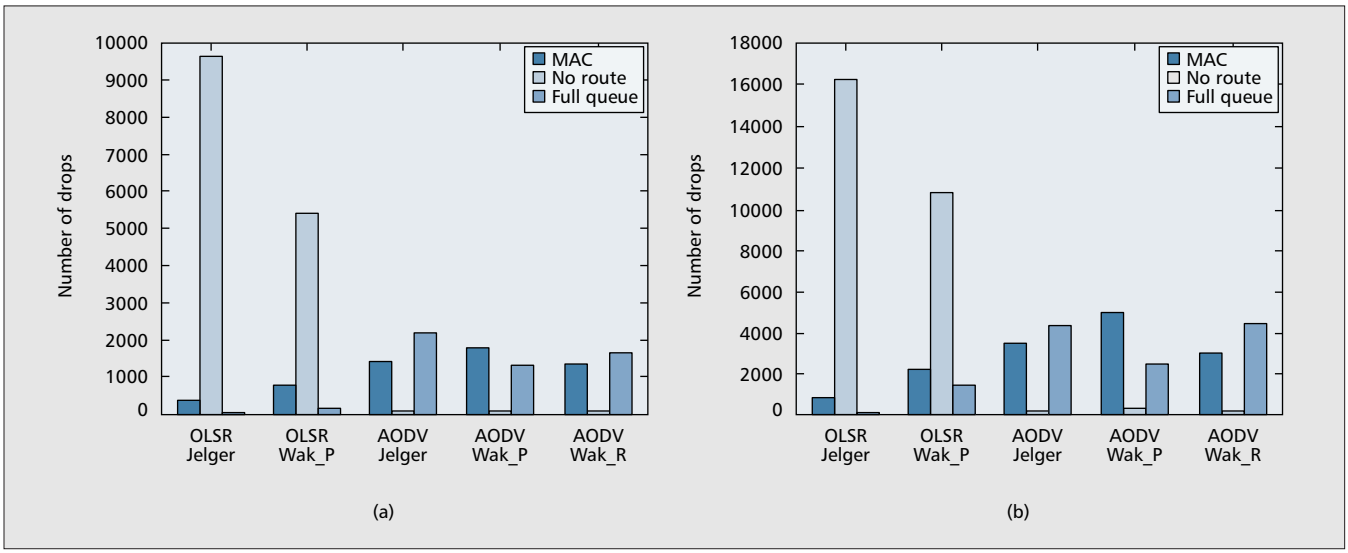

Figure 3. *Cause of packet drops using a pause time of 60 s: a) 10 sources; b) 15 sources.*

amount of control traffic used to propagate GW_INFO messages. However, when mobility is high and there are few data sources, proactive Wakikawa offers the best PDR at the cost of a very high control overhead.

For OLSR, if we compare the graphs for 10 and 15 sources we realize that as the number of sources increases, reactive Wakikawa's performance dramatically worsens. However, Jelger is not as affected by an increasing number of sources.

An interesting aspect shown by Fig. 3 is that with AODV, Jelger produces a higher number of drops due to filling up queues than does Wakikawa. This is due to concentration of data traffic, as explained before. In addition, the MAC layer contention is higher in proactive Wakikawa due to its excessive control overhead. In OLSR there is not as much congestion and contention because drops due to missing routes reduce the traffic load in the network. Therefore, we can conclude from Fig. 3 that there are important interactions between routing protocols and gateway discovery mechanisms that need to be taken into consideration. The same interworking solution may perform completely differently when used with distinct routing protocols.

Normalized Control Overhead

The normalized control overhead is defined as the number of control and data packets sent out to deliver a data packet to the final destination. The optimal normalized control overhead of one unit can only exist when there is no control traffic. The advantage of this metric is that it considers not only control packets, but also the trade-off between packet delivery ratio and control overhead.

As seen in Fig. 4a and 4b, OLSR approaches have higher control overhead due to the additional control messages used by OLSR to those used in AODV. In general, all the approaches increase the control overhead at increasing mobility rates due to the need to react to link

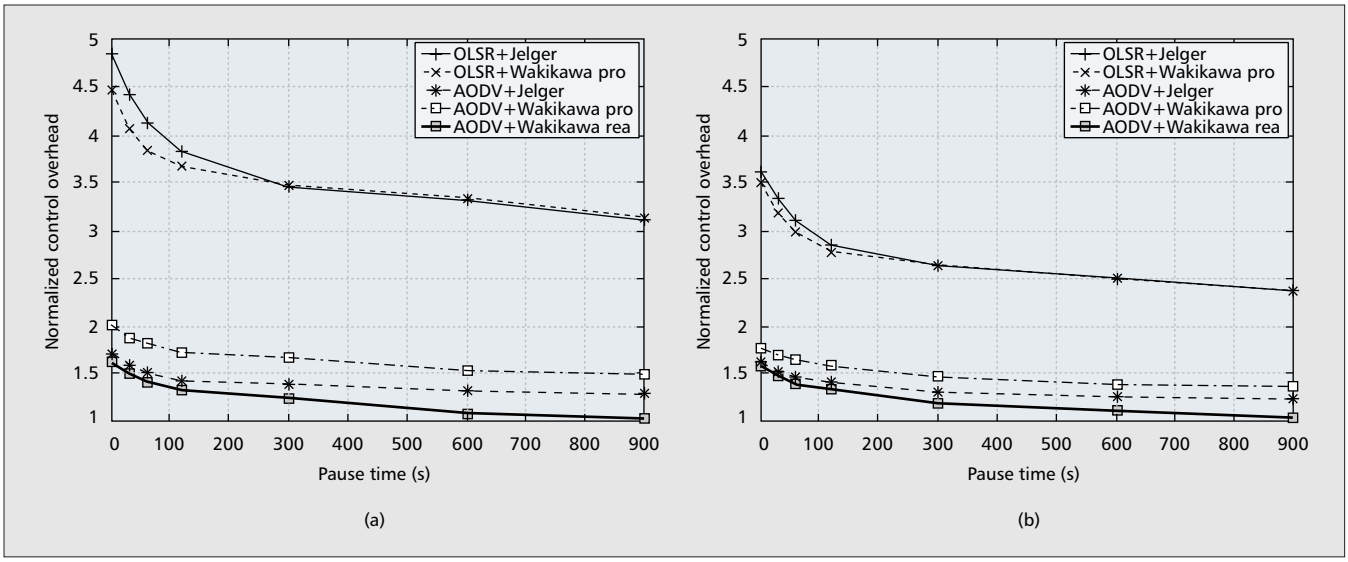

Figure 4. *Normalized control overhead for different pause times: a) 10 sources; b) 15 sources.*

breaks. Surprisingly, the combination of OLSR and Jelger offers higher normalized control overhead than OLSR with proactive Wakikawa when mobility is high. This is because it delivers a lower amount of data packets than proactive Wakikawa, even though its control overhead is lower. This makes the normalized metric offer lower performance. This clearly indicates again the need for quick recovery of routes to gateways when OLSR is used.

Regarding the solutions based on AODV, the results are pretty much was expected. Reactive Wakikawa has the lowest normalized control overhead, whereas proactive Wakikawa has higher overhead. This means that although the PDR achieved by the proactive approach is higher, the control overhead required to achieve that higher PDR is way too much compared to the reactive approach. The AODV-Jelger combination lies somehow in the middle of the other two approaches. The reason is that its control overhead is a little lower than proactive Wakikawa (due to its optimized GW_INFO propagation), whereas its PDR is very similar to proactive Wakikawa.

These results seem to confirm that adaptive gateway discovery algorithms like the ones we proposed in [7] can definitely help very much to enhance the results for reactive ad hoc routing protocols.

Gateway Discovery Overhead

Finally, we evaluate the overhead of the gateway discovery function in each proposal. As seen in Fig. 5, AODV simulations result in higher gateway overhead as the mobility of the network increases. This is due to the increase in the link break rate, which makes ad hoc nodes find a new route to the Internet as soon as their default route is broken. Moreover, AODV generates higher gateway overhead as the number of sources increases. Again, adding more sources means that more nodes need to find new routes to Internet gateways due to mobility. We can clearly see that proactive Wakikawa generates the biggest amount of Internet–gateway messages due to its periodic flooding through the whole network. Reactive Wakikawa shows the minimum gateway overhead thanks to its reactiveness. Jelger sits between the other two, due to its limited periodic flooding. As expected, the gateway discovery overhead for Internet connectivity mechanisms combined with OLSR remains almost unaffected by network mobility. This is due to the fact that Internet connectivity messages are periodically sent out by OLSR without reaction to link breaks, so its gateway control overhead is not heavily affected by mobility. When the number of sources increases, the gateway overhead decreases because the heavier traffic load saturates the network and a lower number of control messages can be injected. Figure 5 shows that Jelger always maintains a lower overhead than proactive Wakikawa due to the restriction of forwarding imposed by prefix continuity. The difference remains almost constant independent of the mobility of the network and the number of sources.

Conclusions and Future Work

In this article we have reviewed the current efforts being performed within the research community to address the issue of hybrid mobile ad hoc networks. There is currently no performance comparison of these mechanisms. This makes it hard for protocol designers within the MANET research community to assess which approaches are better and what their performance trade-offs are. That is why, in addition to the description and comparison of existing specifications, we have conducted simulations to evaluate performance under different scenarios and traffic loads.

We believe that the results in this article of evaluating the most recent solutions and routing protocols under consideration within the IETF MANET working group offer a valuable and timely insight. The simulation results showed that some of the features that seem more interesting from the specifications (e.g., prefix continuity) have some performance limi-

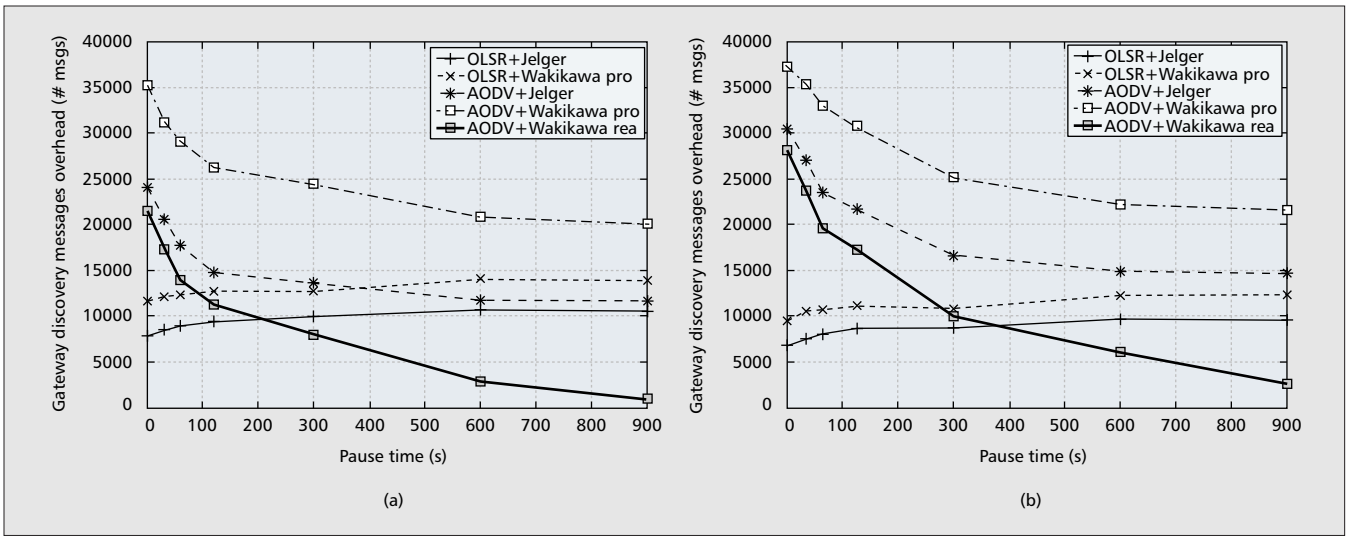

Figure 5. *Gateway discovery overhead for different pause times: a) 10 sources; b) 15 sources.*

tations when used with OLSR that need to be addressed. Another interesting result is that proactive gateway discovery needs a constrained flooding mechanism to avoid the huge amount of overhead associated with the discovery of gateways. We believe this work will help to create some awareness within the MANET research community about the proper direction regarding the design of Internet connectivity mechanisms.

There is still a lot of work to be done on the horizon. For instance, based on the results of this work, we believe that an interesting future research topic is the work on adaptive gateway discovery mechanisms. In addition to gateway discovery and auto-configuration, there are other areas related to interworking with fixed networks in which there is still a lot to do. These include, among others, improved DAD mechanisms, efficient support of DNS, discovery of application and network services, network authentication, and integrated security mechanisms.

Acknowledgment

Part of this work has been funded by the Spanish MCYT by means of the Ramon y Cajal work program, the ICSI Call for Spanish Technologists, and the SAM (MCYT, TIC2002-04531-C04-03) project. The authors are very thankful to the anonymous reviewers whose comments helped very much to enhance this article.

References

[1] S. Singh *et al.*, "Ad Hoc Network Autoconfiguration: Definition and Problem Statement," Internet draft, draft-singh-autoconf-adp-01.txt, July 2005.
[2] J. Broch, D. A. Maltz, and D. B. Johnson, "Supporting Hierarchy and Heterogeneous Interfaces in Multi-Hop Wireless Ad Hoc Networks," Proc. Wksp. Mobile Comp. held in conjunction with IEEE Int'l. Symp. Parallel Architectures, Algorithms, and Networks, Perth, W. Australia, June 1999.
[3] U. Jonsson *et al.*, "MIPMANET: Mobile IP for Mobile Ad Hoc Networks." *IEEE/ACM Wksp. Mobile and Ad Hoc Net. and Comp.*, Boston, MA, Aug. 1999, pp. 75–85.
[4] H. Ammari and H. El-Rewini, "Performance Evaluation of Hybrid Environments with Mobile Gateways," *Proc. 9th Int'l. Symp. Comp. and Commun.*, vol. 1, Alexandria, Egypt, June 2004, pp.152–57.
[5] P. Ratanchandani and R. Kravets, "A Hybrid Approach to Internet Connectivity for Mobile Ad Hoc Networks," *Proc. IEEE WCNC 2003*, vol. 3, New Orleans, LA, Mar. 2003, pp. 1522–27.
[6] J. Lee *et al.*, "Hybrid Gateway Advertisement Scheme for Connecting Mobile Ad Hoc Networks to The Internet," *Proc. 57th IEEE VTC 2003*, vol. 1, Jeju, Korea, Apr. 2003, pp. 191–95.
[7] P.-M. Ruiz, and A.-F. Gomez-Skarmeta, "Adaptive Gateway Discovery Mechanisms to Enhance Internet Connectivity for Mobile Ad Hoc Networks," *Ad Hoc and Sensor Wireless Networks*, vol. 1, no. 1, Mar. 2005, pp. 159–77.
[8] S. Singh *et al.*, "Mobile Multi-gateway Support for IPv6 Mobile Ad Hoc Networks," Internet draft, draft-sinhg-manet-mmg-00.txt, June 2004.
[9] R. Wakikawa *et al.*, "Global Connectivity for IPv6 Mobile Ad Hoc Networks," Internet draft, draft-wakikawa-manet-globalv6-03.txt, Oct. 2003.
[10] C. Jelger, T. Noel, and A. Frey, "Gateway and Address Autoconfiguration for IPv6 Ad Hoc Networks," Internet-draft, draft-jelger-manet-gateway-autoconf-v6-02.txt, Apr. 2004.

Biographies

PEDRO M. RUIZ graduated with M.Sc and Ph.D. (2002) degrees in telematics from the University of Murcia, Spain. In December 2003 he was awarded a Ramon y Cajal research position at the same university. He has worked as a post doctoral researcher at the International Computer Science Institute of the University of California at Berkeley. He has published over 70 refereed papers in international journals and conferences, and participates in the editorial board of IJPEDS and many international program committees. His main research interests include mobile ad hoc networks, wireless sensor networks, adaptive multimedia communications, and distributed algorithms.

FRANCISCO J. ROS received his B.Sc. degree in computer science in 2004 from the University of Murcia. He works currently as a researcher in the Department of Communications and Information Engineering at the same university. His main research interests are routing and autoconfiguration in hybrid mobile ad hoc networks and sensor networks.

ANTONIO F. GOMEZ-SKARMEA received an M.S. degree in computer science from the University of Granada, and B.S. (Hons.) and Ph.D. degrees in computer science from the University of Murcia. Since 1993 he has been a professor in the Department of Communications and Information Engineering at the University of Murcia, and for 2001–2005 chair of the department. Additionally, he has participated in several EU projects, and has published more than 130 international papers in journals and conferences, also being a member of several international program committees. His current research interests include security in mobile computing and networks, ad hoc networks, and distributed systems.