International Journal of Wireless Information Networks, Vol. 10, No. 3, July 2003 (© 2003)

Two-Hop Polling: An Access Scheme for Clustered, Multihop Ad hoc Networks

G. Dimitriadis[1] and F-N. Pavlidou[1]

In multiple channel environments, clustering provides a convenient framework for channel access and bandwidth allocation. Many clustering schemes, however, demand that terminals may communicate directly only if they share a common clusterhead. This requirement deactivates otherwise helpful links; those between nodes that belong to different clusters (intercluster links). Links between nodes that belong to different clusters constitute a distributed gateway. In this paper, we evaluate the importance of distributed gateways for two different clustering schemes and propose a novel access scheme for clustered environments using the link-cluster architecture, called two-hop polling (2HP). Two-hop polling manages to utilize intercluster links, leading to better connectivity and throughput.

KEY WORDS: Ad hoc; distributed gateway; two-hop polling.

1. INTRODUCTION

Ad hoc networks are emerging as a new promising field in wireless communications. Modern computing and communication devices are becoming so miniaturized, and consequently portable, that the traditional methods of networking through wired lines diminish the usability of the devices. Smart mobile phones, personal digital assistants (PDAs), and portable computers are some of the devices that a growing portion of the population uses daily. The most obvious alternative solution that does not undermine the mobility of users is a wireless network.

A special case of wireless networks, an ad hoc network, consists of terminals that are able to organize and configure themselves into a network by means of distributed algorithms. This feature enables them to be deployed in sites that have little or no infrastructure in minimal time. Having no need for a centralized entity that manages and configures the network, ad hoc networks are well-suited for applications ranging from military communications in hostile territories to cases of spontaneous communication in conference halls. Moreover, many ad hoc networks are designed to be multihop. This means that a message can be relayed by one or more intermediate terminals until it arrives to its destination, leading to much higher efficient transmission ranges compared to the physical transmission range that can be achieved by the terminals' radios. By implementing a routing mechanism at each terminal, each node becomes aware of the network topology and forwards data packets to the right direction while the entire process is completely transparent to the user.

The lack of centralized control, combined with the mobility of the terminals and the limited bandwidth, processing, and power resources, poses a severe strain to the routing and medium access protocols. The routing protocols should be able to adapt swiftly to the constantly shifting topology of the network whereas the medium access protocols must coordinate the transmissions of the terminals, in order to ensure fairness, efficiency, and minimal amount of collisions. Furthermore, these tasks should be accomplished with as little consumption of

[1] Telecommunications Division, Department of Electrical and Computer Engineering, Aristotle University of Thessaloniki, 540 06 Thessaloniki, Greece. E-mail: gedimitr@auth.gr; niovi@vergina.eng.auth.gr

1068-9605/03/0700-0149/0

resources as possible. Many protocols proposed for ad hoc networks are based on well-tested schemes for wired networks but slightly modified in order to cope with the demanding environment of an ad hoc network. In this paper, we examine the importance of a case of links in clustered networks (intercluster links or distributed gateways) and introduce an efficient medium access scheme based on polling which takes advantage of them.

Polling in single-channel environments cannot easily be implemented in networks consisting of large populations because of the lack of centralized control and the hidden terminal problem [1]. However, in a multiple channel environment (i.e., using different frequencies or different spreading codes), polling becomes feasible. The network is divided into clusters, providing this way a convenient framework for channel access and bandwidth allocation. Each cluster employs different channels than its neighboring clusters in order to isolate transmissions taking place within it. A node is elected in each cluster to play the role of the local coordinator, called the clusterhead. The clusterhead is responsible for polling each member of its cluster and thus coordinating the transmissions taking place within it. Consequently, this multiple access scheme could be viewed as polling on top of FDMA (Frequency Division Multiple Access) or CDMA (Code Division Multiple Access), for example.

The link-cluster architecture [2] provides a framework for minimizing the channel contention in the network through clustering and electing the nodes that coordinate the transmissions taking place within the clusters (clusterheads). Within this framework, intercluster links do not integrate well because their end nodes are not within range of a single clusterhead, but of two different ones. Taking also into account the fact that the end nodes of intercluster links use different channels (because they belong to different clusters), the utilization of these links becomes even more difficult.

If the medium access scheme demands that two terminals may communicate only if they share a common clusterhead, terminals that are within range but belong to different clusters may not communicate directly, but have to use a route that consists purely of intracluster links. This way, many potent links between nodes are not used; links that can lead to a more robust, better connected network. Intercluster links and their nodes that support them constitute a distributed gateway (DG). The importance of these links ranges from crucial (the sole path from node A to node B contains a distributed gateway) to handy (the shortest path from node A to node B contains a distributed gateway). In this paper, we evaluate and quantify the importance of these links and propose a scheme that utilizes them, called two-hop polling (2HP).

Section 2 discusses the clustering process and presents two well-known algorithms for doing so. Next, the importance of the distributed gateways for these two schemes is evaluated. In section 3, two-hop polling is introduced where we describe the mechanics of the proposed scheme. Afterward, in section 4, the simulation results are presented, and, finally, section 5 concludes the paper.

2. CLUSTERING

2.1. The Link-Cluster Architecture

One of the most convenient ways of scheduling and coordinating transmissions in a multiple channel environment is provided through the use of clusters. A cluster is defined as a subset of the network's terminals having a maximum diameter of two hops. Clusters have a node at their center (clusterhead) and extend to all the nodes that are within one hop from the clusterhead. Consequently, every member of the cluster is one hop away from its clusterhead and at most two hops away from every other member of the cluster. A cluster may take many different forms, but at its simplest form, it can be depicted as a star graph having a clusterhead at its hub.

There are many different ways of organizing the network into clusters and electing the clusterheads. Most algorithms can be divided into two categories, identifier-based [3] and connectivity-based [4]. With identifier-based clustering, a node is elected to be a clusterhead if it has the lowest or highest ID (Identification Number) in its vicinity. With connection-based clustering, clusterhead becomes the node with the most neighbors in its vicinity. Both of these algorithms have centralized and distributed versions. Naturally, in every ad hoc context the distributed versions are used.

Communication between different clusters is provided by gateway nodes; nodes that belong to more than one clusters. Gateway nodes reside at the overlapping region of two or more clusters and consequently are directly connected to two or more clusterheads. Gateways are capable of employing different codes depending on which cluster they mean to transmit to.

Many clustering schemes do not support the use of intercluster links (distributed gateways), but rather demand that all communications take place via intracluster links. Figure 1 presents an example of a network topology where without the usage of distributed gateways, the two clusters would be unable to communicate. Nodes 1 and 2 have been elected as clusterheads by usage of the lowest-ID algorithm while all the other

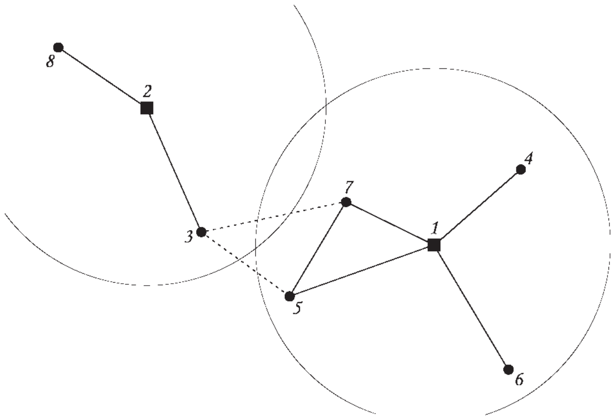

Fig. 1. Clustering Example.

nodes are directly connected either to node 1 or to node 2. Solid lines represent links between nodes sharing a common clusterhead, and dashed lines represent links over distributed gateways; that is, links between nodes that belong to different clusters. Normally, the node pairs {3, 7} and {3, 5} would not be able to communicate directly. In this example, the activation of the distributed gateways is necessary in order to establish communication between the two clusters. Surely, this example seems to be rather improbable, so in the next subsection, we quantify the importance of distributed gateways.

2.2. Evaluation Metrics

Our first results stem from the network graph alone. That is, we do not use a full simulation environment yet, but extract information from the topology graph that is formed by different random configurations. This process is repeated, and the results are then averaged. Specifically, a population of 100 nodes is randomly placed in a square area of 1 km^2, which is actually the configuration used in our network simulator later. The metrics that we evaluated are two: the connectivity and the average hops per route. With the first metric we assessed the ability of the network to cover as large a portion of the node population as possible. That is, the ability of a node to communicate with as many other nodes as possible, no matter in how many hops. With the second, we evaluated the route optimization that can be achieved through the activation of distributed gateways.

In Fig. 2, we present the connectivity of the network versus the transmission range. The connectivity metric is defined as the number of pairs that can communicate; that is, the number of pairs that are connected by a valid route, divided by the number of all possible pairs. The nodes are randomly placed in the area, where the resulting network is clustered by using the lowest-ID algorithm. All links over distributed gateways are removed, and then the connectivity of the resulting graph is measured. On the same initial network, this process is repeated, but now the better-connected nodes are elected to be clusterheads.

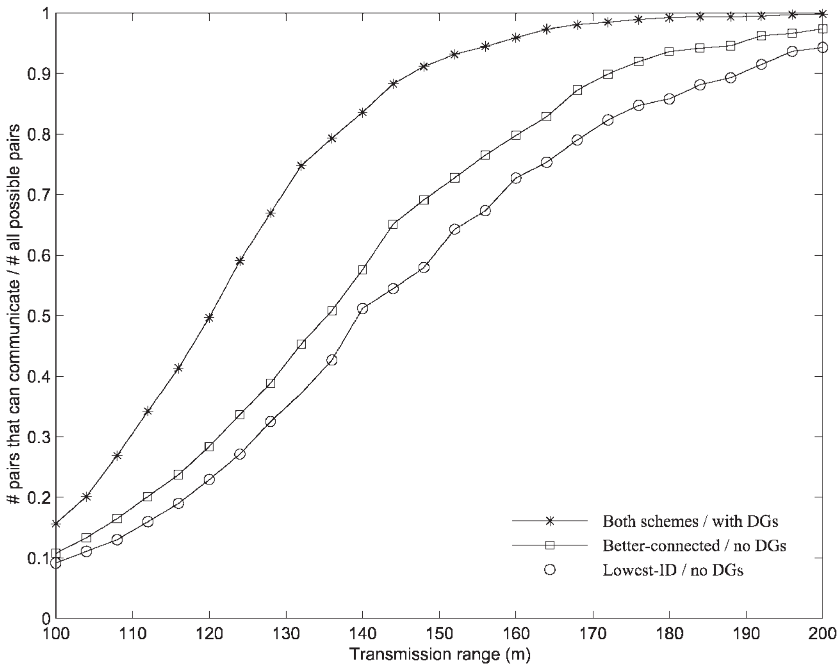

Fig. 2. Connectivity versus transmission range.

Of course, activating the distributed gateways in either clustered version will result in the same graph in which every node pair within range can communicate. The connectivity is measured for this graph also. In Fig. 2, there are two curves that correspond to the connectivity measured for the two clustered versions of the network without using distributed gateways as well as a curve for the same network using all possible links. This last curve corresponds to a clustered network that takes advantage of the distributed gateways, no matter the clustering algorithm.

As we can see, the highest connectivity is achieved when distributed gateways are activated. When not, the network clustered by the lowest-ID algorithm showed a worse behavior than the one that was clustered on a connectivity basis. This is expected because with connectivity-based clustering, the nodes that show the highest number of neighbors in their vicinity become clusterheads. Because all links that have a clusterhead at one of their ends are intracluster, these links are usable by definition. Nodes that have fewer neighbors become ordinary members of a cluster, lowering this way the possibility that an intercluster link may appear. On the other hand, when a network is clustered by using the lowest-ID algorithm, the clusterheads are elected more or less randomly. Consequently, this way the possibilities of an intercluster link appearance are not as low.

Because of the limited radio range, fragmentation may occur, resulting in a network that is effectively partitioned into multiple communication subsets. Clearly, it can be seen that distributed gateways provide a means for otherwise disconnected regions to communicate. Unless the transmission range is very high, where as we can see the three curves tend to unity, the activation of distributed gateways provide the extra links that enhance the network's connectivity.

The activation of intercluster links may also provide shortcuts to packets traveling through the network. In Figs. 3 and 4, the average number of hops per route is evaluated. Of course, since there are fewer routes when the distributed gateways are not used, in this metric participate only the common routes between the two versions of the network: the one that uses distributed gateways and the other that does not. In Fig. 3, the network was clustered with the lowest-ID algorithm whereas in Fig. 4 the network was clustered with the algorithm based on connectivity. In both figures, we can see that the average number of hops per route is lower when the distributed gateways are used. That means that packets reach to their destination faster, resulting in significant savings in bandwidth and energy consumption.

3. TWO-HOP POLLING

Two-hop polling (2HP) is a revised version of polling, especially tailored for the needs of a clustered environment.

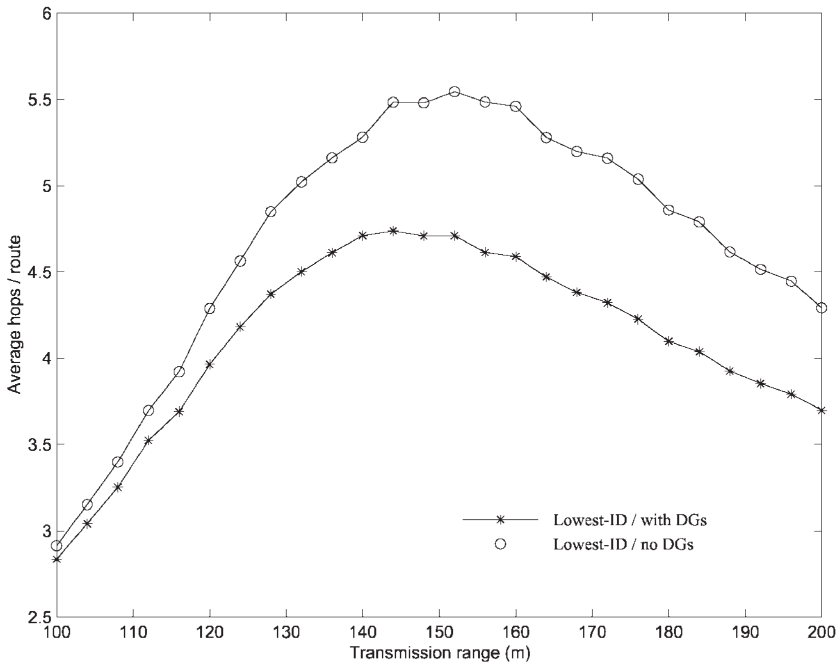

Fig. 3. Average hops per route versus transmission range (Lowest-ID algorithm).

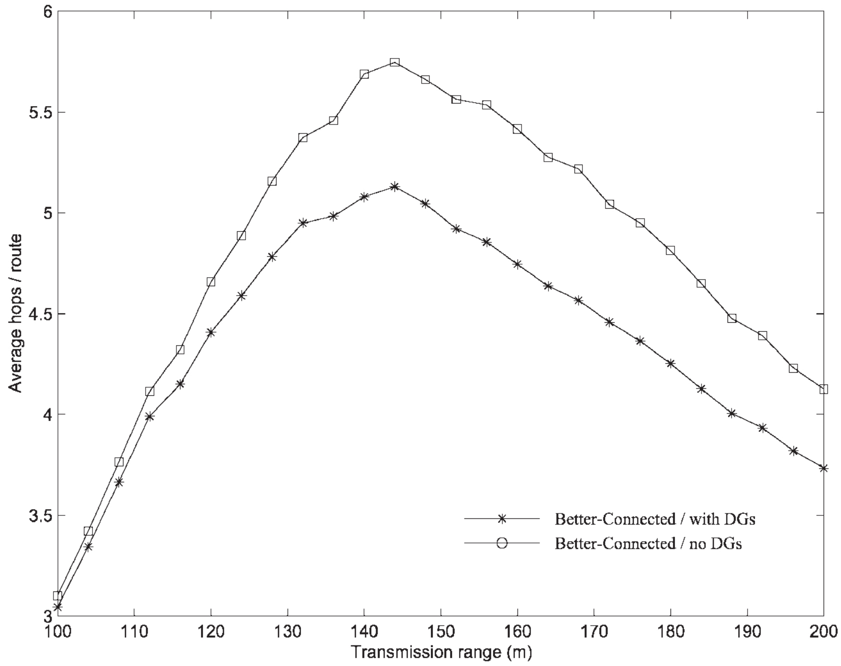

Fig. 4. Average hops per route versus transmission range (Better-connected algorithm).

With this scheme it is possible to utilize intercluster links (distributed gateways) without adding much to the complexity of polling.

With ordinary polling, the clusterhead invites each neighboring terminal to transmit in a round-robin fashion. As we recall from the previous section, intercluster collisions are not possible because adjacent clusters use different spreading codes. 2HP changes the medium access protocol by giving more liberty to nonclusterhead nodes.

Polling can be thought of as a token-based scheme. A token is issued by the clusterhead which circulates in the cluster and passes from every member of it. Terminals that hold the token are authorized to transmit data packets. After they have finished transmitting, they return the token back to the clusterhead, which transmits the token to a different terminal. With two-hop polling, ordinary nodes are not obliged to return the token back to the clusterhead as soon as they have finished transmitting. Rather, they can forward it to another, which resides at the other end of a distributed gateway. This way, links over distributed gateways can be used, and this is how the name of the scheme has been derived. Terminals as far as two hops from the clusterhead may be polled, though indirectly.

Returning to the example of Fig. 1, with ordinary polling the node pairs 3, 7 and 3, 5 would not be able to communicate directly. Using 2HP, node 5 instead of returning the token back to its clusterhead (node 1) transmits the token to node 3, allowing it this way to transmit data either to node 5 or to node 7. The token then returns to the clusterhead over the same route. Similarly, node 3 would poll node 5 and node 7 so that they can transmit to it and then return the token back to its clusterhead (node 2). This can be viewed from another perspective; nodes that participate in distributed gateways become "members" of the clusters that their peer nodes at the other side of the gateway belong to. These nodes that are two hops away from the clusterhead serve traffic that is inbound to the cluster. Thus, continuing our example, node 3 becomes a member of the cluster defined by node 1, and it is polled in order to transmit packets inbound to that cluster. The same applies of course to nodes 5 and 7. They become members of the other cluster as long as inbound traffic is concerned.

After the clusterheads are elected and the clusters are formed, the clusterhead does not know yet whether any members of its cluster are engaged in distributed gateways. As coordinator of the cluster, the clusterhead has to know if and which nodes are engaged in DGs. This information is especially useful in cases where more than one node of a cluster is connected via intercluster links to the same node. The clusterhead cannot allow that the same extracluster node gets polled more than once in the same polling cycle, in order to ensure fairness. So, every clusterhead constantly listens to the common medium. When it detects that a token was transmitted and the destination is not itself (and thus a node over a distributed gateway was polled), the clusterhead marks the IDs of the transmitting node and the destination node. When the token returns back to the clusterhead, it transmits a small control packet (or uses some space within the token), informing every node in the cluster that only the node that just returned the token has the permission to poll the node at the other side of the distributed gateway (as well as every other node that this node has polled over distributed gateways). Continuing the previous example, after node 5 has polled node 3 and returned the token back to its clusterhead (node 1), the clusterhead would inform every node in its cluster that only node 5 has permission to poll node 3, thus prohibiting node 7 from polling by itself node 3.

This way, a token is in the possession of a node that is a member of a different cluster. However, conflict-free communication is guaranteed because, as we recall, communication in different clusters is based on different spreading codes. Throughout this paper, we have assumed that every node transmits using the code of the token's origin cluster. Another problem arises from the fact that the token may get lost. This could happen if the token is sent to a node that is transmitting in a different code or to a node that has just gone out of range. This problem is amended through the use of timeout timers. If the token is not returned in a preconfigured interval, the token sender creates a new token.

It would seem more logical to allow a node to communicate with its partner in a distributed gateway simply by acquiring its cluster's token and then transmitting in a different spreading code. This scheme however poses two problems. First, the clusterhead would not know that the node is transmitting (because it uses a different code) and would go on polling a different node. Multiple nodes that belong to the same cluster could transmit simultaneously, or even worse in our case, multiple instances of a token could simultaneously exist in a cluster. Second, this way it is not guaranteed that the destination node is not itself transmitting, thus being unable to receive. The transmission of a token over a distributed gateway is equivalent to a handshake; it ensures that the node that just polled and all the other members of its cluster are ready to receive.

The benefits of the 2HP scheme are easily understood. First of all, the network becomes more robust and better connected. This means that previously unreachable destinations are now available and that packets are transmitted along better optimized routes. A more subtle benefit arises from the fact that the transformation of a link from intracluster to intercluster (and vice versa) no longer causes harm. In a wireless multihop network, a link between two nodes may switch at random times from intracluster to intercluster. The disappearance of an otherwise valid link will pose problems to the routing mechanism until the routing tables come up to this topology change through the use of routing updates. This abrupt disappearance of a link can lead to loops and, therefore, to a waste of network resources. With 2HP, the shift of a link from intra- to intercluster does not render the link unusable, thus the topology of the network appears to be less dynamic.

On the other hand, 2HP introduces more latency than ordinary polling, mostly because the token passes through more nodes. Moreover, there is greater possibility of a token to be lost in 2HP than in ordinary polling. With ordinary polling, the token may get lost when it is transmitted to a gateway node that is engaged in a transmission to a member of a different cluster. With 2HP the same can also happen when the token is transmitted over a distributed gateway. Finally, a less important disadvantage of the proposed scheme is that 2HP slightly adds to the complexity of the medium access compared to ordinary polling. Therefore, a careful study of pros and cons is necessary for the specific architecture in consideration. Our proposal has extensively been studied through simulation trials.

4. SIMULATION

4.1. Simulation Scenario

Our network model is based on the proposal of Gerla and Tsai in Ref. [4]. Network time is divided into frames. The frame is further divided into two phases: the control phase and the info phase. During the control phase, the nodes transmit in a TDM (time devision multiplexing) fashion in the common code information about their neighbors and their IDs. This way, the distributed version of the lowest-ID algorithm is executed, and every node of the network identifies its neighbors. At the end of the control phase, the clusters have been formed and the network is ready to serve data packets. At the beginning of the info phase, the clusterhead issues a token that circulates in the cluster as is defined by the 2HP scheme.

The simulation program was developed in C++. The network consists of 100 nodes moving randomly at a predefined speed in an area 1000×1000 m^2. The channel capacity is 2 Mbit/s while the packet size is 10 Kbits for data packets, 500 bits for control packets, and 200 bits for tokens. The info phase lasts 100 ms. Traffic load is created as follows: 100 sessions between random nodes are created. Every node in a session generates data packets following a Poisson process with an average time interval of 2.5 s. The routing protocol used is the fisheye state routing protocol (FSR) as proposed in Ref. [5]. The scope radius is 2 hops while routing tables are refreshed every 2 s for in-scope nodes and every 6 s for out-scope nodes. Messages containing routing updates are transmitted as data packets in the info phase, but they are of greater priority than normal data. In all our experiments, packets heading toward an unreachable (at that time) destination were not buffered, but dropped.

4.2. Simulation Results

Our first metric for the evaluation of a network using two-hop polling is the delivery ratio. The delivery ratio of the network is defined as the portion of the packets that asked to be served which actually arrived at their destination. Figures 5 and 6 compare the delivery ratio of polling versus 2HP; that is, they show the effect of the distributed gateways. It is obvious that 2HP shows a much better behavior than its counterpart that supports no distributed gateways. In some cases, 2HP demonstrates an almost double delivery ratio. This is owed primarily to the fact that when the distributed gateways are activated, routes to more destinations can be established. Also, messages containing routing updates propagate over distributed gateways, so they spread more quickly. Moreover, with 2HP the network appears to be less dynamic, as it was mentioned above.

It should be obvious in this point that the curves of Figs. 2 and 5 bear a resemblance. This is not strange; actually, the curves of Fig. 2 represent the absolute maximum the delivery ratio of a real-world network can achieve. If the behavior of the routing protocol was ideal and responded instantaneously to every change of the topology, under the same circumstances (identical node density and transmission range) the curves would be exactly the same.

Figure 7 compares the two schemes in a somewhat more instructive manner. In this figure is depicted the average number of hops each successfully delivered packet makes until it reaches its destination. At low speeds, 2HP shows a much higher hop count than its counterpart. As we recall from Fig. 2, when the distributed gateways are activated, the network becomes better connected and further destinations are reachable. This is also the main reason why 2HP shows a much higher delivery ratio. When intercluster links are not used, routes to nodes that are far away cannot be established and many packets are dropped.

As speed increases, we notice that the curve that corresponds to 2HP declines at first, but beyond 8 m/s the average hop count increases. With ordinary polling, data packets constantly make more and more hops until they

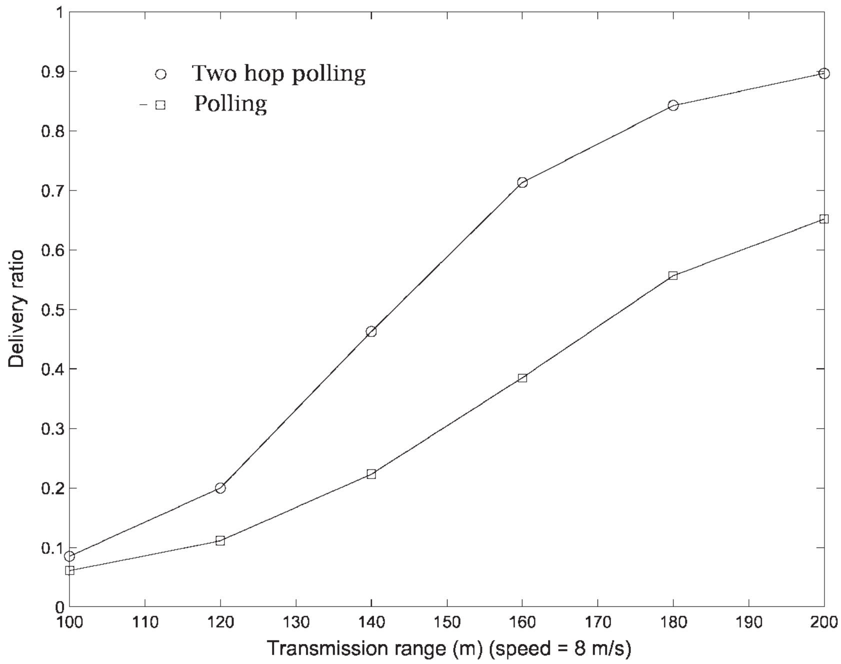

Fig. 5. Delivery ratio versus transmission range.

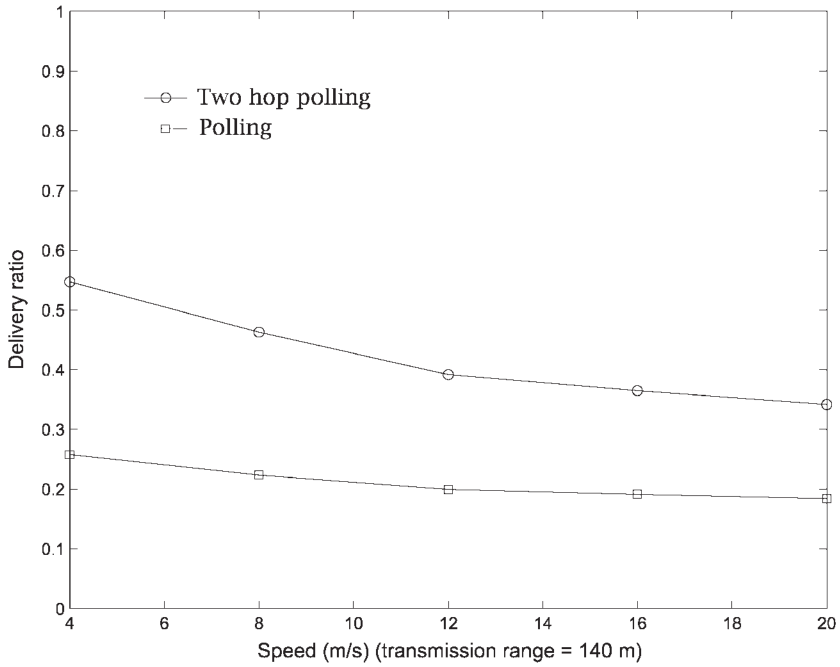

Fig. 6. Delivery ratio versus mobility.

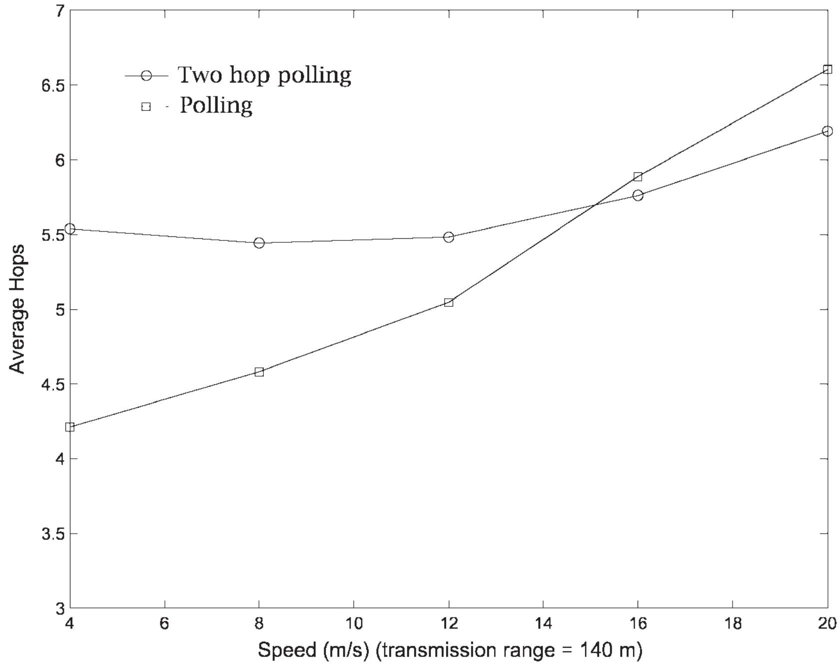

Fig. 7. Average hop count versus mobility.

reach their destination as speed increases. Of course, this phenomenon cannot be accounted to the fact that packets reach more distant destinations. In environments of high mobility, distant destinations become unreachable, because the routing mechanism cannot keep track of the constantly shifting topology. Bearing in mind from Fig. 6 that the delivery ratio declines, we come to the conclusion that this increase of the packets' hop count is an indication of loops.

Loops are being formed when two or more network nodes have an opposing view of the network topology, which results in packets traveling in circular routes, resulting in a great waste of network resources. Loops are formed because of inconsistencies in the routing tables, especially when a routing update has not propagated to every node in the network. FSR is based on the link state algorithm so long-term loops cannot exist. However, loops as transient phenomena may be formed, and they last until a routing update is propagated throughout the network, eliminating them. Bearing in mind from Ref. [5] that in FSR, routing updates are propagated one hop at a time, rather than being flooded; this process may take some time.

In our network model it was provided that a packet that made more than 100 hops was discarded because it clearly was involved in a loop. However, there is a possibility that packets involved in a loop may ultimately reach their destination before they exceed the 100-hops limit. These are the "lucky" packets that find a proper route toward their destination because a routing update eliminated the loop they were caught in before they exceeded the hop limit. These are the packets that caused the increase in hop count that affected both versions of polling. In Table I, we present the number of packets that were discarded because they exceeded the 100-hops limit. Totally, 90,690 packets asked to be served by the network.

As we see here, the version of the network that employed polling is more susceptible to loops than the one that uses 2HP. This is also implied by the average hop results. At high speeds, the scheme that does not use intercluster links shows a much higher average hop count, even though only packets headed toward near destinations are generally successfully delivered. This behavior is a consequence of two things. Messages containing routing updates propagating via intercluster links cover bigger portions of the network in less time. Thus, loops have a shorter time span. Also, as previously mentioned, when DGs are activated, the network appears to be less dynamic.

In our last figure, we compare the two schemes on an average delay basis. As we see in Fig. 8, with 2HP, data packets take more time to arrive at their destination, no matter the mobility. Thus, in delay terms, activating the distributed gateways results in worse performance. This behavior is owed to two facts. First, in the ordinary polling case, many packets are dropped, so the network effectively works under lighter load than that using 2HP. Second, and this is an inherent disadvantage of 2HP, the token passes

Table I. Number of Packets Discarded Because They Were Caught in Loop

	4 m/s	8 m/s	12 m/s	16 m/s	20 m/s
Polling	739	1316	1733	2058	2412
2HP	353	756	992	1397	1626

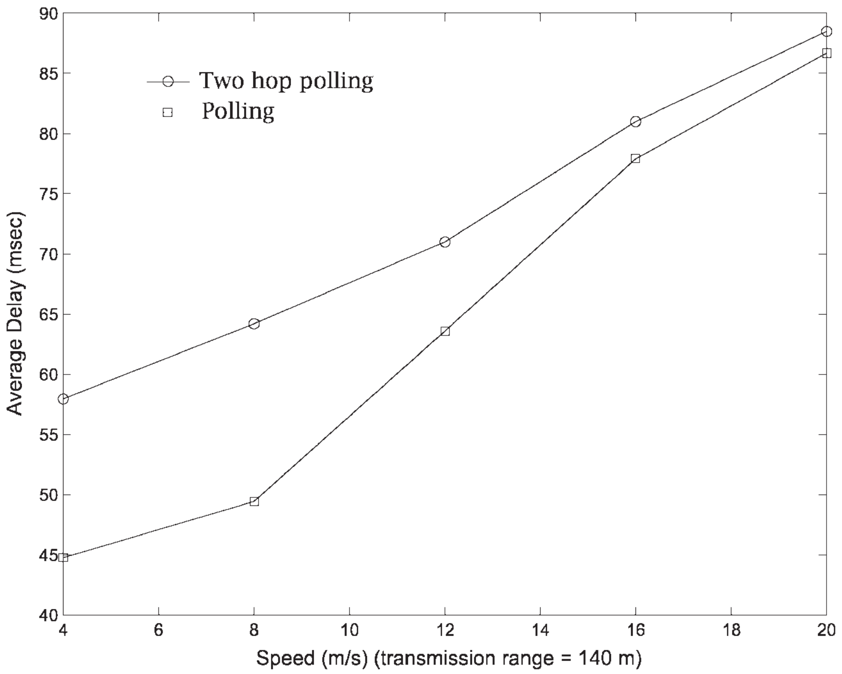

Fig. 8. Average delay versus mobility.

through more nodes in each polling cycle, resulting this way to nodes waiting longer to acquire the permission to transmit their data. However, as we recall from Fig. 7, with 2HP, packets reach more distant destinations, making more hops. This greater hop count translates to longer delays until the packets arrive at their destinations.

As speed increases, both curves follow an increasing trend. The appearance of loops results in a waste of network resources which hinders the transmission of data packets. Also, the packets that were temporarily caught in loops but eventually arrived at their destination strongly bias the results toward higher delay values (same effect as with the average hop count).

5. CONCLUSIONS

In this paper, we have described and analyzed the importance of intercluster links. Simulation studies on the network graph and on a full simulation environment showed that distributed gateways should be used whenever possible. To achieve this, we have proposed a scheme based on polling that does not add much to the complexity of multiple access. 2HP improves the network efficiency by improving the connectivity and lessening the probability of loop appearances. In most cases, 2HP managed to outperform ordinary polling.

REFERENCES

1. F. A. Tobagi and L. Kleinrock, Packet switching in radio channels: Part II—the hidden terminal problem in carrier sense multiple-access and the busy-tone solution, *IEEE Transactions in Communications*, Vol. 23, No. 12, pp. 1417–1433, 1975.
2. D. J. Baker and A. Ephremides, A distributed algorithm for organizing mobile radio telecommunication networks, In *Proceedings of the Second International Conference on Distributed Computer Systems*, Versalles, France, pp. 476–483, 1981.
3. A. Ephremides, J. Wieselthier, and D. Baker, A design concept for reliable mobile radio networks with frequency hopping signaling, In *Proceedings of IEEE*, Vol. 75, No. 1, pp. 56–73, 1987.
4. M. Gerla and J. Tsai, Multicluster, mobile, multimedia radio network, ACM-Baltzer *Journal of Wireless Networks*, Vol. 1, No. 3, pp. 255–265, 1995.
5. A. Iwata, C-C. Chiang, G. Pei, M. Gerla, and T.W. Chen, Scalable Routing Strategies for Ad-Hoc Wireless Networks, *IEEE Journal on Selected Areas in Communications*, Vol. 17, No. 8, pp. 1369–1379, 1999.
6. F-N. Pavlidou and G. Dimitriadis, Adaptive cluster link utilization based on a new two-hop approach, *Proc. 3GIS (3rd Generation Infrastructure and Services) Symposium*, Athens, Greece, 2001.
7. IEEE Standard for Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, Nov. 1997, P802.11.
8. D. Bertsekas and R. Gallagher, *Data Networks*, Prentice-Hall, Englewood Cliffs, NJ, 1987.
9. C. Perkins, *Ad Hoc Networking*, Addison-Wesley, Boston, MA, 2001.

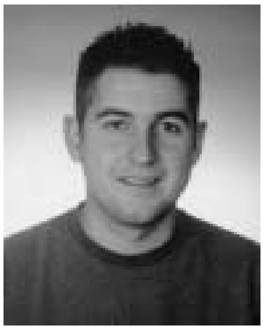

Gerasimos Dimitriadis received his Diploma in electrical and computer engineering from the Aristotle University of Thessaloniki, Greece, in 2001. He is currently working toward his Ph.D. degree in the same department. His research interests include medium access as well as routing in multihop wireless networks.

Foteini-Niovi Pavlidou received the Ph.D. degree in electrical engineering from the Aristotle University of Thessaloniki, Greece, in 1988 and the Diploma in mechanical-electrical engineering in 1979 from the same institution. She is currently an associate professor at the Department of electrical and computer engineering at the Aristotle University and is engaged in teaching for the under- and postgraduate programs in the areas of mobile communications and telecommunications networks. Her research interests are in the field of mobile and personal communications, satellite communications, multiple access systems, routing and traffic flow in networks, and QoS studies for multimedia applications over the Internet. She is involved with many national and international projects in these areas (Tempus, COST, Telematics, IST), and she has been chairing the European COST262 Action on "Spread Spectrum Systems and Techniques for Wired and Wireless Communications." She has served as a member of the TPC in many IEEE/IEE conferences and she has organized/chaired some conferences such as the "IST Mobile Summit 2002," the 6th "International Symposium on Power Lines Communications-ISPLC2002," the "International Conference on Communications-ICT1998," and so forth. She is a permanent reviewer for many IEEE/IEE journals. She has published about 80 papers in refereed journals and conferences. She has served as guest-editor on special issues such as "Mobile Ad Hoc Networks (MANETs): Standards, Research, Applications" in the *International Journal of Wireless Information Networks* and "Power Line Communications and Applications" in the *International Journal on Communications Systems*. She is a senior member of IEEE, currently chairing the joint IEEE VT&AES Chapter in Greece.