

Address Aggregation in Mobile Ad hoc Networks

Christopher F. Shiflet and Elizabeth M. Belding-Royer

Department of Computer Science
University of California, Santa Barbara
{cshiflet,ebelding}@cs.ucsb.edu

Charles E. Perkins

Communication System Laboratory
Nokia Research Center
charliep@iprg.nokia.com

Abstract—Current mobile ad hoc network (MANET) routing protocols operate with networks as flat address spaces, treating all nodes in the network as independent peers. However, these nodes are not always independent; they can form groups due to physical or environmental restraints. By taking advantage of these natural groupings, there is an opportunity to optimize the performance of MANET routing in these settings. This paper examines address aggregation in MANETs, focusing primarily on its ability to reduce routing overhead. We specifically describe the work necessary to implement address aggregation in the Ad hoc On-demand Distance Vector routing protocol (AODV). We found the benefits of aggregation to be quite significant, reducing overhead and delay by 42% to 82% and 30% to 62%, respectively.

I. INTRODUCTION

Mobile ad hoc networking (MANET) is a relatively new field of research that has been growing tremendously over the past few years. During this time, a number of routing schemes have been developed, with current research continuing to refine them. The reduction of routing overhead has been a key concern in creating these routing protocols and still remains an active area of research. Without better techniques to mitigate the effects of routing overhead, it seems that the scalability of MANETs is limited.

Current MANET routing protocols deal with networks as flat address spaces, treating all nodes in the network as peers. Every node is required to participate in an equal capacity. As such, each node must act as a router, maintaining individual routes to other nodes. Wired networks, on the other hand, do not generally treat all nodes as peers. Since few topological changes occur, static routers can be configured for even large groups of computers. This is accomplished elegantly through the concept of network and host bits. In Internet Protocol (IP) version 4, there are 32 bits in each address. A mask is used to indicate how many leading bits compose the network prefix. All addresses whose leading bits are the same as those specified by the prefix are part of the same sub-network, or subnet. This is the basic concept that underlies routing in the Internet today, namely Classless Inter-Domain Routing (CIDR) [4].

Although MANETs typically experience much more topological change than their wired counterparts, nodes are still not completely independent. They may, for example, be grouped due to physical or environmental constraints. Nodes might be located within an automobile, airplane, bus, cruise ship or train, to name a few. These groups of nodes create an opportunity for aggregation similar to that found in wired

networks. By forming subnets, larger quantities of nodes can be represented by a single route in lieu of multiple independent routes. This reduction in routing data results in a lessened need for control overhead, which can be especially beneficial in MANETs with their characteristically lower bandwidth. It also allows for a general reduction in the amount of routing table entries at each node, circumventing scalability issues inherent to the routing table's size.

As deployment of mobile networks, especially ad hoc networks, continues to increase, there is a growing potential for hierarchical applications. Military applications have already been seeking this functionality to more tightly couple nodes with similar objectives. Consumer applications may benefit as mobile nodes become more pervasive, for instance coupling many wireless devices within a car or grouping users within a cruise ship or airplane. In the latter example, an airline might send updated flight information to passengers aboard an aircraft. On an airplane, users desire uninterrupted connectivity throughout their flight. By leveraging MANET routing protocols, an airplane can dynamically maintain connectivity through any available data path, for instance, connecting through an 802.11 hot spot while docked at the airport gate. By grouping the nodes in such an aircraft into a subnet, the amount of routing overhead needed to route each node its flight information is dramatically reduced. Now a single route request and reply cycle takes place for the subnet, whereas an individual discovery cycle would have previously been necessary for each node.

In this paper we implement and examine the usefulness of address aggregation in MANETs, focusing primarily on its ability to reduce routing overhead. We specifically explain the work necessary to implement address aggregation in the Ad hoc On-demand Distance Vector routing protocol (AODV) [1]. Since the Internet Protocol (IP) was inherently designed to support hierarchical routing, the core operations of the routing protocol require only slight modification to support CIDR. The AODV RFC already includes specifications for all constructs necessary for the support of aggregation, including the network prefix field in both route replies and the routing table [7]. All necessary changes are related to the implementation of AODV rather than its specification, such as adding support for longest match first routing table lookups. Through these modifications, AODV is given the capability to address multiple nodes with a single route and consequently service multiple route requests with a single route reply.

The rest of this paper is organized as follows. First, we discuss the motivation for aggregating nodes in MANETs, followed by an explanation of the changes necessary to support aggregation. We then present our experimental setup, report our findings and conclude.

II. MOTIVATION

Many important networking technologies have progressed toward a hierarchical structure over time. Take for instance ARPANet, the precursor to the modern Internet. Early on, ARPANet consisted of only a few computers using a flat address space for identification. Each computer stored a route to all of the other computers, maintained by a simple routing protocol. As ARPANet grew larger, evolving into the NSFnet and later the commercial Internet, the number of nodes, and consequently routes, increased rapidly. Since every node was responsible for its own routing, this resulted in an increasingly heavy routing load for each computer in the network. In order to reduce the routing load, computers were aggregated into subnets that could be handled as a single entity for routing purposes. For example, a large university could have many computers that connect to the Internet through a single point. Rather than maintain individual routes for all of these computers, aggregation allows them to be addressed by their point of connection, also called their border or gateway router. This significantly reduces the amount of routing information in the network.

One benefit of reduced routing information is that communication capacity can be freed in the network. The maintenance of routes requires nodes to exchange topological information. This data is sent throughout the network, consuming part of the capacity that would otherwise be available for use by data traffic. As one might expect, aggregating routes can free a significant percentage of network bandwidth in large networks. Whether or not MANETs grow to extremely large sizes, their more restricted resource set will benefit from even small reductions in overhead.

In addition to lowering control overhead, aggregation can reduce the size of the routing table for a given network. Because subnet routes address multiple nodes, it is possible to represent connectivity to more nodes in fewer routing table entries, resulting in decreased routing table size. Generally speaking, a smaller routing table results in lower look up latency, which also improves routing performance. It is also generally accepted that routing table size is a significant limiting factor in terms of routing scalability.

In the next section, we explain the modifications necessary to support subnets in the AODV routing protocol, while the remainder of the paper quantitatively demonstrates the ability of aggregation to provide greater scalability in MANETs.

III. SUBNET SUPPORT

A number of nuances must be considered while implementing subnets in AODV. At its essence, aggregation requires a designated node within a subnet to become responsible for the routing of the remaining subnet nodes. This node,

called the subnet router, must consequently support AODV-specific routing semantics for the entire subnet, most notably the maintenance of a subnet sequence number. Furthermore, the formation and leadership of subnets is an open problem that must be addressed by any aggregated routing protocol. In the remainder of this section, possible solutions and the specific operation of subnets within our implementation of AODV are described in detail.

A. Subnet Membership and Leadership Selection

Subnet formation is an interesting and difficult problem. Because nodes are mobile, it can be challenging to group them in an effective manner. It is also necessary to properly configure the addresses of subnet nodes since aggregated nodes are determined by the relationship of their addresses. A dynamic system similar to MANETconf [5] could serve as an effective mechanism for maintaining subnet membership and addressing, where a node within a given subnet is configured to provide addresses appropriate for membership in its own subnet upon request. More centralized, client/server approaches could also prove effective. Alternately, subnets can be configured in a static manner by assigning the appropriate addresses, alleviating the need for a complicated grouping system.

In addition to grouping nodes, it is also necessary to determine leadership within the subnet. Similar to membership concerns, choosing a subnet leader can also be handled dynamically or statically. The only stipulation upon the selection of a leader is that there may be only one leader in each subnet. If, per chance, more than one leader appears in a subnet, it is necessary for one of the leaders to abdicate. There are a number of possibilities for determination of the leader. A leadership election is a strong candidate for a dynamic configuration technique. Specific events cause an election to occur, such as detecting the lack of a subnet leader or the existence of multiple subnet leaders. Upon commencement of the election process, nodes advertise information about themselves that is pertinent to their leadership abilities; for example, their processing speed or available power. After a period of time, the most adept node discovered during the election assumes the role of leader. Static leadership designation is also possible and much simpler than dynamic methods. A user designates a given node to act as leader for the subnet and the process is complete. No overhead in terms of bandwidth usage or time is required.

Although there are a number of interesting and possibly beneficial methods by which subnets could be dynamically formed, this specific implementation uses static methods to define both membership and leadership. Each node's IP address and routing options must be correctly configured by the user for it to operate as part of a subnet. Despite the typically dynamic nature of MANETs, statically configured subnets can still be both useful and desirable. As an example, situations can arise where multiple nodes are physically grouped, possibly in a car, airplane or other vehicle. In these contexts, a resource-rich node, such as a computer powered by the vehicle's

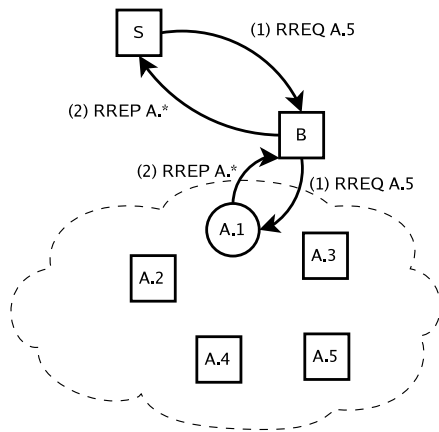


Fig. 1. In step (1), node S sends a RREQ for node A.5. The subnet router, A.1, sends a RREP for the entire subnet back to S in step (2).

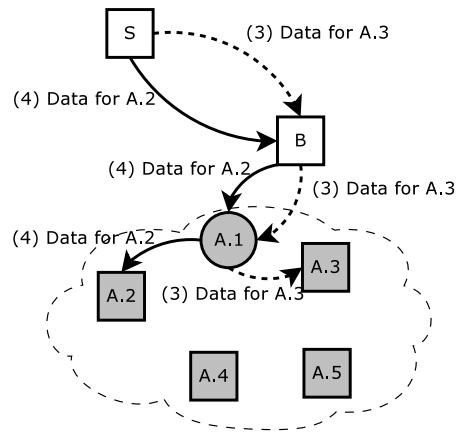


Fig. 2. In step (3), node S sends data to A.3 using the information obtained through step (2). In step (4), node S sends data to node A.2.

electrical system, can be configured to act as the subnet router. This removes much of the routing load from other nodes in the subnet. Statically configured subnets are also useful when deploying AODV to handle routing in transit networks, such as rooftop or mesh networks, since it is easy to statically determine subnets and their routers in this context.

B. Route Request/Reply Cycle

Due to AODV's reactive nature, a node typically executes a route discovery procedure when it desires to send data to an unknown destination. For single destinations, a route discovery proceeds in the following manner. First, the source node issues a RREQ for its desired destination. This RREQ is subsequently flooded throughout the MANET. Each flooding node creates a reverse route back to the source as the RREQ propagates. Once the destination node receives a RREQ matching its own address and adds the reverse route to its routing table, it creates a RREP containing the routing information it obtained from the RREQ. This RREP is subsequently unicast back to the source along the established reverse route. When the source node receives the RREP, it adds the enclosed route, known as the forward route, to its routing table and may begin communication with the destination node.

With the addition of subnet support to AODV, the route request and reply procedures change slightly. The RREQ is created and sent in the same manner, as seen in step 1 of figure 1, but the subnet router is now responsible to reply to RREQs for any node within its subnet. In order to accomplish this, the subnet router compares the destination of each RREQ with its own address and network mask. The subnet router then replies if the network bits of its address and the destination address are the same, as shown in step 2 of figure 1. The prefix field of the RREP is used to indicate the range of hosts that are reachable through this route [7]. Upon reception at the requester, this reply provides reachability information for every node in the subnet, indicated by the darkened nodes in figure 2. Steps 3 and 4 in figure 2 illustrate that data may be sent to multiple nodes within the subnet without the need for

additional route request and reply cycles. When using AODV without subnet support, each node to which data is sent causes the source to initiate another route discovery cycle.

C. Router Operation

By definition, a subnet router handles route requests reception of a request for a member of its subnet, the router issues a route reply providing a route for every node in the subnet, as described previously. Member nodes are not allowed to create such a reply. Instead, they must forward requests for their subnet on to their subnet router. Because subnet routes always pass through the router, a route to each node in the subnet is maintained by the router at all times. One way this can be done is by requiring subnet nodes to remain within a single hop of the router so that connectivity is maintained through hello messages.

In order to enable aggregation in AODV, most changes must be made in the subnet router. The router is responsible for all routes to the subnet as a whole; therefore, it is crucial that the router correctly maintain a sequence number for the subnet. The use of sequence numbers within AODV precludes the existence of routing loops by allowing nodes to distinguish between fresh and stale routing information. In order to support proper sequence numbers for the subnet, one option is for the router to use its personal sequence number as the sequence number of the subnet. Because the router responds with routing information pertaining to the whole subnet whenever it replies, the subnet router responses essentially take the place of the individual responses of the router. Under this scheme, it is important that members of the subnet do not respond on behalf of the subnet because this would require the subnet sequence number to be maintained in a distributed fashion throughout the entire subnet. Due to this inherent difficulty, member nodes must forward RREQs to their subnet router, allowing the router to reply.

As a byproduct of the subnet router replying to all RREQs for its subnet, all resultant routes pass through the router. Consequently, the router must maintain routes to all nodes




Fig. 3. This comparison of total control overhead shows that aggregation slows the growth of control overhead substantially.

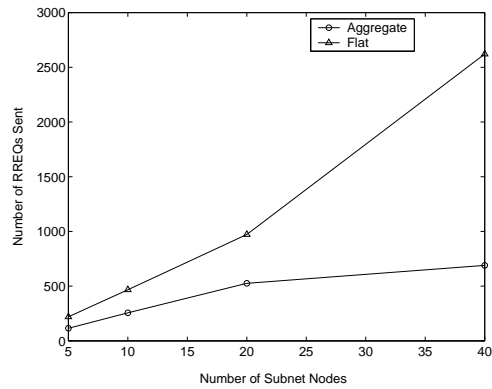


Fig. 4. This figure displays the RREQ component of control overhead, that parallels the trend found in overall overhead.

within its subnet. By doing so, data may arrive at the router destined for any of its nodes and the router will be able to forward it appropriately. This requirement also allows the router to detect when data has been sent to a non-existent or otherwise unreachable node, in which case it issues an Internet Control Message Protocol (ICMP) host unreachable message.

Up to this point, protocol interactions external to the nodes have been discussed, but the internal structure of the router also requires modification. Because the network now supports aggregation, a routing table look up must consider which nodes are encompassed by each entry. The only complication of this system is the case in which an address matches a given network, but is not routable through that network. In order to enable communication in these instances, routers must choose the longest network prefix that matches a given address. With this technique two routes with overlapping network prefixes can co-exist in the same routing table without hampering communication.

Collectively, these modifications embody all that is necessary for the support of address aggregation in AODV. The next section examines the quantitative advantage that can be provided by this support in terms of its ability to reduce routing overhead and delay.

IV. EXPERIMENTAL SETUP

As a motivating example for this experiment, recall the example of an airplane mentioned earlier. This airplane contains multiple nodes of an AODV network. Because of their constrained movement, nodes group together to form a subnet. This is a straightforward mechanism by which airlines might provide fully-routable Internet access in the most cost effective manner. AODV's dynamic nature enables nodes to access the Internet through a hot spot near the airport terminal while on the ground and seamlessly transition to using a satellite-based route once in the air. By doing so, the less expensive terrestrial link can be utilized when available while still maintaining the ability to provide connectivity through the entirety of the flight. Because nodes within an airplane are physically confined to the same space, they provide a perfect opportunity for aggregation.

To test the performance advantage provided by aggregating the nodes in our airplane scenario, we have compiled a series of simulations using the ns-2 network simulator [3]. The scenario simulates a situation in which an airline might send an instant message to some of the passengers aboard an airplane detailing flight schedule changes.

Our scenarios are configured with 100 independent mobile nodes where one of these nodes is a subnet router for a subnet ranging from 5 to 40 nodes. It is quite conceivable to imagine an airplane with 40 or even more passengers accessing the Internet [8]. These values also give us a good range over which we can identify the trend of efficiency improvement. Independent nodes move at speeds ranging from 0 to 10 m/s following the random waypoint mobility model. They are configured not to pause, but to remain in motion throughout the simulation. This environment is more difficult than required by our hypothetical situation; however, it is useful to shed light on other more mobile applications. All simulations are conducted using the AODV-UU implementation [6] with modification to support aggregation. All simulations run in an area of 1000 m x 1000 m where a single node outside of the subnet sends a short burst of data to every node in the subnet. The source node sends 3 packets of 512 bytes to each node within the subnet at a rate of one packet every 0.25 seconds.

V. RESULTS

To quantitatively gauge the affects of address aggregation we consider total routing overhead, routing overhead due to route requests, packet delivery ratio and delay as our primary metrics. Routing overhead is the number of AODV control packets sent through the course of the simulation, including route requests, route replies and route errors. The routing overhead caused by route requests is also considered independently because of its large role in the overall control overhead. Packet delivery ratio is the number of user data packets received divided by the number that were sent. Finally, delivery delay is the average time that elapses between sending a user data packet and its arrival at the destination, considering only properly received packets.




Fig. 5. Both aggregated and flat networks delivered over 90% of the data packets offered to the network. However, the delivery ratio in the aggregated networks is slightly better than that of the flat networks.

In simulation, it was found that overhead is substantially reduced in networks with support for aggregation while also providing a slight improvement in delivery ratio. Figure 3 clearly shows that the control overhead in the aggregated networks grows much slower than that of the flat networks. This trend is due to the reduction in route requests achieved through using aggregated routes, shown in figure 4. Since route requests are responsible for the majority of routing overhead in networks running AODV [2], their reduction causes the overall overhead to be reduced in a similar manner. Even at the smallest subnet size of five nodes, the number of route requests was reduced by 42%, ranging up to as high as 82% in subnets of 40 nodes.

In addition to reduced overhead, aggregated networks were able to deliver a slightly higher percentage of data packets than their flat counterparts as subnet sizes increased (figure 5). In subnets of 5 nodes, the delivery ratios are very similar, at 0.93 in aggregate networks and 0.94 in flat networks. As the subnets grew in size though, the delivery ratios in the aggregate networks became superior to those in the flat networks by 2% to 3%. This is caused by reduced contention for network resources in the aggregate networks since the routing overhead in these networks was reduced.

Delivery delay was also reduced in the aggregate networks by 30% to 61% (see figure 6). However, this trend was not so much dependent upon the size of the subnet as the fact that the flat networks required a route discovery cycle before sending data to each new node. The aggregate networks only needed to execute this cycle once.

VI. CONCLUSION

In this paper, we have demonstrated the motivation and modifications necessary to support address aggregation in MANETs. Specifically, support for aggregation was added to AODV utilizing static subnet member and leader designation




Fig. 6. Packet delivery delay remains fairly constant across different subnet sizes, reduced by 30% to 61% in the aggregate networks.

without modifying the pre-established specifications for prefix advertisement in the AODV RFC. With this implementation, we studied the characteristics of aggregation in MANETs by comparing against the corresponding original implementation and found that we were able to significantly reduce routing overhead in instances where nodes can be effectively aggregated. Overall routing overhead was reduced by 42% to 82%. Delay was also reduced by 30% to 61%. Considering the limited resources of MANETs, this finding provides significant motivation for future use of address aggregation in these networks. We plan to continue our research by determining the best methods for grouping nodes into subnets and evaluating their relative effectiveness. We also plan to explore different methods of leader selection, quantitatively evaluating both their costs and benefits.

REFERENCES

- [1] E. M. Belding-Royer and C. E. Perkins. Evolution and Future Directions of the Ad hoc On-Demand Distance Vector Routing Protocol. *Ad hoc Networks Journal*, 1(1), July 2003.
- [2] S. R. Das, C. E. Perkins, and E. M. Royer. Performance Comparison of Two On-demand Routing Protocols for Ad hoc Networks. In *Proceedings of the IEEE Conference on Computer Communications (INFOCOM)*, pages 3–12, Tel-Aviv, Isarel, March 2000.
- [3] K. Fall and K. Varadhan. ns Manual. <http://www.isi.edu/nsnam/ns/doc/>, 1999. The VINT Project.
- [4] V. Fuller, T. Li, J. Yu, and K. Varadhan. Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy. *RFC 1519*, September 1993.
- [5] S. Nesargi and R. Prakash. MANETconf: Configuration of Hosts in a Mobile Ad Hoc Network. In *Proceedings of IEEE Infocom 2002*, New York, USA, June 2002.
- [6] E. Nordstrom and H. Lundgren. AODV-UU Implementation from Uppsala University. <http://user.it.uu.se/~henrik/aodv/>.
- [7] C. Perkins, E. Belding-Royer, and S. Das. Ad hoc On-Demand Distance Vector (AODV) Routing. *RFC 3561*, July 2003.
- [8] J. Wallace. Boeing debuts in-flight, broadband Internet. *Seattle Post-Intelligencer*, January 2003. <http://www.msnbc.com/local/PISEA/104318.asp?0cv=BB10&cp1=1>.