

Mobility-Adaptive Protocols for Managing Large Ad Hoc Networks

Stefano Basagni

Center for Advanced Telecommunications Systems and Services (CATSS)

Erik Jonsson School of Engineering and Computer Science

The University of Texas at Dallas

E-mail: basagni@utdallas.edu

Damla Turgut

Sajal K. Das

Center for Research in Wireless Mobility and Networking (CReWMaN)

Department of Computer Science and Engineering

The University of Texas at Arlington

E-mail: {turgut, das}@cse.uta.edu

Abstract- In this paper, we propose a new protocol for efficiently managing large ad hoc networks, i.e., networks in which all nodes can be mobile. We observe that, since nodes in such networks are not necessarily equal in that they may have different resources, not all of them should be involved in basic network operations such as packet forwarding, flooding, etc. In the proposed protocol, a small subset of the network nodes is selected based on their status and they are organized to form a *backbone* (whence the name “backbone protocol” or simply *B-protocol* to our proposed solution). The B-protocol operates in two phases: first the “most suitable” nodes are selected to serve as backbone nodes, then the selected nodes are linked to form a backbone which is guaranteed to be connected if the original network is. The effectiveness of the B-protocol in constructing and maintaining in face of node mobility and node/link failure a connected backbone that uses only a small fraction of the nodes and of the links of the original networks is demonstrated via simulation. The obtained results show that both the selected backbone nodes and the links between them in the backbone are considerably smaller than the nodes and the links in the flat network.

I. INTRODUCTION

Networks comprised of wireless nodes all of which can be mobile—often termed *ad hoc networks*—have recently gained increasing popularity. The use of ad hoc networking technology is now shifting from the tactical/military scenario (widely deployed for over three decades) to all those situations in which a wired infrastructure is not viable.

Recent advances in processors, memory and radio technology have enabled ad hoc networks with potentially very large number of small, lightweight, and low-cost nodes. A typical application for this “larger” kind of ad hoc networks concerns *distributed microsensing*: Each node of the network is capable

of monitoring a given surrounding area (sensing), and coordinating with the other nodes in the wireless network to achieve a larger sensing task. These ad hoc *sensor networks* are an emerging technology of current interest.

Problems in ad hoc networking range from the definition of new routing protocols to techniques for Quality of Service provisioning, energy conservation, and privacy and security issues. The main difference with respect to solutions proposed for wired networking and cellular wireless networks is that in the ad hoc scenario all the nodes can be mobile, and thus no fixed infrastructure can be taken into account. With this fundamental challenge in mind, several distributed protocols have been proposed for ad hoc networks, which include solutions for multipoint communication such as routing [1], multicast [2] and broadcast [3, 4]), node/resource discovery [5], and so on. Many of these protocols rely on basic network services such as flooding (e.g., for route discovery) resource location, user tracking and geographic messaging that, being network-wide, induce significant bandwidth and energy overhead which are often unbearable. In order to implement these services efficiently, *scalable* protocols are needed that meet the critical requirements of ad hoc networks, namely, adaptiveness to the mobility of a potentially very big number of nodes, and minimizing the bandwidth and energy overhead associated with the transmission of network information.

The main obstacle to the realization of such protocols is due to the fact that existing solutions consider each node of the network equally suitable for any task. This imposes a “flat” vision of the network, where nodes that are considered equally endowed are willing to forward packets for any other node, independently of the current local status of the nodes (traffic congestion, battery life, memory overflow, etc.). In this framework, every node and every inter-node link are equally involved in the execution of a distributed task.

In this paper, we define a new protocol for the management

of critical network operation in support of basic network services protocols. Our solution is based on the observation that an efficient management of network resources can be obtained by deploying only a subset of the network nodes. In particular, those nodes whose *local status* allows them to guarantee reliable communication *a)* among themselves, and *b)* with any other node, will be selected to form a network dynamically superimposed over the flat network. Our proposed *backbone protocol*, called *B-protocol*, sets up and maintains a connected network (the *backbone* or simply the *B-network*) in face of node mobility and node/link failures. The B-network can then be used to convey at each node critical and time-sensitive network management information from all over the network, with minor overhead and in a timely manner. For instance, the location of a node or a specific resource/user could be determined with a simple query broadcasted over the B-network. Route discovery, and multipoint communication protocols could be easily implemented in a similar way.

The proposed B-protocol comprises two major tasks. 1) *B-nodes selection*, to select the backbone nodes (*B-nodes*). These nodes are in charge to “serve” all the other non selected nodes (we term these nodes F-nodes, i.e., nodes that belong to the flat network and not to the B-network). 2) *B-links establishment*, where backbone links (*B-links*, i.e., links among the B-nodes) are established so that the resulting B-network is always connected.

The task of selecting the B-node is performed at each node based on a node’s own *weight*, i.e., a real number ≥ 0 that each node constantly computes based on what is most critical to that node for the specific network application (e.g., node mobility, its remaining battery life, and its *connectivity degree*, i.e., the number of its neighbors). The highest the weight of a node, the more suitable that node is for being a B-node. Once a node *b* determines its role as a B-node, all its neighbors become the F-nodes served by *b*. B-nodes selection will be adaptive to node mobility, and in general to changes in its local status (as expressed by the node’s weight). B-links establishment determines the inter-B-nodes links to be established in order for the network to be connected.

The B-nodes selection protocol does not impose any limitation on the number of neighbors that each B-node can serve. For dense networks, i.e., network in which each node has a large number of neighbors, this can lead to non negligible management overhead, increased delays, decreased local throughput, and other similar problems. Therefore, we propose a modification of the B-protocol by imposing a realistic limitation on the number $k < n$ of F-nodes that each B-node can serve (n is the number of the nodes in the network). This solution can be useful when, due to technical restrictions or efficiency considerations, it is better for a B-node not to be overcharged with serving too many F-nodes. This is the case, for instance, of the Bluetooth technology [6], where it is best for each *master* to serve at most a fixed number of *slaves* (this number is 7 according to the current Bluetooth specifications).

We demonstrate the effectiveness of the proposed B-protocol and of its described modification through simulations in networks with up to 1000 nodes. We observe that the number of B-nodes in the B-network is considerably smaller than the total number, n , of nodes in the network (it is at *most* 15% of n) and that, independently on how the B-links are implemented, the number of the B-links is just a fraction of the number of the links in the flat network. When limitations are imposed on the number of the F-nodes that each B-node can serve we obtain similar results.

As a final note, we mention that there exist solutions in the ad hoc networks literature to connect a number of selected nodes to form a backbone. These solutions are mainly based on the concept of “spine,” which is a connected structure made up of nodes that are neighbors in the flat network (for details and further references, see [7]). Such structures are difficult to maintain in face of mobility and/or node/link failures; to the best of our knowledge, no investigation has been made that demonstrates that the number of “spine-nodes” is consistently smaller than the total number of nodes in the network, especially when the number of nodes is large.

The rest of the paper is organized as follows. Section II describes the B-protocol. The section terminates with the list of properties that are satisfied by our protocol. The following Section III demonstrates the effectiveness of our protocol in building a backbone of a small fraction of the nodes and of the links of the flat network. Simulation results are presented both for the case in which each B-node can serve any $k < n$ F-nodes, and for the case in which k is instead bounded by a constant (we consider two cases for $k \leq 3$ and $k \leq 7$). Section IV concludes the paper.

II. B-PROTOCOL DESCRIPTION

In this section we describe the protocol that implements the B-nodes selection and the corresponding B-link establishment that dynamically maintain the B-network on top of the flat network. For details and the pseudo-code of the protocol procedure the reader is referred to [8].

A. B-nodes selection

The protocol for the selection of the B-nodes is executed at each node in such a way that just by knowing its own identifier (ID) and weight, as well as the IDs, the weights and the role (either B-node or F-node) of the one-hop neighbors, a node decides autonomously if it is going to be a B-node or an F-node.

In the following, let us illustrate the protocol operations with an example (see Figure 1 where, for the sake of discussion, we have identified nodes with their weights).

As soon as a node starts its operation it executes a routine to identify the the neighbors and their weights, and consequently decides if it is going to be a B-node or an F-node. Specifically, a node is going to be a B-node if it has the highest weight among all its neighbors (that are not served by another B-node). Thus, nodes 8, 10 and 12 in the network depicted in Fig-

ure 1 (left) will declare themselves as B-nodes by broadcasting a corresponding message. On receiving this message, node 6 immediately decides to be served by node 8, while nodes 2 and 3 need to wait for a message from node 9 whose weight is bigger than 8. Node 9, however, having received from B-node 10 that it is going to be part of the backbone, decides to be served by 10. The same happens to node 7 which has received a message from B-node 12. Now, nodes 2 and 3 have received the message stating that node 9 will be served by B-node 10, therefore they decide to be served by B-node 8.




Figure 1: A flat network (left) and the corresponding B-network (right). The squared nodes are the B-nodes. The “thicker” links are the B-links.

An important part of the B-node selection protocol has to do with the choice of the weights, dealing with their variation in time, and coping with node mobility. These tasks have been introduced and discussed in [9, 10, 11], and the corresponding techniques apply to the B-nodes selection process in a completely similar way.

B. B-links establishment

Given the B-nodes as selected through the application of B-nodes selection, here we define the rules for connecting them in such a way that if the (flat) network is connected, the resulting B-network is connected as well.

These rules are stated by the following theorem, first proven in [12].

Theorem 1 *Given a set B of network nodes such that no two of them are neighbors and every other node has a link to a node in B , then a connected backbone is guaranteed to arise if each node in B establishes links to all other nodes in B that are at most three hops away. Moreover, these links are all needed for the deterministic guarantee in the worst case, in the sense that if any of them is left out then it is not true anymore that the arising backbone is connected for any underlying flat network.*

We notice that a B-node does not need to know neither the identity nor the number of B-nodes at most three hops away. In order to guarantee that the message will reach the intended destination through the backbone, a B-node has only to make sure that a message sent on the B-network will reach three hops away. This implies that the B-nodes connection protocol is not only adaptive to mobility, but also extremely bandwidth and energy efficient, since there is no need to keep at each node any information regarding the B-nodes three hops away. The

application of this rule to the simple network of Figure 1 (left) produces the backbone depicted in Figure 1 (right).

C. Properties of the B-protocol

We conclude the description of the B-protocol by listing some of its unique properties.

1. To run B-nodes selection locally, each node in the (flat) network needs to know *only* its one hop neighbors. This is the minimum amount of information possible needed for the selection to produce the set of the B-nodes, which induces the minimum possible overhead.
2. B-links establishment is run at each B-node *only*, with no knowledge of either the identity or the number of the surrounding B-nodes. A B-node needs only to send packets to all and only other B-nodes at most three hops away. Again, this induces the minimum possible overhead.
3. Every B-node serves a number of F-nodes each of which is at most one hop away from the B-node. The weight-based B-nodes selection protocol guarantees that all the F-nodes are served by one neighboring B-node. This “covering” guarantees timely and cost-effective communication to/from an F-node from/to the serving B-node.
4. No two B-nodes are neighbors (in the flat network). This guarantees a well-scattered set of B-nodes which can cover all the remaining nodes.
5. B-nodes selection is based on the node current status (expressed by a node’s weight) that may vary in time, rather than invariants like a node’s unique identifier. This always guarantees that the “best” nodes in the network serve as B-nodes.
6. The B-network is always connected (provided that the underlying flat network is connected). This guarantees that what is reachable via the flat network is also more effectively reachable via the B-network.
7. Our solution takes into account different technologies and mechanisms that can be used to link the B-nodes among themselves to form the B-network. Thus, when nodes enabled with power control are deployed, or in the case a node can use both omni and directional antennas, our solution adapts to the fact that two B-nodes can talk directly to each other. When this is not possible, we extend the original protocol to map the “virtual” links among the B-nodes over at most three physical links that involve also F-nodes.

III. SIMULATIONS RESULTS

We have simulated the B-protocol to demonstrate its effectiveness in building a B-network made up of some of the “best” nodes of the flat network (here “best” is according to the criteria defined in the previous sections). The results show that

not only the the number of B-nodes is just a small fraction of the number n of the network nodes, but also that the number of B-links is very small as well.

We used a simulator of an ad hoc network, implemented in C++.¹ The $n = 100, 200, \dots, 1000$ nodes of the ad hoc network can freely move around in a rectangular region (modeled as a grid) according to the following mobility model. (To ease the modeling, the node movements are discretized to grid units with a grid unit = 1 meter.) Each time it moves, a node determines its direction randomly, by choosing between its current direction (with 75% probability) and uniformly among all other directions (with 25% probability). The node then moves in the chosen direction according to its current speed. When a node hits a grid boundary, it bounces back into the region with an angle determined by the incoming direction.

The fixed transmission range of each node (250m) and the grid sides (in meters) have been chosen to obtain good network connectivity. Independently of the number of nodes, more than 98% of the time, after network topology changes, the network was connected. Each link is modeled by a FCFS queue with service time as the packet transmission time characterized by a bandwidth of 1 Mbps.

Each protocol packet contains the time-stamped, node identified weight of the sending node. All packets are intended for the one-hop neighbors only (i.e., no packet is forwarded further).

The measures investigated concern:

- The fraction (%) of B-nodes that form the backbone, in the three cases when a B-node has no limitation on the number $k < n$ of F-nodes that it can serve (“unbounded k ”), and when instead $k \leq 3$ and $k \leq 7$. The first bound for k has been chosen since we observed that, on average, the network degree (maximum number of neighbors for each node) lies between 5 and 7, and thus we wanted to investigate the case in which the restriction really imposes a limitation on the B-nodes selection protocol. (When we have to select k among $> k$ neighbors, we choose those with the smaller weight.)

In Figure 2 we show the fraction of the B-nodes involved in the B-network as the percentage with respect to the number n of the nodes in the flat network. In the unbounded case, only at most 15% of the network nodes are selected to serve as B-nodes. The weight-based B-nodes selection process described in the previous section guarantees that these are the best suitable nodes for this role. Basically the same result is obtained when imposing at most 7 F-nodes per B-node (the “Bluetooth restriction”). As expected, things changes when $k \leq 3$. Since each node has, on average, a number of neighbors that ranges between 5 and 7, the protocol is forced to select only 3 nodes (based on their weight), resulting in a greater number of B-nodes. The number of B-nodes is however not bigger than twice the number of B-nodes in the unbounded case.

¹ Currently, our study is limited to network-layer details, thus no link- or physical-layer are modeled.




Figure 2: Number of B-nodes (%) with respect to the number of the network nodes).

- The fraction (%) of backbone links, in the three cases mentioned above, when a direct link between B-nodes at most three-hops away can be established without involving intermediate F-nodes (the direct link could be implemented using, for instance, power control, or a directional antenna).

As depicted in Figure 3, the number of B-links is never more than 15% of the number of the links in the flat network. This holds even in the case in which a B-node can serve no more than 3 F-nodes (case $k < 4$ in the picture), in which a higher number of B-links is expected. It is interesting to observe that when k is unbounded, or strictly bounded by 8, the larger is the network (in terms of number of nodes), the smaller is the fraction of B-links (when $n \geq 500$ they are less than 5% of the links of the flat network).




Figure 3: Number of B-links (%) when a physical link between any two B-nodes can be established directly.

- The fraction (%) of backbone links, in the three cases mentioned above, when a direct link between B-nodes at most three-hops away *cannot* be established. In this case we have

implemented the virtual link among two B-nodes by a corresponding physical path with at most three links.

Figure 4 shows that the number of the B-links is never more than 25% of the number of the links in the flat network. The increase with respect to case 2 is justified by the fact that now every link between two B-nodes is mapped onto either two or three physical links, that are considered B-links. However, the fraction of the B-links is still considerably smaller than the number of the link in the flat network. As in the previous case, we notice that when k is unbounded and $k < 8$, for networks with larger number of nodes (500 and up), the number of B-links is less than 5% of the number of links in the flat network.




Figure 4: Number of B-links (%) when a link between B-nodes is implemented by a physical path with at most three hops.

All the simulations run for a time long enough to achieve a confidence level of 95% with a precision within 5%.

IV. CONCLUSIONS

This paper presented the B-protocol for ad hoc networks, a novel scalable protocol which, by constructing and maintaining a backbone (B-network) made up of selected nodes and links of an ad hoc networks, can be used to efficiently implement network management operations and several other applications (routing, user tracking, geocasting, etc.). The main strength of the B-protocol stems from the fact that, by generating the minimum possible overhead for the construction and maintenance of the B-network, it always selects the B-nodes that are best suited to carry network management information, without affecting the performance of less endowed nodes. We have observed, through the use of extensive simulations on ad hoc networks with up to 1000 nodes, that the number of B-nodes and B-links is just a small fraction of the total nodes and links in the flat network, respectively. This demonstrates that, beyond being an excellent solution for dealing with mobility and node/link failures, the B-protocol is the ideal solution for very large ad hoc networks.

REFERENCES

- (1) E. M. Royer and C.-K. Toh. A review of current routing protocols for ad hoc mobile wireless networks. *IEEE Personal Communications*, 6(2):46–55, April 1999.
- (2) E. M. Royer and C. E. Perkins. Multicast using ad-hoc on-demand distance vector routing. In *Proceedings of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking, MobiCom'99*, pages 207–218, Seattle, WA, 15–20 August 1999.
- (3) S. Basagni, D. Bruschi, and I. Chlamtac. A mobility transparent deterministic broadcast mechanism for ad hoc networks. *ACM/IEEE Transactions on Networking*, 7(6):799–807, December 1999.
- (4) S.-Y. Ni, Y.-C. Tseng, Y.-S. Chen, and J.-P. Sheu. The broadcast storm problem in a mobile ad hoc network. In *Proceedings of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking, MobiCom'99*, pages 151–162, Seattle, WA, 15–20 August 1999.
- (5) R. Castaneda and S. R. Das. Query localization techniques for on-demand routing protocols in ad hoc networks. In *Proceedings of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking, MobiCom'99*, pages 186–194, Seattle, WA, 15–20 August 1999.
- (6) J. C. Haartsen. The bluetooth radio system. *IEEE Personal Communications*, 7(1):28–36, February 2000.
- (7) B. Das and V. Bharghavan. Routing in ad-hoc networks using minimum connected dominating sets. In *1997 IEEE International Conference on Communications. ICC'97.*, pages 376–380, Montreal, Que., Canada, 8–12 June 1997.
- (8) S. Basagni, D. Turgut and S. K. Das. A Scalable Backbone Protocol for Managing Large Ad Hoc Networks Technical Report UTDCS-07-01, Department of Computer Science, The University of Texas at Dallas, February 2001.
- (9) M. Chatterjee, S. K. Das, and D. Turgut. An on-demand weighted clustering algorithm (WCA) for ad hoc networks. In *Proceedings of IEEE Globecom 2000*, pages 1697-1701, San Francisco, CA, November 27–December 1 2000.
- (10) S. Basagni. Distributed clustering for ad hoc networks. In A. Y. Zomaya, D. F. Hsu, O. Ibarra, S. Origuchi, D. Nassimi, and M. Palis, editors, *Proceedings of the 1999 International Symposium on Parallel Architectures, Algorithms, and Networks (I-SPAN'99)*, pages 310–315, Perth/Fremantle, Australia, June 23–25 1999. IEEE Computer Society.
- (11) S. Basagni. Distributed and mobility-adaptive clustering for multimedia support in multi-hop wireless networks. In *Proceedings of the IEEE 50th International Vehicular Technology Conference, VTC 1999-Fall*, volume 2, pages 889–893, Amsterdam, The Netherlands, September 19–22 1999.
- (12) I. Chlamtac and A. Faragó. A new approach to the design and analysis of peer-to-peer mobile networks. *Wireless Networks*, 5(3):149–156, May 1999.