

TITAN: On-Demand Topology Management in Ad Hoc Networks

Cigdem Sengul
sengul@uiuc.edu

University of Illinois Urbana-Champaign

Robin Kravets
rhk@cs.uiuc.edu

I. Introduction

An ad hoc network is a multi-hop wireless network that is established by a group of mobile nodes without depending on any infrastructure. Due to the disconnected nature of such mobile nodes, a fundamental problem in ad hoc networks is energy-efficient operation to extend the lifetime of the nodes and the network. A promising strategy is to reduce the power consumption of the wireless interface since it is a significant contributor to the overall energy consumption. Essentially, while traffic load defines energy consumption by the wireless interface during active communication [1, 2], idle-time energy dissipation dominates total system energy consumption in the presence of low to moderate traffic [3, 4]. To this end, current approaches allow nodes to switch to a power-save mode where they spend most of their time in a low-power sleep state. However, allowing all nodes to operate in power-save mode imposes additional delay on all communication and can severely limit the capacity of the network as load increases [4]. To compensate for these limitations, some nodes can stay in active mode and serve as stable relays in the network to support low delay and high throughput [3, 4, 5]. Since the choice of nodes that remain active determines both energy consumption and communication quality, the main challenge to any idle-time energy conservation protocol is selecting the set of active nodes through which all traffic flows.

Approaches for selecting active nodes fall into two classes: proactive and reactive. The proactive approach, known as *topology management* [5], builds a forwarding backbone of active nodes, but does not tie the choice of backbone nodes to traffic in the network and so, requires these nodes to stay awake even if they are not participating in routing. A reactive approach to this problem is *on-demand power management (ODPM)* [4], which allows nodes to stay in power-save mode as long as they are not used for routing. However, ODPM does not use any knowledge about the power-management mode of nodes when choosing routes and may result in unnecessary activation of some nodes.

Given the limitations of current approaches, we propose TITAN (Traffic-Informed Topology-Adaptive Network) that combines the benefits of reactive and proactive approaches. From reactive approaches, TITAN allows current network traffic to drive the choice of forwarding nodes. Additionally, as in proactive approaches, once a node is chosen as a forwarding node, it is favored over power-saving nodes for future routes. Using a cross-layer approach by basing decisions for maintaining a forwarding backbone on information from routing and MAC layers, TITAN provides on-demand topology management. Essentially, each node independently decides to join the backbone as route requests flow in the network based on its current power-management mode and local neighborhood information. Initial simulation results show that TITAN successfully maintains a forwarding backbone on-demand, allowing nodes that are not required for forwarding to stay in power-save mode. The improvement in energy savings results from eliminating control overhead to build and maintain the forwarding backbone. Additionally, due to its on-demand nature, TITAN saves energy by adapting to traffic and allowing active nodes to disconnect from the forwarding backbone if they are not actively forwarding. Furthermore, TITAN achieves high communication performance (i.e., high delivery ratio and acceptable delay), while providing significant energy savings.

The remainder of the paper is organized as follows. Section II presents the protocol design of TITAN. Section III shows the effectiveness of our protocol via simulations. Finally, Section IV concludes with future directions.

II. TITAN: On-Demand Topology Management

To support on-demand topology management, the choice of nodes on a forwarding backbone should be driven by the network traffic. Therefore, a forwarding backbone in TITAN is defined as a set of *active* nodes that serve as a source, destination or relay for *active* flows in the network. The overhead of main-

taining such a forwarding backbone should be kept to a minimum to ensure that any energy savings is not compromised by the energy cost of maintenance. To this end, TITAN provides implicit and reactive topology management by tying the decisions about nodes in the backbone to routing choices in the network and the current power-management mode of nodes. Essentially, each node in TITAN independently decides how to participate in route set-up. Once a route is selected, TITAN reactively selects all nodes along that route to join the backbone by transitioning these nodes into active mode. Given this design, a node's decisions as to how to participate in routing impact the chance of that node being selected in a route and so the chance of joining the backbone. Additionally, the novelty of TITAN comes from its ability to implicitly direct traffic to routes through current nodes on the forwarding backbone. Furthermore, TITAN dynamically adapts the forwarding backbone to the current traffic allowing nodes to connect and disconnect from the backbone based on routing decisions.

TITAN is designed to work with an on-demand routing protocol (e.g., DSR [6] or AODV [7]) where the source initiates a route discovery by flooding the network with Route Requests (RREQs). While no changes are made to the routing protocol, TITAN impacts a node's decision as to when to forward a RREQ. These decisions are based on two criteria. First, TITAN aims to ensure that only one node in a given area is active. Therefore, each node monitors the power management mode of the nodes in its neighborhood. If there is already an active node in the neighborhood, a power-saving node defers forwarding the RREQ to allow the backbone node to respond first. Second, although TITAN aims to wake up as few new nodes as possible for a new flow, bounding the amount of time a power-saving node backs off enables the choice of shorter routes through power-saving nodes. Therefore, a power-saving node should still participate in route discovery as determined by how long the node defers the RREQ. However, to reduce the number of redundant RREQs in the network, a power-saving node that has heard a Route Reply (RREP) for a particular flow cancels sending the buffered RREQ. Assuming the destination sends a RREP to only the first RREQ, this design ensures that the backbone nodes that forward the RREQ immediately dominate the route discovery process.

Forwarding backbone maintenance in TITAN is realized by three cooperating mechanisms: 1) a *back-off decision mechanism*, 2) a *back-off scheduling mechanism* and 3) and *neighbor discovery*. Using the back-

```

FORWARD-RREQ-PSM ()
1   $r_i = \text{UNIFORM-RAND}(0, 1)$ 
2  if  $\alpha_i^* \geq 0$  and  $r_i < p_i(*)$ 
3    then Buffer RREQ for one beacon interval
4    else Send RREQ

 $(*)p_i = \begin{cases} 1 - \frac{1}{\delta_i^*}, & \text{if } \alpha_i^* = 0 \\ 1 - \frac{1}{\delta_i^* \alpha_i^*}, & \text{otherwise.} \end{cases}$ 

 $\delta_i^*$  and  $\alpha_i^*$  are the number of all neighbors
and active neighbors of node  $i$  not counting the
node that sent the RREQ.

SEND-BUFFERED-RREQ( $rreq$ )
1  if CHECK-RREP( $rreq.dest \rightarrow rreq.src$ )
2    then Cancel RREQ
3    else Send RREQ

```

Figure 1: Pseudo-code for TITAN integrated with IEEE 802.11 PSM sleep scheduling protocol.

off decision mechanism, a power-saving node defers forwarding a RREQ probabilistically based on the number of active and all neighbors. Essentially, if a power-saving node already has at least one active neighbor, it backs off. Additionally, a node that does not have any active neighbors still backs off based on the number of its neighbors. To achieve this behavior, each power-saving node uses a simple increasing function of active and all neighbors to back-off from forwarding a RREQ (see p_i in Fig. 1). The neighborhood information is obtained via neighbor discovery, which passively monitors neighbors to determine their presence and power management mode. Finally, the back-off scheduling mechanism determines the length of the back-off interval for power-saving nodes that decide to back-off. Essentially, the amount of delay introduced via back-off scheduling impacts when a RREQ arrives at its destination, hence discovery of routes. Using these three mechanisms, TITAN impacts the decisions about which nodes are selected for the backbone, while relying on ODPM to manage transitions between active and power-save modes.

Our prototype of TITAN is built on DSR, ODPM and IEEE 802.11 PSM [8] for sleep coordination. While TITAN in essence is not limited to any sleep coordination mechanism to perform back-off scheduling, different possibilities for lengths of back-off intervals are mainly determined by how power-saving nodes schedule their on-off times. The reason for our choice of IEEE 802.11 PSM is three-fold. First, although not specifically designed for ad hoc networks, IEEE 802.11 PSM is the standard protocol for power

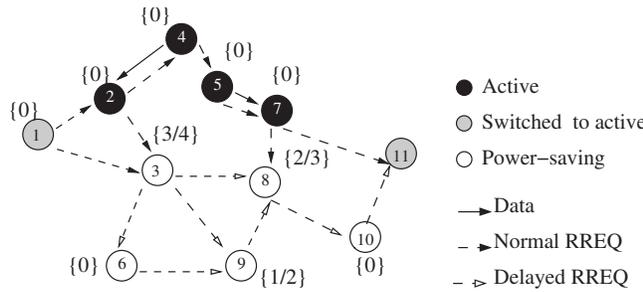


Figure 2: Example network with probability assignments based on the back-off decision mechanism. The RREQ propagation is shown for the case when both node 3 and 8 back off from forwarding.

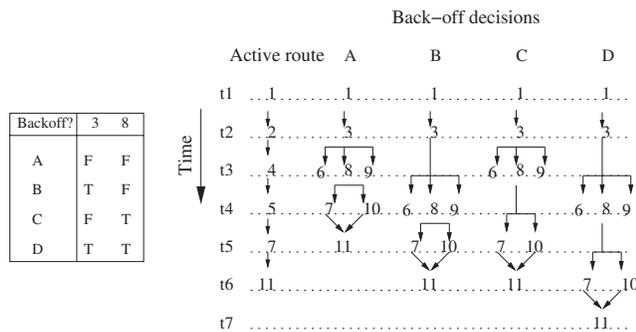


Figure 3: Timeline of RREQ propagation based on back-off decisions. Active route is found in condition D, power-saving route is found in condition A, and a power-saving or active route may be found in conditions B and C.

management. Second, it has a complete solution for broadcast communication in power-saving networks compared to [9]. Third, it does not assume the existence of any additional hardware such as a wake-up signaling radio as in [10].

Each power-saving node runs the distributed algorithm in Fig. 1, which implements the back-off decision and scheduling mechanisms. Fig. 2 illustrates a simple example of forwarding backbone maintenance in TITAN. In the example network, nodes 2, 4, 5, 7 are the active nodes that form the current forwarding backbone. When node 1 sends a RREQ for node 11, the active nodes send RREQs as normal, while power-saving nodes 3, 6, 8, 9, 10 may back off. Based on RREQ propagation, TITAN may discover three possible routes for flow $1 \rightarrow 11$: r_1 : 1 - 2 - 4 - 5 - 7 - 11 and $r_2(3)$: 1 - 3 - 8 - 7(10) - 11. While r_1 consists of only active nodes, r_3 consists of only power-saving nodes. The RREQ propagation in TITAN is affected by the RREQ back-off probability assignments at each node (shown in parentheses in Fig. 2). Back-off decisions

taken by power-saving nodes 3, 6, 8, 9 and 10 determine the time a RREQ reaches the destination. Based on these probabilistic decisions, TITAN may find either a shortest route or an active route. For example, TITAN finds r_2 or r_3 if both nodes 3 and 8 do not back off. In this case nodes 3 and 8 join the backbone, while node 10 may join the backbone if it captures the channel before node 7. TITAN finds r_1 if both nodes 3 and 8 back off, in which case the forwarding backbone stays the same. However, if only node 3 or node 8 backs off, r_1 , r_2 or r_3 contend for the channel at a meeting point (see Fig. 3). The winner of the channel determines if new nodes should join the backbone. Therefore, TITAN provides a natural selection between shorter routes with power-saving nodes and longer routes with active nodes. Essentially, as long as the time to traverse the longer routes does not exceed the accumulated delay from power-saving nodes in shorter routes, the destination replies to a longer route with active nodes. However, if this is not the case, the destination may reply to a shorter route with power-saving nodes, which switch to active mode once they are selected as relays.

III. Performance Evaluation

The goal of our evaluation is to show that TITAN does not degrade communication quality and conserves energy compared to ODPM and a topology management protocol such as Span [5]. We also evaluate the performance when all nodes are active (Active).

The effectiveness of TITAN as an on-demand topology management protocol can be evaluated by its impact on forwarding backbone maintenance, energy conservation and communication performance. The metric of interest for characterizing the forwarding backbone is the forwarding backbone size, which is defined as the average number of active nodes in a unit time interval. The performance in terms of saving energy is evaluated using energy goodput (bit/J), which is defined as the ratio of total bits transmitted to total energy consumed (i.e., the energy spent for data communication including the routing and MAC layer overhead). The computation of total bits transmitted uses data packets only (i.e., does not include control packets). Finally, we use the data delivery ratio to measure communication performance. Data delivery ratio quantifies the packet loss rate and is calculated as the ratio of data packets delivered to the destinations to data packets sent by the sources.

We implemented our prototype of TITAN in the ns-2 network simulator using the CMU wireless ex-

tension [11]. In our simulations, all nodes communicate with half-duplex wireless radios that conform to IEEE 802.11-based wireless radios with a bandwidth of 2Mbps and a nominal transmission radius of 250m. We use the energy model in [5] with transmit, receive, idle and sleep powers as 1.4W, 1W, 0.83W and 0.13W respectively. Our simulation results represent an average of five runs with identical traffic models, but different randomly generated network topologies.

We compare TITAN to Span using the implementation provided by Chen et al [5]. However, the current implementation of Span is coupled with a geographical routing protocol. Therefore, we evaluate the performance of Span with geographical routing and ODPM and TITAN with DSR. The simulations with Active also uses geographical routing to eliminate control overhead from DSR. Additionally, we do not simulate mobility in the network to avoid any control overhead from mobility. Therefore, both Active and Span protocols are at an advantage in terms of routing overhead compared to TITAN and ODPM. Although an exact comparison between Span and TITAN is not possible due to the use of different routing protocols, the simulation results in this section still provide an understanding of how each protocol performs in terms of energy conservation and communication quality.

In our simulations, 100 forwarding nodes and 10 source and 10 destination nodes are simulated using the Span-topology as described in [5] in a $1000m \times 1000m$ static network. Based on the Span-topology [5], source and destination nodes are placed, uniformly at random, on each of two 50 meter-wide full-height strips located at the left and right sides of the network. A source on the left side must send to a destination on the right side and vice versa. The initial positions of the 100 forwarding nodes are chosen uniformly at random in the entire network. The traffic is CBR (Constant Bit Rate), and the start time for each flow is determined randomly between 20s and 120s. Each simulation runs for 600s.

We first evaluate the forwarding backbone maintained by each protocol. Based on Span [5] simulations, source and destination nodes are not counted as a part of the backbone. Therefore, not including source and destination nodes, simulation results show that TITAN uses approximately 20% fewer nodes on average compared to Span, while the difference between TITAN and ODPM is more significant (see Fig. 4). The number of nodes involved in the forwarding backbone in TITAN is reduced because of two reasons: (i) active nodes forward RREQs ear-

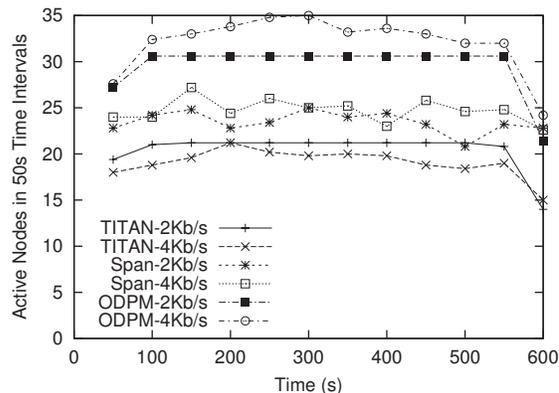


Figure 4: Number of active nodes vs. time in Span-topology, static network. TITAN achieves the smallest forwarding backbone through implicit topology management.

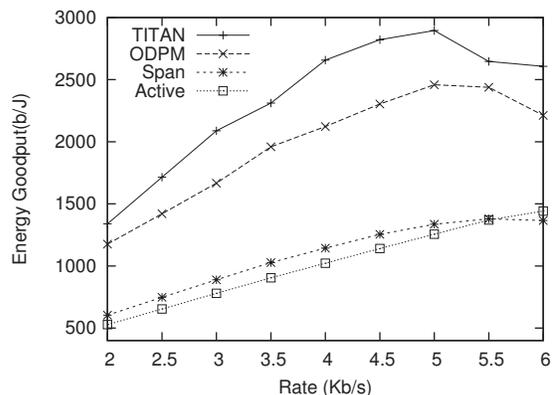


Figure 5: Energy Goodput vs. traffic load in Span-topology, static network. While TITAN provides the highest energy savings, Span's performance degrades as the traffic increases.

lier than power-saving nodes and (ii) the destination only replies to the first RREQ. Additionally, as the network traffic decreases towards the end of the simulation runs, the size of the backbone in ODPM and TITAN decreases. This is due to active nodes disconnecting from the forwarding backbone and switching to power-save mode as they are no longer required to forward traffic. This behavior is not observed in Span due to its proactive operation. Furthermore, TITAN maintains a forwarding backbone comparable to Span, although TITAN does not use three-hop connectivity information that is available to SPAN for backbone maintenance. TITAN uses information from routing and MAC layers and makes more educated decisions about which nodes are necessary in the network.

Next, we evaluate TITAN in terms of saving en-

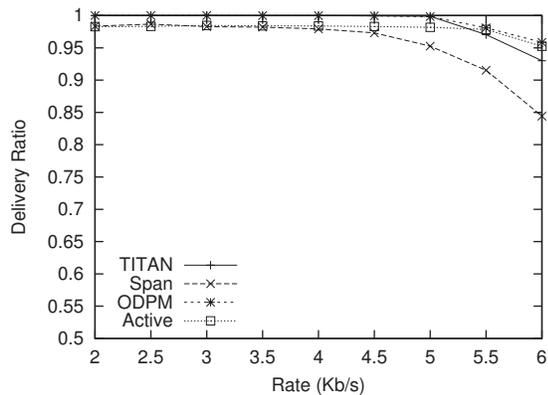


Figure 6: Delivery Ratio vs. traffic load in Span-topology, static network. Despite using fewer nodes for forwarding, TITAN does not reduce capacity.

ergy. The energy consumption of nodes that joined the backbone because they are either a source or destination is not included in the calculations. (Similar performance trends with lower energy goodput values are observed when source/destination power consumptions are accounted for). While TITAN is able to achieve 90-120% higher energy goodput than Span, the energy savings compared to ODPM is 13-25% (see Fig. 5). Essentially, the energy spent in Span for coordination messages to determine active nodes is significant. These results confirm our expectations that TITAN is able to save energy by building and maintaining a forwarding backbone implicitly and reactively. However, Span simulations do not use Span-specific improvements for IEEE 802.11 (e.g., advertised traffic window). When these improvements are used, TITAN continues to provide 10-20% higher energy goodput compared to Span. Therefore, even though TITAN simulations do not utilize such improvements, TITAN achieves the best performance by reducing the energy consumption while maintaining high delivery ratios (see Fig. 6).

IV. Future Directions

Conserving energy in ad hoc networks is challenging due to the trade-off between keeping nodes in power-save mode and maintaining efficient and effective communication. Topology management protocols try to address this challenge by identifying redundant nodes that may power down their radios at the cost of additional control overhead to build and maintain a forwarding backbone of active nodes. In this paper, we propose TITAN, which differs from current topology management protocols in the sense that it

does not require any knowledge of location or coordination among nodes to determine the nodes that should stay active. Initial results from our simulation studies show that TITAN achieves the efficiency of a topology maintenance protocol without any explicit forwarding backbone maintenance and verify that on-demand topology management is the right approach for ad hoc networks. For our future work, we plan to incorporate load balancing into TITAN to distribute the network traffic more fairly among nodes. Since the nodes in the forwarding backbone are always active and all traffic is tunneled to active nodes, these nodes drain their batteries faster. If necessary measures are not taken, the early death of forwarding nodes may shorten the total lifetime of the network. An obvious solution is that nodes with low remaining battery power delay RREQs longer to reduce the probability of joining the forwarding backbone. Additionally, we plan to extend TITAN to adapt to the changes in network load.

References

- [1] L. M. Feeney and M. Nilsson, "Investigating the energy consumption of a wireless network interface in an ad hoc networking environment," in *IEEE INFOCOM*, April 2001, pp. 1548–1557.
- [2] R. Kravets and P. Krishnan, "Application-driven power management for mobile communication," *Wireless Networks*, vol. 6, no. 4, pp. 263–277, 2000.
- [3] Y. Xu, J. Heidemann, and D. Estrin, "Geography-informed energy conservation for ad hoc routing," in *7th Annual International Conference on Mobile Computing and Networking (MobiCom)*, July 2001, pp. 70–84.
- [4] R. Zheng and R. Kravets, "On-demand power management for ad hoc networks," in *IEEE INFOCOM*, March 2003, pp. 481–491.
- [5] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris, "Span: An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks," in *7th Annual International Conference on Mobile Computing and Networking (MobiCom)*, July 2001, pp. 85–96.
- [6] D. B. Johnson, D. A. Maltz, and J. Brosch, *Ad Hoc Networking*, chapter DSR: The Dynamic Source Routing Protocol for Multi-Hop Wireless Ad Hoc Networks, pp. 139–172, Addison-Wesley, 2001.

- [7] C. E. Perkins, E. M. Royer, and S. R. Das, “Ad hoc on-demand distance vector (AODV) routing,” IETF RFC 3561, 2003.
- [8] IEEE 802 LAN/MAN Standards Committee, “Wireless LAN medium access control (MAC) and physical layer (PHY) specifications,” IEEE Standard 802.11, 1999.
- [9] R. Zheng, J.-C. Hou, and L. Sha, “Asynchronous wakeup for ad hoc networks,” in *4th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, June 2003, pp. 35–45.
- [10] C. F. Chiasserini and R. R. Rao, “Combining paging with dynamic power management,” in *IEEE INFOCOM*, April 2001, pp. 996–1004.
- [11] “Network simulator-ns2 and CMU Monarch wireless and mobility extensions to ns,” <http://www.isi.edu/nsnam/ns> and <http://www.monarch.cs.cmu.edu>.