

Singapore Management University

## Institutional Knowledge at Singapore Management University

---

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

---

3-2003

### Comparing the Routing Energy Overheads of Ad-Hoc Routing Protocols

Sorav BANSAL

*IBM India Research Laboratory*

Rajeev SHOREY

*IBM India Research Laboratory*

Archan MISRA

*Singapore Management University, archanm@smu.edu.sg*

Follow this and additional works at: [https://ink.library.smu.edu.sg/sis\\_research](https://ink.library.smu.edu.sg/sis_research)



Part of the [Software Engineering Commons](#)

---

#### Citation

BANSAL, Sorav; SHOREY, Rajeev; and MISRA, Archan. Comparing the Routing Energy Overheads of Ad-Hoc Routing Protocols. (2003). *WCNC 2003: IEEE Wireless Communications and Networking Conference: 16-20 March, New Orleans*. 1155-1161.

Available at: [https://ink.library.smu.edu.sg/sis\\_research/704](https://ink.library.smu.edu.sg/sis_research/704)

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email [cherylids@smu.edu.sg](mailto:cherylids@smu.edu.sg).

# Comparing the Routing Energy Overheads of Ad-Hoc Routing Protocols

Sorav Bansal<sup>†</sup>, Rajeev Shorey  
IBM India Research Laboratory,  
Block 1, Indian Institute of Technology,  
Hauz Khas, New Delhi 110016, India  
Email: srajeev@in.ibm.com

Archan Misra  
IBM T J Watson Research Center,  
19 Skyline Drive,  
Hawthorne, NY 10532, USA  
Email: archan@us.ibm.com

*Abstract*— We use simulations to study the comparative routing overheads of three ad-hoc routing protocols, namely AODV, DSDV and DSR. In contrast to earlier studies, we focus exclusively on the energy consumption and not on other metrics such as the number of routing packets. In particular, we study the ‘range effects’ of the three protocols, i.e., how changes to the transmission power and transmission radius affect the overall energy consumed by routing-related packets. Due to the broadcast nature of the wireless medium, the energy spent in packet receptions is almost as important as the transmission power; using the number of transmissions as an indicator of the routing overhead can thus be fairly misleading. Our studies show that the energy overhead of the three protocols varies with the transmission power in distinct and non-obvious ways.

## I. INTRODUCTION

While a wide variety of routing algorithms have been proposed for ad-hoc wireless networks, comparative analyses of their performance (e.g., [1], [2]) have largely focused either on:

- Performance metrics associated with the *data* packets.
- The routing overhead generated by the routing algorithms, defined purely from the *number* and *size* of routing-related messages and packets.

This is understandable, since traditional networks have always operated under the assumption that the data traffic load is usually much more significant than the routing overheads.

There are, however, several future ad-hoc networking scenarios where a significant fraction of the total communication energy is spent on merely maintaining routing/reachability information. In such scenarios, the actual data transfer phases may be fairly infrequent and sporadic, and the amount of transferred data may be fairly small. For example, push-based notification applications may generate only small amounts of periodic data to pervasive devices that connect to the Internet via an ad-hoc cloud (an architecture discussed in [3]). ‘Presence’-based applications (e.g., instant messaging) are another class of popular applications that are specially relevant to mobile and pervasive computing scenarios, and that can be characterized by sporadic bursts of small amounts of data. In such environments, it is thus necessary, at least as a first step, to explore in isolation the energy overheads associated with route establishment and maintenance algorithms.

<sup>†</sup> Sorav Bansal is currently a graduate student in the CS department at Stanford University, USA. Email: sbansal@stanford.edu

In this paper, we investigate the comparative performance of several ad-hoc routing protocols from the viewpoint of *routing energy-efficiency*. Our study focuses on three of ad-hoc routing algorithms, namely AODV, DSDV and DSR [4]. More precisely, we study how the performance of these routing algorithms, in terms of the energy spent in transmitting and receiving route requests and updates, varies with changes to the transmission range (power levels) of the constituent nodes. We also explore the dependence of this energy consumption on various other properties of the ad-hoc environment, such as the mobility rate of the ad-hoc nodes, their relative density and the arrival rate and holding times of data sessions. Our studies are based on extensive simulations performed with mobile nodes using the IEEE 802.11 MAC [7] protocol. Unlike previous efforts, we do not simply concentrate on the number and size of the routing-related signaling messages, but instead on the total energy spent in their transmission. Clearly, it is possible for higher-powered transmissions to spend more energy than lower-powered ones, even if a lower number of bits were transmitted at the higher power level. Moreover, we also explicitly consider the energy spent by nodes in packet *reception*. We shall see that the number of uniquely generated routing packets is a particularly misleading indicator of energy overheads, since each transmission, even if unicast, leads to the consumption of energy at *all* nodes that lie within the transmission range. This is especially important for 802.11-based networks, since empirical evidence [8] shows that commercial 802.11 cards spend similar energy in receiving or sending packets.

It should be intuitively clear that the energy overhead of route establishment and maintenance will depend on the power level, or more directly, the transmission radius of the individual nodes, for both the *distance-vector* and *link-state* families of routing protocols. In both families of routing protocols, a larger transmission radius clearly leads to greater energy consumption at the sender. Moreover, a larger transmission radius also implicitly leads to a greater number of one-hop neighbors. On the other hand, a larger transmission radius effectively decreases the diameter (in terms of hop count) of the ad-hoc network and the average hop count of the session paths. For protocols where connectivity changes in one portion of the network eventually ripple

throughout the entire network, a larger transmission radius may conceivably result in a reduction in the total number of independent routing updates generated due to a single connectivity change. Finally, a larger transmission radius also decreases the frequency of link breakages that occur due to node mobility, and thus lowers the signaling overhead spent in re-establishment of routes with broken links.

The rest of the paper is organized as follows. Section II lists the various parametric constants for each protocol used in our simulations, and provides a qualitative explanation of how some of the protocol-specific features are likely to affect the energy efficiency. Section III provides the details of our simulation environment and our choice for various simulation parameters. While section IV presents our simulation studies for varying transmission power levels and changing mobility rates, section V discusses additional simulation results that study the effect of variation in the packet arrival rates. Finally, section VI concludes the paper, highlighting the limitations of this study and the roadmap for future work.

#### A. Related Work

There have been several studies on the relative performance of ad-hoc routing protocols [9], [1], [2]. Note, however, that the earlier studies did not focus explicitly on the energy consumed by the ad-hoc routing protocols. In [2], the routing overhead was computed based on the number of packets and bytes used for routing-related signaling, rather than the energy level associated with each transmission and the power spent by receivers in this broadcast medium. Feeney and Nilsson [5] investigate the energy consumption of a IEEE 802.11 wireless network interface operating in the ad hoc mode. While they study the energy consumed in the MAC operations, our paper investigates the energy consumed due to routing protocol overheads. Kravets and Krishnan [6] present a transport protocol which suspends communication for select short periods of time, thus achieving power savings.

The effect of varying transmission ranges on ad-hoc routing performance were considered indirectly in [12] and directly in [13]. [12] studied the performance of the data packets, including the packet delivery ratio, as a function of the average number of one-hop neighbors, which is itself dependent on the transmission range of individual nodes. On the other hand, [13] studies how variations in the transmission range affect the capacity of multicast streams using AODV's multicast routing mechanism.

## II. OVERVIEW OF AD-HOC ROUTING PROTOCOLS

A detailed overview of various ad hoc routing protocols is presented in [4]. In this section, we evaluate some of the interesting ways in which changes in the transmission range of individual nodes in an ad hoc network can affect their energy consumption.

#### A. AODV-specific Effects on Routing Energy

Table I lists the values for the various AODV-related constants that were used in our simulation studies. As with earlier studies, link-layer notification was used to explicitly alert the AODV protocol of link breakages; AODV HELLO packets are thus used only as a backup mechanism. It is to be noted that all the parametric constants that are used in the simulation are based on the default parameters in ns version 2.1b8 [10].

TABLE I  
PARAMETRIC CONSTANTS FOR AODV

Timeout for active route	50 secs
Lifetime in RREP sent by destination node	60 secs
Initial timeout for maintaining reverse route information	10 secs
Flood record time (for ignoring duplicated RREQs)	6 secs
Default value for forwarding latency on a hop	30 msecs
Maximum number of attempted RREQs	3
Default initial time before retrying RREQ	60 msecs
Link-layer breakage detection	yes

By inspecting the functional behavior of AODV, we can predict some qualitative effects of varying the transmission radius and mobility rates on the energy consumed in routing overheads.

a) As an on-demand protocol, route requests are generated only when data packets are generated or when ongoing data sessions encounter intermediate link failures. Accordingly, AODV's energy consumption is expected to be fairly low in scenarios with negligible node mobility, and increase proportionately with an increase in the incidence of link breakages.

b) Since AODV does not cache multiple routes, the failure of an intermediate link requires all sessions currently using that path to issue new RREQs. Such RREQs often ripple through the system until an alternative route is discovered; moreover, previous work [2] has shown that such RREQs are the dominant component of messages exchanged by the AODV protocol. The amount of signaling required to re-establish routes upon link breakage is thus usually higher in AODV, than in alternative protocols such as DSR which maintain multiple simultaneous routes.

c) Nodes in AODV maintain next-hop information for a specific destination as long as they are used by one or more active flows with packets intended for that destination. In the absence of traffic, such forwarding information is deleted after the expiration of an inactivity timer, and must be re-established for future data packets. AODV's routing overhead should thus be fairly independent of the packet inter-arrival times, as long as the inter-arrival gap is small enough to ensure, on average, the initiation of a new packet before an inactivity-related timeout. AODV's energy consumption should however exhibit a noticeable increase whenever the interval between successive packets becomes slightly larger than the inactivity-related timeout, since

most packets will now trigger the re-establishment of forwarding paths.

### B. DSDV-specific Effects on Routing Energy

Table II shows the various constants and their values used in our simulation of DSDV.

TABLE II  
PARAMETRIC CONSTANTS FOR DSDV

Route update period	15 secs
Max. number of updates missed before broken link	3
Initial triggered update settling time	6 secs
Weighted settling time factor	0.875
Route advertisement aggregation time	1 sec

It is relatively easy to see the following effects of DSDV behavior on the energy spent in routing overheads.

- As a classical distance-vector protocol, DSDV is responsible for periodically announcing its routing table to all one-hop neighbors. Since DSDV routing tables contain a list of next-hop entries for every node in the ad-hoc network, the size of this routing update is essentially independent of a node's transmission range or power level.
- Since DSDV is not an on-demand protocol, and requires the periodic transmission of reachability information, the *ambient* energy consumption of DSDV (for very low packet arrival rates and negligible mobility) should be much higher than that of alternative on-demand protocols. However, for the same reason, DSDV's energy consumption can be expected to be less sensitive to higher mobility rates (frequent link breakages) or variations in the packet arrival rates (since no caches are used). Moreover, since DSDV does not require the original sender to re-establish an entire route in response to the breakage of an intermediate link, DSDV may prove to have a slower rate of increase in its energy consumption for increasing mobility rates than alternative protocols such as AODV and DSR.

### C. DSR-specific Effects on Routing Energy

TABLE III  
PARAMETRIC CONSTANTS FOR DSR

Min. time between retransmitted RREQ (Exponentially backed off)	500 msecs
Min. time between gratuitous replies for same route	1 secs
Timeout for non-propagating search	30 msecs
Timeout for cached routes	300 secs

Unlike the other protocols considered here, DSR's routing overheads are not confined to control messages alone; the use of source-routing implies a variable routing-dependent overhead in the data packets as well. Some of the DSR-specific features that impact the actual energy spent in the routing process include:

- Due to its use of source routes, DSR naturally results in longer headers (containing the IDs of all intermediate nodes)

in data packets. Since a larger transmission radius effectively decreases the average hop count of a path, we can expect that the size of the data packet itself would decrease with an increase in the transmission power levels. However, this paper *completely neglects the energy associated with the transmission of data packets* and accordingly does not consider this overhead.

b) DSR supports the maintenance of multiple routes to a specific destination in the routing cache of a sender. Accordingly, the failure of a route may not automatically lead to the generation of new RREQs; rather, intermediate nodes will switch to alternative routes.

c) It has been observed [2] that DSR routing packets are dominated by ROUTE REPLIES, which are unicast and thus incur the RTS-CTS overhead of the 802.11 MAC.

d) To aggressively cache available routes, DSR nodes typically operate in promiscuous mode and snoop on ongoing transmissions. DSR thus leads to energy spent on processing packets at all nodes within earshot of a sender, even when the transmission consists of a *data* packet and the promiscuous node is not the intended recipient. (This energy consumption is related to packet *processing* at higher layers of the protocol stack and should be distinguished from the energy spent in packet *reception*, which could actually be zero if nodes not interested in a reception could disable their radio interfaces). Due to our exclusive focus on the routing overheads, we do not, however, attempt to capture this processing overhead.

## III. SIMULATION DETAILS

We have used the ad-hoc routing protocol implementations included as part of the ns-2.1b6 package [10]. Our simulation results are based on the study of 30 nodes distributed in a square grid of  $670 \times 670$  meters<sup>2</sup> and using the 802.11 MAC protocol. Each of the  $\binom{30}{2}$  (source, destination) pairs in the simulation generate packets according to a Poisson arrival process, i.e., the inter-packet times for each (source, destination) pair were exponentially distributed. Our node mobility model is based on the *random waypoint* model enumerated in [1]. Our energy measurements do not include the energy spent in either transmitting or receiving any *data* packet.

We performed each simulation run for a total of 1300 secs, with the measurements of the energy divided into two distinct phases. The *transient* phase included the energy spent in routing in the interval (0, 300) secs (when on-demand protocols would always need to establish a route for the first time), while the *steady-state* phase measured the energy expended in the interval (300, 1300) secs. While packets were generated from 0 secs (to initiate the route establishment process), the energy overheads for our steady-state computations were not measured until 300 secs had elapsed. We have also separately studied the energy consumed in the first 300 secs (the transient phase). While the results for the transient phase are available in the detailed version of this paper [14], they are relatively less interesting and are not discussed in this paper.

To make our study as realistic as possible, we decided to in-

clude the 802.11 MAC overheads (for routing packets) in our energy computations. Most of the routing control messages use link-layer broadcast; in 802.11, such broadcast packets are not subject to either contention resolution or the 4-way handshake. In contrast, some of the routing messages are unicast; such messages are obviously subject to the 4-way RTS-CTS based handshake mechanism, and thus incur slightly higher overhead.

To study the effect of the reception energy on the overall routing overhead, we had to assign values to both the transmit power  $P_t$  and the receive power  $P_r$ . While we vary  $P_t$  (and thus change the transmission range) in our simulations,  $P_r$  is an implementation-specific constant and should not vary with  $P_t$ . Empirical results [8] show that  $P_r$  and  $P_t$  are usually comparable in current 802.11 cards; the best cards usually result in  $P_r$  being about  $\frac{2}{3}$ rd of  $P_t$ . We thus adopted the following strategy for determining  $P_t$  and  $P_r$ . While the transmission range of the wireless nodes was varied within the range (225,500) meters in our simulations, we set  $P_r$  to half the value of  $P_t$  chosen for the 225 meter (smallest) range. We observed that setting the transmission range to less than 225 meters caused the graph to occasionally become disconnected. More specifically, we assumed a transmission power of 0.0271 W to result in a transmission range of 225 meters; according  $P_r$  was set to 0.0135 W in all our simulations. Note that in the simulations, we use an attenuation factor of  $K = 2$ . To study the effect of the transmission overheads alone, we also repeated the simulations by setting  $P_r = 0$ , thereby completely neglecting the receiver energy consumption. The exact absolute values of  $P_r$  and  $P_t$  are relatively unimportant, since we are interested in the comparative performance of the various ad-hoc routing protocols.

#### IV. PERFORMANCE STUDIES FOR VARYING RANGE AND NODE SPEED

We first study the routing overhead of the various protocols as the transmission range of the nodes is varied, while the packet arrival rate is held constant. We experimented with different values of  $v_{max}$ , which we call the mobility rate, in the set (0.01 m/s, 5 m/s, 10 m/s, 20 m/s and 30 m/s). All the plots in this section correspond to a Poisson packet arrival process (for each node pair) with a mean inter-arrival time of 15 secs. It is important to emphasize that our primary focus is on *the relative qualitative behavior of the three protocols, since the precise quantitative behavior will depend on the precise choice of various protocol-specific constants.*

##### A. Steady State Behavior

We first study the comparative energy consumption of the three protocols in the steady state. In Figure 1, Figure 2 and Figure 3, we plot the individual energy consumption of the three protocols versus the transmission range for different maximum node speeds (mobility rate) with  $P_r = 0$ . Clearly, the three protocols exhibit significantly different behavior in terms of routing overhead. While these plots depict only the energy spent in packet transmissions, they do not directly correlate with the

*number* of routing packets, since the transmission energy per packet is not constant.

Figure 1 shows that the energy overhead of DSDV increases with the transmission range and is almost insensitive to change in the node mobility rates. Due to its pro-active nature (periodic generation of routing updates), DSDV's energy consumption is relatively insensitive to the higher mobility rates (frequent link breakages) than either AODV or DSR. Furthermore, unlike AODV or DSR, DSDV does not require the original sender to re-establish an entire route in response to the breakage of an intermediate link, and hence has a slower rate of increase in energy consumption with increasing mobility rates. It would also appear, somewhat counter-intuitively, that the energy consumption of DSDV is lower than that of AODV or DSR. This is really due to the absence of MAC-layer notification in our DSDV implementation; a node detects the loss of connectivity to a neighboring node only after 45 secs. For the node speeds considered here, link breakages occur far more frequently; accordingly, AODV and DSR generate mobility-driven RREQs much more often. Accordingly, the DSDV packet delivery ratio is very low, since most packets attempt to use links that are already broken.

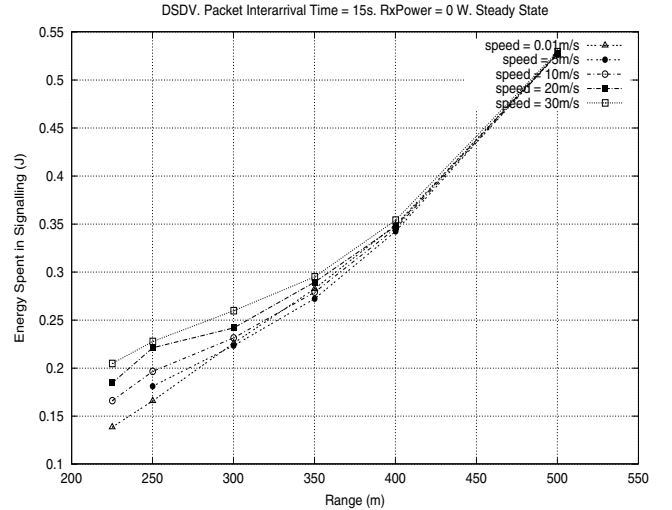


Fig. 1. Steady-State Routing Overhead,  $P_r = 0$ , DSDV

In contrast to DSDV, the on-demand nature of both AODV and DSR makes their energy consumption more sensitive to variation in the mobility rates (see Figure 2 and Figure 3), with higher mobility rates expectedly resulting in greater energy consumption. Since the packet inter-arrival times in this set of plots is small enough to ensure that a cached route usually never expires (see Table I, Table III), the bulk of the AODV transmission overhead consists of RREQs generated due to mobility-driven route failures. Accordingly, a larger value of  $R$  significantly decreases the frequency of such link breakages and hence, lowers the overall energy consumption. Since this decrease happens even though a larger  $R$  implies greater transmission power spent on a single transmission, it is apparent that the *number of*

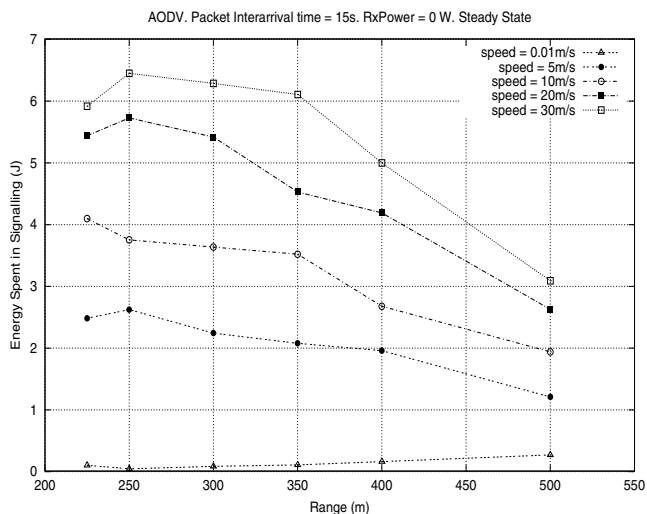


Fig. 2. Steady-State Routing Overhead,  $P_r = 0$ , AODV




Fig. 3. Steady-State Routing Overhead,  $P_r = 0$ , DSR

packet transmissions decreases very rapidly with an increase in the transmission range. The AODV transmission energy overhead typically exhibits a monotonically decreasing behavior as the transmission range is increased.

We also see from Figure 2 and Figure 3 that the energy consumption of DSR is markedly different from AODV, and exhibits a unimodal behavior, with the energy consumption peaking at an intermediate transmission range. As the transmission range initially increases, the static component of the energy overhead increases, while the frequency of link breakages does not exhibit a corresponding decrease. Accordingly, as  $R$  is initially increased, the transmission energy overhead also increases. As the transmission range is further increased, the rate of link breakages begins to exhibit a sharp decrease, effectively leading to a decrease in the mobility-related component of the signaling energy. Moreover, a larger value of  $R$  also leads to

greater caching efficiency, since an overhearing node is now within the transmission range of a greater number of nodes. Accordingly, the overall signaling energy begins to decrease with an increase in the transmission radius. As expected, this inflexion point, or the value of  $R$  at which the energy consumption begins to exhibit a decrease, is larger for higher values of mobility rates. Accordingly, DSR operation is most efficient at either low transmission range (low transmission power) or at very high transmission ranges (high transmission power). The energy overheads of DSR and AODV are seen to be fairly comparable, especially since we've ignored the source-routing overhead for DSR data packets.

We now consider the more realistic (at least for 802.11-based networks) scenario where the individual nodes expend additional energy on receiving every transmission that they overhear in the broadcast medium. In Figure 4, Figure 5 and Figure 6, we plot the energy consumption of the three routing protocols when  $P_r = 0.013$  W, i.e., when the energy overhead includes the energy spent in packet receptions. On comparing these plots with those in Figure 1, Figure 2 and Figure 3, the significance of the packet reception component in the overall routing overhead is immediately apparent. To begin with, the absolute values of the energy consumed in routing is now often larger by an order of magnitude. For example, the AODV energy consumption now varies between 0.1-95 Joules (as opposed to the corresponding value of 0.05-6.7 Joules for  $P_r = 0$ ).

In contrast to the corresponding plot in Figure 1, DSDV clearly shows a decrease in the overall energy consumption with increasing  $R$ . This apparent contradiction can be explained by noting that an increase in the transmission energy can occur although the number of transmissions itself decreases (since each transmission consumes more energy). Clearly, although a larger  $R$  does lead to the consumption of reception energy at a greater number of receivers (one-hop neighbors), the reduction in the number of transmissions reduces the overall energy spent in packet receptions. Most interestingly, the total DSDV energy overhead begins to increase again beyond  $R = 400$  meters in this plot; beyond this value, the reduction in the number of update packets is not rapid enough to compensate for the increase in the energy spent by neighboring nodes in receiving a single packet.

In AODV (see Figure 5), the overall energy consumption decreases with increasing transmission radius  $R$  across all mobility rates. A larger  $R$  always implies a larger value of the number of one-hop neighbors of a particular transmitter, and hence a larger overall consumption of reception energy for a single transmission. However, we can see that the reduction in the actual number of AODV packet transmissions is dramatic enough to counteract the corresponding increase in the number of one-hop neighbors.

In contrast to AODV, the total energy consumption in DSR (Figure 6) does not exhibit a monotonic behavior. As with the corresponding plot in Figure 3, the total energy exhibits a maximum for an intermediate value of  $R$  and is lower for both smaller and larger transmission ranges. Clearly, for values of  $R$  larger




Fig. 4. Steady-State Routing Overhead,  $P_r = 0.013$  W, DSDV

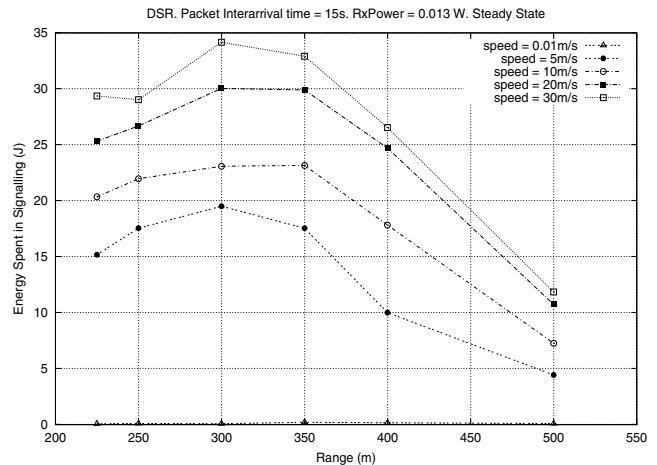


Fig. 6. Steady-State Routing Overhead,  $P_r = 0.013$  W, DSR

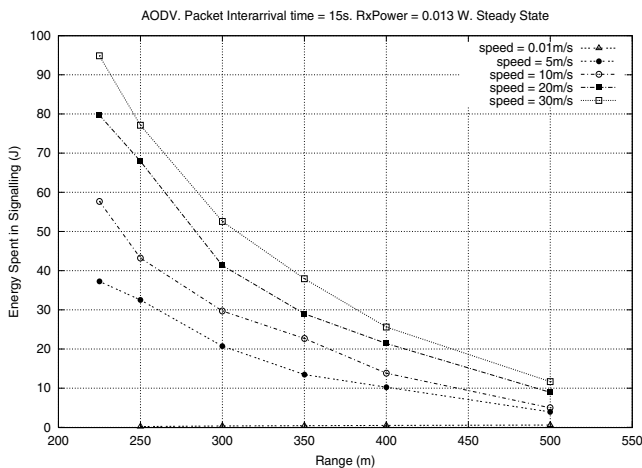


Fig. 5. Steady-State Routing Overhead,  $P_r = 0.013$  W, AODV

than the inflexion point, the reduction in the *number* of packet transmissions leads to a much smaller value of the total packet reception energy, even though the number of neighboring nodes increases with increasing  $R$ . This decrease in the packet reception energy swamps any possible increase in the packet transmission energy— accordingly, this inflexion point occurs earlier in the DSR plot in Figure 6 than the corresponding plot in Figure 3 (for example, while the maximum energy consumption for speed = 10 m/s occurs for  $R = 350$  meters in Figure 6, it occurs at  $R = 400$  meters in Figure 3.)

## V. PERFORMANCE STUDIES FOR VARYING PACKET ARRIVAL RATES

The previous section clearly showed that estimating the energy overhead of an ad-hoc routing protocol from the packet *transmission* activity (either in terms of energy or the number of such packets) can be very misleading, since it is the packet reception energy component that usually dominates the overall energy consumption. We now consider the effect of varying the

packet inter-arrival rates. Changes in the mean time between packets should clearly affect the on-demand routing protocols that cache routes from previous packet transfers for a finite duration. The plots in this section all correspond to a maximum node speed  $v_{max}$  of 10 m/s. For reasons of space, we principally focus on the steady-state total (transmission+ reception) energy of the three protocols, i.e., obtain plots with  $P_r = 0.013$  W.

Figure 7 and Figure 8 plots the steady-state energy consumed by the AODV and DSR routing protocols respectively for different values of the packet inter-arrival time. We omit the DSDV plot; since DSDV is a pro-active protocol, the routing-related energy consumption is independent of the arrival rate of the data packets.

Figure 7 shows that the total AODV steady-state energy decreases with increasing transmission range for all values of the packet inter-arrival time. In general, we would expect a smaller packet inter-arrival time would normally lead to the generation of larger number of RREQs and hence, a higher routing overhead: indeed, a larger inter-arrival time (less frequent packet arrivals) does generally result in a lower routing overhead. On the other hand, since AODV routes are cached for a maximum of 50 secs, a significantly large inter-arrival time also increases the probability of cache timeouts and could lead to a greater number of RREQ messages. Accordingly, we can see that, when the transmission range is large, the greater frequency of cache timeouts and consequent RREQ generation causes AODV to have a higher energy consumption for a large (30 sec) value of the packet inter-arrival time than for the smaller values of 10, 15 and 20 secs. We confirmed this phenomenon by plotting the AODV transmission energy alone ( $P_r = 0$ ), and noticing a sharp increase in the transmission energy when the inter-arrival time exceeds 30 secs.

Figure 8 shows that the steady state energy overheads of DSR are similar to AODV, in that the DSR overhead is higher for smaller inter-arrival times. (The unimodal nature of the plots has




Fig. 7. Steady-State Routing Overhead with  $P_r = 0.013$  W, AODV




Fig. 8. Steady-State Routing Overhead with  $P_r = 0.013$  W, DSR

already been explained in the previous section.) Unlike AODV, the DSR energy consumption for a inter-arrival time of 30 sec is not larger than smaller inter-arrival times of 15 and 20 secs. This is due to the larger value of the DSR cache timeout (300 secs), which almost eliminates the likelihood of inactivity-based cache timeouts, even when the average inter-arrival time is as large as 30 secs.

Our simulations thus reveal that, in both AODV and DSR, a larger packet inter-arrival time usually leads to less frequent generation of RREQ and hence a lower energy consumption. However, this packet inter-arrival time should not be too low; excessive gap between successive packets can lead to either cache timeouts or the use of incorrect (stale) routing entries, both of which can significantly increase the routing overhead.

## VI. CONCLUSION

In this paper, we have studied the comparative signaling energy overhead of three popular ad-hoc routing protocols, namely

DSDV, AODV and DSR. Simulation studies showed that the three protocols exhibited significant differences in the way in which their energy overhead depended on the transmission range. Most importantly, the simulations demonstrated the *critical importance of including the packet reception energy in computing the total energy cost*, and the need to distinguish between the total transmission energy and the total *number* of packet transmissions. While a larger  $R$  usually resulted in lower energy consumption for all three protocols, the energy consumption in DSR registered an initial increase with increasing  $R$ .

We have also studied the dependence of the routing energy on the packet inter-arrival times. Contrary to the general expectation, the total energy for on-demand protocols did not always decrease with larger inter-arrival times— if the inter-packet gap became fairly large, cache timeouts and the potential for the use of stale routing entries could lead to greater routing overheads.

This work generates several interesting avenues for future work. We would clearly like to study additional ad-hoc routing protocols, especially those belonging to the link-state family. Moreover, we would also like to extend our energy framework to include the energy overhead of the data packets as well, thereby enabling us to study the tradeoffs between various metrics (such as packet delivery ratio or forwarding latency) in a unified framework.

## REFERENCES

- [1] J. Broch, D. Maltz, D. Johnson, Y. Hu and J. Jetcheva, "A Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols", Proceedings of ACM Mobicom, 1998.
- [2] S. Das, C. Perkins and E. Royer, "Performance Comparison of Two On-Demand Routing Protocols for Ad-Hoc Networks", Proceedings of IEEE INFOCOM 2000, March 2000.
- [3] Y. Lin and Y. Hsu, "Multihop Cellular: A New Architecture for Wireless Communications", Proceedings of IEEE INFOCOM'00, March 2000.
- [4] C. E. Perkins (Editor), "Ad Hoc Networking", Addison-Wesley, 2001.
- [5] L. Feeney and M. Nilsson, "Investigating the Energy Consumption of a Wireless Network Interface in an Ad Hoc Networking Environment", In Proceedings of IEEE Infocom'2001, April 2001.
- [6] R. Kravets and P. Krishnan, "Application-Driven Power Management for Mobile Communication", ACM/URSI/Baltzer Wireless Networks (WINET), Volume 6, Number 4, 2000.
- [7] IEEE Computer Society LAN MAN Standards Committee, Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, IEEE Std. 802.11-1997, IEEE, 1997.
- [8] "Power Consumption of Wireless Networking PC Cards", <http://www.synack.net/wireless/consumption.html>.
- [9] V. Park and S. Corson, "A Highly Adaptive Distributed Routing Algorithm for Mobile Wireless Networks", Proceedings of IEEE INFOCOM '97, March 1997.
- [10] "The ns-2 network simulator", <http://www.isi.edu/nsnam/ns>.
- [11] S. R. Das, R. Castaneda, J. Yan and R. Sengupta, "Comparative Performance Evaluation of Routing Protocols for Mobile, Ad Hoc Networks", 7<sup>th</sup> International Conference on Computer Communications and Networks (IC3N), October 1998.
- [12] E. Royer, P. Melliar-Smith, L. Moser, "An Analysis of the Optimum Node Density for Ad Hoc Mobile Networks", Proceedings of IEEE International Conference on Communications (ICC), June 2001.
- [13] E. Royer and C. Perkins, "Transmission Range Effects on AODV Multicast Communication", to appear in ACM Mobile Networks and Applications, preprint.
- [14] S. Bansal, R. Shorey and A. Misra, "Comparing the Routing Energy Overheads of Ad Hoc Routing Protocols", Detailed Manuscript, <http://www.cse.iitd.ernet.in/srajeev/publications.htm/wcnc-fullpaper.pdf>