Energy-Efficient MANET Routing: Ideal vs. Realistic Performance

Thomas Kunz
Systems and Computer Engineering
Carleton University
1125 Colonel By Drive
Ottawa, Ont., Canada K1S 5B6
tkunz@sce.carleton.ca

***Abstract*—Nodes in a Mobile Ad Hoc Network are typically battery-powered. Energy consumption is therefore an important metric to consider in designing routing protocols for such networks. In this paper, we present some results from our work on integrating energy-efficiency aspects into a standard MANET routing protocol, OLSR. In particular, we are exploring the impact of nodes having only inaccurate/imprecise knowledge of the energy levels of other nodes. We use two different energy-efficient variants of the OLSR protocol and simulate a wide range of scenarios. Our results show that for all mobility scenarios, traffic loads, and protocol variations, the achievable protocol performance is negatively affected by the imprecise information available. The loss in performance can be as high as 10%, emphasizing the need to collect more precise routing-related information.**

***Index Terms*—Energy-efficient, MANET, OLSR, Routing, Ideal Knowledge, Knowledge Accuracy**

I. INTRODUCTION

A Mobile Ad Hoc Network (MANET) is defined by the MANET Working Group as "an autonomous system of mobile routers (and associated hosts) connected by wireless links - the union of which forms an arbitrary graph". Because of the antenna's limited transmission range, the nodes in the network may act as a router to forward packets to other nodes, and then a routing protocol is needed. The main characteristics of a MANET are:

- Packets may need to be forwarded by several nodes to reach the destination.
- Dynamic topology due to the nodes' mobility or nodes leaving/joining the network, which causes packet loss and route change.
- Resource constrains: wireless medium bandwidth, device's battery, processing speed and memory.

To obtain the correct network topology, frequent control message exchanges between nodes are required; on the other hand, these control messages will consume valuable wireless bandwidth resources. This tension poses a challenge for developing routing protocols. Existing MANET routing protocols basically can be classified as proactive (table-driven), reactive (on-demand) and hybrid. An example of a proactive is the Optimized Link State Routing Protocol (OLSR) [1], and example of a reactive routing protocol is the Ad hoc On-Demand Distance Vector (AODV) Routing Protocol [2]. Detailed reviews and performance comparisons of these protocols can be found in early papers such as [3][4], which focused on the routing protocol performance in terms of Packet Delivery Ratio, packet latency, control message overhead, path optimality, etc.

Nodes within a mobile ad hoc network communicate via wireless interfaces. This complicates their operation in a number of ways. One of the major issues is the limited energy, which is usually supplied by node batteries that network nodes possess. Amongst other sources, energy is consumed by wireless interfaces in several modes of operation, and for many radio technologies, wireless communication is the dominant energy cost for a node. Nodes expend energy in packet sending and receiving modes. However, it has been also shown that nodes consume energy through their wireless interfaces while idle [5]. As a radio is in idle mode for the majority of time, idle energy consumption can easily dominate the overall nodal energy consumption [6][7]. As wireless interfaces also typically have a much more low-powered "sleep" mode, in which they cannot receive packets, many proposals have been suggested in the literature to allow nodes to stay in "sleep" mode for extended periods of time. These solutions range from pure MAC-layer solutions such as the power management functionality in IEEE 802.11 [8] to solutions that combine routing and MAC layer functionality [9][10] and proposals that are designed as a separate network layer component intended to complement any routing protocol [11][12]. We therefore assume in the remainder of this paper that the problem of idle energy consumption is appropriately addressed by one of the above proposals, in particular one that is independent of the specific routing protocol.

Once idle energy consumption has been addressed, further performance gains are possible if the MANET routing protocol takes energy into consideration when making routing decisions. Early proposed routing protocols typically selected hop count as routing metric, optimizing the number of retransmissions for each packet. For wireless interfaces with fixed transmission power (such as IEEE 802.11), such a strategy may, at first glance, also appear optimal from an energy perspective. However, as explored in [13], minimum-hop-count routing tends to choose paths through the "center" of the network, resulting in at least the following two problems:

- Many packets are contending with each other, leading to high collisions, retransmissions, and ultimately packet

978-1-4244-2202-9/08/$25.00 © 2008 IEEE

loss.

– Nodes in the more central area of the network disproportionally relay more data packets, causing them to deplete their energy first. This will lead to many longer routes, and in the extreme case may cause some nodes to become disconnected from the rest of the network.

Recognizing these issues, researchers have started to explore energy-efficient MANET routing protocols for a number of years. An early example of an energy-aware routing protocol is [14], and a survey of that early work can be found in [15]. However, such protocols typically were completely new protocols, and were not picked up by the larger community. As best-effort unicast MANET routing protocols such as OLSR, DSR, and AODV, were standardized in the IETF, more recent work has focused on adding enery-awareness to these routing protocols, see for example [16].

This paper reports some of our work on making OLSR an energy-efficient routing protocol. In particular, in earlier work we identified two modifications to the protocol that have the potential to increase network performance under constrained energy operations, assuming that each node accurately knows the remaining energy levels of every other node. In reality, though, nodes learn about this and other network-related information through the exchange of topology information. Therefore, the energy level information at each node is, at best, an approximation of the actual energy levels of nodes in the network, specifically in scenarios with high mobility and/or high traffic volumes. In this paper, we are exploring the impact of inaccurate/imprecise information on the performance of the energy-efficient modifications, providing one of the few quantifications of this impact. The paper is organized as follows. The next section describes the core operation of OLSR in more detail, Section III reviews related work. Section IV discusses modifications to OLSR to increase its performance under limited available energy and the modifications to propagate energy-level information throughout the network. Section V presents a summary of our extensive simulation results based on NS2 and the Hipercom OOLSR implementation. We conclude the paper in Section VI with a brief summary and outlook on future work.

II. Brief Review of OLSR

To put the related work in context, we first briefly describe the protocol we are working with, the Optimized Link State Routing Protocol (OLSR) [1]. OLSR is an optimization of the pure link state algorithm. The key concept used in the protocol is that of multipoint relays (MPRs). Each node selects a set of its neighbor nodes as MPRs. Only nodes selected as MPRs are responsible for forwarding control traffic, intended for diffusion into the entire network. MPRs thus provide an efficient mechanism for flooding control traffic by reducing the number of (re-)transmissions required. While determining the optimal (minimal) number of MPRs for a given topology is NP-hard, the OLSR RFC contains a simple heuristic that this typically results in fairly small MPR sets, while still allowing the path determination algorithm to find minimum-hop paths.

Nodes selected as MPRs also have a special responsibility when declaring link state information in the network. Indeed, the only requirement for OLSR to provide shortest path routes to all destinations is that MPR nodes declare link-state information for their MPR selectors.

Nodes that have been selected as multipoint relays by some neighbor node(s) announce this information periodically in their control messages. Thereby a node announces to the network that it has reachability to the nodes which have selected it as an MPR. In route calculation, the MPRs are used to form the route from a given node to any destination in the network.

Due to its proactive nature, OLSR works with a periodic exchange of messages. The key messages are Hello and TC messages. Hello messages are periodically exchanged to inform nodes about their neighbors and their neighbors' neighbors and are 1-hop broadcast messages. The 2-hop neighborhood information is then used locally by each node to determine MPRs. In contrast, TC messages are flooded through the network to inform all nodes about the (partial) network topology. At a minimum, TC messages contain information about MPRs and their MPR selectors.

To control the protocol overhead, OLSR defines a number of parameters. The Hello-interval parameter (default: 2 seconds) represents the frequency of generating a Hello message. Increasing the frequency of generating Hello messages leads to more frequent updates about the neighborhood and hence a more accurate view of the network.

The TC-interval parameter (default: 5 seconds) represents the frequency of generating a TC message. TC messages are one of the major sources of overhead in MANETS, as they are flooded throughout the network, but they facilitate the topology discovery process. Since nodes learn about the whole topology by exchanging TC messages, the more frequently nodes generate TC messages, the more recent the knowledge nodes have about the topology.

The MPR-coverage parameter (default: 1) allows a node to select redundant MPRs. The criterion for selecting MPRs is that all 2-hop neighbors must be reachable through at least one MPR node. Nodes should select their MPR set to be as small as possible in order to reduce protocol overhead. Redundancy of the MPR set affects the overhead through affecting the amount of links being advertised, since a node will be selected by more neighbor nodes as an MPR. The parameter also affects the amount of nodes advertising links, since more nodes will be selected as MPRs, reducing the efficiency of the MPR flooding mechanism. On the other hand, redundancy in the MPR set ensures that more nodes advertise reachability for a node.

The TC-redundancy parameter (default: 0) specifies, for the local node, the amount of information that may be included in the TC message. A TC-redundancy of 0 specifies that the advertised link set of the node is limited to links to its MPR selectors. A TC-redundancy of 1 specifies that the advertised link set of the node is the union of links to its MPR selectors and to other MPRs. A TC-redundancy of 2 specifies that the advertised link set of the node is the full neighbor link set. The TC-redundancy parameter affects the overhead through affecting the amount of links being advertised as well as the amount of nodes advertising links.

III. Related Work

Many energy-efficient routing protocols have been developed, modifying the routing metric to take energy costs into account. Representative examples are [17][18]. These protocols either explore routes with maximum bottleneck residual energy at one of the intermediate nodes, minimize the total end-to-end transmission energy for a packet, or a (weighted) combination of both. The objective is typically to reduce energy consumption and to increase nodal lifetime, resulting in increased network lifetime and performance. However, minimizing transmission energy only differs from shortest-hop routing if nodes can adjust transmission power levels, such that multiple short hops are more advantageous, from an energy perspective, than a single long hop. As we are interested in running our protocol over an IEEE 802.11-based MAC, we do not have access to this capability. We therefore are only focusing on routing mechanisms that avoid nodes with low residual energy levels.

In addition, as also motivated in the introduction, we are more interested in enhancing standard MANET protocols, rather than inventing new ones. Examples are modifications to DSR [16], or AODV [19][20]. With regard to OLSR, a few researchers have explored the suitability of the protocol for QoS routing [21][22]. To achieve the best results, these QoS protocol variants typically modify both the MPR selection criteria and the path determination algorithm. [23][24][25] discuss energy-efficient variants of the OLSR. Unlike the QoS-variants, these three papers differ significantly in the necessary modifications. [23] solely explores modifications to the MPR selection mechanism, choosing MPRs based on their residual energy levels, rather than on their coverage of 2-hop neighbors, as suggested by the original protocol RFC [1]. [24] modifies the path determination algorithm, selecting paths based on the residual energy level of intermediate nodes. And finally [25] combines both MPR selection and path determination to again select paths with maximum bottleneck residual energy level. All papers publish simulation results that show that, for certain scenarios, taking energy constraints into consideration can increase nodal lifetime and network performance.

None of the above papers discuss in detail how nodes collect residual energy levels, how accurate this information is, or what the impact of any inaccuracy may be on the routing protocol performance. In general, we found very few studies that systematically explored the problem of state information accuracy, either for best-effort routing or QoS routing. For example, Ge et al. [21] developed a QoS version of the OLSR protocol, based on link bandwidth as QoS metric. This QoS OLSR protocol searches for paths with maximum bottleneck bandwidth, propagating changes in the link bandwidth to all nodes when the available link bandwidth exceeds a certain threshold. The paper evaluates the performance of this QoS OLSR model under different bandwidth change threshold values and compares it to the basic OLSR protocol performance. The main conclusion from the paper is that the availability of more accurate state information throughout the network, via more frequent updates, improves the performance of QoS routing, albeit at the cost of higher protocol overheads. However, this work does not investigate quantitatively the accuracy of the QoS metric (link bandwidth).

[26] and [26] investigated the impact of extending topology knowledge (by varying the MPR-coverage and TC-redundancy parameters. Both focus on having more accurate information at the topology (network) level and how it affects the routing protocol performance. In other words, they study the effect of tuning the OLSR protocol parameters on accuracy in terms of network status (existing nodes and links) and not in terms of state information available at nodes and links.

The closely related body of work to the problem of QoS routing in the presence of inaccurate information is a set of papers aimed at exploring state-aggregation issues and their impact on routing performance in large networks such as [27]. These papers emphasize on developing good aggregation techniques that minimize inaccuracy in network state information, while allowing substantial reductions in the amount of state data. Since we are dealing with relatively small networks, our work investigates the impact of collecting fine-grained data (at the level of individual nodes) on the accuracy of state information as an upper bound on achievable accuracy.

[29], [30], and [31] investigate the impact of inaccuracies, in the available network state and metric information, on the path selection process for flows which require QoS guarantees. In particular [31] evaluates the impact of inaccurate state information on the performance and overheads of QoS routing and showed, similar to [21], that relatively coarse-grained update policies reduce overheads while also reducing the performance of QoS routing.

In this paper, we are exploring modifications to OLSR based on [23][24][25]. We evaluate the impact of these modifications on the network performance under a wide range of scenarios. We have previously identified two promising variations of the OLSR protocol that perform superior to the original OLSR protocol for a wide range of scenarios. In these modifications, we assumed that each node had instantaneous access to other node's residual energy levels. In this paper, we modify our OLSR implementation to propagate residual energy values through protocol control messages. Nodes use this information (rather than the omniscient information we utilized in previous work), and we compare the achievable protocol performance. To the best of the author's knowledge, this is the first study that systematically quantifies the performance impact of inaccurate/imprecise knowledge on the performance of a MANET routing protocol, energy-efficient or otherwise.

IV. MODIFICATIONS TO OLSR

We made essentially two modifications to OLSR: similar to [23], we modified the MPR selection criteria, and similar to [24][25], we modified the path determination algorithm.

MPRs are a subset of the 1-hop neighbors that provide access to all 2-hop neighbors of a node. Reducing the number of MPRs each node selects is key to the OLSR optimizations. Determining the minimal MPR set is NP-hard, but the OLSR RFC [1] suggests a simple heuristic to approximate this minimal set. In essence, the heuristic iteratively adds 1-hop neighbors with connectivity to the maximum number of 2-hop neighbors to the MPR set until all 2-hop neighbors are covered. Our modification to this strategy is to iteratively add 1-hop neighbors with maximal residual energy level to the MPR set until all 2-hop neighbors are covered. As routes ultimately are build from MPRs (except potentially the first and last node), this will avoid nodes with low residual energy levels, unless they are the source or destination.

The path determination algorithm in the original OLSR protocol is essentially a Dijkstra shortest-path algorithm. Nodes learn through TC messages a partial network topology, over which they then perform shortest-path routing to populate the routing table. It can be shown that the shortest path in this partial topology has the same length as the shortest path determined over the complete topology. Our modification to this algorithm is in changing the weight associated with each link. Rather than assigning each link the same constant weight of 1, we assign it the reciprocal value of the sending node's residual energy level. Again, this will penalize routes that traverse nodes with low residual energy level. In addition, using the reciprocal value has the advantage that the algorithm does not need "artificial" thresholds to determine when a node's residual energy level is low and the link should therefore be assigned a higher weight. As a node's energy level depletes, the path determination algorithm tries increasingly "harder" to avoid such nodes.

In our initial work, as we used a simulator, we simply allowed nodes to access the residual energy levels when they selected MPRs or made routing decisions, as described above. In a more realistic setting, nodes need to learn about residual energy levels through message exchange, so we added this capability to our protocol. There are two basic ways in which residual energy levels can be propagated throughout the network. Either we define a new message type to carry that information, or we include it in the OLSR protocol messages (Hello and TC messages) to be available to other nodes in the network. With the first approach, a new message type has to be defined and exchanged. This will incur a potentially large overhead in the network since more messages will be exchanged. In addition, these messages will include a lot of redundant information and the same gain can be obtained by including the residual energy information in the OLSR protocol messages. Therefore, we chose that second approach.

Through the exchange of OLSR control messages, each node accumulates information about the network. This information is stored according to the OLSR specifications. However, to store the residual energy levels associated with a node, a new field was added to the neighborhood information base and to the topology information base maintained by the protocol. To populate these fields, the message format of Hello and TC messages was extended as well.

Extended Hello messages are broadcast to all one-hop neighbors. They contain not only a list of addresses of neighbors, but also the most recent energy level associated with those neighbors from the sender node's perspective. In addition to that, the message also contains the remaining energy level of the sender node itself at the time the message is generated. The other fields are loaded according to the OLSR specifications.

Extended TC messages are broadcast and retransmitted by the MPRs in order to diffuse topology information into the entire network. TC messages contain not only a list of addresses of a node's MPR selectors, but also the energy levels associated with those nodes from the originator node's perspective. In addition to that, the message also contains the remaining energy level of the originator node at the time the message is generated.

A consequence of these modifications is that, for a given node X, a single node may have multiple different residual energy levels stored in various internal databases. Also, as a result of message delays and message losses, a node may learn old information about node X from other, intermediate nodes, at a later point in time. To disambiguate between these entries and to determine the most recent values, we associate a timestamp with each data point and modify the control messages and local repositories accordingly. As energy level is a monotonically decreasing value, a simpler solution would be to always believe/use the smallest value. However, this would not work if we were to explore other, non-monotonic, QoS metrics such as queue lengths (for load-balanced routing) or link bandwidth (as in [21]). No clock synchronization will be required if these timestamps are only used to compare values originating from the same node (to determine the most recent value). However, if we were to, in addition, use these timestamps to analyze the information for age or propagation delay, some relatively coarse-grained clock synchronization among nodes becomes necessary.

V. SIMULATION RESULTS

A. Simulation Setup

We extensively conducted simulation experiments to evaluate the performance of our various OLSR versions under different traffic loads and mobility scenarios. All simulations were performed with the NS2 simulator with the OOLSR implementation of OLSR provided by the Hipercom project (NS2 version 2.27 with OOLSR version 0.99.15)[32]. The common and fixed simulation parameters are summarized in Table 1. We simulated a MANET with 50 nodes in a 1000 x 1000 meter square area. Each node is equipped with an IEEE 802.11 wireless interface with a transmission range of 250m, and buffers outbound packets in a priority queue of size 50 that drops packets at the queue end in case of overflow. We deliberately choose a rather low initial energy value to generate scenarios where nodes deplete their battery and die. As we discussed in the first section, we assume that idle energy consumption is addressed with any of the existing complementary approaches. Therefore, and to emphasize the possible gains due to routing decisions, we ignored idle energy consumption by setting it to 0.

All results reported here are the averages for at least 10 simulation runs. During each run, 3 sources send 128 bytes CBR packets at a given fixed intervals to three different receivers. The simulations are done with 4 different intervals to study the effect of low, medium and high traffic rates. The intervals are 0.2, 0.14, 0.09, and 0.04 seconds – an interval of 0.2 means a node will send a packet every 0.2 of a second (i.e. 5 packets per second). At the shorter intervals, in particular for 0.04, a significant number of packets (data and control packets) will be lost due to network congestion, testing the performance under stressful overload scenarios. These may not be representative of long-term sustained traffic, but often occur in MANETs due to the link bandwidth limitations and the bursty nature of data traffic.

Table 1: Simulation Parameters

Simulator Parameters	
Propagation model	TwoRayGround
Network Type	IEEE 802.11
Transmission Range	250 m
Mobility Model	Static Network
Queue Length	50
Interface Queue	DropTail/PriQueue
Scenario Parameters	
Topology area	1000*1000
Number of nodes	50
Simulation time (secs)	200 seconds
Energy Model Specifications	
Initial Energy (Joules)	15
Transmission Power (Watt)	1.4
Receiving Power (Watt)	1.0
Idle Power (Watt)	0.0

We also varied the mobility patterns. In the static scenarios, nodes are randomly placed into the network, but do not move during the simulation. In the low mobility scenario, nodes move around the simulation area based on the RWP mobility model, with a maximum speed of 2 m/s and a pause time of 10 seconds (which could for example model a pedestrian mobility pattern). In the high mobility scenario, nodes move at speeds of up to 20 m/s or 72 km/h and no pauses, which corresponds more to vehicular movements at relatively high speeds.

In previous work, we explored a number of OLSR variants, and found that the best overall performance across different traffic loads and mobility scenarios is obtained by using the new path determination algorithm, combined with the original MPR selection criteria. For many scenarios, in particular as node mobility increases, changing the MPR selection criteria (in addition to the new path determination algorithm) is beneficial as well. So we use these two protocol variants: *Modified Routing* refers to the version that uses the original MPR selection criteria, but uses the new path determination algorithm, whereas *Modified MPR/Routing* combines both the new MPR selection and the new path determination algorithm. We use two different versions of each protocol, the *ideal* version (where a node has access to the actual, current residual energy level of remaining nodes when selecting MPRs and determining paths) and the *realistic* version, in which a node needs to rely on the residual energy levels it learned through protocol messages as described above.

B. Performance Metric

RFC 2501 [33] describes performance metrics for the evaluation of routing protocols and has been widely adopted. Due to space limitations, we focus on only one parameter here, and that is the network's ability to deliver data packets. This is traditionally expressed as Packet Delivery Ratio or PDR, measured as the ratio of the number of packets delivered to a traffic sinks relative to the number of packets sent by the traffic sources. However, in our scenarios, where nodes die due to the depletion of their energy source, this metric can be rather misleading. Unlike some earlier evaluations such as [9][10], all our nodes are energy-constrained, including traffic sources and sinks. In this case, a routing protocol that would assure that traffic sources deplete their energy first would then achieve high PDR metrics, even though the network itself only delivered very few packets in total. We therefore evaluate the various protocol variants by counting the total number of packets successfully delivered.

C. Evaluation

Figures 1 to 3 summarize the results for the three mobility scenarios, plotting traffic load against protocol performance for the various protocol variants introduced above. In all three figures, the solid lines indicate the performance for the *Modified Routing* variation of OLSR, and the dashed lines show the performance for the *Modified MPR/Routing* variant. The square symbol indicates the performance of the *ideal* protocol version, whereas the diamond shows the performance of the *realistic* protocol version. For the static scenarios, shown in Figure 1, there is a clear difference between the ideal and realistic version, with the *realistic* version delivering from close to 250 packets at low traffic rates to roughly 1000 packets fewer than the *ideal* version for the *Modified Routing* variant.

For low mobility scenarios, the total number of packets delivered drops, see Figure 2. This is consistent with many other MANET protocol performance studies that found that mobility has a detrimental impact on overall performance. Similar to the static scenarios, the *ideal* version clearly outperforms the *realistic* version for all traffic loads.

Figure 3 summarizes our results for the high mobility scenarios. The total number of delivered packets dropped further, but the comparative performance stays the same: the *ideal* protocol versions outperform the *realistic* versions for all traffic loads. At the lowest traffic load, the *ideal* version of the

Modified Routing variant delivers over 1040 more packets than the *realistic* version, for a performance gain of 10.7%. Similarly, the *ideal* version of the *Modified MPR/Routing* variant delivers over 740 additional packets for a performance gain of 8.4%, compared to the *realistic* version.

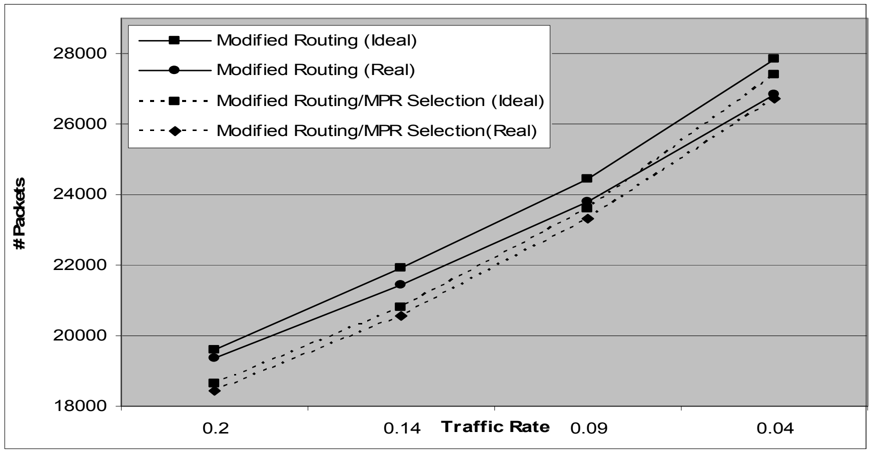

Figure 1: Protocol Variant Performance for Static Network Scenarios

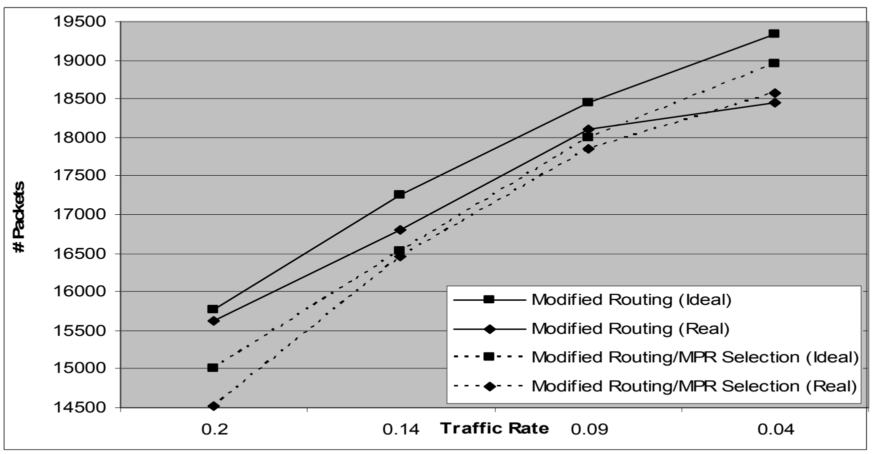

Figure 2: Protocol Variant Performance for Low Mobility Network Scenarios

We performed some initial analysis for the *realistic* protocol versions, exploring the accuracy and timeliness of the collected routing information (residual energy levels for other nodes in the MANET). In general, the collected routing information is less accurate as the traffic load and mobility patterns increase. The inaccuracy is directly related to the age of the collected information: the older, the more inaccurate. For example, under high traffic load, a node's knowledge about other node's residual energy levels often is more than 20 seconds old (i.e., the most recent information, received either via a Hello or a TC message, has a timestamp more than 20 seconds lower than a node's current time). Such information is clearly inaccurate, causing a node to make relatively poor routing decisions.

VI. CONCLUSIONS

In this work, we studied different OLSR protocol variants that aim to increase node lifetime and network performance, taking a node's residual energy levels into account to avoid routing through nodes with low residual energy. We focus particularly on the impact this information being inaccurate; comparing the protocol performance for *realistic* and *ideal* versions. In the *ideal* versions, a node has access to the current, accurate, routing-relevant parameters (here residual energy level). In the *realistic* version, a node collects this information through protocol control messages, which results in the information being somewhat out-of-date and potentially quite inaccurate. We evaluated our protocol variants under a range of different scenarios, varying traffic load and mobility pattern. For all traffic loads and mobility scenarios, the *realistic* protocol versions showed poorer performance than the *ideal* versions, delivering in some cases as much as 1000 fewer packets, for a relative performance loss of about 10%. These results show, among others, that testing a routing protocol in a simulator, assuming ideal knowledge about the relevant routing parameters, may not yield realistic results.

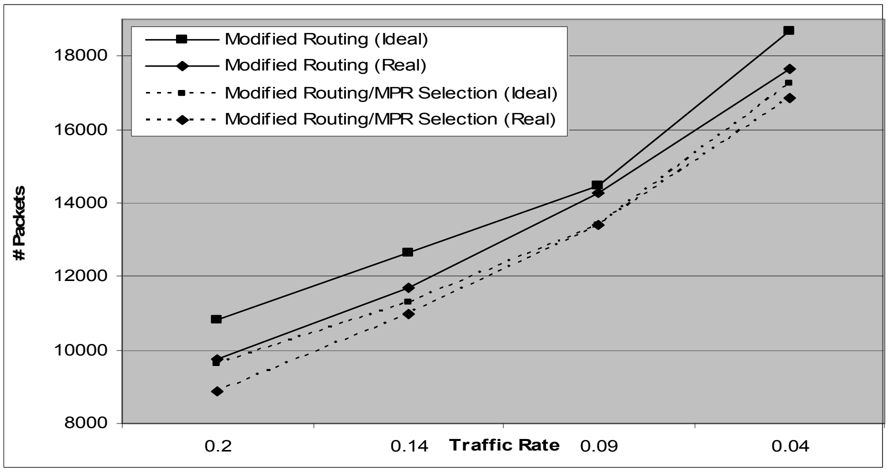

Figure 3: Protocol Variant Performance for High Mobility Network Scenarios

This work is only a starting point. We have not yet explored whether changing the OLSR protocol control parameters would significantly impact the protocol performance, for example by increasing the frequency of Hello or TC messages (which may potentially provide more recent and therefore accurate routing-related information). We also have not yet analyzed in-depth the amount of inaccuracy incurred in the *realistic* protocol variants, and how to improve this accuracy, using probing or triggered updates, as suggested in some of the reviewed related work. Note however that all these mechanisms will also increase the protocol overheads and consume nodal energy. Finally, we may consider other ways a node may collect/derive (more) accurate nodal energy levels, for example through a combination of guessing and prediction.

VII. ACKNOWLEDGEMENTS

This work was supported by the National Sciences and Engineering Research Council Canada.

REFERENCES

[1] T. Clausen, P. Jacquet, "*Optimized Link State Routing Protocol (OLSR)*", RFC 3626, Oct. 2003. http://www.ietf.org/rfc/rfc3626.txt
[2] C. E. Perkins et al., "*Ad Hoc On-Demand Distance Vector (AODV) Routing*," RFC 3561, July 2003. http://www.ietf.org/rfc/rfc3561.txt
[3] E. Royer and C.-K. Toh, "*A Review of Current Routing Protocols for Ad Hoc Mobile Wireless Networks*", IEEE Personal Communications Magazine, April 1999, pp. 46-55.
[4] S. Das et al., "*Simulation-Based Performance Evaluation of Routing Protocols for Mobile Ad Hoc Networks*," Mobile Networks and Applications, 2000, 5(3):179-189.
[5] L. M. Feeney, "*An Energy Consumption Model for Performance Analysis of Routing Protocols for Mobile Ad Hoc Networks*", Mobile Networks and Applications, 6(3):239-249.
[6] Y. Gaddallah, T. Kunz, "*Energy consumption in ad-hoc routing protocols: Comparing DSR, AODV and TORA*", Proc. 1st Int. Conf. on Ad-Hoc Networks and Wireless, Sept. 2002, pp. 161-176.
[7] M. Fotino et al.., "*Evaluating Energy Consumption of Proactive and Reactive Routing Protocols in a MANET*", Proc. First Int. Conf. on Wireless Sensor and Actor Networks, Sept. 2007, pp. 119-130.
[8] IEEE 802.11 Standard, LAN MAN Standards Committee of the IEEE Computer Society, 1999.
[9] Y. Xu et al., "*Geography-informed Energy Conservation for Ad Hoc Routing*", Proc. ACM Mobile Comp. and Networking Conf., July 2001, pp. 70-84.
[10] B. Chen et al., "*Span: An Energy-Efficient Coordination Algorithm for Topology Maintenance in Ad Hoc Wireless Networks*", Proc. ACM Mobile Comp. and Networking Conf., July 2001, pp. 85-96
[11] R. Zheng, R. Kravets, " *On-demand Power Management for Ad Hoc Networks*", Proc. INFOCOM, April 2003, pp. 1-11.
[12] Y. Gadallah, T. Kunz, "*A protocol independent energy saving technique for mobile ad hoc networks*", Special Issue on Mobile and Wireless Networking, International Journal of High Performance Computing and Networking (IJHPCN), 4(1-2):39-50, 2006.
[13] T. Kunz, "*On the inadequacy of MANET routing to efficiently use the wireless capacity*", Proc. IEEE Int. Conf. on Wireless and Mobile Computing, Networking and Communications (WiMob 2005), Aug. 2005, pp. 109-116.
[14] S. Singh et al., "*Power-aware Routing in Mobile Ad Hoc Networks*", Proc. ACM Mobile Computing and Networking Conference, Oct. 1998, pp. 181-190.
[15] C. E. Jones et al., "*A Survey of Energy Efficient Network Protocols for Wireless Networks*", Wireless Networks, Kluwer Academic Publishers, 2001, pp. 343-358.
[16] Y. Luo et al., "*An energy-efficient DSR routing protocol based on mobility prediction*", Proc. Int. Symposium on a World of Wireless, Mobile and Multimedia Networks, June 2006.
[17] X. Jin et al., "*A RED based minimum energy routing algorithm for wireless ad-hoc networks*", Proc. 2005 Int. Conf. on Wireless Communications, Networking and Mobile Computing, Sept. 2005, pp. 757 – 761.
[18] D. Kim et al., "*Routing Mechanisms for Mobile Ad Hoc Networks Based on the Energy Drain Rate*", IEEE Transactions on Mobile Computing, Vol. 2, No. 2, Jan. 2003, pp. 161-173.
[19] S. Senouci, M. Naimi, "*New routing for balanced energy consumption in mobile ad hoc networks*", Proc. 2nd ACM Int. Workshop on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks, Oct. 2005, pp. 238 – 241.
[20] S. Jung et al., "*Energy efficiency of load balancing in MANET routing protocols*", Proc. Sixth Int. Conf. on Software Eng., Artificial Intelligence, Networking and Parallel/ Distributed Computing, 2005 and First ACIS Int. Workshop on Self-Assembling Wireless Networks, May 2005, pp. 476 – 483.
[21] Y. Ge et al., "*Proactive QoS routing in ad-hoc networks - but at what cost?*" Proc. 2nd Int. Conf. on Ad-Hoc Networks and Wireless, Oct. 2003, pp. 60-71.
[22] K. Kowalik et al., "*Making OLSR Aware of Resources*", Proc. Int. Conf. on Wireless Communications, Networking and Mobile Computing, 2007, Sept. 2007, pp. 1488-1493.
[23] N. Ghanem et al., "*New energy saving mechanisms for mobile ad-hoc networks using OLSR*. Proc. 2nd ACM Int. Workshop on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks, Oct. 2005, pp. 273 – 274.
[24] Z. Guo, B. Malakooti, "*Energy Aware Proactive MANET Routing with Prediction on Energy Consumption*", Proc. 2007 Int. Conf. on Wireless Algorithms, Systems and Applications, August 2007, pp. 287-293.

[25] A. Benslimane et al., "*Energy Power-Aware Routing in OLSR Protocol*", Proc. First Int. Mobile Comp. and Wireless Communication Conf., Sept. 2006, pp. 14-19.

[26] T. Clausen et al., "*Investigating the Impact of Partial Topology in Proactive MANET Routing Protocols"*, Proc. 5th Int. Symposium on Wireless Personal Multimedia Communications, Oct. 2002, Vol. 3, pp. 1374-1378.

[27] P.E. Villanueva-Pena et al., "*Extending network knowledge: Making OLSR a Quality of Service conductive protocol"*, Proc. Int. Conf. on Communications and Mobile Computing, Vancouver, Canada, July 2006, pp. 103-108.

[28] Q. Han and N. Venkatasubramanian, "*Information collection services for QoS-Aware mobile applications*," IEEE Transactions on Mobile Computing,, May 2005, 5(5):518-535.

[29] G. Apostolopoulos et al., "*Quality of Service Based Routing: A Performance Perspective*", Proc. ACM SIGCOMM, Vancouver, Canada, Oct. 1998, Vo. 28, No. 4, pp. 17-28.

[30] R.A. Guerin, A. Orda, "*QoS routing in networks with inaccurate information: theory and algorithms"*, IEEE/ACM Transactions on Networking, June 1999, 7(3):350-364.

[31] A. Shaikh,, J. Rexford and K. G. Shin," *Evaluating the impact of stale link state on quality-of-service routing"*, IEEE/ACM Transactions on Networking, Apr. 2001, 9(2):162-176.

[32] Hipercom Project, *OOLSR Implementation of the OLSR Optimized Link State Routing Protocol*, http://hipercom.inria.fr/oolsr/

[33] S. Corson, J. Macker, "*Mobile Ad hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations*", RFC 2501, January 1999. http://www.ietf.org/rfc/rfc2501.txt