

Cross-Layer Performance Evaluation of Sensor Networks: Routing over Energy Efficient MAC

Hsing-Jung Huang^a, Ling-Jyh Chen^d, Polly Huang^{a,b,c}

^aDepartment of Electrical Engineering

^bGraduate Institute of Communication Engineering

^cGraduate Institute of Networking and Multimedia
National Taiwan University

^dInstitute of Information Science
Academia Sinica

r93921046@ntu.edu.tw, ccljj@iis.sinica.edu.tw, phuang@cc.ee.ntu.edu.tw

ABSTRACT

Delivering data in Wireless Sensor Networks (WSNs) is very challenging in that an ideal solution has to not only provide a reasonable data throughput, but also take into consideration a variety of constraints, such as battery life, computation capabilities, wireless interference, node mobility, and etc. As numerous solutions have been proposed in the past years, most designs have focused on the functionalities within a particular layer, i.e., routing protocol design in the network layer or MAC enhancements in the data link layer. However, the overall performance depends also on the interaction of mechanisms across layers. A comprehensive cross-layer performance evaluation is essential to a better understanding of the overall system performance when deploying a new protocol. In this case study, we simulate the interaction between a more recently proposed routing protocol and a number of MAC protocols. The results show that the performance of the routing protocol varies as different MAC schemes are employed. In particular, the state-of-the-art MAC protocol might not perform the best with the routing protocol in study. No single combination dominates the other combinations in terms of all metrics in all cases.

Categories and Subject Descriptors

C.4 [Computer Systems Organization]: PERFORMANCE OF SYSTEMS—*Design studies, Performance attributes*

General Terms

Design, Performance

Keywords

Cross-Layer Design, MAC, Routing, Sensor Network

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PM²HW²N'08, October 31, 2008, Vancouver, BC, Canada.
Copyright 2008 ACM 978-1-60558-239-9/08/10 ...\$5.00.

1. INTRODUCTION

The energy efficiency and reliability requirements impose great challenges on the design of communication protocols for WSNs. Although there have been many proposals for routing and MAC layers, as well as a significant number of innovations in each of these areas, very few works have addressed the interaction of protocols across different layers. Many works propose changes in a specific layer and assume a certain implementation for the other layers. For example, it has been confirmed that IEEE 802.11 is an inadequate MAC for WSNs due to inefficient use of energy. The majority of energy consumed by a IEEE 802.11 radio is for idle listening. Previous works that do not take this factor into consideration when evaluating their routing protocols might be biased.

Application developers are in an awkward situation, evaluations of individual protocols might not be comparable. The interaction between different protocols at various layers of the network stack remains largely unknown. Because of these limitations, it is hard to select an appropriate combination of protocols for the applications at hand.

There has been a number of related works on cross-layer interaction in ad hoc networks. As noted in [1], MAC protocols selected for simulation studies are critical to the performance of routing protocols, and this aspect must be considered when comparing the performance of routing protocols. In [2], the authors also suggest that routing layers affect MAC layers and vice versa. It is not meaningful to consider a routing or a MAC protocol in isolation. Thus, the relative performance of routing protocols needs to be examined when using MAC protocols designed specifically for WSNs.

In this work, we simulate and analyze the performance of magnetic diffusion [3], a multi-path routing protocols, in combination with IEEE 802.11 [4], S-MAC [5], B-MAC [6], and Z-MAC [7]. We discover that the relative performance of the routing protocol's two different forwarding modes changes when interacting with different MAC protocols. No single combination dominates the other combinations in terms of all metrics in all cases. Looking into the details of the mechanisms, we find the interaction between the routing and MAC layers might result in surprising performance degradation. Lastly, the traffic workload affect also significantly the cross-layer network performance.

2. STUDIED ROUTING PROTOCOL: MAGNETIC DIFFUSION

Magnetic diffusion [3], referred to as MD, is based on a simple, and yet effective idea. Consider the data sink as a magnet and the data as metallic nails. The data is attracted to the sink because of the magnetic field, just as the nails are attracted to the magnet. The magnetic field is established by setting up appropriate magnetic charges on the sensor nodes within the magnetic influence of the data sink. The strength of the charge is determined by the hop distance or delay to the sink. The data is propagated based on the magnetic field from low-charge to high-charge nodes. Forwarding data in this fashion results in multiple shortest paths, which provide a higher delivery rate without compromising delay.

MD operates in two phases: the interest broadcast and data propagation. The magnetic field is established in the interest broadcast phase, such that the data can be disseminated towards the sink in the data propagation phase. In the subsections below, we describe in more detail the MD operation and the implementation options.

2.1 Interest Broadcast

We first introduce two data entities involved in the interest broadcast phase: the interest and the entry. The *interest* is a message used to establish the magnetic field to facilitate the proper flow of data toward the sink. The *entry* contains information about where a node records the data type and the magnetic charge in each node.

When a sink wants to collect data, it sends an interest message to its neighbor nodes. After receiving an interest message for the first time, a node creates an entry for the interest. Then, the node decrements the magnetic charge of the interest by one. The node records the data type and the magnetic charge in its entry and then forwards the interest message to its neighbors. The decremented magnetic charges from the sink to the source guide the flow of data in the reverse direction, mimicking the movement of metallic nails from low-charge to high-charge points in a magnetic field.

The sink node will broadcast interest messages periodically to re-establish proper magnetic charges in the network. This provides an environment for robust data dissemination, especially in a dynamic network.

In Fig. 1(a), the sink node broadcasts the interest to its neighbor nodes. In Fig. 1(b), nodes A and B receive the interest, create corresponding entries, set the data type and proper magnetic charges in the entries, and then propagate the interests to their neighbors. In Fig. 1(c), nodes C, D, and E receive the interests from nodes A and B and repeat the actions taken by A and B in Fig. 1(b). Each node may receive and discard duplicate interests or interests with weaker magnetic charges.

2.2 Data Propagation

Next, we describe how the data is disseminated throughout a network. We have two implementation strategies for data propagation. One is *gradient-based* and the other is *broadcast-based*.

2.2.1 Gradient-based Strategy

In the interest broadcast phase, when a node receives an interest from its neighboring nodes with a stronger magnetic




Figure 1: An example of interest propagation.




Figure 2: An example of data propagation with the gradient-based mechanism.

charge, the node establishes a gradient toward the interest-sending node. This gradient directs the data through the network from the source to the sink.

When a node senses data, it checks if it has an entry matching the data type. If it does, it sends the data to the nodes indicated by the existing gradients. The forwarding process continues until the data reaches the sink.

Fig. 2(a) shows the gradients established by the interest broadcast in 1. In Fig. 2(b), the source sends data to the two neighbor nodes indicated by the gradients from the source. In Fig. 2(c), the receiving nodes continue to send data to their neighbor nodes based on their gradients until the sink receives the data.

2.2.2 Broadcast-based Strategy

In the broadcast-based mode, the nodes do not establish gradients in the interest broadcast phase. Instead, the magnetic charge is included in the data being disseminated. The receiving node can tell from the charge carried in the data where the data is from and whether to forward further downstream.

More specifically, when a node receives data, it checks if it has any matching entry. If it does, it compares the magnetic charge in the entry with the magnetic charge of the data. If the former is greater than the data, it sets the magnetic charge of the data to that in the entry and then broadcasts the data. This means the data is sent from the node whose magnetic charge is lower than the intermediate node, and the intermediate node repeats this process. If a




Figure 3: An example of data propagation with the broadcast-based mechanism.

node receives duplicate data, or data whose magnetic charge is greater than that in the entry, the data will be discarded.

In Fig. 3(a), the magnetic charge of every node is established; the sink's charge is 7. In Fig. 3(b), when the source wants to send data to the sink, it will broadcast the data to its neighbor nodes. In Fig. 3(c), the nodes with charge strength 6 broadcast the data because the magnetic charge of the nodes is greater than that of the data. Thus, the sink receives the data.

3. STUDIED MAC PROTOCOLS

3.1 IEEE 802.11

The 802.11 [4] uses CSMA/CA to handle multiple access and prevent collisions in a shared medium. A node senses the medium before a transmission; if the medium is busy, then it defers. If the medium is free for a specified time (called DIFS, Distributed Inter Frame Space) then the node is allowed to transmit.

802.11 uses a simple handshake mechanism between senders and receivers to address the hidden terminal problem. The sender first sends an RTS packet to the intended receiver who then replies with a CTS packet. There is a duration field in the RTS and CTS packets to indicate the amount of time for the following transaction, so that neighboring nodes know how long they should back-off in order not to interfere with the transmission. The sender starts transmitting data after it receives the CTS packet, and receiver replies with an ACK packet after each data packet arrives. Transmissions in 802.11 follow the sequence: RTS-CTS-DATA-ACK.

In 802.11, exponential back-off is used to resolve contention problems. Before transmitting a packet, a node waits for a random period of time chosen between zero and its contention window size. If the medium is free after waiting that amount of time, the node transmits immediately. If the medium is busy, the transmission is deferred, and the process will be repeated when the channel is free. If a collision happens during the transmission, the node will increase the contention window size exponentially and repeat the transmission process when the channel is free.

3.2 S-MAC

The goal of the S-MAC [5] protocol is to reduce energy consumption. In wireless sensor networks, nodes might spend a significant amount of time on idle listening and thus waste energy. To reduce energy consumption, S-MAC lets nodes go into the sleep mode periodically. Each node goes to sleep for a period, and wakes up to listen for a period. This constitutes a periodic listen and sleep schedule for nodes.

In S-MAC, a node needs to choose a schedule and inform its neighboring nodes. The schedule of a node is determined as follows. First, the node listens for a certain amount of time. If it does not hear a schedule from any other node, it randomly chooses a schedule and immediately broadcasts it to the neighbors. If a node receives a schedule from a neighbor before choosing its own schedule, it follows that schedule by setting its schedule to be the same and rebroadcasts this schedule. If a node receives a different schedule after it selects and broadcasts its own schedule, it adopts both schedules. Due to clock drift, synchronization is needed among neighboring nodes. Each node broadcasts a synchronization packet periodically, and receivers adjust their timers accordingly.

S-MAC is a contention-based protocol like 802.11. A node performs randomized carrier sensing before initiating a transmission. Unicast packets follow the sequence: RTS-CTS-DATA-ACK. Broadcast packets are sent without RTS/CTS. In the transmitted packet, there is a duration field that indicates how much longer the transmission will last. The neighboring nodes then know how long they should sleep to avoid overhearing and thereby save energy.

3.3 B-MAC

B-MAC [6], which is also a contention-based protocol, provides power management via low power listening (LPL). A node maintains a listening duty cycle separated by a specific period of time called the check interval. The node wakes up to check activity every check interval. If activity is detected, the node stays awake and receives the incoming packet. However, if the medium is clear, the node goes back to sleep. To support LPL, each transmission is preceded by a preamble as long as the check interval so the intended receiver is definitely aware of the transmission and receives the incoming packet. This scheme shifts the load from receivers to senders, and saves a substantial amount of energy when the traffic load is light.

B-MAC provides a flexible interface that allows reconfiguration of the system parameters on the fly according to the state of the network. It also gives control to the upper layers. For example, link-level retransmission is disabled in default settings of B-MAC. The higher layers are allowed to enable these schemes if necessary. The idea is to facilitate cross-layer optimization while simultaneously preserving the layered architecture.

3.4 Z-MAC

Z-MAC [7] is built on top of B-MAC, and also uses LPL to save energy. It attempts to combine the strengths of TDMA and CSMA to achieve high channel utilization under both low contention and high contention. In Z-MAC, a higher overhead is incurred initially because of the slot assignment algorithm. The design philosophy is that the high initial overhead is amortized over a long period of network operation, and compensated for by improved throughput and energy efficiency.

First, each node gathers its two-hop neighbor information, which is used as input to the slot assignment algorithm. After the slot assignment, each node decides its own local frame and broadcasts its frame size and slot number to its two-hop neighborhood. Thus, each node knows about the slot and frame information of its one-hop and two-hop neighbors.

Unlike TDMA, a node may transmit during any slot in Z-MAC, though each node has its own time slot. Before a transmission, a node performs randomized carrier sensing. If the medium is free, it transmits the packet. Otherwise, the process will be repeated when the medium is not busy. However, the owner of a slot always has priority over non-owners in accessing the medium. The priority is implemented by adjusting the contention window size such that the owner always has the chance to transmit earlier than non-owners. In this way, non-owners can still compete for transmission when the slot is not being used by its owner; thus, channel utilization is enhanced.

Let us consider the details of transmission control. To address the hidden terminal problem, Z-MAC defines two modes: low contention level (LCL) and high contention level

Table 1: Simulation Settings

Nodes	50
Size	160 x 160
Radio range	40 m
Bandwidth	20kbps
Tran power	60 mW
Recv power	45 mW
Idle power	45 mW
Sleep power	0.33 mW
Data rate	10 sec
Periodic Interest	120 sec
Duty cycle of S-MAC	10%
Check interval of LPL	50 ms
Simulation time	1200 sec

(HCL). A node is in HCL only when it receives an explicit contention notification (ECN) message from a two-hop neighbor within a certain period. In LCL, a node can compete for transmission in any slot, but in HCL, only the owner of a slot and its one-hop neighbors are allowed to compete for the medium. The ECN messages are used to notify two-hop neighbors not to compete with the owner of a slot when contention is high. Each node makes a local decision to send an ECN message based on its local estimation of the contention level.

Z-MAC must be synchronized to maintain the TDMA-like frame. Therefore, a synchronization packet is sent at specific times. In the worst case, where nodes are completely unsynchronized, it falls back to CSMA.

4. PERFORMANCE EVALUATION

TABLE 1 summarizes the parameters used in the experiments. To study the performance of different combinations of routing and MAC protocols, we run the experiments in ns-2 simulator. In order to get the average behavior, there are 10 distinct runs for each setup. In each run, there are 50 nodes placed at random in a 160m by 160m square. We use one source and one sink randomly selected from the 50 nodes. Unless otherwise specified, we use the default settings of S-MAC, B-MAC, and Z-MAC. Most of the parameters of the simulation environment are based on the capability of Mica2, which is the development platform of the above MAC protocols.

We choose several metrics to measure the performance of different combinations of routing and MAC protocols: **Energy** measures the amount of power consumed by the network. The metric is very important in a resource-constrained wireless sensor network. **Reachability** measures the probability that the sink will receive a data packet successfully. This metric is important for medical applications in that the sensor data is mission critical, and its loss could be a matter of life-or-death. **Latency** measures the data transmission time from the source to the sink. For medical applications, the metric represents the timeliness and temporal reliability of the data.

We examine the performance of combinations of MD and four MAC protocols: Thus, we have eight protocol combinations: MDB with 802.11 (MDB-802), MDG with 802.11




Figure 4: Energy consumption of all combinations of routing and MAC protocols




Figure 5: Energy consumption of all combinations of routing and MAC protocols with increasing numbers of sources

(MDG-802), MDB with S-MAC (MDB-SMAC), MDG with SMAC (MDG-SMAC), MDB with B-MAC (MDB-BMAC), MDG with B-MAC (MDG-BMAC), MDB with Z-MAC (MDB-ZMAC), and MDG with Z-MAC (MDG-ZMAC).

4.1 Energy Consumption

Fig. 4 shows the energy consumption of all combinations of routing and MAC protocols. The energy consumption of MDG-802 is a bit higher than that of MDB-802 because MDG sends more data packets over the network. A similar result can be observed in S-MAC, B-MAC, and Z-MAC.

Recall that to avoid overhearing unicast packets, idle nodes in S-MAC are turned off. This mechanism helps conserve energy for nodes that are not involved in the communication. Therefore, the energy consumption of MD using S-MAC is significantly lower. Furthermore, the amount of packets required to transmit in MDG is higher than that in MDB. This explains the slightly higher energy consumption in MDG-SMAC.

Relative to B-MAC, Z-MAC is more sophisticated. However, our simulation results show that MD-ZMAC consumes more energy than MD-BMAC. This is due to a substantial amount of energy consumed by Z-MAC in the setup phase. Although the start-up cost will gradually take a lower proportion as the system runs longer, the combinations with B-MAC quantitatively outperforms the other combinations in this set of experimental setups.

Let us further analyze the energy consumption of combinations of routing and MAC protocol combinations for different numbers of sources. For the ease of understanding,




Figure 6: Reachability of all combinations of routing and MAC protocols

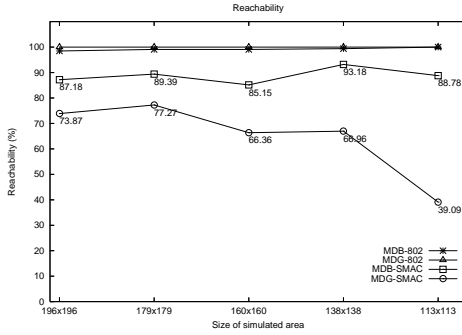


Figure 7: Reachability of MDB-802, MDG-802, MDB-SMAC, and MDG-SMAC for simulated areas of various size

we plot only the results of S-MAC, B-MAC, and Z-MAC in Fig. 5. The energy consumption of MDG-BMAC and MDG-ZMAC steadily increases as the number of sources grows. This is because B-MAC and Z-MAC use LPL to save energy so that the sender must send a long preamble before a transmission. This scheme shifts the load from the receiver to the sender. Hence, the energy consumption increases quickly when the traffic load is high. In contrast, the energy consumption of MDB-BMAC and MDB-ZMAC is unchanged because the overhead of MDB does not increase as much as that of MDG. This result indicates that the choice of routing protocols is also critical to power consumption when combined with the MAC protocols that use LPL to save energy. The energy consumption of S-MAC is relatively stable compared to the other two MAC protocols. The difference is due to the fixed schedule of S-MAC, which causes nodes to spend most of their energy on idle listening.

4.2 Reachability

Fig. 6 shows the reachability of all combinations of routing and MAC protocols. In S-MAC, nodes might have accumulated a number of packets during a long sleep interval, but these packets can only be sent at the beginning of a listening interval. Hence, the contention is quite high and collisions often occur. This results in the lower reachability observed in MDB-SMAC and MDG-SMAC.

Intuitively, the gradient-based mechanism (MDG), which is capable of collision detection and retransmission, should perform better than the broadcast-based mechanism (MDB). MDG indeed performs better than MDB when combined




Figure 8: Reachability of all combinations of routing and MAC protocols with increasing numbers of sources

with 802.11, B-MAC, and Z-MAC. However, MDB outperforms MDG when using S-MAC. S-MAC assumes that, after a scheduling phase, each node knows about the existence of all of its neighbors; thus, all neighboring nodes have the same schedule. Therefore, a node cannot send packets to an unknown node because it thinks the schedule of that node is unclear. Unfortunately, in a dense network with a lot of collisions, a node might have the schedule of an unknown neighbor because it can receive that schedule from other neighboring nodes. This causes a strange situation, in that nodes cannot send a unicast packet to unknown neighbors, even if they have the same schedule. However, S-MAC sends multiple copies of a broadcast packet, one for each schedule. Therefore, the neighbors can receive the packet, even if the sender is not aware of their existence. This explains the difference in the reachability of MDB-SMAC and MDG-SMAC. This example shows the complex interaction between the routing and MAC layers. The relative performance of routing protocols varies when combined with different MAC protocols.

To verify our observation about the strange situation, we change the size of the simulated area, but keep the number of nodes the same. In other words, the density of the nodes changes from low to high. The reachability of different protocol combinations is shown in Fig. 7. The plot shows that the reachability of MDG-SMAC gradually decreases as the density of the nodes increases, while that of MDB-SMAC remains the same.

Fig. 8 shows the reachability of all combinations of routing and MAC protocols for different numbers of sources. MDG-802 and MDG-ZMAC outperform other combinations under low and high traffic loads respectively. The reachability of MDB declines as the number of sources increases which indicates that MDG performs better than MDB when the level of contention is substantially higher. Recall that Z-MAC utilizes the topology information to avoid two-hop collisions. Because of retransmission and the control of two-hop contention, MDG-ZMAC achieves the highest reachability. This plot clearly shows the strength of Z-MAC under a high traffic load. The substantial energy cost involved in the setup phase is compensated for by increased energy efficiency later.

4.3 Latency

Fig. 9 shows the cumulative probability of data delivery latency for all combinations of routing and MAC protocols.

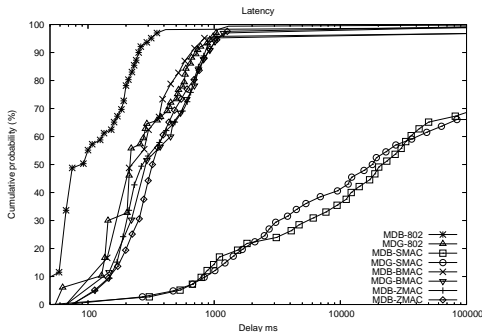


Figure 9: Cumulative probability of data delivery latency for all combinations of routing and MAC protocols

Clearly MDB-802 has the least latency, while protocols that use S-MAC have the most. To save energy, S-MAC delays packets until the beginning of a listening interval; hence, the delay is much longer than in the other MAC protocols. MDB-802 has the least delay because it broadcasts data and bypasses the RTS/CTS handshake. MDB-SMAC, however, performs worse than MDG-SMAC. As noted above, S-MAC sends multiple copies of a broadcast packet, one for each schedule. Thus, in a dense network, nodes might have accumulated a number of packets, so newly received data cannot be sent immediately. This observation also shows that the relative performance of the routing protocols might change when run on different MAC protocols.

B-MAC performs slightly better than Z-MAC because Z-MAC ensures that a transmission does not cross over the boundary of a slot. In contrast, B-MAC allows nodes to transmit at any time, as long as there is no contention. Another reason is because, in HCL of Z-MAC, nodes can only transmit at slots that are not owned by two-hop neighbors; thus, they tend to wait for a long time. The advantage of MDB is not obvious when working with B-MAC and Z-MAC, because a node has to send a long preamble before a transmission. Therefore, the difference in delay is small.

Note that MDB-BMAC is comparable to MDG-802. In MDB-BMAC, the delay of all data packets is shorter than 1 second. Additionally, about 80% of the data is delayed for less than 500 ms. This result implies a surprising fact: it is possible that we can achieve a short delay even when using an energy-efficient MAC in wireless sensor networks.

4.4 Selection Guidelines

Table 2 summarizes our comparison of all combinations of routing and MAC protocols. Based on the experiment results, we derive the following selection guidelines.

- For wireless sensor networks with restrict energy constraints, the combination of MDB and B-MAC is the most plausible solution. It achieves a balanced performance in terms of energy efficiency, reliability, and data latency.
- For a system that requires absolute reliability, MDG with Z-MAC outperforms all other combinations in an environment where the sensor nodes are battery-powered. If a mains supply is available, MDG with 802.11 is the best choice.

Table 2: Best protocol combinations for different application requirements and deployment environments

	Battery	Wired
Energy	MDB B-MAC	
Reachability	MDG Z-MAC	MDG 802.11
Latency	MDB B-MAC	MDB 802.11

- For applications that requires timely data delivery, the combination of MDB and B-MAC is a better option in an environment where is difficult to get a continuous power supply. If a mains supply is available, MDB with 802.11 achieves the best performance in terms of latency.

We find that the choice of routing and MAC protocols is highly application dependent and deployment environment dependent. Take the elevator application in [8] as an example. The sensor nodes are battery-powered and the application requirements are high reachability and timely delivery of data. MDG with Z-MAC provides the highest reliability, but the delay is too long. MDB with B-MAC has the best performance in data delivery latency but it is unreliable when the traffic load is high. However, we observed that there are only two sources in the network, and it is rare that two sources send data at the same time. The traffic is light and the contention is not competitive in this application. Fortunately, we know that contention-based mechanisms work well under a low traffic load. Thus, we choose the combination of MDB and B-MAC to serve the elevator application. The system works quite well. This example illustrates how to select from alternative combinations of sensor networking protocols. Application designers have to evaluate the trade-off between different configurations and choose the most appropriate one according to the application requirements, the traffic load, and the type of traffic. The performance results and the system guidelines we present here will help application developers select the best combination of routing and MAC protocols for the application at hand.

4.5 Discussion

Based on the results of our experiments and our observations, we note the following points.

- An MAC protocol that uses a fixed duty cycle on radio, such as S-MAC, is not flexible enough to handle traffic changes in a network, and expends a lot of energy on idle listening. Nodes are only allowed to send packets accumulated during a long sleep period at the beginning of a listening period. This makes the contention quite competitive and causes a lot of collisions. B-MAC and Z-MAC use LPL provided by the hardware to enable power management. This scheme works much better than the fixed duty cycle. However, we observe that the energy consumption of this scheme increases rapidly when the traffic load is high.

- The relative performance of routing protocols might change when run over different MAC protocols. Thus, the evaluation of protocols should be done in totality. Additionally, we found that both the MAC protocol and the routing protocol are critical to the energy consumption and network performance.
- No single combination dominates the other combinations in terms of all metrics in all cases. The traffic load and the type of traffic dramatically affect network performance. Thus, the choice of protocol combination is both application and deployment environment dependent.
- Retransmission enhances network reliability. The control of two-hop contention effectively reduces the number of collisions. In wireless sensor networks, packets are usually delayed when an energy efficient MAC protocol is applied. This makes the contention competitive, and increases the probability of collisions. These issues need to be addressed when designing the next generation of MAC protocols for wireless sensor networks.

5. CONCLUSION

We have presented performance comparisons of the MD routing protocols when combined with 802.11, S-MAC, B-MAC, and Z-MAC. We found that the relative performance of the routing protocols changes when they are run on different MAC protocols, and vice versa. Thus, it is not meaningful to discuss routing or MAC protocols in isolation. The interaction between the routing and MAC layers might cause unpredictable behavior in the network. Additionally, the application workload and the type of traffic dramatically affect network performance. Application designers should consider a trade-off between different combinations of network protocols, and select the most appropriate one according to the application requirements, the traffic load, and the type of traffic.

We also demonstrate that the reliability of a system drops dramatically when two-hop contention is competitive. This issue needs to be addressed when designing the next generation of MAC protocols for wireless sensor networks.

We believe that our experience provides a more thorough understanding of the interaction between various network layers, which should be of value in the design of future protocols for wireless sensor networks.

6. REFERENCES

- [1] E. Royer, S. Lee, and C. Perkins, "The effects of mac protocols on ad hoc network communications," in *IEEE Wireless Communications and Networking Conference*, 2000.
- [2] C. Barrett, M. Drozda, A. Marathe, and M. Marathe, "Characterizing the interaction between routing and mac protocols in ad-hoc networks," in *ACM International Symposium on Mobile Ad Hoc Networking and Computing*, 2003, pp. 92–103.
- [3] H. J. Huang, T. H. Chang, S. Y. Hu, and P. Huang, "Magnetic diffusion: Scalability, reliability, and qos of data dissemination mechanisms for wireless sensor networks," in *Computer Communications Special Issue on Wireless Sensor Networks: Performance, Reliability, Security, and Beyond*, in press, 2006.
- [4] "Wireless lan medium access control (mac) and physical layer (phy) specification," in *IEEE Std. 802.11-1999 edition*.
- [5] W. Ye, J. Heidemann, and D. Estrin, "An energy-efficient mac protocol for wireless sensor networks," in *INFOCOM*, 2002.
- [6] J. Polastre, J. Hill, and D. Culler, "Versatile low power media access for wireless sensor networks," in *Proceedings of the Second ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2004.
- [7] I. Rhee, A. Warriier, M. Aia, and J. Min, "Z-mac: a hybrid mac for wireless sensor networks," in *ACM Sensys*, 2005.
- [8] S. Y. Lau, T. H. Chang, S. Y. Hu, H. J. Huang, L. de Shyu, C. M. Chiu, and P. Huang, "Sensor networks for everyday use: The bl-live experience," in *IEEE International Conference on Sensor Networks, Ubiquitous, and Trustworthy Computing (SUTC 2006)*.