Comparative Factor Analysis and Experimental Validation of Low Power MAC Performance

Wenjie Zeng, Jing Li, Anish Arora
Computer Science and Engineering
The Ohio State University
Columbus, OH, USA
{zengw, ljing, anish}@cse.ohio-state.edu

ABSTRACT

We analyze and experimentally validate the throughput capacity of duty cycled wireless networks in general and state-of-the art MAC protocols in particular. This enables a comparison of the energy efficiency of the protocols, which differ by multiples from up to 17x (for low traffic) to 2x (for high traffic), the latency of these protocols, and the tradeoff between these two metrics. We also evaluate the impact of specific protocol design factors on performance. Notably, we find that one extant protocol has the best energy efficiency and is latency-competitive across a broad range of network configurations and traffic. [1]

Categories and Subject Descriptors

C.2.5 [**Computer-Communication Networks**]: Network Protocols; C.2.2 [**Computer-Communication Networks**]: Local and Wide-Area Networks—*experimental evaluation, performance measures*

1. INTRODUCTION

In low-power wireless networks, MAC protocols not only coordinate packet communications to ensure reliable delivery —which is the primary objective for wall-powered wireless networks— they also coordinate the sleep/wakeup of nodes to control the energy cost and/or latency overhead. Notwithstanding a decade of productive research in low power MACs, optimizing MAC throughput, energy efficiency, and/or latency performance still remains as a challenge.

We attribute the difficulty of optimizing MAC performance to several inter-related reasons: (i) The capacity of a wireless network in a duty-cycled setting is not easily computed, and methods for its estimation have not received much attention. (ii) The coordination strategy of MACs has a substantial effect on the maximum throughput that can be achieved at any given duty cycle[2] of sleep/wakeup, but this effect is inadequately characterized; thus, given a MAC protocol and a traffic load, we do not know how to formally choose the least duty cycle at which the protocol can reliably deliver the traffic.[3] (iii) Likewise, the coordination strategy of MACs has a substantial effect on the latency, but we lack methods to choose duty cycles for a given traffic load (and, conversely, to choose traffic loads for a given duty cycle) to control both latency and energy efficiency.

Metric	Comparison
Throughput	RI-MAC > X-MAC [1]
	BoX-MAC > X-MAC [2]
	X-MAC > SCP-MAC [3]
	SCP-MAC > B-MAC [4]
	Crankshaft > SCP-MAC [5]
	SCP-MAC > Crankshaft (at high load) [5]
Energy Efficiency	O-MAC > B-MAC [6]
	RI-MAC > X-MAC [1]
	BoX-MAC > X-MAC [2]
	X-MAC > SCP-MAC [3]
	SCP-MAC > B-MAC [4]
Latency	RI-MAC > X-MAC [1]
	BoX-MAC > X-MAC [2]
	X-MAC > SCP-MAC [3]
	SCP-MAC > Crankshaft [5]
	SCP-MAC > B-MAC [4]
	SCP-MAC > Crankshaft (at high load) [5]

Table 1: Known MAC protocol performance comparisons

To appreciate that performance characterization of state-of-the-art MAC protocols is inadequate, consider the synopsis of performance comparisons from the literature given in Table 1, which are (with one exception) for relatively low traffic loads. Note that even for low traffic it is unclear which protocol performs best and whether their performance gaps are significant. It is not known whether the best protocol remains the same as traffic load changes substantially. With regard to (i), we do not know how far the maximum throughput of all of these protocols at various duty cycles are from what is achievable in theory.

It is notable that some of the comparisons do not attempt to parameterize each protocol so as to optimize the performance of the protocol under test. This may be rationalized as is not trivial to experimentally establish whether the duty

[1]Last updated February 25, 2011

[2]By duty cycle, we mean the percentage of the time that a node is up, including when the node is receiving, transmitting or idling.

[3]In general, the duty cycle of each node depends upon its traffic load, local channel utilization, and interference. If the MAC protocol allows it, the selection the duty cycle should not be the same for all nodes.

cycle achieved is optimal for a given traffic and network configuration. One reason for the difficulty is that each MAC typically only controls part of duty cycle, for instance the receiver's wakeup interval and/or the sender's transmission rate. The cumulative duty cycle is a derivative quantity that depends on a number of parameters, which makes its predication difficult.

The more important issue is that it is not well known how to calculate for a wireless network the maximum throughput that a given MAC protocol can achieve at a given duty cycle. By the same token, it is not known how to determine the minimum duty cycle at which the protocol can reliably deliver a given traffic load. Thus, it has been onerous for experimenters to compare the maximum throughput and energy efficiency of protocols across many network configurations. The net effect is that the question of which protocol to use for optimizing energy efficiency (and possibly latency also) is open, as is the tradeoff is between energy efficiency and latency.

Contributions of this paper. In this paper, we analyze the achievable throughput capacity of MACs in general and of state-of-the-art protocols in particular. We experimentally study the maximum throughput, energy efficiency, and latency of these protocols across diverse network configurations. Our study includes consideration of ultra-low loads at which the MACs are essentially always-off as well as high loads at which these MACs become always-on. Further, we investigate the tradeoff between energy efficiency and latency. Moreover, we deconstruct the comparative performance of the protocols in terms of central MAC design factors.

Duty cycled capacity. In contrast to previous studies in wireless networks, we consider capacity in a duty-cycled wireless network, wherein a node only wakes up its radio to communicate for a fraction of time. Our capacity analysis shows that *for random wireless networks, there is a bound on the duty cycle (which is usually less than* 100%*) beyond which the maximum throughput of the network does not increase.* Equally importantly, the capacity that is achievable in an n-node general network at a duty cycle of ψ increases not with $O(\frac{\psi}{\sqrt{n \log n}})$ but with $O(\sqrt{\frac{\psi}{n}})$ till the capacity limit is reached. And for the case of one hop MAC traffic, the capacity that is achievable increases linearly in ψ and independently of n till the capacity limit is reached.

We note that there is a close relationship between capacity and energy efficiency. MAC energy efficiency —the ratio of the useful communication energy expended to the total energy expended— depends on the ratio of goodput to the overall duty cycle. Now, if a MAC is able to deliver its traffic load reliably and goodput is equal to the throughput (i.e., duplicates packets are not received), then maximizing the energy efficiency essentially reduces to maximizing the throughput for a given duty cycle or, equivalently, minimizing the duty cycle for a given traffic load.

Performance. We evaluate four representative state-of-the art protocols: BoX-MAC [2], RI-MAC [1], SCP-MAC [4], and O-MAC/Crankshaft [6, 5]. Both analytically and experimentally, we find that for almost all traffic, *O-MAC achieves reliable delivery at the lowest duty cycle and has the highest capacity/energy efficiency.* RI-MAC in turn is better than BoX-MAC up to a rather high traffic load, beyond which BoX-MAC is better.

At ultra-low duty cycles, the capacity/energy efficiency gap is the highest. As the duty cycle increases the gap shrinks. The energy efficiency gap ranges from 13.8x to 2.2x between O-MAC and RI-MAC, and 16.9x to 2.1x between O-MAC and BoX-MAC in our experiments; capacity gaps are comparable.

We note that as the duty cycle becomes 100%, these protocols degenerate into unslotted CSMA protocols (with the exception of SCP-MAC which degenerates to slotted CSMA). The performance gap at 100% duty cycle is dominated by the diverse overheads in their respective protocols. As expected, the most minimal and simplest of the protocols, BoX-MAC, has least overhead, with O-MAC being a fairly close second and RI-MAC with substantially higher overhead. Keep in mind though that none of these protocols are intended for use at 100% duty cycle from the perspective of energy efficiency, since the throughput capacity of CSMA is known to be low [7].

In terms of latency, we find that: (a) the best latency for each protocol (regardless of traffic load) is achieved at 100% duty cycle; of course, if MACs choose to minimize latency this way, they can suffer a substantial loss of energy efficiency; (b) given a traffic load, the protocol which is least efficient has the lowest latency; informally, this follows from the fact that the least efficient protocol wakes up the most often to communicate; and (c) interestingly, *if a more efficient protocol is operated at a higher-than-need-be duty cycle so as to match the lowest latency achieved (by the least efficient protocol), its resulting energy efficiency is still better than that of the least efficient protocol.* In other words, the protocol to use for improving latency is the same protocol to use for optimizing energy efficiency.

Outline of the paper. Section 2 briefly recalls the design of four state-of-the-art MAC protocols in terms of design factors that primarily impact performance. Section 3 presents our analysis of MAC capacity limits as a function of duty cycle as well as the (significantly lower) limits achieved by the four protocols. Section 4 experimentally corroborates the analytic results for three of these protocols, and also compares their energy efficiency and latency performance. Section 5 qualitatively analyzes the impact of protocol mechanisms on the experimental performance. Section 6 discusses our future work and makes concluding remarks.

2. MAC FACTORS AND EXEMPLARS

In this section, we identify key protocol design factors that materially impact MAC throughput capacity, energy efficiency and latency, and distinguish how state-of-the-art protocols differ in terms of these factors.

Coordination centricity. The responsibility of coordination is centered on the sender or the receiver. In sender centric protocols, the sender chooses the moment of transmission, whereas in receiver centric protocols, the receiver

Protocol	Synch.(S) / Async.(A)	Sender(S)/ Receiver(R) Centric	Data Probes(D) Preambles(P)
X-MAC	A	S	P
BoX-MAC	A	S	D
RI-MAC	A	R	P
SCP-MAC	S-Regional	S	P
O-MAC	S-Local	R	–

Table 2: Factor selection in MAC protocols

chooses the moment of transmission. The moments may be chosen in an independent or coordinated fashion among nodes.

Level of synchrony. The more minimal MACs require no time synchronization between sender and receiver. Unlike asynchronous MACs, some others require global (or at least regional) time synchronization whereby the coordinated nodes wakeup at the beginning of the same slot in each periodic *frame* of time. Yet others require local time synchronization between receivers.

Probing. Asynchrony implies that once a sender is ready to transmit it performs some activity (waiting or probing) until the receiver wakes up independently. Two forms of probing are popular: preambles, in which the probe has no data, and data probes, which contain the sender data payload and may or may not contain control information. Some synchronous MACs use probing as well to notify the synchronized receivers of incoming traffic.

Table 2 identifies the factor choices made in four representative protocols, whose core protocols we describe next.

X-MAC [8] is an extension of the canonical asynchronous B-MAC protocol [9] that adopts the Low Power Listening (LPL) mechanism, where receivers independently and periodically poll the channel for activity using low power. Each sender wakes up its receiver by sending it a preamble that is at least as long as the receiver's frame length. By embedding destination information in the sender preambles, X-MAC reduces overhearing energy loss by allowing non-intended receivers to return to sleep earlier. Also, by inserting short gaps in between preambles, the sender avoids sending a continuous preamble and initiates data transmission upon receiving an acknowledgement from the intended receiver after some preamble.

BoX-MAC [2] further refines X-MAC by sending data packets instead of preambles repeatedly, thus eliminating the sender's energy cost in sending the data packet after the preamble. At the receiver, BoX-MAC conserves energy by adopting an LPL-like channel activity detection mechanism instead of the more costly preamble detection, which consumes an order of magnitude more energy than the former.

RI-MAC [1] is receiver-centric: in contrast to the previous two protocols, the rendezvous is initiated by the receiver. Receivers periodically broadcast a preamble to their neighbors. Receivers choose their periods independently and their wakeups are thus likely to not overlap, thereby attempting to avoid contention among senders for different receivers. Once senders have data to send they keep their radio active in receive mode and contend for the channel upon receiving a preamble from their intended receiver.

SCP-MAC [4] is an extension of the canonical regionally synchronous S-MAC protocol, wherein all nodes in a region wake up simultaneously in each frame. Since receivers poll the channel for activity in an aligned fashion, sender preambles become short *wakeup tones* that are sent just before receiver polls. Such synchronized polling not only reduces the energy cost in sending the long preambles, but also improves channel utilization during the wakeup slots.

O-MAC [6] is receiver-centric and locally synchronous: receivers use pseudo-random wakeup schedules to avoid waking up simultaneously and communicate the seed for their schedule to neighboring senders. Senders with pending data thus wakeup just before their intended receiver does. The sharing of the seeds makes the use of probes unnecessary.

We note that the selection of mechanisms for carrier sensing, for packet dwelling (where a sender that wins channel contention is allowed to send queued packets back-to-back), time synchronization, neighborhood discovery, and frame length (and therefore duty cycle) adaption all have secondary impact on performance. Section 5 addresses aspects of this impact.

3. ANALYSIS OF DUTY-CYCLED MAC PERFORMANCE

In this section, we analyze how throughput capacity and energy efficiency scale in duty cycled networks, by extending the Gupta-Kumar optimal scheduler for general random networks to accommodate duty cycling. We then focus on the case of MAC where traffic flows are only of 1-hop (in both general networks and clique networks). Finally, as a basis for corroborating our results in a controllable setting, we analyze and numerically compare the capacity and energy efficiency of the four representative MACs.

Recall the well known Gupta-Kumar [10] result that the per node throughput capacity, λ, of random wireless networks has an asymptotic tight bound of $\lambda = \Theta(\frac{W}{\sqrt{n \log n}})$ in the Protocol Model. At first glance, one may hypothesize that the capacity of a network whose duty cycle fraction is ψ, where $\psi \in [0, 1]$, would simply be $\lambda = \Theta(\frac{W}{\sqrt{n \log n}} \cdot \psi)$. We show however that the network can actually achieve a substantially higher capacity than the ψ proportion via careful scheduling of the wakeup and communication times of nodes.

3.1 System Model

Consider a random network where n nodes are uniformly and independently placed in a unit square. Each node, X_i, $i \in 1, ..., n$, sends data to a random destination node. The transmission range r_n, where $r_n = \Omega(\sqrt{\frac{\log n}{n\pi}})$, and the traffic are the same for all nodes. Each node has a maximum bandwidth of W bps and wakes up on average to communicate for ψ fraction of time.

We adopt the *Protocol Model* [10] which postulates a geometric condition for successful transmission: a transmission from X_i to X_j is successful if for all other nodes X_k, $k \neq i, j$, that are concurrently transmitting over the same channel,

Symbol	Meaning	Value
W	maximum transmission rate	250 Kbps
q	uniform probability of backoff	1/16
u	beacon length relative to data length	0.4
E_{bit}	energy for sending a bit	0.217 μW
E_{radio}	energy consumed by radio per second	54.3 mW
n	total # of nodes in network	-
λ	throughput capacity	-
ψ	radio duty cycle	-
r_n	communication range	-
Δ	guard zone factor for interference-free communication	-

Table 3: Model parameters

the following inequality holds:

$$|X_k - X_j| \geq (1 + \Delta)|X_i - X_j|. \quad (1)$$

The circle of radius of $(1 + \Delta)|X_i - X_j|$, $\Delta > 0$, centered at a receiver delineates the guard zone within which there is no destructive interference. Table 3 summarizes the notation alongwith representative values for subsequent numerical comparisons.

3.2 Capacity of the Optimal Scheduler

Our main capacity result for multi-hop traffics in general duty-cycled networks is:

THEOREM 1. *The throughput capacity of a duty-cycled random wireless network in the Protocol Model is*

$$\lambda = \Theta(\frac{W}{\Delta}\sqrt{\frac{\psi}{n}}) \quad (2)$$

until it reaches its network capacity of $\Theta(\frac{W}{\Delta^2\sqrt{n \log n}})$ *when* $\psi \geq \frac{32}{\Delta^2 \log n}$.

Two comments are in order about the result. For one, below the limiting duty cycle, the capacity scales better than linearly in ψ and better than the inverse of $\sqrt{n \log n}$. For two, the factor $\sqrt{n}$ reflects the average number of hops between source and destination, and $\sqrt{\psi}$ reflects the cumulative duty cycle to forward the source traffic along the route.

Proof. We only sketch the proof for reasons of space; a complete proof is in [11]. Let Z be the expected distance between a node and its destination. On average, each source is then $\frac{Z}{r_n}$ hops away from the destination, and the bit rate the network needs so as to accommodate its traffic is at least $n\lambda\frac{Z}{r_n}$, where $0 < Z < 1$. In an optimal schedule, at most $n/2$ senders can be concurrently active in any given slot. Suppose each sender on average wakes up for t out of T slots, the maximum number of potential transmissions in the network is $\frac{nt}{2}$. On the other hand, the maximum number of simultaneous transmissions the network can support is no more than $\frac{16}{\Delta^2 r_n^2 \pi}$[11]. Therefore, the number of achievable transmissions during a period of T is $\frac{16}{\Delta^2 r_n^2 \pi}T$. As long as the number of potential transmissions does not exceed the network capacity, i.e.,

$$\frac{nt}{2} \leq \frac{16T}{\Delta^2 r_n^2 \pi} \Rightarrow \frac{n}{2}\psi \leq \frac{16}{\Delta^2 r_n^2 \pi} \Rightarrow \psi \leq \frac{32}{n\Delta^2 r_n^2 \pi} \leq \frac{32}{\Delta^2 \log n}, \quad (3)$$

an optimal scheduler can accommodate the traffic, which leads to Eq. (4):

$$n\lambda\frac{Z}{r_n} \leq W \cdot \min\{\frac{n\psi}{2}, \frac{16}{\Delta^2 r_n^2 \pi}\}. \quad (4)$$

Consider the inequality in (4) with respect to the second term in the *min* function. Since r_n is asymptotically larger than $\sqrt{\frac{\log n}{n\pi}}$, based on Eq. 4 we can derive:

$$\lambda \leq \frac{16W}{n\pi\Delta^2 r_n Z} \leq \frac{c_1 W}{\Delta^2\sqrt{n \log n}}. \quad (5)$$

Now, the first term in the *min* function represents throughput before the maximum capacity shown in Eq. (5) is reached. By plugging in the constraint $r_n^2 \leq \frac{32}{\Delta^2 n\pi\psi}$ derived from Eq. (3), we complete the proof for the theorem with:

$$\lambda \leq \frac{W\psi r_n}{2Z} \leq \frac{c_2 W}{\Delta}\sqrt{\frac{\psi}{n}}. \quad (6)$$

Results for 1-hop traffic and clique networks. Since we are interested in MAC performance, we now restrict our analysis for the case where the communication is for senders and receivers that are one hop apart. We derive capacity of 1-hop traffic for both general and clique network.

For 1-hop traffic, Z/r_n has a constant value of 1, which leads to the following corollary on the throughput of 1-hop traffic in general networks:

COROLLARY 1. *The throughput capacity of a duty-cycled random wireless network with 1-hop traffic is*

$$\lambda = \Theta(W\psi) \quad (7)$$

until it reaches its network capacity of $\Theta(\frac{W}{\Delta^2 \log n})$ *when* $\psi \geq \frac{32}{\Delta^2 \log n}$.

For the special case of clique networks, the maximum number of simultaneous transmissions reduces from $\frac{16}{\Delta^2 r_n^2 \pi}$ to 1, which leads to our next corollary:

COROLLARY 2. *The throughput capacity of a duty-cycled clique network with 1-hop traffic is*

$$\lambda = \Theta(W\psi) \quad (8)$$

until it reaches its network capacity of $\Theta(\frac{W}{n})$ *when* $\psi \geq \frac{2}{n}$.

Notably, below the limiting duty cycle, nodes in clique networks achieve a throughput that is independent of n.

3.3 Capacity of Extant MACs

So that we can corroborate our analysis with properly controlled experimentation, we analyze the maximum throughput of extant MACs for the case of clique networks. In contrast to an optimal scheduler, which guarantees that only one node transmits in an interference region, the representative MACs only ensure that the probability of a successful communication for any node during a slot is τ, where

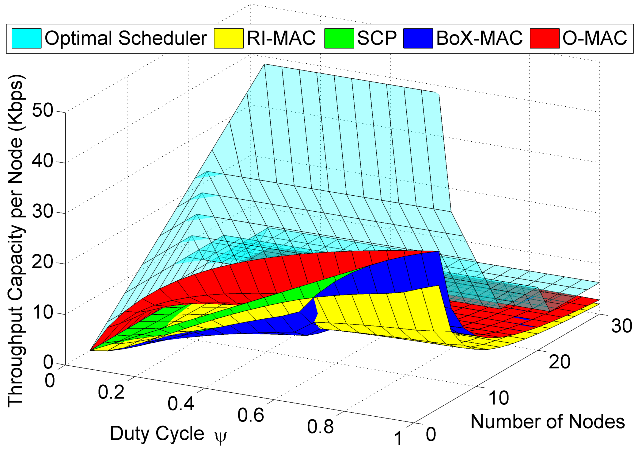

(a)

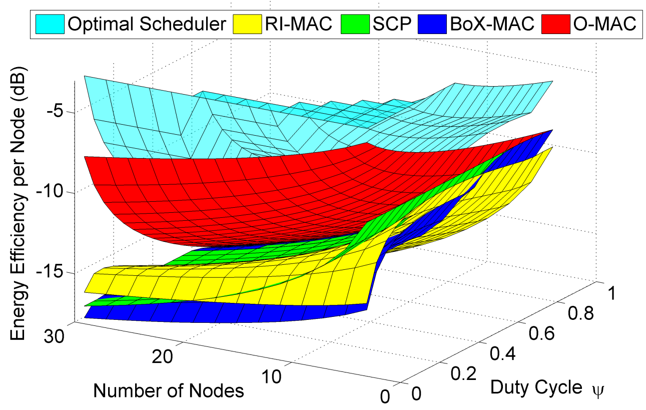

(b)

Figure 1: Comparison of throughput capacity and energy efficiency in clique networks

$\tau \in [0, 1]$. Hence, the expected number of successful transmissions for $n/2$ senders is $n\tau/2$. Accordingly, the capacity of these MACs is

$$\lambda \leq W \cdot \min\{\frac{\tau}{2}, \frac{1}{n}\}. \quad (9)$$

We now present a framework for calculating τ for CSMA-based MAC protocols, which subsumes the representative protocols. In CSMA, typically, when a node attempts to transmit a packet, it first randomly selects one out of C_s contention slots and monitors the channel until that slot to ensure that no other transmission is occurring within its communication range. To avoid overloading the word slot, we henceforth refer to a contention slot as a timeslice. If any transmission is detected before the chosen timeslice, the sender withdraws its transmission attempt; otherwise, it immediately starts the data transmission after the timeslice. Let the probability of selecting any timeslice be q and $\hat{\epsilon}$ be the expected number of contenders in each node's communication range. The probability for a node to successfully

MAC	Duty Cycle Constraint	$\hat{\eta}$ $(= \hat{\epsilon})$
SCP-MAC	$p_t\psi_r + \psi_r = 2\psi$	$p_t\eta$
O-MAC	$p_t\psi_r + \psi_r = 2\psi$	$p_t\psi_r\eta$
BoX-MAC	$(\frac{1}{2\psi_r} + 1) \cdot p_t\psi_r + \psi_r = 2\psi$	$(\frac{p_t}{2} + p_t\psi_r)\eta$
RI-MAC	$(\frac{1}{2\psi_r} + 1) \cdot p_t\psi_r + \psi_r = 2\psi$	$(p_t + u)\psi_r\eta$

Table 4: Capacity framework parameter for protocols

access the channel, denoted as p_a, is thus

$$\begin{aligned} p_a &= q + q(1-q)^{\hat{\epsilon}-1} + q(1-2q)^{\hat{\epsilon}-1} + ... + q \cdot q^{\hat{\epsilon}-1}, \\ &= q \sum_{i=0}^{1/q-1} (1 - iq)^{\hat{\epsilon}-1}. \end{aligned} \quad (10)$$

Let the expected number of contenders in the interference range be $\hat{\eta}$. Of course, the transmission is guaranteed to succeed when there is no other transmission within the interference range of the receiver. However, in general the probability of successful transmission in any given slot when data is available is equal to $p_a(1-p_a)^{\hat{\eta}-1}$. Thus, the total probability of successful transmission in any slot is

$$\tau = p_d \cdot p_a(1 - p_a)^{\hat{\eta}-1}, \quad (11)$$

where p_d indicates the probability of transmitting data.

Eq. (11) serves as the parameterized framework for analyzing the four representative MAC schedulers. We assume that q follows the same probability distribution for all MAC protocols. In a clique, the total number of contenders within a communication range, denoted by η, is equal to that within the interference range. The total duty cycle of any sender-receiver pair is 2ψ, out of which a node spends ψ_r, $0 \leq \psi_r \leq 1$, in receiving mode to account for the time over which a node wakes up, polls the channel, possibly receives a packet, and goes to sleep. A corresponding sender may choose to send data with probability p_t once the receiver is known to be awake, which leads to the equation $p_d = p_t \cdot \psi_r$. We use the representative values for the constant parameters of Table. 3 and the key constraints subject to which we optimize the capacity for each MAC in Table 4.

Fig. 1(a) shows the MATLAB simulation throughput capacity results at different duty cycles for a network size ranging from 4 to 30 nodes. We observe that of the four protocols, O-MAC approximates the optimal scheduler best, although the performance gap decreases at high duty cycles. At low density, SCP-MAC outperforms RI-MAC as inter-receiver contention —contention caused by traffic destined to different receivers— is low and the synchrony in SCP-MAC substantially reduces the overhead in probe detection. As inter-receiver contention increases with density, RI-MAC takes over in performance. At full duty cycle, all MAC protocols converge to a pure CSMA scheme except for RI-MAC whose use of probes becomes a major constraint.

3.4 Energy Efficiency of Extant MACs

Provided that a MAC can schedule all source traffic within its throughput capacity, its energy efficiency, denoted by e, is the following:

$$e = \frac{\lambda \cdot t \cdot E_{bit}}{\psi \cdot t \cdot E_{radio}}, \quad (12)$$

where E_{bit} is the energy cost of sending one data bit, E_{radio} is the energy consumption rate for active radio, and t is the period of time considered.

Fig. 1(b) shows the MATLAB simulation results on energy efficiency of different MACs with the same configurations as in Fig. 1(a). The energy efficiency of all protocols is higher at low network density. As duty cycle increases, the efficiency of receiver-centric protocols decreases while the efficiency of BoX-MAC increases. Among the representative MACs, O-MAC remains the most energy efficient protocol under all configurations, with a maximum gap of $10dB$ over BoX-MAC and SCP-MAC, and a gap over RI-MAC ranging from $3dB$ to $8dB$.

4. EXPERIMENTAL PERFORMANCE COMPARISON

In this section, we experimentally corroborate our theoretical analysis of performance as well as study the tradeoff between energy efficiency and latency. Our evaluation spans O-MAC, RI-MAC and Box-MAC[4] in a wide range of network configurations. We seek to interpret the results in term of the primary coordination design factors: synchrony, centricity, and probing.

Our findings are based on two sets of experiments. In the first set, which is based on clique networks, we compare the three protocols in terms of energy efficiency, delivery ratio, and delay over a wide range of traffic loads, frame lengths (wake-up intervals) in different network densities. We begin by fixing the network density (and hence interference) and study the impact of traffic. Then, we fix the traffic and study the impact of network density. Next, we experiment with different traffic loads and densities to study their combined effect. To verify that our results are robust with respect to different environments, we repeat some experiments in a qualitatively different testbed setting than our office testbed.

The second set of comparative experiments is based on multi-hop networks, which accommodate consideration of hidden terminals. We repeat a number of experiments conducted in the first set to compare the results.

Energy measurement. To monitor the energy cost of sensor nodes, we implemented a software module in the TinyOS CC2420 driver that measures the cumulative duty cycle and is used by all three MAC protocols. The radio is considered active upon detecting the oscillator stabilization signal and inactive upon detecting the radio stop signal. The small power difference between the transmission mode (52.2 mW) and reception mode (56.4 mW) is ignored in the computation of energy efficiency; the average, 54.3 mW, is used as the power of an active radio. The useful energy cost is computed as the product of the energy cost per packet, denoted as E_{packet}, and the total number of unique packets transmitted, where E_{packet} is computed as:

$$E_{packet} = 40 \times 32 \times 54.3 \approx 69 \ \mu W \ , \qquad (13)$$

as each packet is 40 bytes long, including the 802.15.4 header, and each byte takes approximately 32 μs to transmit. The total energy cost is computed by the product of the average radio power, the duty cycle, and the experiment duration. Sequence numbers are attached to data packets to filter out duplicate packets. The measurement of the total energy cost begins when all flows have started transmitting data[5].

Experiment duration	15 minutes
Length of data packet	40 bytes
Carrier sensing window	0-2 ms
Number of nodes	twice the number of flows
Maximum retransmission	5
Transmission buffer	32 packets
Sender frame length	100 seconds

Table 5: Values of parameters shared by all protocols in experiments

Implementation and configuration. Our experiments evaluated existing implementations of O-MAC and Box-MAC[6] in TinyOS-2.1 and of RI-MAC in TinyOS-2.0.2. All experiments were performed on the TelosB platform. We configured each MAC as follows: packets were buffered in a FIFO queue of length 32 and retransmission was enabled for up to 5 times per packet. To decouple the sender duty cycle from the receiver's and since traffic is one-way, sender frame length was set to a negligible constant, namely, 100 seconds. Table 5 summarizes the shared parameters we used in this experiment.

For BoX-MAC, we reduced the default receiver dwell time from 100 ms to a less conservative 20 ms, which is about twice the time to account for carrier sensing, radio buffer loading, SIFS, and acknowledgement; this change was necessary as we found that with the original dwell time the receiver ran at 100% duty cycle whenever the data interval was less than 100 ms.

For O-MAC, we let the implementation use a hold time to deal with lack of alignment of receiver and sender wake-ups, use Disco [12] for neighbor discovery, and use the clock skew estimation in FTSP [13] for local time synchronization. Time synchronization information in O-MAC is piggybacked on the acknowledgements once the initial neighborhood discovery is finished.

4.1 Capacity and Energy Efficiency

Impact of traffic on a single link. We begin with experiments that measure the maximum throughput of different MAC protocols over a wide range of duty cycles. We use only a single sender and a single receiver to separate the concerns of contention handling and overhearing loss from achievable throughput. Since the total duty cycle is a derivative (and not fully controllable) metric, we choose to search for the lowest duty cycle by experimenting with 12 traffic loads ranging from 1 packet every 60 second to 125 packets per second.

[4]SCP evaluation is omitted given the results in [3] and in Section 3.

[5]O-MAC requires an initial neighbor discovery period to acquire its initial time synchronization. In a static network, the overhead of this period would be amortized over time and thus negligible. So, we exclude this period in calculating the O-MAC energy cost.

[6]The BoX-MAC version we tested is included as the default

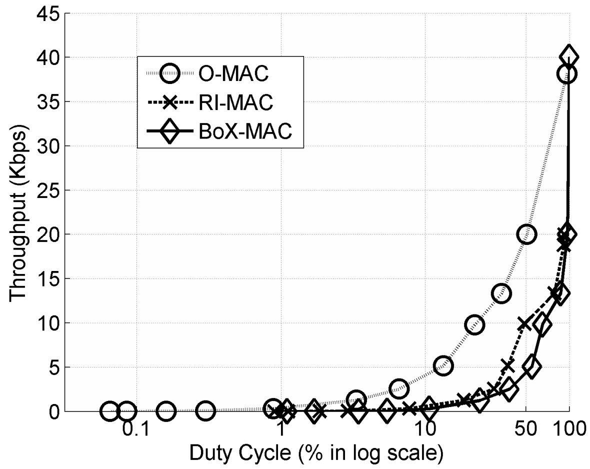

Figure 2: Comparison of capacity at different duty cycles.

Since we seek to minimize the duty cycle at which each protocol accommodates a given traffic load, we optimize the protocols as follows. For asynchronous protocols that use LPL, we recall that for each periodic traffic load with a fixed number of senders, there is an optimal duration for receiver sleep-wakeup [4]. Specifically, the optimal frame length is $O(\sqrt{T_{data}})$ where T_{data} is the data interval. For BoX-MAC, we therefore search for the lowest duty cycle for a given traffic, by varying the frame length over a set of 6 different values whose average value is $\sqrt{T_{data}}/C$, where the constant C depends on the sensor platform. Empirically, we found that $C = 7$ was a best approximation for the optimal frame length in most experiments. For RI-MAC, analogous reasoning applies (with sending preambles and receiver polling inverted to receiver probes and sender polling). We empirically found that the constant C for RI-MAC was 10.

For O-MAC, since a sender's overhead is decoupled with its receiver's frame length, the optimal duty cycle is always achieved with the longest frame length. We conservatively chose the 6 frame length values with the first one being $T_{data}/2$ and each successive value being half that of its previous one. Note that we could further improve the energy efficiency of O-MAC by increasing the frame length until the time synchronization overhead became dominant. (A more detailed discussion of the time synchronization overhead for ultra-low traffic is in Section 5.3.)

Fig. 2 shows our maximum throughput results (note that the duty cycle axis is in log scale). The growth of maximum throughput with respect to duty cycle is almost linear in O-MAC, whereas that of the two asynchronous protocols is slow in the low duty cycles and continues to gain speed as duty cycle increases. At the lowest of the traffic we tested, the minimum duty cycle of asynchronous protocols was at least 0.9% whereas that of O-MAC was 0.065%, which translates to a gap in energy efficiency of 16.9x between O-MAC and BoX-MAC and a gap of 13.4x between O-MAC and RI-MAC. At higher traffic, the gap between O-MAC and the other two decreases because the overhead in their coordination strategies becomes less in proportion. The maximum throughput of RI-MAC is slightly better than BoX-MAC before the former reaches its capacity limit at 20 $Kbps$, i.e., 62 packets per second, after which the maximum throughput remain the same even if the sender and the receiver run at a higher duty cycle. The maximum throughput of BoX-MAC overtakes O-MAC at very high traffic load, i.e, at 125 packets per second, BoX-MAC has 100% delivery ratio but O-MAC does not.

MAC protocol for TinyOS-2.1

Summary: Synchrony essentially obviates the need for preamble transmission and preamble detection. At low traffic, synchronous MACs outperform asynchronous ones. As traffic increases, receiver-centric, asynchronous protocols start to underperform compared to their sender-centric counterparts, because of the additional overhead in our their use of data-free probes. At ultra-high traffic, the overhead for synchronous protocols is essentially wasted as synchronized slots are not really needed; in the extreme, BoX-MAC overtakes O-MAC. For the most part, however, O-MAC performs the best.

Impact of density. We now consider experiments that study the impact of contention between concurrent traffic flows. Each traffic flow consists of a unique source and a unique destination. The network consists of a clique of 20 TelosB motes in an indoor lab testbed. The traffic and frame length is kept fixed across these experiments: each source generates one packet per second, and the receiver frame length is set to 1 second for all MACs. The number of flows is varied; for each number of flows, we repeat the experiment twice with each experiment lasting for 15 minutes and record the average value. (To study the impact of probes, we also tested a variation of O-MAC, O-MAC with probes, in which a receiver broadcasts a probe at the beginning of each receiving slot and the sender transmits the data packet upon receiving the probe.)

All of the protocols deliver 100% of the generated packets for each flow configuration. As Fig. 3(a) shows, O-MAC achieves a sender duty cycle upper bound of 1.14% in all configurations, as compared to 52.4% and 54.5% in RI-MAC respectively. Fig. 3(b) shows the receiver duty cycle. Across all configurations, O-MAC achieves a receiver duty cycle upper bound of 0.71%, as compared to 1.96% in RI-MAC and 3.1% in BoX-MAC respectively. The 20 ms dwell time in BoX-MAC makes its receiver run at a slightly higher duty cycle than is needed in the receiver-centric protocols. Another reason for the lower receiver duty cycle for O-MAC and RI-MAC is their adaptive dwelling scheme, which we will discuss further in Section 5.1.

The latency results in Fig. 3(c) show different trends between the two coordination centricities. While the latency for O-MAC and RI-MAC remains flat as the number of flows increases, the latency in BoX-MAC increases. Finally, Fig. 3(d) shows the energy efficiency of the tested protocols. The energy efficiency of O-MAC is 15 dB (31x) better than BoX-MAC and RI-MAC.

One might argue that testing all MAC protocols with the same frame length is unfair to some; in this case, the 1-second frame length is longer than is optimal for both BoX-MAC and RI-MAC. One way to improve the efficiency of

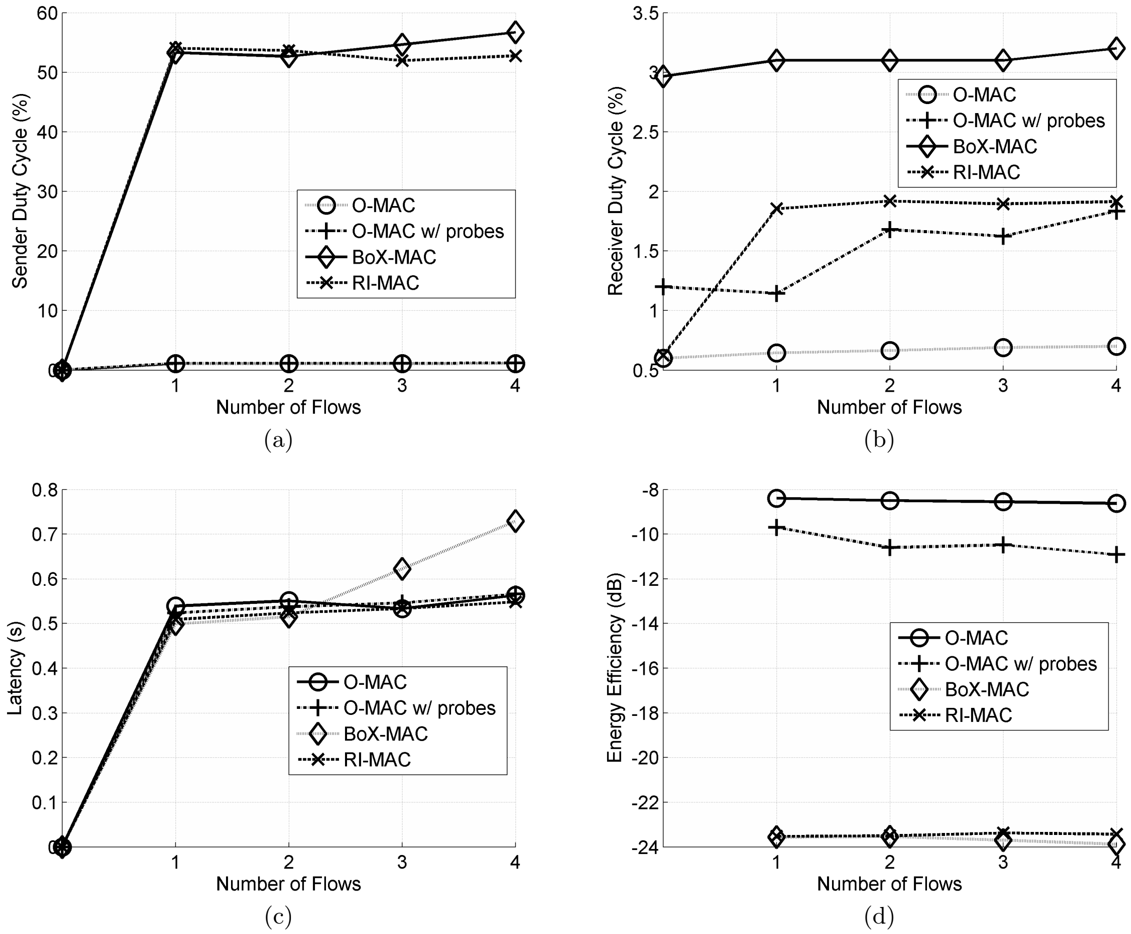

Figure 3: Performance of O-MAC, O-MAC w/ probes, RI-MAC, and BoX-MAC with contending flows at low traffic (1 packet/second per sender)

the asynchronous MAC protocols would be to adapt the receiver's frame length to the traffic rate, which will save energy in either preamble transmission or beacon detection. As we will see shortly, however, O-MAC still significantly outperforms the asynchronous protocols even when they are respectively operated at their most energy efficient point.

Summary: Receiver centric protocols are more efficient at managing inter-receiver contention than sender centric ones, provided the overall traffic is within their capacity limits. Receiver wakeups tend to not overlap in the former, hence contention between flows eschews inter-receiver contention and is only for intra-receiver flows. In contrast, the latter suffer from inter-receiver contention as the number of flows increases because of their long preambles and the resulting high channel occupancy. O-MAC consistently performs the best in energy efficiency despite increase in the contention level.

Impact of both traffic and density. We next consider experiments that let us study the impact of both traffic rate and density on delivery ratio, latency, and energy efficiency. For each MAC protocol, four numbers of flows (from 1 to 4), five traffic rates (1, 4, 8, 16, and 32 packets per second), and six different frame lengths (25, 50, 100, 200, 500, and 1000 *ms*) are considering, leading to a total of 120 experiment configurations per MAC. Each experiment is conducted once for a duration of 15 minutes.

For a given number of flows and traffic rate, we select the best energy efficiency by searching over the different frame lengths. (This redresses the previously discussed shortcoming of the density experiments.) In all the configurations, both BoX-MAC and RI-MAC never achieve their optimal energy efficiency at a frame length on the boundary, i.e., 25 or 1000 *ms*, verifying that the interval of frame lengths we have considered is sufficient for searching for the optimal efficiency.

In terms of delivery ratio, the outermost contour line in Fig. 4(a) shows that RI-MAC reaches its capacity limit whenever the product of number of flows and the traffic rate per flow increases above 40 packets per second. The advantage of using of probes in RI-MAC thus turns out to be limited. Both O-MAC and BoX-MAC achieve 100% duty cycle in almost all test cases: with 4 flows and 32 packets per second, the delivery ratio of O-MAC and BoX-MAC respectively reduce to 84% and 95%, and with 1 flow and 125 packets a

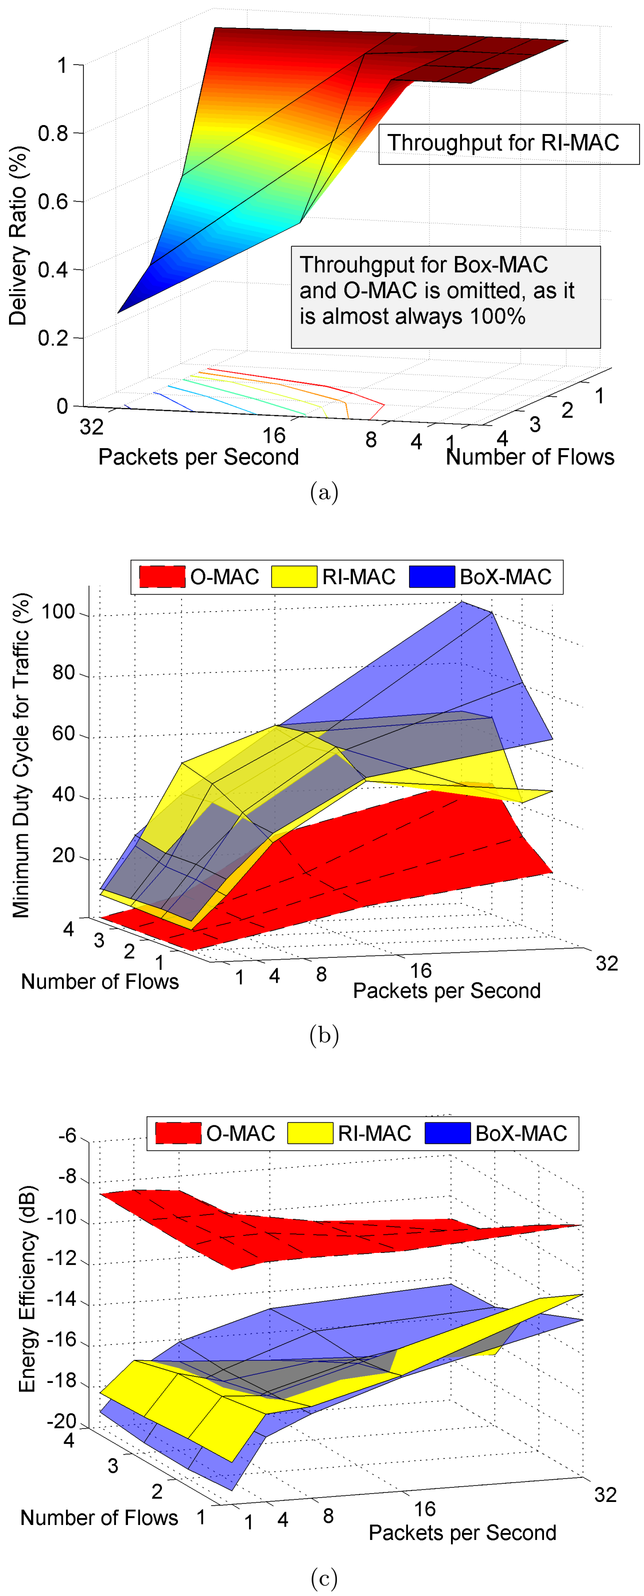

Figure 4: Performance of O-MAC, RI-MAC, and BoX-MAC at diverse flow and traffic configurations

second, the delivery ration of O-MAC is 95%.

Fig. 4(b) and Fig. 4(c) show the lowest duty cycles and the best energy efficiencies of the three MAC protocols in the 360 experiments. When the delivery ratio is 100%, energy efficiency grows linearly with the inverse of duty cycle, as predicted by our analysis. In general, the energy efficiency of all protocols is higher when the number of flows (and thus contention) is lower. In terms of energy efficiency, for OMAC, the gap over BoX-MAC ranges from 3 *dB* to 13 *dB*, i.e., from 2x to 20x, and over RI-MAC ranges from 3 *dB* to 10 *dB*, i.e., from 2x to 9x.

Summary: Receiver-centricity combined with synchrony best deals with maximizing throughput and minimizing contention. Across almost all tested configurations, O-MAC is most energy efficient. In terms of trends, we see: (i) As the product of traffic load and flow count decreases, the receiver-centric, synchronous protocol is least affected and the sender-centric, asynchronous protocol is most affected. (The latter suffers from its use of preambles in the presence of increased inter-receiver contention.) (ii) As the product increases, the receiver-centric, asynchronous protocol is most affected and the sender-centric, asynchronous protocol the least affected. (Growth in inter-receiver contention is less effectively handled by receiver-centric protocols and they behave more like asynchronous protocols; as compared to BoX-MAC, the use of preambles as opposed to data probes hurts RI-MAC first and the use of time sync hurts O-MAC at ultra-high traffic. BoX-MAC senders makes more effective use of the channel by immediately sending arriving packets in the form of data probes.)

Impact of environment. For the purpose of establishing the robustness of our results to environment, we next repeat the experiment of Fig. 3 in a publicly available testbed [14] that has about 400 TelosB motes deployed and possibly concurrent experiments running in different radio frequencies. The results are essentially identical. The largest difference in duty cycle is 8% when the sender duty cycle of BoX-MAC increases from 54.48% to 58.9% in the 4-flow case, in which case its latency also increased by 96 *ms* from 694 *ms* to 790 *ms*, which is probably due to sporadic interfering wireless activities.

4.2 Latency and Energy Efficiency

It has been argued in the literature that synchronous MACs tend to achieve higher efficiency at the cost of latency [3, 5], in particular when the traffic is low. In this section, we show that this observation in insufficient in that (receiver-centric) synchronous protocols can be better in both energy efficiency and competitive in latency.

Based on the results in Section 4.1, we study whether the latency drawback is due to synchrony by examining data points where O-MAC runs at a duty cycle higher than necessary and achieves a latency comparable to the lowest latency achieved by the asynchronous protocols when they are optimized for energy efficiency. To separate the concerns of synchrony from the effects of congestion backoff, dwell time, and carrier sensing, for each MAC protocol we select 70 experiment configurations where the per flow traffic rate is at most 8 packets per second (at higher traffic loads, these

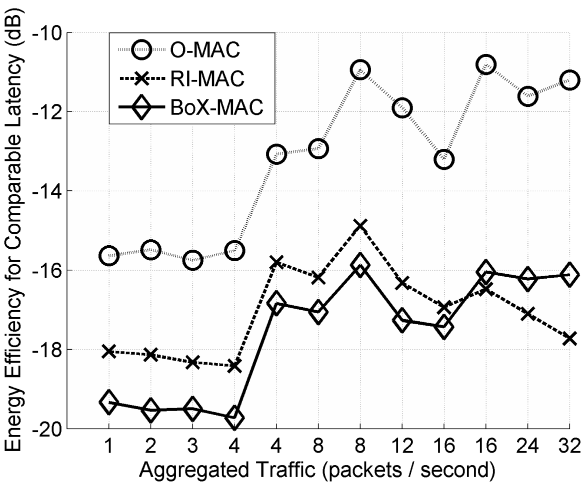

Figure 5: Comparison of energy efficiency at T^*, where T^* is the best latency that would be achieved among the protocols if they were optimized for energy efficiency.

mechanisms have nontrivial impact). We consider two latencies to be comparable if they are within 20% of each other.

Fig. 5 shows our results. Each data point is obtained as follows: for a given aggregated traffic rate, i.e., the product of the traffic rate per flow and the number of flows, we find the minimal latency achieved by the three protocols when their energy efficiency is optimized[7]. Then, from the data points where O-MAC runs at different frame lengths for that aggregated traffic, we select the data point where O-MAC achieves comparable latency. (Incidentally, in 6 out the 12 selected results, the latency achieved by O-MAC is lower than the minimal latency of the other two.) As we can see, the energy efficiency of O-MAC is at least $2.5x(4dB)$ better than BoX-MAC and RI-MAC even when O-MAC sacrifices efficiency to make its latency comparable.

Summary: Even though receiver-centric, synchronous protocols have the best energy efficiency when MACs are optimized for efficiency, asynchronous protocols have best latency at these operating points. This essentially follows from the difference in frame lengths between asynchronous and synchronous protocols. The former achieve optimal energy efficiency at a frame length of $O(\sqrt{T_{data}})$ [4], while the latter achieve optimal energy efficiency at a frame length lower bounded by $O(T_{data})$ (because otherwise the receiver waste energy in idle listening).

Nevertheless, we see that O-MAC can trade its energy efficiency for latency, achieving a latency comparable to that of asynchronous protocols while still outperforming them in efficiency.

4.3 Comparison in Networks with Hidden Terminals

We conclude with experiments to verify whether the comparative performance of the MACs shows the same trends in a multi-hop network, in which sedners can be hidden terminals with respect to one another. In terms of experiment setup, we found it nontrivial to create, maintain, and measure hidden terminals in networks while still ensuring reliable channels between the respective senders and receivers. So we emulated the desired topology by forcing the radio to always return a clear channel after carrier sensing in the context of hidden terminals. Also, transmissions were issued with the STXON strobe instead of the STXONCCA strobe in the CC2420 radio so that no CCA sampling was done prior to the transmission. The number of flows ranged from 1 to 4 and each source generated 8 packets per second to ensure a non-negligible collision probability.

[7]Identical aggregated traffic rates occur when we have configurations such as two flows with 4 packets per second versus one flow with 8 packets per second.

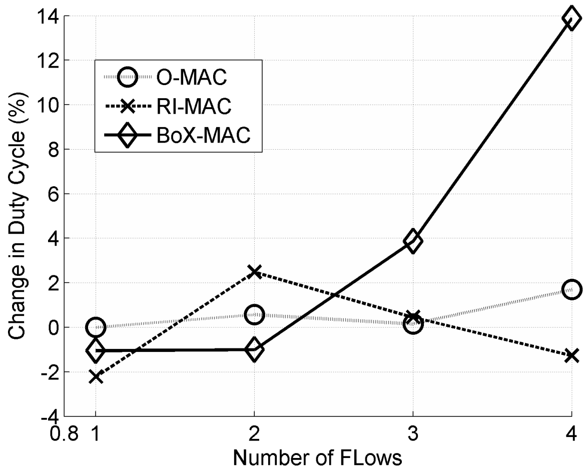

Figure 6: Comparison of percentage change in duty cycle when senders become hidden terminals with respect to one another.

Fig. 6 shows the percentage change in duty cycle when senders became hidden terminals to one another. The change in percentage in O-MAC and RI-MAC remains negligible in all cases. The duty cycle of BoX-MAC start to be the same as the clique network results and then increases by as high as 14% as the number of flows reaches 4.

Summary: With hidden terminal senders, avoiding interfering concurrent traffic becomes critical for energy efficiency. Receiver centricity avoids concurrent transmissions across different flows, which substantially reduces the number of collisions. In contrast, sender centricity along with asynchrony does not explicitly randomize the transmission of different flows and requires that senders transmit long preambles. As a result, it suffers from the increased probability of collisions when the number of concurrent flows is large. We expect the efficiency for sender centric, synchronous protocols, such as SCP-MAC, will also suffer in this case since they explicitly increase the number of contending senders by synchronizing the channel polling time of receivers.

5. IMPACT OF AUXILIARY MECHANISMS

In this section, we qualitatively study the impact of mechanisms that complement that primary coordination strategy factors. We begin with addressing duty cycle adaption. Next, we address the sorts of impact that carrier sensing has on different coordination strategies. Lastly, we address the issue of time synchronization, where needed, and compare our findings with previous results.

5.1 Duty Cycle Adaption

Nodes adjust duty cycle by changing either the wakeup frame length or the *dwell time*, i.e., the minimum radio-on time before a receiver concludes that no sender is currently transmitting to it. Adapting the former has a larger impact on MAC protocol performance than the latter. Moreover, adapting the former is related to carrier sensing, which we will discuss in the Section 5.2.

Frame length selection. As shown in Section 4.1, energy efficiency is sensitive to the selection of frame length. Note that between Fig. 3(a)-(b) and Fig. 4(b) we see that the duty cycle of BoX-MAC and RI-MAC respectively improve by an order of magnitude for each traffic and flow configuration. The receiver incurs idle listening overhead when the frame length is shorter than need be, and the sender wastes energy on increased contention and, where applicable, on probes when the frame length is longer than need be. In either case, energy efficiency decreases.

Fig. 7 plots the experimental results for O-MAC when there is one traffic flow with traffic loads of 8 and 16 packets per second. In contrast to BoX-MAC and RI-MAC, for O-MAC the energy efficiency still increases after its frame length exceeds the data interval. However, this increase is minimal compared to the negative impact on latency. In other words, *although one can further trade latency for improved efficiency in O-MAC, operating O-MAC at a frame length close to the data interval strikes a good balance between energy efficiency and latency.*

Dwelling. All three protocols use dwell time to allow senders to transmit queued packets. In both BoX-MAC and RI-MAC, as soon as a sender finishes transmitting to a receiver, all other pending senders for that receiver can immediately contend to transmit their queued packets. Pending senders wait for a congestion backoff period before contending to transmit again, and the dwell time length is chosen accordingly. Of the two, the receiver centric RI-MAC performs more efficient dwelling based on its better control of the contention period. We discuss this further in Section 5.2.

O-MAC achieves even more efficient dwelling by limiting the dwell time to the winning sender, which allows senders that lose contention to avoid potentially long congestion backoffs. O-MAC also embeds a flag in data packets that signals the termination of dwell time. Before a sender transmits the last packet in its queue, it sets this flag in the outgoing packet, which allows its receiver to go back to sleep immediately after processing that packet.

5.2 Carrier Sensing

For sender-centric protocols, carrier sensing is always required as (i) long preambles in asynchronous protocols substantially increase contention, and (ii) synchronized channel polling essentially batches all senders in a synchronized region to contend simultaneously.

In contrast, receiver-centric protocols try to approximate time-division scheduling of receivers and so eschew carrier

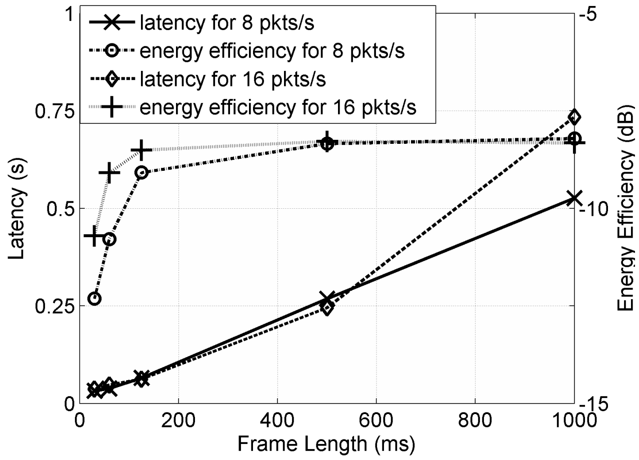

Figure 7: Impact of frame length on O-MAC

sensing for dealing with contention among senders to different receivers. In other words, carrier sensing is only for contention among intra-receiver senders and not inter-receiver senders. Since receivers can locally estimate their number of senders, they can adaptively choose the contention window size to improve efficiency.

A tradeoff exists between carrier sensing duration and frame length. By reducing the frame length, a receiver can reduce the average number of contending senders every time it wakes up. However, this tradeoff favors receiver-centric protocols as the duty cycle of each receiver needs to account only for the traffic that is intended for it. On the other hand, for sender-centric protocols, receivers need to account for the aggregated traffic in the neighborhood.

5.3 Time Synchronization

Synchronous protocols exchange time information in every *synchronization period*. Ye et al have shown that there exists an optimal synchronization period, along with a corresponding guard time for transmissions to compensate the synchronization error [4]. Dutta et al have shown that the constraint in periodic synchronization makes the duty cycle lower bound of any synchronous protocol an order of magnitude higher than the theoretical lower bound to handle the given amount of traffic [15].

However, both these analyses uses clock skews as the basis for computing the synchronization period and guard time. The FTSP protocol compensates for the relative clock skews between nodes [13]: in addition to computing the clock offsets between two nodes, each node estimates the clock rate of a remote node via linear regression. Such compensation drastically reduces the error in local estimation of the global time even when the synchronization period is long. It has been shown that the synchronization error is less than $40\mu s$ on Mica2 platform even when the synchronization period is 30 minutes, which translates to a frequency skew of $11 \times 10^{-3} ppm$ based on the well-known equation:

$$t_{error} = 2 \cdot r_{skew} \cdot T_{ts}$$

, where t_{error} is the synchronization error, r_{skew} the frequency skew, and T_{ts} the time synchronization period. This overhead is negligible even at ultra low duty cycles and if time sync information is not piggybacked onto data packets, keeping in mind its accompanying efficiency benefits and especially when only local time synchronization is needed.

6. CONCLUSIONS

As traffic decreases (and therefore the duty cycle decreases), MAC contention becomes less of an issue whereas avoiding overhearing and idle listening become more prominent. Our evaluation confirms that time division approaches, which are realized in a light-weight manner by receiver-centric protocols such as O-MAC and RI-MAC, are best suited for avoiding overhearing in this case. Of these two protocols, we have found that O-MAC is more energy efficient due to its use of local time synchronization between sender and receiver — which can be achieved with low overhead even at very low duty cycles without requiring global time synchronization— that reduces the sender wakeup time.

As traffic increases, MAC contention increasingly becomes more of an issue compared to overhearing and idle listening. Using probes that contain data (as in BoX-MAC) versus without (as in RI-MAC) is more efficient, so the capacity of BoX-MAC grows faster than RI-MAC and eventually overtakes it.

As the number of flows increases, the resulting mutual interference decreases the throughput capacity and the energy efficiency, as we might expect. At higher traffic loads, we have found that this decrease is most pronounced for RI-MAC as its data free probes increasingly collide; it has a slight impact for O-MAC because its slot overhead is higher than BoX-MAC, for which it is least.

Overall, we have found that across most traffic and flow densities, O-MAC has highest throughput capacity and energy efficiency, suggesting that the combination of the factors of receiver-centricity and synchrony is important for its overall performance advantage. Moreover, it can be used without loss of latency performance compared to the other protocols we studied.

Finally, although our analysis and experiments provide a comparative analysis of state-of-the-art MAC protocols in terms of capacity and energy efficiency, the absolute performance numbers of all of these protocols are significantly lower than achievable limits. This indicates that there is still significant room for both protocol and engineering efforts to improve the MAC layer for low-power networks.

7. REFERENCES

[1] Y. Sun, O. Gurewitz, and D. Johnson, "RI-MAC: A receiver-initiated asynchronous duty cycle MAC protocol for dynamic traffic loads in wireless sensor networks," in *Proc. of the 6th Intl. Conf. on Embedded Network Sensor Systems*, pp. 1–14, ACM, 2008.

[2] D. Moss and P. Levis, "BoX-MACs: Exploiting physical and link layer boundaries in low-power networking," tech. rep., Technical Report SING-08-00, Stanford University, 2002.

[3] K. Klues, G. Hackmann, O. Chipara, and C. Lu, "A component-based architecture for power-efficient media access control in wireless sensor networks," in *Proc. of the 5th Intl. Conf. on Embedded Networked Sensor Systems*, (New York, NY, USA), pp. 59–72, ACM, 2007.

[4] W. Ye, F. Silva, and J. Heidemann, "Ultra-low duty cycle MAC with scheduled channel polling," in *Proc. of the 4th Intl. Conf. on Embedded Networked Sensor Systems*, pp. 321–334, 2006.

[5] G. Halkes and K. Langendoen, "Crankshaft: An energy-efficient MAC-protocol for dense wireless sensor networks," *Wireless Sensor Networks*, pp. 228–244, Apr 2007.

[6] H. Cao, K. Parker, and A. Arora, "O-MAC: A Receiver Centric Power Management Protocol," in *Proc. of the 14th IEEE Intl. Conf. on Network Protocols*, 2006.

[7] C. Lin and M. Gerla, "Adaptive clustering for mobile wireless networks," *Selected Areas in Communications, IEEE Journal on*, vol. 15, no. 7, pp. 1265–1275, 1995.

[8] M. Buettner, G. Yee, E. Anderson, and R. Han, "X-MAC: A short preamble mac protocol for duty-cycled wireless sensor networks," in *Proc. of the 4th Intl. Conf. on Embedded Networked Sensor Systems*, pp. 307–320, 2006.

[9] J. Polastre, J. Hill, and D. Culler, "Versatile low power media access for wireless sensor networks," in *Proc. of the 2nd Intl. Conf. on Embedded Networked Sensor Systems*, pp. 95–107, 2004.

[10] F. Xue and P. Kumar, *Scaling laws for ad-hoc wireless networks: An information theoretic approach.* Now Publishers Inc., 2006.

[11] "Anonymized technical report," `https://sites.google.com/site/anonymoustechnicalreport/tech-report`.

[12] P. Dutta and D. Culler, "Practical asynchronous neighbor discovery and rendezvous for mobile sensing applications," in *Proc. of the 6th Intl. Conf. on Embedded Network Sensor Systems*, pp. 71–84, ACM, 2008.

[13] M. Maróti, B. Kusy, G. Simon, and Á. Lédeczi, "The flooding time synchronization protocol," in *Proc. of the 2nd Intl. Conf. on Embedded Networked Sensor Systems*, pp. 39–49, ACM, 2004.

[14] "Kansei testbed," `http://kansei.cse.ohio-state.edu`.

[15] P. Dutta, D. Culler, and S. Shenker, "Procrastination might lead to a longer and more useful life," *Proc. of HotNets-VI*, p. 7, 2007.