The impact of wakeup schedule distribution in asynchronous power save protocols on the performance of multihop wireless networks

Laura Marie Feeney*†, Christian Rohner†, Bengt Ahlgren *

*Swedish Institute of Computer Science, Kista, Sweden
†Uppsala University, Uppsala, Sweden

Abstract— **By definition, the operation of an asynchronous power save protocol permits an arbitrary distribution of nodes' wakeup schedules. This wakeup schedule distribution creates an uncoordinated pattern of times at which nodes will attempt to transmit. Intuitively, we would expect that some patterns will be more (or less) favorable than others for a given traffic pattern.**

We investigate the impact of this wakeup pattern on network capacity and present simulation data showing that the capacity associated with the best wakeup patterns is significantly larger than that of the worst. This result not only gives insight to the behavior of such protocols, but also acts as a feasibility study showing the potential benefit of mechanisms by which nodes adapt their wakeup schedules to obtain improved performance.

I. INTRODUCTION

Ad hoc networks are intended to operate in a self-organizing fashion. For the lowest layer communication protocols, this means that nodes may not enjoy centralized synchronization or scheduling of channel access. For network layer services, this means nodes must cooperatively forward traffic for each other to maintain network connectivity.

Because nodes do not *a priori* know when they will be called on receive and forward traffic, they must be prepared to do so at any time. Unfortunately, listening to the wireless channel consumes significant energy, requiring the use of power saving mechanisms that allow the network interface to spend as much time as possible in a low energy consumption sleep state. Such a mechanism must provide a way for nodes to cooperatively buffer traffic for their sleeping neighbors and arrange appropriate rendezvous times to exchange traffic.

A variety of power save protocols have been proposed for CSMA/CA-based ad hoc networks: Some proposals are based on synchronous operation, using periodic broadcast of traffic announcement messages ([1]), often in conjunction with some form of clustering ([2]). To avoid the overhead of synchronization, a number of asynchronous protocols have also been proposed. A few are based on probabilistic mechanisms ([3], [4]), but most are based on wakeup schedules that ensure some deterministic overlap between nodes' wake intervals (Figure 1(a)).

Such asynchronous power save protocols generally require three elements. (1) A common, periodic wakeup schedule, which each node follows independently of its neighbors. This wakeup schedule is defined such that the nodes' wake intervals overlap in some deterministic way, regardless of the phase difference between them. Such wakeup schedules are often based on quorum scheduling or similar structures ([5], [6], [7]). (2) A neighbor discovery mechanism, which allows a node to make use of these structures to rendezvous with a neighbor. (3) A traffic scheduling mechanism that allows nodes to predict or negotiate intervals during which a neighbor will be awake to receive traffic. In this paper, we consider the simple majority power save protocol [5].

Because neighbors can only exchange traffic when they are both awake, the distribution of the nodes' wakeup schedules creates some spatial and time pattern, which we refer to as a *wakeup pattern*, in the nodes' attempts to access the channel. Given periodic wakeup schedules and assuming relatively long-lived flows and slowly changing topologies, such patterns will exist over moderate time intervals. A natural question is whether, for a given topology and traffic flows, there are particularly favorable or unfavorable wakeup patterns.

The answer is obvious in trivial cases: For example, if nodes A and B are receiving traffic from nodes C and D, minimal overlap between their wakeup schedules reduces contention and interference, even though nodes C and D cannot communicate directly (Figure 1(b)).

The answer is less obvious in more realistic scenarios: The shared multihop channel is characterized by complex patterns of inter- and intra-flow contention, extending over multiple hops. (Many multihop link scheduling problems are computationally hard [8]). It may be that, in practice, the cumulative effect of these interactions tends to work against the possibility of obtaining significant overall advantage.

A positive answer would be particularly interesting, as it would suggest the potential benefit of improving network performance by manipulating the wakeup pattern. The current work may therefore be viewed as a feasibility study in this regard.

To answer our question, we formulate the following experiment: We fix the topology, the routing, and the offered load (set of randomly chosen source-destination pairs), then measure the flow capacity of the network for each of a large number of randomly generated wakeup patterns. We choose

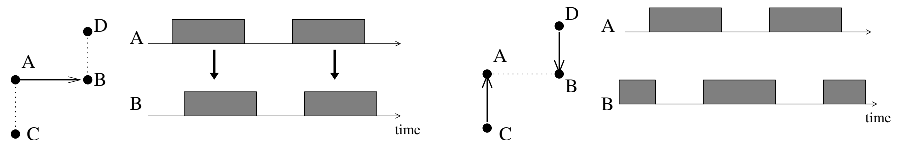

(a) Nodes alternate between wake(dark) and sleep (light) states:The two nodes have the same period, but the offset is arbitrary.

(b) Interfering transmissions are separated by schedules with minimal overlap.

Fig. 1. Distributions of wakeup schedules.

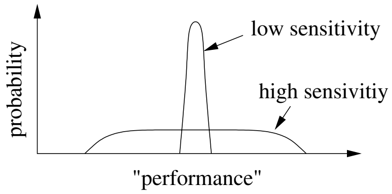

Fig. 2. In our experiment, we would like to distinguish between these cases.

the flow capacity metric because it reflects overall network performance in a practical sense.

These measurements will define some probability distribution (Figure 2). A very narrow probability distribution would imply that capacity is not too much affected by the wakeup pattern: what is good for some flows is bad for others and it is difficult to obtain some overall advantage. Conversely, a flat probability distribution would suggest that the capacity is very sensitive to the wakeup pattern. Because the most interesting implications of this work are related to the feasibility of improving performance by managing the wakeup pattern, our goal is effectively to distinguish between these two cases.

We have developed a simulation tool that allows us to efficiently determine the number of feasible CBR flows in a network for a large number of wakeup patterns. The data show that the flow capacity measurements exhibit a narrow central distribution with long tails. We conclude that there is substantial difference between the performance of best and worst wakeup patterns, but that such patterns are relatively rare. The result suggests the potential benefit of developing techniques for nodes to adapt the wakeup pattern to obtain better performance.

To the best of our knowledge, no previous work has investigated the whether the performance of an asynchronous power save protocol is sensitive to the pattern of the nodes' wakeup schedules. The primary contribution of this work is to show that the wakeup pattern generated by the operation of an asynchronous power save protocol has a significant effect on the network capacity. The result suggests the potential benefits of developing mechanisms in which the wakeup pattern is adapted to obtain better performance. The design of such a mechanism is out of scope of this work.

II. NETWORK MODEL

The experiment above requires that we compute flow capacity for each of many wakeup patterns, for each of many topologies and network configurations. This large number of simulation runs makes it impractical to use more detailed discrete event simulators, such as ns-2[9].

However, our immediate goal is not protocol design or optimization. Our goal is to measure the sensitivity of the flow capacity to wakeup pattern – to distinguish between the two hypothetical curves in Figure 2. That is, we are interested in the variation in performance, rather than in the absolute value of a performance metric.

We therefore develop a network model that is well suited for this kind of rapid exploration of the performance space. The key simplification is that we model only the "steady state" operation of the network, rather than its evolution over time. This approach requires three closely related assumptions:

First, the model considers only CBR traffic and periodic power save protocols, which need to have compatible (integer multiple) periods (e.g. a wakeup schedule with a period of 100ms and CBR flows of 10 packets per second.) This assumption allows us to evaluate a single period, which is assumed to be representative of the steady state performance.

Second, we ignore transient overhead traffic (e.g. routing, admission control, and upper layer protocols), assuming that it is sufficiently infrequent that only a small proportion of periods are affected. It is possible to represent constant periodic overhead, but there is no reason to, as we are not interested in protocol optimization.

Third, we assume variation in channel access time is small compared to the total transmit time for the frame. This is a particularly strong assumption for a CSMA-based MAC layer, which does not provide such a structure. However, the CBR flow capacity metric implies a system in which there is some soft admission control mechanism (e.g. [10]) and that the network is thus unlikely to be heavily overloaded.

The model is described in more detail below and in [11].

Each node has fixed transmit radius and an interference radius that is 40% larger than the transmit radius (the so-called "protocol model"). Nodes are assumed to communicate without error to any node within their transmit radius, while nodes within the interference radius of a transmitter sense the channel as busy.

We assume that the MAC layer operates without error to prevent packet loss due to collisions. Each transmission occupies the channel at all nodes within interference range of the transmitter and within communication range of the receiver. This rule roughly models the RTS/CTS operation of IEEE 802.11.

The time that the transmission occupies the channel is intended to reflect the complete channel access process, with an implicit assumption that the variation in channel access time (e.g. random backoff) is small relative to the total transmission time. The transmission time is 2.2 ms, which corresponds with a short (ca 137-byte) IEEE 802.11b frame transmitted at 11Mbps or a longer frame transmitted on a newer IEEE 802.11 interface.

TABLE I

CONNECTIVITY (MIN-MAX) AND MEAN PATH LENGTH (STDDEV). AVERAGE OVER 50 TOPOLOGIES.

large square	nodes	connected pairs (%)	mean path length	**small square**	nodes	connected pairs (%)	mean path length
	50	39% (11-92)	3.9 (1.2)		15	89% (19-100)	2.1 (0.4)
total area = 1.0	75	79% (27-100)	5.5 (1.1)	total area = 1.0	20	96% (47-100)	2.1 (0.3)
6.3 x 6.3 hops	100	94% (49-100)	5.3 (0.6)	2.6 x 2.6 hops	30	99% (63-100)	2.0 (0.2)
	125	98% (78-100)	5.0 (0.3)		40	99% (81-100)	2.0 (0.1)

large rectangle	nodes	connected pairs (%)	mean path length	**small rectangle**	nodes	connected pairs (%)	mean path length
	75	58% (21-100)	5.2 (1.6)		20	82% (31-100)	2.4 (0.6)
total area = 1.0	100	86% (33-100)	6.3 (1.2)	total area = 1.0	35	98% (48-100)	2.6 (0.3)
3.2 x 13 hops	125	95% (39-100)	6.5 (0.8)	1.3 x 5.2 hops	50	99% (49-100)	2.6 (0.2)
	150	99% (45-100)	6.3 (0.4)		65	100% (100)	2.5 (0.1)

We model the simple majority power save protocol[5]. The wakeup schedule is based on the observation that, if every node is awake slightly more than half of each period, its awake interval will overlap with that of each of its neighbors, regardless of the pattern of their wakeup schedules (Figure 1). Neighbors use these periods of guaranteed overlap to learn the phase difference between their wakeup schedules and thus to predict when it is possible for them to exchange data traffic. The period of the power save protocol is 100ms, with a 55% duty cycle (i.e. 55ms awake, 45ms asleep).

We chose this power save protocol, despite its limited energy saving, for its simplicity and because we expect its high duty cycle to make it resilient to variation.

A transmission is feasible if the transmitter and receiver are both awake and have a free channel for an interval corresponding to the transmission time. The transmission is assigned to the interval that meets these criteria.

The traffic consists of CBR flows with the same period as the power save protocol: 10 packets per second. A flow is feasible if each transmission along its route is feasible in one period. That is, the source and each forwarding node must be able to transmit once each period, so that one packet enters and one packet leaves the network in each period. (Note that the latency for a given packet may be greater than one period.) Shortest path routes are computed at initialization time and not changed.

III. SIMULATION EXPERIMENT

We consider four scenarios: large and small square and large and small rectangle, each with a variety of node densities (Table I). A topology is generated according the scenario parameters, with nodes deployed on a rectangular field, their x- and y-coordinates uniformly distributed random variables over the size of the field. Nodes are stationary.

The shortest path routes for all pairs are computed and an ordered set of *candidate* flows (offered load) is randomly selected from the set of connected source-destination pairs. The size of the candidate set is chosen such that most, but not all, flows in the set are feasible (see below).

To determine the *baseline* flow capacity of the topology, we evaluate the feasibility of each flow in the candidate set with the power save protocol turned off. This determines the baseline set of flows, which are known to be feasible in the absence of any constraints imposed by a wakeup pattern.

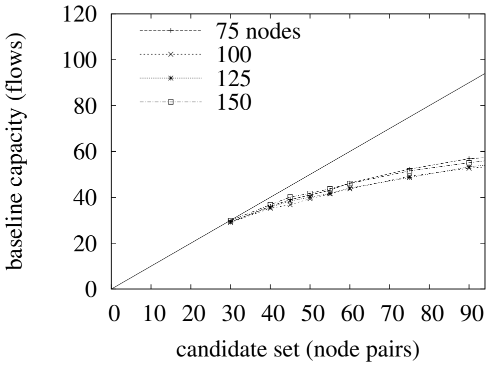

Fig. 3. Baseline flow capacity (no power save protocol) as a function of candidate set size (large rectangle scenario, averaged over 50 topologies).

We then apply the 55ms on/45ms off wakeup schedule described above to each node, such that the phase difference between the nodes' schedules is randomly and uniformly distributed over the period. The feasibility of each flow in the baseline set is evaluated and the total flow capacity computed. The process is repeated for a large number of wakeup patterns to obtain the probability distribution of the flow capacity.

We earlier referred to the candidate set size. If the set is too small, all of the flows are feasible. The network is underutilized and unlikely to exhibit significant capacity variation due to the operation of the power save protocol. If candidate set is too large, there is a "cherry picking" effect: we are effectively searching through many candidates for feasible flows. Such an offered load is also unrealistic in the sense that the network is badly under-dimensioned, since most user flows are rejected.

We would therefore like the size of the candidate set to result in an offered load such that the network is mildly congested: most, but not all, flows in the offered load are feasible. The candidate set size should also reflect a reasonable offered load (flows per user) in the context of the scenario.

To determine the size of the candidate set, we compute the baseline (no power save protocol) capacity of a given topology for each of several candidate set sizes (Figure 3).

TABLE II

SIMULATION PARAMETERS

timing	period	100ms
	resolution	500 slots/period
power mgmt	duty cycle	55% (55ms on, 45ms off)
pkt size	transmit time	2.2ms
	(137byte @ 11Mbps IEEE 802.11)	
traffic	CBR traffic	10 pkts / sec
candidate set	small sq.	40 flows
	small rect.	40 flows
	large sq.	50 flows
	large rect.	45 flows
experiments	350 wakeup schedules / topology	
	4 scenarios × 50 topologies	

When the size of the candidate set is small, then all of the flows are feasible and the capacity is equal to the size of the candidate set. As the size of the candidate set increases, the network becomes more congested. Some flows are infeasible and the baseline capacity grows more slowly than the size of the candidate set. Eventually, no more feasible flows can be found, no matter how large the candidate set.

Using this data, we (visually) choose the candidate set size at which 85-90% of the flows are feasible in the absence of any power save protocol. An extremely precise estimate is not needed, we just want to define experimental parameters such that the network is likely to be moderately congested. The candidate set sizes used for each scenario are shown in Table II. Depending on the scenario, the size of the candidate set reflects a load of 0.25–2 flows per node.

IV. SIMULATION RESULTS

We first present a subset of the raw data in some detail and then present integrated data and compare the baseline performance with that obtained by the "best" wakeup pattern.

A. Sampling the data

In this section, we examine in detail an instance of the large rectangle scenario, with 100 nodes and a mean path length of about six hops.

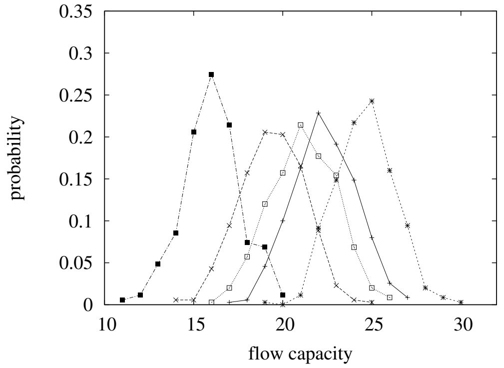

Fig. 4. Raw probability distribution (8 topologies)

Figure 4 shows the probability distributions for the flow capacity of five topologies (randomly selected from the 50 topologies evaluated for this configuration). Each curve represents a single topology, showing the probability (y-axis) of observing a given flow capacity (x-axis), based on flow capacity measurements for 350 different wakeup patterns.

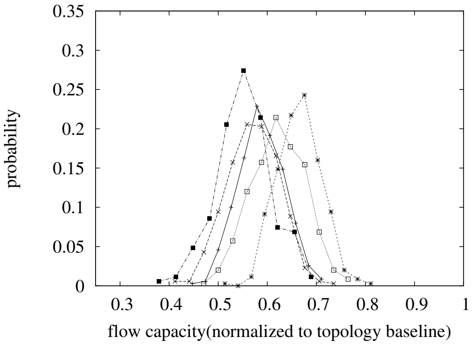

Fig. 5. Normalized probability distribution

In general, we observe a roughly symmetric normal distribution. This is not too surprising, given that the underlying random variable – the overlap between nodes' wakeup schedules – is (almost) uniformly distributed. Note the considerable variation in absolute flow capacity among the various topologies.

Figure 5 shows the same result, with flow capacities normalized to the baseline (no power save) capacity of their associated topology. The variation between topologies persists: in some topologies the best wakeup patterns obtain over 80% of the baseline capacity, while in others they obtain only about 65%.

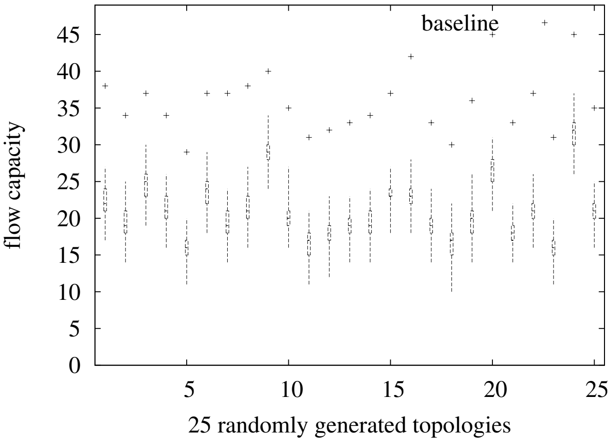

Fig. 6. Quartile representation

Plotting the flow capacity measurements as quartiles loses some of the detail present in the full distribution, but makes it possible to present more topologies in a single plot (Figure 6). We observe that the two middle quartiles account for only a small part of the total variation, while the extremes are significantly larger. We further note that the best wakeup patterns generally obtain about 70-75% of the baseline capacity.

TABLE III

MEDIAN FLOW CAPACITY: MEAN OF THE DISTANCE OF EACH QUARTILE FROM THE MEDIAN, RELATIVE TO THE MEDIAN (COMPUTED FOR EACH TOPOLOGY). EACH MEAN IS COMPUTED OVER 50 TOPOLOGIES (STDDEV IN PARENTHESES).

	nodes	median capacity	inner quartile relative to median		min - max relative to median	
small square	15	20.2 (3.6)	-0.06 (0.02)	+0.07 (0.02)	-0.31 (0.05)	+0.25 (0.05)
	20	20.8 (2.7)	-0.06 (0.02)	+0.06 (0.02)	-0.28 (0.04)	+0.23 (0.04)
	30	23.3 (2.7)	-0.05 (0.02)	+0.05 (0.02)	-0.25 (0.05)	+0.21 (0.04)
	40	24.7 (2.8)	-0.05 (0.02)	+0.05 (0.01)	-0.22 (0.04)	+0.19 (0.03)
small rectangle	20	22.3 (5.7)	-0.06 (0.02)	+0.06 (0.02)	-0.27 (0.07)	+0.22 (0.05)
	50	22.7 (3.2)	-0.05 (0.02)	+0.05 (0.02)	-0.22 (0.05)	+0.19 (0.05)
	35	22.1 (3.6)	-0.06 (0.02)	+0.05 (0.02)	-0.25 (0.05)	+0.21 (0.04)
	65	24.1 (3.2)	-0.04 (0.01)	+0.04 (0.01)	-0.21 (0.04)	+0.17 (0.03)
large square	50	24.2 (7.3)	-0.05 (0.02)	+0.06 (0.02)	-0.24 (0.05)	+0.24 (0.07)
	75	23.4 (5.1)	-0.05 (0.02)	+0.06 (0.02)	-0.25 (0.05)	+0.23 (0.05)
	100	23.6 (3.2)	-0.06 (0.02)	+0.05 (0.02)	-0.23 (0.04)	+0.23 (0.05)
	125	26.6 (2.7)	-0.05 (0.02)	+0.05 (0.02)	-0.22 (0.03)	+0.20 (0.04)
large rectange	75	22.5 (6.2)	-0.06 (0.02)	+0.06 (0.02)	-0.25 (0.06)	+0.24 (0.07)
	100	21.7 (4.8)	-0.06 (0.02)	+0.06 (0.02)	-0.25 (0.05)	+0.23 (0.05)
	125	21.3 (3.4)	-0.06 (0.02)	+0.06 (0.02)	-0.25 (0.05)	+0.23 (0.05)
	150	22.4 (3.1)	-0.05 (0.02)	+0.05 (0.02)	-0.24 (0.04)	+0.21 (0.04)

B. Summarizing the data

We present the integrated quartile data for each configuration and compare the baseline performance with that of the best observed wakeup pattern for each topology.

a) Quartile analysis: For each topology, we compute the distance from the median flow capacity to each quartile. Each distance is computed relative to the median for that topology to eliminate the inherent variation in capacity between topologies. This relative inter-quartile distance is then averaged over the 50 topologies simulated in each scenario. The results are shown in Table III. (Note that the stddev reflects the variation between topologies and not the wakeup schedule dependent variation within a topology.)

We see that the inner quartiles lie within a range about ±5% of the median. In absolute terms, this means that 50% of wakeup patterns will support only one flow more or less than the median flow capacity of 20-25 flows – a negligible difference. The max-min excursions are much larger: the "best" wakeup schedule patterns support some 20-25% more flows than the median, while the "worst" patterns support some 20-25% fewer flows. This represents a loss or gain of four or five flows relative to the median capacity of 20-25 flows.

b) Maximum capacity: We also compare the capacity of the "best" wakeup pattern and the baseline capacity (no power save protocol). Figure 7 shows, for each scenario, the proportion of topologies (y-axis) in which the best wakeup pattern obtains a flow capacity that exceeds some percentage (x-axis) of the baseline capacity for that topology.

Observe that for all topologies, the ratio between the best flow capacity and baseline capacity is larger than the 55% duty cycle. Even in large dense networks, in a majority topologies, the best wakeup pattern obtains 75-85% of the baseline capacity. In over 70% of observed topologies, the best wakeup pattern obtains at least 70% of the baseline capacity obtained in the absence of the power save protocol.

V. DISCUSSION

In the previous section, we showed that there is considerable variation in flow capacity, but we did not show that this variation is an underlying property of the network. One possible explanation for our result is that variation in the number of feasible flows simply reflects variation in the path-length of the admitted flows.

We call this the total "total transmissions hypothesis": If a long path-length flow happens to be feasible in some wakeup pattern, the flow capacity only increases by one, even though many transmissions are required. Conversely, if the long path-length flow is not feasible, then two (or more) short path-length flows may be able to use its transmission times, increasing the flow capacity by two (or more).

The model of an offered load consisting of flows of varying path length is realistic. However, the previous results would not be meaningful with regard to finding wakeup patterns that improve network performance, since it is trivially possible to increase the number of admissible flows by preferentially admitting short path-length flows.

If the total transmission hypothesis were true, we would expect that the total number of transmissions in the network – the number of feasible flows times their mean path length – should be roughly constant for each wakeup pattern.

Figure 8 shows the mean distance of each quartile from the median, calculated as described in the previous section. Generally speaking, the inner quartiles (50% of the observed flow capacities) vary ±10% from the median, while the extremes vary $\pm 30 - 40\%$, which is even larger than the variation in the flow capacity. In short, there does not seem to be any evidence for a pattern of small variation around the natural total capacity of each topology – and thus little support for the total transmission hypothesis.

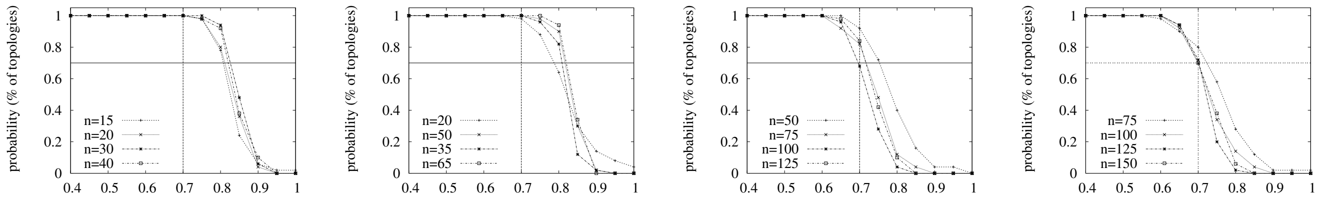

(a) small square (b) small rectangle (c) large square (d) large rectangle

Fig. 7. Maximum flow capacity (relative to baseline)

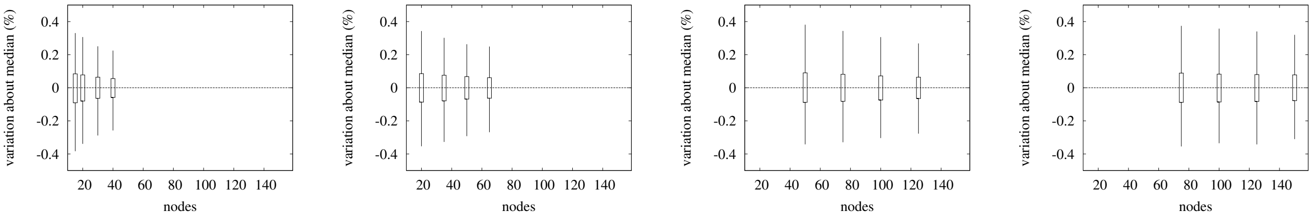

(a) small square (b) small rectangle (c) large square (d) large rectangle

Fig. 8. Quartiles: Variation in the total number of transmissions. The variation in the total number of transmissions is large, suggesting that flow capacity does not simply reflect the variation in path length among the flows in the offered load.

VI. CONCLUSION AND FUTURE WORK

We investigated the impact of the wakeup pattern created by the operation of an asynchronous power save protocol on the capacity of a multihop wireless network. To do this, we defined a simple model of network operation that is suitable for rapid exploration of the performance space and allows us to efficiently compute flow capacity for a large number of wakeup patterns.

We determined that, for a given topology and offered load, the flow capacity (and total number of transmissions) in a network can vary significantly depending on the wakeup pattern. The simulation results show that the probability distribution of flow capacity measurements has a narrow central variation and long tails, with the maximum and minimum observed capacities varying some ±20-25% from the median. Although the study of more complex power save protocols remains future work, we expect the general result to hold for the more restrictive wakeup schedules associated with these protocols.

The value of this result is not just that it provides insight into the behavior of such protocols, but more importantly, that it suggests that wakeup schedule adaptation may significantly improve network performance.

Because asynchronous power save protocols function correctly for any wakeup pattern, nodes can adapt their wakeup schedules to obtain more favorable patterns. Unfortunately, the relative rarity of good patterns suggests that purely randomized approaches may not be very effective and that more sophisticated techniques will be required. In this regard, we note that the problem bears some similarities to the difficult problem of computing assignments for spatial reuse TDMA([12]), albeit at a much coarser granularity and at a higher network layer. This topic is the focus of future work.

ACKNOWLEDGMENTS

This work was partly supported by the Winternet program funded by the Swedish Foundation for Strategic Research.

REFERENCES

[1] IEEE, 1999. http://www.ieee.org.

[2] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris, “Span: An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks.,” *Wireless Networks*, vol. 8, no. 5, pp. 481–494, 2002.

[3] M. J. McGlynn and S. A. Borbash, “Birthday protocols for low energy deployment and flexible neighbor discovery in ad hoc wireless networks.,” in *MobiHoc*, pp. 137–145, ACM, 2001.

[4] P. Hurni, T. Braun, and L. M. Feeney, “Simulation and evaluation of unsynchronized power saving mechanisms in wireless ad hoc networks,” in *4th Internationl Conference on Wired/Wireless Internet Communications WWIC 2006*, (Bern, Switzerland), May 2006.

[5] L. M. Feeney, “A QoS aware power save protocol for wireless ad hoc networks,” in *Proceedings of the First Mediterranean Workshop on Ad Hoc Networks (MedHocNet 2002)*, 2002.

[6] J.-R. Jiang, Y.-C. Tseng, C.-S. Hsu, and T.-H. Lai, “Quorum-based asynchronous power-saving protocols for ieee 802.11 ad hoc networks.,” *MONET*, vol. 10, no. 1-2, pp. 169–181, 2005.

[7] R. Zheng, J. C. Hou, and L. Sha, “Asynchronous wakeup for ad hoc networks,” in *MobiHoc*, pp. 35–45, ACM, 2003.

[8] A. Sen and M. L. Huson, “A new model for scheduling packet radio networks,” *Wireless Networks.*, vol. 3, no. 1, pp. 71–82, 1997.

[9] http://www.isi.edu/nsnam.

[10] G.-S. Ahn, A. T. Campbell, A. Veres, and L.-H. Sun, “Supporting service differentiation for real-time and best-effort traffic in stateless wireless ad hoc networks (SWAN).,” *IEEE Trans. Mob. Comput.*, vol. 1, no. 3, pp. 192–207, 2002.

[11] L. M. Feeney, “Impact of phase distribution in asynchronous communication protocols,” Tech. Rep. T2005:15, Swedish Insititute of Computer Science, Dec. 2005.

[12] J. Grönkvist, “Distributed scheduling for mobile ad hoc networks - a novel approach,” in *Proceeedings of the 15th IEEE Int’l Symposium on Personal, Indoor and Mobile Radio Communications PIMRC*, 2004.