The feasibility of leveraging a power save protocol to improve performance in ad hoc networks

Laura Marie Feeney * Christian Rohner †

We discuss how the operation of an asynchronous power save protocol affects the scheduling of transmissions in an ad hoc network. We present simulation results that show that network capacity is sensitive to the random distribution of wakeup schedules. Finally, we discuss the implications of this result with respect to improving performance and speculate about potential new cross-layer architectures.

Power save protocols reduce energy consumption in ad hoc and sensor networks by maximizing the amount of time each node's wireless interface spends sleeping. In a CSMA/CA environment, nodes cannot predict when they will receive (or forward) traffic, so they must follow some schedule that determines when they wake up to receive any pending traffic.

Some protocols are designed to operate asynchronously. In this case, each node follows a periodic wakeup schedule with an unknown offset relative to its neighbors. The wakeup schedule is defined so as to ensure deterministic overlap between the wake intervals of neighboring nodes (e.g. via quorum scheduling[1]). Nodes use these overlapping wake intervals to rendezvous with their neighbors and exchange traffic.

The (random) distribution of the nodes' wakeup schedules may be expected to affect the network performance. For example, a node pair that forwards a lot of traffic will be able communicate more efficiently if their common wake interval has minimal overlap with that of an interfering node pair.

To examine this intuition more carefully, we studied how the network performance depends on the wakeup schedule distribution. Conceptually, we measure the network "performance" for one topology and many different wakeup schedule distributions. The measurements define a probability distribution that indicates how sensitive the performance is to the pattern of overlap intervals defined by the wakeup schedule. Figure 1 shows two hypothetical extremes.

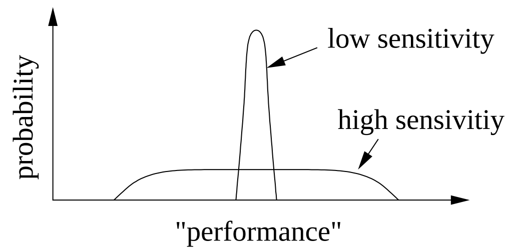

Figure 1: A narrow distribution suggests that wakeup schedule distribution has little impact on performance.

Actual simulation results suggest that the network performance – as represented by CBR flow capacity – exhibits relatively little variation, but there do exist distributions with considerably better and worse performance.

In our experiment, we define a simple wakeup schedule in which each node is awake slightly more than half of each period. We fix a topology and an offered load (CBR flow between randomly selected multi-hop source-destination pairs). This metric reflects the impact of the wakeup distribution, as well as contention and interference. The offered load is such that the network is slightly under-dimensioned: most, but not all, flows are admissible. Each flow in the offered load is evaluated in turn; the flow capacity is the total number of admitted flows. For each topology, we generate a large number of wakeup schedule distributions and determine the flow capacity of each. Figure 2 shows an extract from the simulation data.

To evaluate flow admissibility, we use a simplified model of network operation. To avoid the overhead of modeling time varying behavior, we consider only the "steady state" operation of a static network. To do this, we require that the periodicity of the wakeup schedule and the pe-

*Ph.D. Student, Uppsala University and Research Staff, Swedish Institute of Computer Science. Contact: lmfeeney@sics.se.

†Post-Doctoral Researcher, Uppsala University. Contact: christian.rohner@it.uu.se

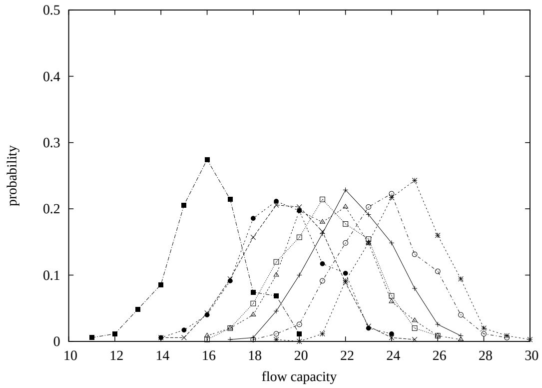

Figure 2: Sample of performance distribution for 8 different topologies. The absolute capacity of the topologies varies, but the distributions have roughly similar shapes.

riodicity of the CBR flows are compatible and model only a single period. Transient traffic such as routing and admission control is ignored. We also assume an ideal MAC that operates without transmission error and without large variation in channel access times. This allow us to evaluate flow admissibility using simple matrix operations.

The model is clearly less realistic than a conventional simulator such as ns-2, but it allows us to efficiently simulate (statistically) many wakeup schedule distributions for each of many topologies. Moreover, we are more interested in qualitative behavior, i.e. to distinguish between the hypothetical curves in figure 1, than in quantitative capacity measurements.

The simulation results show that the flow capacity is roughly normally distributed with wakeup schedule variation. Over a wide range of scenarios, half of all distributions account for a variation of only $\pm 5\%$ about the median flow capacity - a negligible difference given the model. The extremes, however, vary $\pm$ 20-25% about the median, suggesting opportunities to obtain considerable advantage.

The simulation results also show that total number of transmissions varies as much as the flow capacity. This confirms that the variation in flow capacity is the result of variation in the wakeup schedule and not just an artifact of selection of flows from the offered load.

These results are interesting because they imply that network capacity can be improved by adjusting the nodes' wakeup schedule distribution. Local adjustment of the offset between node's wakeup schedules is trivial, because the mechanism is intended to operate asynchronously and therefore supports any wakeup distribution. It is not yet clear how to obtain a global improvement, given the relative scarcity of "good" distributions indicated by these results.

Alternatively, consider the problems of intra-flow contention and interference along a multihop flow (e.g. [2]). A "good" wakeup schedule distribution would ensure that the wake intervals for potentially conflicting transmissions did not overlap , if necessary by adjusting the offsets between the nodes' wakeup schedules. The CSMA/CA would usually sense the channel as clear, reducing the time wasted in contention.

Given the known difficulty of computing, much less instantiating, an optimal multihop TDMA slot assignment, it is not clear how such "good" distributions could be obtained. Our preliminary results suggest that trivial randomization is not likely to be sufficient. An better strategy is likely to be adaptive, although care must be taken to avoid problematic feedback loops. A more sophisticated simulation environment is needed to properly investigate these issues and is the subject of ongoing and future work.

More speculatively, we hypothesize that future cross-layer MAC architectures might make use of these flexible timing structures: Despite their inefficiencies, CSMA-based MAC's (e.g. IEEE 802.11) are widely used in ad hoc networks because of their natural distributed, asynchronous operation. The slot schemes associated with TDMA MAC's can be more efficient, but have proven challenging to implement in the ad hoc environment. The timing patterns defined by the wakeup schedule distribution create coarse-grained, variable-length "slots", which might be used for higher-layer scheduling or traffic management. Moreover, heuristic methods could be applied, since the channel would be further protected by the underlying CMSA.

[1] Jiang, J.-R., Tseng, Y.-C., Hsu, C.-S., and Lai, T.-H. Quorum-based asynchronous power-saving protocols for ieee 802.11 ad hoc networks. *MONET 10*, 1-2 (2005), 169–181.

[2] Sanzgiri, K., Chakeres, I. D., and Belding-Royer, E. M. Determining intra-flow contention along multihop paths in wireless networks. In *BROADNETS* (2004), pp. 611–620.