Performance of Energy-Conserving Access Protocols Under Self-Similar Traffic

Jason Redi^a and Dimiter Avresky^b
^{a} BBN Technologies, GTE, Cambridge, MA, *redi@bbn.com*
^{b} Boston University, Department of Electrical and Computer Engineering, Boston, MA, *avresky@bu.edu*

***Abstract* - Since widespread commercial use of wireless technology is still many years off, there is very little data to determine what constitutes an accurate model for wireless traffic. Many recent studies have shown that landline network traffic does not follow a traditional Poisson model but instead exhibits self-similar behavior. Previous papers have introduced and analyzed the performance of fundamental classes of energy conserving protocols under traditional Poisson models. In order to evaluate these protocols under a distribution which might more accurately depict future wireless network traffic, we use simulation and proven models to consider the trade-off between energy and delay under a self-similar arrival distribution.**

I. INTRODUCTION

As new devices and services proliferate, there is a general lack of knowledge about what the interarrival distributions of different types of traffic will be on future wireless networks. The very recent deployment of these networks, changing user expectations as well as billing and end-user device issues have kept the industry from definitively pinning down the current traffic profile, let alone understanding enough to predict future use. However, what is well established through the market is that end users typically expect the same applications, power and connectivity in their laptop computers as they have on the computer which sits on their desk. Therefore, a realistic estimate of expected traffic on future wireless networks is a mirror of today's landline network traffic patterns.

Traditional network modeling and analysis has relied on the assumption of exponential interarrival rates and Poisson traffic models. However, the ground-breaking work in [7] used long, high resolution traces of Ethernet packets for the first time to show that arrival rates of IP packets on a LAN exhibit not Poisson, but self-similar behavior. This work has been subsequently repeated many times in the literature to show that self-similar behavior is common throughout many types of traffic including telnet , FTP [8], World-Wide-Web [1], ISDN, and signaling (SS7) traffic [3]. This presents a fundamental challenge to the networking community as self-similar processes have dramatically different characteristics than traditional Poisson traffic models. In particular, self-similar processes can be described by their characteristic of being scale-invariant, meaning that a perfectly self-similar process on average looks exactly the same regardless of the time scale observed. When posed in the context of traffic load, this means that we see extended bursts of activity and inactivity regardless of whether we look at millisecond, second, minute, or hour averages. This is in dramatic contrast to Poisson processes which have exponentially distributed, independent interarrival times such that over long time scales, Poisson processes can be estimated well by a constant mean rate.

It has been suggested that many theoretical protocols and systems need to be reevaluated under this different type of traffic distribution before practical implementations are constructed. In particular, given the limited energy and bandwidth availability in wireless networks, insuring predictable and efficient performance is of paramount importance. Research papers and commercial systems have introduced energy conserving access protocols which attempt to reduce the amount of time a mobile node needs to have its receiver on (to reduce energy), while not substantially adding to the access delay for packets queued at the base station. These types of protocols can be classified into three distinct groups similar to those described in [2]. Protocols such as these are very important to the success of a wireless device and are representative classifications of the variety of energy-conserving access protocols used in devices from pagers to wide-area wireless data modems. Since the wireless network transceiver can typically use 15-30% of the power of a typical mobile computer [5], and the battery's energy not expected to substantial increase in potential in the near future, it is of great importance to emphasize energy-conservation of the wireless modem when considering future mobile communication.

These energy-conserving access protocols have been analytically and experimentally evaluated in the framework of Poisson arrivals. To further consider the use of these protocols in a practical system, this paper addresses the energy vs. delay performance of these protocols under self-similar packet arrivals. We adapt theoretically proven models to our network model to determine if these classes of energy-conserving access protocols continue to give substantial savings in energy while keeping additional delay costs within reasonable bounds.

In the next section we present a definition of self-similarity and discuss a proven simulation model. Next, we summarize three classes of energy-conserving access protocols protocols as an introduction to the simulations. We then consider the results and discuss their implications.

II. SELF-SIMILAR MODELS

Self-similar processes can briefly be described by their characteristic of being *scale-invariant*. A perfectly self-similar process on average looks exactly the same regardless of the time scale observed. A wealth of literature on self-similar phenomena has been indexed in [9].

The degree of self-similarity of a process is typically described by the *Hurst* parameter (H). H is between 0.5 and 1.0, where 0.5 represents non self-similar behavior and the closer to 1.0, the more long-range dependent the process is.

A continuous time stochastic process $X(t)$, is then self-similar if $a^{-H}X(at)$, has exactly the same second-order statistics as $X(t)$ for any real $a > 0$ and Hurst parameter $0.5 \leq H \leq 1.0$. The idea is that a direct scaling of time results in a related scaling of the series regardless of what scale is chosen.

A discrete-time definition of self-similarity follows accordingly. If we have a wide-sense stationary process X_t, then the autocorrelation of the function depends only upon the lag k. Let $X_t^{(m)}$ be aggregated versions of the same time series, in other words $X_k^{(m)} = 1/m(X_{km-m+1} + \ldots + X_{km})$. The process is then exactly self-similar if the second-order properties of $X_t^{(m)}$ are the same as those of X_t:

$$Var(X_t^{(m)}) = \frac{Var(X_t)}{m^\beta} \tag{1}$$

$$R_{X_t^{(m)}}(k) = R_{X_t}(k) \tag{2}$$

Where $H = 1 - \beta/2$.

A slightly weaker definition is the following: a process is *asymptotically self-similar* if its aggregated version has the following properties:

$$Var(X_t^{(m)}) = \frac{Var(X_t)}{m^\beta} \tag{3}$$

$$R_{X_t^{(m)}}(k) = R_{X_t}(k) \text{as } m \rightarrow \infty \tag{4}$$

In other words, as we average over larger and larger blocks of size m we loose some small-scale transients in the signal and are left with another process with the same properties. This is the point of divergence from traditional traffic models. When averaged over the long term, a Poisson process will tend to a constant with some added noise, and result in an autocorrelation of 0 and a variance which decreases as $1/m$. The self-similar process exhibits long-range dependence in that autocorrelations can be seen over a very wide range of lags (k).

A. *Modeling Self-Similar Traffic*

Due to the time-dependent nature of self-similar processes, it is currently not possible to analyze with any sufficient detail the effects of self-similar processes in queueing system analysis. However, it has been shown through trace data [11] and also proven [10] that concise modeling of self-similar traffic can be achieved by the aggregation of a large number of ON-OFF packet train sources. The ON-OFF states are strictly alternating where ON represents a state where packets are generated according to some regular rate, and OFF represents a state where no packets are generated. The length of time in which each train spends in either the ON or OFF state should be selected according to a distribution which has long-range dependence. That is, a distribution where the amount of time spent in the state can be very large with a non-negligible probability. The Pareto distribution $(F(x) = 1 - x^{-\alpha}$, with $1 < \alpha < 2)$ has been found to fit very well to the actual distribution of observed packet trains. The long-term correlations which result from using this distribution are synonymous with the *Noah effect*, and is the main point of departure from traditional traffic models. Traditional models such as Poisson distributions show short scale correlations while the Noah effect suggests long scale correlations. As discussed in [11], values of $\alpha = 1.90$ and $\alpha = 1.25$ for the ON-time and OFF-time Pareto distributions match well with observed data and are therefore used for our simulations of 1000 individual packet trains.

III. NETWORK MODEL AND ENERGY-CONSERVING PROTOCOLS DESCRIPTION

For purposes of evaluation of energy-conserving protocols, we consider a single cell system where a base station communicates with N nodes through a radio channel of bandwidth B. The communication is packet-oriented. We assume the time to be slotted and the base station's transmissions to be synchronized to the beginnings of slots. The packet length c is constant, and exactly one packet can be transmitted during one slot. In this model, we do not explicitly treat transmission errors.

We define an *access protocol* as consisting of two components: a *transmission scheduling strategy* at the base station which in each slot selects a packet for transmission from the arrival queue, and a *wake-up schedule* at each node which determines the slots in which the node is awake. In general, the transmission scheduling strategy can take into account different parameters: the number of packets in the queue, the packets' ages, as well as the wake-up schedules of their destinations. In the protocols discussed here the "oldest packet" criterion is generally adopted to help meet the application delay requirements. We next review three classes of protocols for constructing efficient wake-up schedules as originally described in [2]: grouped-TDMA protocols, directory protocols, and pseudo-random protocols.

A. *Grouped-TDMA Protocols*

Very similar to the technique used by pagers, nodes are divided into m disjoint groups, with the cardinality of each group differing by at most one node, and assign (reserve) each slot of a

TDMA cycle to a unique group. This increases the average energy consumption per slot over traditional (one node per slot) TDMA by the cardinality of the groups, but decreases the average delay since there is a greater probability that a node will be awake soon after a packet for it has arrived at the basestation. The optimal selection of this group size is important for receiving the best energy and delay performance of this class of protocols.

B. *Directory Protocols*

Variants of the directory scheme are used by a number of wide area wireless data services as well as 802.11. In our version of the directory protocol, the base station always waits for a group of k packets in the queue to accumulate. The base station then transmits a list or directory of the k packet destinations before transmitting the packets in the subsequent slots. The nodes are all awake during the transmission of the directory, and can therefore schedule their wake-up slots to coincide with the broadcast of their packets. When there is no group being currently transmitted, the nodes wake up periodically every v slots in order to give the base station an opportunity to start the transmission of a new group.

The choice of the parameter k depends on the load and must take into account the trade-off between the increase in the delay due to a larger k and the energy savings from more infrequent broadcasting of directories. Additionally, the parameter v should depend upon k and the load. A system with small value of v and a low load will have the nodes waking up frequently until enough packets have accumulated at the base station, while a large value of v will incur an increase in delay before the start of group's transmission.

C. *Pseudo-Random Protocols*

The pseudo-random protocols are a new class of protocols based on deterministic (pseudo-random) schedules which preserves the power of randomization for fairness, while providing the advantages of determinism, i.e., the base station's ability to predict nodes' state in each slot. In this class of protocols all nodes run the same pseudo-random number generator and determine their state (awake or asleep) at each slot based on a probability p and the stored state of the random number generator. In order to avoid a complete overlap of the wake-up schedules, the pseudo-random generator of each node is initialized using a unique seed, which is known at the base station. Therefore, by using the same pseudo-random number generator it is possible for the base station to determine the schedules of the nodes it wants to transmit to. The base station can initiate changes in the value of p as a function of the load, the number of nodes, etc.

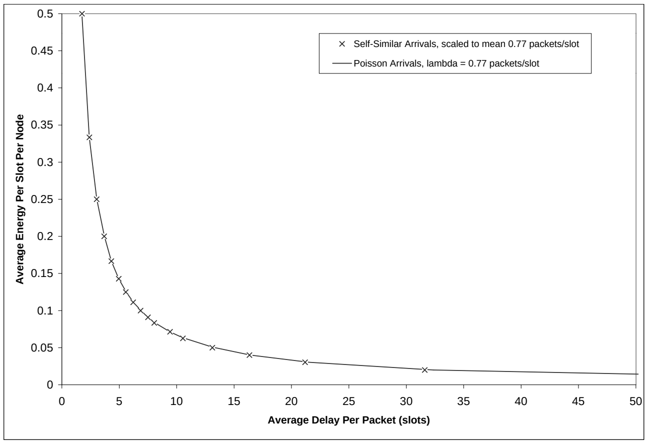

Figure 1: Grouped-TDMA Performance, Gaussian Destination Distribution ($\mu = \frac{N}{2}, \sigma^2 = 10N$)

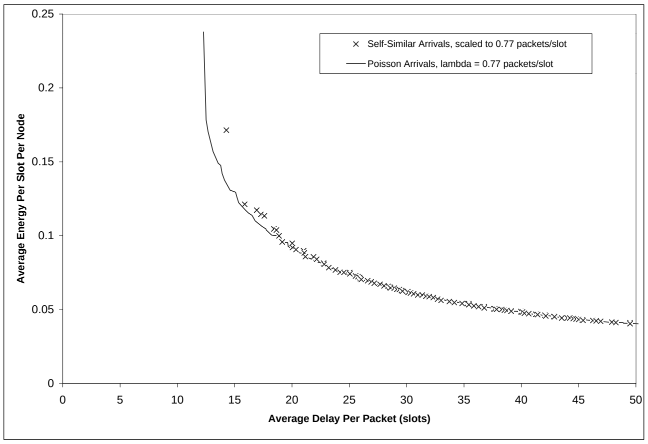

Figure 2: Directory Protocol, Gaussian Destination Distribution ($\mu = \frac{N}{2}, \sigma^2 = N$)

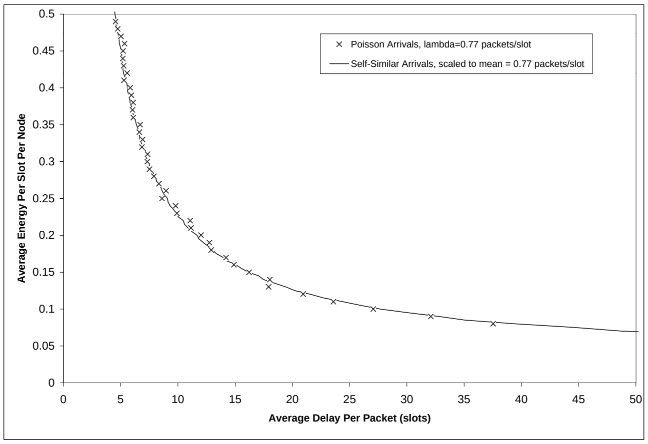

Figure 3: Pseudo-Random Protocol, Gaussian Destination Distribution ($\mu = \frac{N}{2}, \sigma^2 = N$)

IV. RESULTS AND DISCUSSION

Figures 1 through 3 depict the performance of the protocols under two different types of arrival processes. The individual points on the plots depict the energy vs. delay characteristics of the protocols when 100,000 packets are generated for 100 nodes by a self-similar arrival process, while the solid line represents the performance under a Poisson arrival process for the same simulation parameters. The self-similar arrival process consists of 1000 individual packet train sources, each generating packets for a single node. The ON and OFF times, as well as the packet generation rate of the self-similar process have been scaled so the average arrival rate of packets is 0.77 packets per slot, well within the 1 packet per slot capacity bound of our model. We have verified that this rescaling indeed still shows self-similar characteristics, though understandably the emergent self-similar behavior shown in the typical variance and R/S tests begins to occur at a larger aggregated block size.

In the research which experimentally determined network traffic to be self-similar, it was also determined that a large percentage of the traffic was addressed to a small number of destinations. This is an intuitive result as some machines may be routers, servers or simply used by someone with a greater number of network-based applications. It gives validation to the previous interest in protocol performance under clustered traffic shown in the original energy-conserving protocol papers. Therefore, we began with simulations with uniform traffic loads per each node to identify how the burstiness affects the protocols, and then progressed to a Gaussian distribution representing a possibly more typical traffic load. The simulations with uniform traffic load are not shown in this paper, but show nearly identical results.

It is easy to see from the simulation results that self-similar traffic does not have a substantial effect on the average performance of the protocols. This can be intuitively explained by considering the underlying processes that created the self-similar distribution and keeping in mind that a performance change will only occur due to an increased backlog of packets at the queue. Each individual node has relatively few traffic streams associated with it. In each of these streams, single packets arrive according to a fixed rate. For example, in the presence of a uniform traffic distribution with an aggregate rate of 0.77, a single node's packet queue sees 10 individual packet streams where the "ON" state generates arrivals approximately once every 500 slots. It is therefore an extremely rare occurrence for two packets to arrive within a short amount of time for destinations that have with overlapping wakeup schedules (or the same destination). Although the traffic process seen by the base station is self-similar, the protocol performance is mostly determined by the states of the individual mobile node queues at the beginnings of each slot.

Since all of the energy-conserving protocols are based on a cyclical schedule which is very short in comparison to simulation length, we would expect that the performance of the protocols would be most affected with changes in the short-term statistical structure of the traffic. As can be seen visually in previous figures, the short term performance of the self-similar traffic is very similar to that of the Poisson distribution. If we were to allow packet losses due to buffer overflows or arrivals which consisted of more than one packet we may see effects on the average delay characteristics on the protocols by backlogs at the buffers. However, adding these features to our model would change the performance of the Poisson models as well and not necessarily add to the determination of energy-conserving abilities of the protocols. It is also expected that we should be able to handle some of the occasional burstiness simply through dynamic handling of the protocol parameters.

It should be noted that we are interested in considering the *average* delay and energy performance in these plots. Figure 4 depicts the average and maximum queue lengths at the base station for the simulation of Figure 3. These aggregated queue lengths are a reflection of the overall delay performance. We see that a change in traffic arrivals does not change the average queue length, though the maximum queue length has a much wider variation for the self-similar distribution. This is due to the larger-term variation in the arrival process, though since we are within the one arrival per slot capacity bound of the system, the average rates are the same. This reasoning can also be applied to explain why the top-most self-similar data points in Figure 2 do not match the Poisson data points as well as other protocols. These points occur for the directory protocol when $k = 4$. Since this means that, on average, a directory will be broadcast every 5th slot, so for an average arrival rate of 0.77 packets/slot, we are just at the edge of capacity of the system. A long-term rise in the average rate would create backlogs at the queue for extended periods.

We can generally state that that the self-similarity is manifest on a time-scale which is much greater than that of the protocol so that the long-term variations in the arrival process to not substantially effect the protocol performance except in extreme situations. Literature such as [4] has pointed out that the time-scale of self-similarity vs. the protocol under study significantly effects whether long range correlations create changes in performance. In a related work, [6] showed that buffer occupancy of VBR traffic is not affected by the long range dependence if the busy time is not large, and these queues can still be accurately modeled by Markov processes.

V. CONCLUSIONS

The future of wireless packet data traffic cannot be easily estimated due to the constantly changing nature of applications and end-user devices. However, it is reasonable to believe that wireless data services will have demands placed on them which are similar to landline networks. Many landline traffic types and network models exhibit self-similar traffic distributions which show a different statistical properties from traditional Poisson traffic modeling.

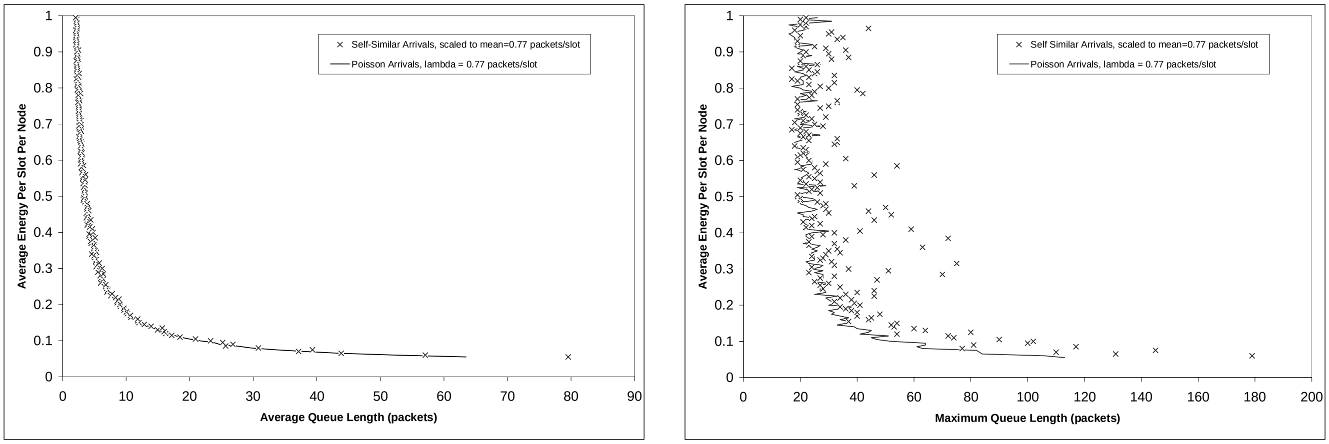

(a) Average Queue Length, Pseudo-Random Protocol, Gaussian Destination Distribution

(b) Maximum Queue Length, Pseudo-Random Protocol, Gaussian Destination Distribution

Figure 4: Comparison of the Average and Maximum Base Station Queue Length Under Poisson and Self-Similar Arrival Processes

We have verified through simulation that self-similar traffic arrivals do not have a substantial effect of the energy vs. delay characteristics of these energy-conserving protocols. The reasons for this are based on small number of multiple packets arriving at the same time for a single queue and the dependence of the protocol only on a small subset of queued packets in each time slot.

The energy-conserving protocols proposed are not complete networking protocols and must be combined with other network and link level layers for a complete mobile system. Although other layers of the networking protocol may be affected by self-similar traffic, it is insightful and adds to the utility of these protocols to know they will not degrade in performance with a change in traffic arrival statistics to what is seen in landline networks.

REFERENCES

[1] Crovella, M.E. and Bestavros, A. Self-similarity in world wide web traffic: evidence and possible causes. *IEEE/ACM Transactions on Networking*, 5(6):835–46, December 1997.

[2] Chlamtac, I., Petrioli, C., and Redi, J. Energy-conserving access protocols for identification networks. *IEEE/ACM Transactions on Networking, February 1999*.

[3] Duffy, D.E., McIntosh, A.A., Rosenstein, M., and Willinger, W. Statistical analysis of CCSN/SS7 traffic data from working CCS subnetworks. *IEEE Journal on Selected Areas in Communications*, April 1994.

[4] Grossglauser, M. and Bolot, J.-C. On the relevance of long-range dependence in network traffic. *Computer Communication Review*, October 1996.

[5] Harris, E.P., Depp, S.W., Pence,E., Kirkpatrick, S., Sri-Jayantha, M., and Troutman, R.R. Technology directions for portable computers. *Proceedings of the IEEE*, April 1995.

[6] Heyman, D.P. and Lakshman, T.V. What are the implications of long-range dependence for VBR-video traffic engineering? *IEEE/ACM Transactions on Networking*, June 1996.

[7] Leland, W.E., Taqqu, M.S., Willinger, W., and Wilson, D.V. On the self-similar nature of ethernet traffic (extended version). *IEEE/ACM Transactions on Networking*, Feb 1994.

[8] Paxson, V. and Floyd, S. Wide-area traffic: The failure of poisson modeling. *Computer Communication Review*, 1994.

[9] Taqqu, M. A bibliographical guide to self-similar processes and long-range dependence. In E. Eberlein and M. S. Taqqu, editors, *Dependence in Probability and Statistics*, pages 137–165. Birkhauser, 1985.

[10] Taqqu, M.S., Willinger, W., and Sherman, R. Proof of a fundamental result in self-similar traffic modeling. *Computer Communication Review*, April 1997.

[11] Willinger, W., Taqqu, M.S., Sherman, R., and Wilson, D.V. Self-similarity through high-variability: statistical analysis of ethernet lan traffic at the source level. *IEEE/ACM Transactions on Networking*, February 1997.