

SPARTA : Scheduled Power and Rate Adaptation

Ramana Rao Kompella and Alex C. Snoeren
University of California, San Diego
{ramana,snoeren}@cs.ucsd.edu

Abstract

Experience has shown that the power consumption of wireless devices is often dominated by their communication patterns. Shannon's law suggests that the energy needed to transmit a fixed amount of information is convex with respect to the rate of the transmission. In other words, the energy required to transmit a packet can be reduced by transmitting the packet at a lower rate. Our approach uses this observation to dynamically vary the transmission power and rate of transmission on a per-packet basis to sustain the input traffic rate, thus saving transmit power over the naive scheme of sending packets at the maximum possible transmit rate. We call this Scheduled Power and Rate Adaptation or SPARTA in short. Previous approaches either adopted static Transmission Power Control or operate in power savings mode to conserve battery resources. We show by simulations that SPARTA achieves upto 50% power savings compared to other popular previous schemes for wireless LANs.

We cannot, however, lower the transmission rate lower than the input rate. Doing so would result in building up of the queues leading to system instability. In SPARTA, we estimate the input traffic rate and choose the transmission rate as close to the input rate as possible. In general, it is not possible to predict the input traffic rate since it depends entirely on the set of applications that a user is running. We estimate the input traffic rate based on a lookahead buffer, that collects packets in an interval and schedule those packets at an optimal rate in the next interval in a pipelined fashion. This introduces a small delay into the overall latency of the TCP packets. By tuning the lookahead, one can control the amount of impact on the Round Trip Time of the TCP connection.

We simulated an FTP transfer of 10MBytes from a mobile host to a server. Our results indicate that data could be transferred with 30% less power as compared to the case when data is transmitted at the highest possible power level. Further, by combining SPARTA with power savings schemes such as PSM and BSD, SPARTA consumes 50% less power than when there is no power management. Further, we also show that the time required to transmit the file is only 10% higher than the case when only PSM or BSD are enabled without SPARTA. This shows that our algorithm does not effect TCP throughput by a great amount. Also we show that by tuning the delay Lookahead value, one can trade-off energy for throughput. Higher values of the lookahead results in increased energy efficiency but also suffers from lower throughput.

In multi-node scenarios, however, sending a packet arbitrarily slowly can cause overall loss in throughput. All the nodes use a distributed coordination algorithm (both in ad-hoc mode as well as in infrastructure mode) to estimate load on the network and transmit at a rate optimal for global throughput. More specifically, the nodes snoop on the wireless channel to record the transmission time of a packet from a particular node. This transmission time is used to estimate the load on the network assuming that all the nodes are well behaved.

In the infrastructure mode, each node indicates to the base station its current load and in-turn obtains the total load on the network from the base station. This is factored in to calculate the appropriate transmission rate for this particular node to achieve power savings while being fair to other nodes. While this scheme has inefficiencies, (for example, it takes time for the distributed algorithm to converge to a final estimate, and during that time, some nodes might send more traffic than their fair share), it would still provide a reasonable estimate on the overall load on the network allowing each node to tune their transmission speeds to achieve power savings while not hurting overall throughput.

SPARTA can be easily incrementally deployed in an existing network with non-conforming nodes already present. In the ad-hoc mode, the conforming nodes use a pseudo-header to indicate that they follow the protocol, whereby only the information from these nodes is used for the purposes of estimation of the overall load on the network. The rest of the nodes are appropriately extrapolated to factor in their load on the network. In the infrastructure mode, the base station can identify which nodes are conforming, depending on which suitable extrapolation can be used in the calculation of the overall load on the network.