Hot-Spot Congestion Relief and Service Guarantees in Public-Area Wireless Networks

Anand Balachandran
University of California, San Diego
anandb@cs.ucsd.edu

Paramvir Bahl
Microsoft Research
bahl@microsoft.com

Geoffrey M. Voelker
University of California, San Diego
voelker@cs.ucsd.edu

Wireless LAN administrators are often called upon to deal with the problem of sporadic user congestion at certain popular spaces ("hot-spots") within the network. To address this problem we describe and evaluate three algorithms that provide hot-spot congestion relief while maintaining users' pre-negotiated bandwidth agreement with the network. The goal of these algorithms is: (1) to accommodate more users by dynamically providing capacity where it is needed, when it is needed; (2) to improve network utilization; and (3) to guarantee at least a minimum amount of bandwidth to users.

An important aspect of designing public-area wireless networks is capacity planning, making the best use of the available network resources to derive the best return on investment while satisfying user service demands. Recent studies of deployments of wireless LANs in public areas have shown that user service requirements are highly dynamic in terms of both time of day and location, and that user load is often distributed quite unevenly among wireless access points (APs). Users tend to localize themselves in particular areas of the network, either due to power-related constraints (i.e., the availability of a fixed power supply), or due to proximity of other services (e.g. gate areas near departing and arriving flights in an airport). These user concentrations create "hot-spots" in the network and complicates the capacity-planning problem, making it difficult to accommodate heavy, concentrated load in different parts of the network without significant, and costly, over-engineering.

Recently, various vendors of wireless LAN products have incorporated load-balancing features in the latest release of network drivers and firmware for APs and wireless PC cards. APs supporting this feature maintain a measurement of the load in their respective cells (number of associated users) and provide it to the mobile stations, which in turn use this information during association with the network. However, these techniques do not take into account explicit user service (QoS) requirements and are local in scope, distributing users only across available overlapping cells.

We propose that both the network and its users should explicitly and cooperatively adapt themselves to changing load conditions depending on their geographic location within the network. When a user requests service from the network in an overloaded region, the network tries to adapt itself to handle the user service request by readjusting the load across its APs. If the network cannot adapt itself to handle the user's request, it provides feedback to the user about where the user can move to get the service requested. As a result, overall network utilization increases, and users get the QoS they request, either transparently or by explicitly moving to specific locations within the network.

In this poster, we describe and evaluate three algorithms for providing hot-spot congestion relief while maintaining pre-negotiated user bandwidth agreements with the network. With *explicit channel switching*, the network trades off signal strength with load by forcing a mobile node to switch from an overloaded AP to a neighboring lightly loaded AP. With *network-directed roaming*, the network balances load by giving explicit feedback to users about nearby locations the user can roam to in order to get the requested service. Lastly, with *wireless cell breathing* the network balances the load, by adjusting AP transmit power among neighboring cells to reconfigure cell boundaries. Reconfiguring cell boundaries can re-target mobile nodes in overlapping cell regions from an overloaded AP to a lightly loaded AP.

We evaluate the benefit of the load-balancing algorithms on network utilization, using simulations of a public-area wireless network. The simulation results show that our algorithms perform well in a variety of user configurations. We use a parameter called balance index to evaluate the extent of balance achieved between the cells in the network. Our algorithms improve the balance index by over 30% in comparison to existing schemes that offer little or no load balancing. Based upon these results, we conclude that public-area wireless networks would benefit greatly from the use of these algorithms.