A Bandwidth Advisor for Wireless Networks

Glenn Judd and Peter Steenkiste
Computer Science Department
Carnegie Mellon University
http://www.cs.cmu.edu/~aura

Introduction

Many network intensive applications can benefit substantially from reasonable estimates of future available bandwidth. As a simple example, consider a user, John, that is about to board a flight that is leaving shortly. Using his palmtop equipped with the Aura ubiquitous computing system, John has just finished editing some very large documents that he needs to email over the airport's wireless network. Just before he begins transmission, a 747 begins deplaning passengers at a gate adjacent to where John is waiting; these deplaning passengers are heavily loading the network in John's cell. John's Aura client contacts the bandwidth advisor and learns that there is insufficient bandwidth for John to complete his transmission before his plane leaves. Fortunately, Aura learns from the bandwidth advisor that a nearby cell has ample bandwidth which enables John to move into that cell briefly and complete his transmission before catching his plane.

The utility of bandwidth advisors has lead to efforts such as Remos and the Network Weather Service which have investigated providing bandwidth advice in wired network settings. This research builds on this previous work by developing a bandwidth advisor that is specifically tuned for wireless networks, and is suited to the needs of mobile applications.

Utilization Prediction

The first component of the bandwidth advisor we developed is a service capable of predicting bandwidth utilization within each cell of the network. We began by collecting bandwidth utilization, via SNMP, from an active wireless network: CMU's campus-wide wireless LAN. Initial inspection of the data confirmed many of our intuitive suspicions: many cells were almost entirely idle while others were heavily utilized; utilization tends to change at hourly intervals and be correlated with the same interval in other weeks. Analysis of autocorrelation of cell utilization over time revealed that cells with high utilization exhibited significant amounts of autocorrelation; this gave us reason to believe that linear predictive models could be useful in predicting future utilization.

We then examined the performance of the following three predictive models at predicting one sample ahead in our traces:
PPREV – predicts future values to be the same as the most recent value observed.
BM – uses an evenly weighted average of the several previous observations.
ARFIMA – computes future values as a dynamically weighted average of past values.

Our results show PPREV and ARFIMA to be superior to BM. Both PPREV and ARFIMA achieve similar error rates, but in differing ways: ARFIMA does a good job predicting background traffic, but does so at the cost of missing higher values; PPREV does a much better job at predicting higher values, but does so at the cost of more overpredictions.

Throughput Prediction

The second component we developed takes as input a utilization estimate for a cell as well as a list of devices currently occupying the cell. It uses this information to estimate the potential maximum throughput a new flow could expect in this cell.

Currently, the accuracy of our throughput estimates is limited by the inability of the access points to provide information about station connection rates as well as how much of the observed traffic was sent at each rate.

Future Work

Our future work will examine refining our utilization predictions by dynamically selecting a utilization predictor for a query based on past performance on similar queries. We will also investigate adding functionality to access points in order to provide sufficient information to make accurate throughput predictions. Lastly, we will examine the possibility of more proactive management of network bandwidth via cooperation between network aware applications and access points.