Research Statement

Vivek Shrivastava

As wireless devices continue to grow in offices and enterprise environments, enterprise wireless LANs (WLANs) have emerged as an important part of an enterprise network. Over the years, network administrators have developed various sophisticated tools and management strategies to run their enterprise networks. Examples include various traffic shaping entities that classify traffic and ensure bandwidth usage limits for different classes. While such solutions in the wired enterprise are relatively well defined, it is not clear whether such tools can be easily extended to the wireless component.

A major challenge in managing bandwidth resources in the wireless domain is the existence of multiple wireless collision domains, each of which interact with a few others. More importantly, the popular 802.11 standards use significant decentralization in resolving medium access across these collision domains. Thus, even if it is possible to specify usage policies of various application classes across the entire enterprise WLAN, implementing such policies and ensuring appropriate sharing of wireless resources to incumbents is a fairly hard problem — a hardness that stems from independent actions of medium access by multiple entities make global policies hard to design and enforce. I believe that centralization of WLANs can be leveraged in solving this problem by providing a fine grained control on different available resources in the system. The thesis of this research is that centralized management of the data and control planes for 802.11 based enterprise WLANs is both necessary and desirable in achieving efficient management of these networks.

Let me explain this further with a simple example. Consider two Access Points (APs), each with one wireless client, located in a manner that the APs cannot sense the presence of each other, i.e., they are outside mutual carrier sense ranges. However, the clients are in carrier sense range of each other (and hence, in mutual interference range). Let us assume that the traffic is downstream along both these links, i.e., from the APs to the clients. It is easy to demonstrate that in such a scenario, depending on traffic patterns, one AP-client traffic can completely starve the other AP-client traffic, inspite of deploying existing QoS mechanisms, such as the 802.11e standards. Efficiently managing the common resource — the wireless medium — under completely decentralized access mechanisms is hard. The goal of this work is to study the challenges in efficiently managing wireless spectrum and designing robust traffic engineering techniques in an enterprise-wide wireless network setting. I exploit the inherent centralized architecture of an enterprise WLAN where a central control element is wired to all the Access Points (APs) in the wireless network. Such a central network element (Controller [1]) provides a natural platform to centrally configure and globally optimize channels and power levels at the APs. Further, the Controller, has a clear view of all downstream traffic [2] that will be transmitted across the entire enterprise wireless LAN. If the Controller can infer something about different interference domains in the wireless network, it can perform data plane optimizations like centralized scheduling, to meet desired objectives of traffic engineering.

Based on these observations, my work investigated both control plane and data plane mechanisms to enhance the performance of enterprise WLANs. As a part of our efforts in the data plane, we have explored the design of a centralized scheduler that leverages the vantage point of the Controller in intelligently scheduling downlink traffic. On the control plane side, I designed and implemented an interference detection and power control mechanism to aid fine grained management of wireless resources by the Controller. In designing these mechanisms, I have assumed no client support, so that they can be implemented by any individual WLAN vendor, by applying software updates to the Controller and the APs. I believe that this property makes my solutions attractive for immediate deployment and adoption by wireless vendor, allowing them to have greater practical impact on the design of enterprise wireless solutions. Below, I discuss my most important contributions in building enterprise wireless solutions, both for the control plane and the data plane.

Centralized Data Plane

While it is natural for control plane mechanisms in enterprise WLANs to be centralized in nature, it is not immediately obvious whether data plane mechanisms, i.e., channel contention and access for competing transmitters, should also be centralized. However, in order to facilitate the convergence of different services like voice, data and video in a single WLAN, there needs

[1] I refer to the central control element as Controller in the rest of this document.

[2] About 80% of total enterprise traffic is downstream (from the Internet to the wireless clients).

to be control on the data plane of the wireless medium to enforce multiple service levels, bandwidth contracts, traffic shaping and spectrum utilization. A major challenge in this direction is for the system to scalably identify all potential interfering link groups within the enterprise WLAN, keeping in mind that mobility will lead to frequent changes to these groups. A second challenge arises in the issue of time-scales of operation. Although any centralized scheduling solution can operate and adapt at the order of a few milliseconds or more, some interference effects that occur on wireless channels needs adaptations at the sub millisecond level. Below, I discuss how my proposed solutions overcome these challenges:

Interference estimation in enterprise WLANs

As discussed above, accurate, fast and scalable mechanism for detecting potentially interfering links is critical for realizing efficient data and control plane mechanisms. Such a tool for interference estimation can enable WLAN managers to improve the network performance by dynamically adjusting operating parameters like channel of operation and transmit power of access points, but also diagnose and potentially proactively fix problems. Prior work on interference estimation employ active probing techniques and suffer from three main problems: a) it incurs moderate to significant measurement overhead and cannot be employed to continuously obtain interference information as they evolve over time, b) it offers limited visibility into the root cause of interference, c) it often requires specific client support.

Motivated by these observations, I designed and implemented a Passive Interference Estimator (PIE) that can dynamically generate fine grained interference estimates across an entire WLAN. The most attractive feature of PIE is that it imposes no measurement traffic, and yet provides an accurate estimate of WLAN interference as it changes with client mobility, dynamic traffic loads and varying channel conditions. Our experiments conducted on on two different testbeds show that PIE is able to not only provide high accuracy but also operate beyond the limitations of prior tools, providing a true solution to performance diagnosis and real time WLAN optimization, as manifested through its use in multiple WLAN optimization applications, namely channel assignment, transmit power control, and data scheduling. This work is currently under submission.

Implementing centralized scheduling

Enterprise WLANs have made a dramatic shift towards centralized architectures in the recent past. The reasons for such a change have been ease of management and better design of various control and security functions. The data path of WLANs, however, continues to use the distributed, random-access model, as defined by the popular DCF mechanism of the 802.11 standard. While theoretical results indicate that a centrally scheduled data path can achieve higher efficiency than its distributed counterpart, the likely complexity of such a solution has inhibited practical consideration. In this paper, we take a fresh, implementation and deployment oriented, view in understanding data path choices in enterprise WLANs. I have performed extensive measurements to characterize the impact of various design choices, like scheduling granularity on the performance of a centralized scheduler, and identify regions where such a centralized scheduler can provide the best gains.

My detailed evaluation with scheduling prototypes deployed on two different wireless testbeds indicates that DCF is quite robust in many scenarios, but centralization can play a unique role in 1) mitigating hidden terminals — scenarios which may occur infrequently, but become pain points when they do and 2) exploiting exposed terminals – scenarios which occur more frequently, and limit the potential of successful concurrent transmissions. Motivated by these results, we design and implement CENTAUR – a hybrid data path for enterprise WLANs, that combines the simplicity and ease of DCF with a limited amount of centralized scheduling from a unique vantage point. This work appeared in Mobicom 2009, where is received promising reviews and received the best paper award.

Centralized control plane

The centralized Controller is at a unique vantage point where it has a global view of the entire WLAN. This can be used to identify different wireless contention domains and could also be leveraged to implement smart centralized mechanisms that configure the power levels and wireless channels of the APs to enforce administrator defined system wide policies. I have explored the following two control plane mechanisms:

Measurement-driven Power Control

In order to provide multiple service levels in enterprise WLANs, we investigated the feasibility of using fine grained power control on a per client basis that can enable the centralized controller to provide desired service (bit-rate) to any particular client by adjusting the power levels for it's corresponding AP. However, we observe that multipath, fading, shadowing and external interference from wireless devices, make the implementation of fine grained power control challenging in practical settings.

The initial focus of our work was to determine the right set of power levels that are practically usable in indoor power control mechanisms.

My work suggests that a few, 3-5, discrete power level choices is sufficient to implement any robust power control mechanism in typical indoor WLAN environments. Through our work, we also build an empirical model that determines these appropriate number and choices of power values that is adequate in any setting. On the basis of this work, I implemented a joint power-rate adaptation mechanism and showed that using the model allows the Controller to quickly converge to the desirable power settings, providing significant gains in client mobility scenarios. Our findings from this work was accepted for publication in IMC 2007 and also won the first prize in the ACM Student Research Competition at ACM Mobicom 2006.

Channel Management

As a part of our research in centralized control plane for enterprise WLANs, we have worked on the idea of exploiting partially overlapped channels in 802.11 standards to improve throughput in wireless networks. I showed that by efficiently assigning partially overlapped channels to dense access point deployments in WLANs, we can get a factor of 3 improvement in throughput (published in ACM Sigmetrics 2006). As a part of this work, we have also provided a detailed interference and throughput model for 802.11 standards, parametrized by the channel separation and applicable to the host of 802.11 standards (a/b/g).

I am currently studying channel hopping in centralized WLANs, where the controller can facilitate very fine grained synchronization between Access Points (APs) and enable fast switching of channels, which is not only useful for fairness properties but also provides robust security against malicious jamming attacks. I have evaluated the performance of channel hopping mechanisms in unplanned wireless hotspots, where we show that channel hopping enables efficient spectrum sharing with minimal coordination (published in ACM Mobicom 2006). Summarizing, we find that careful interference modeling between overlapped channels, coupled with channel hopping mechanisms can enhance throughput and fairness properties of WLANs and our prototype implementation shows that both these mechanisms can be efficiently implemented using a centralized controller.

Other work

Beyond this, I believe that my work has equipped me with requisite tools to investigate problems that are outside my core interests. My projects have spanned both wireless networks like city-wide wireless mesh, centralized WLANs and wired networks like production data centers and peer-to-peer systems. Next, I briefly describe my work in these related fields.

Context aware VM migration in production datacenters

Driven by practical requirements of a production data center, we propose the concept of context based virtual machine (VM) migration. The underlying observation is that VM migration in a production data center must take into account the context of the application which the VM is running to ensure migration is useful and serves the overall goals of the data center. Examples of such contexts are 1. Application dependencies between different VMs 2. Spatial and temporal locality of resource requirements 3. Application specific load patterns. Such context allows us to choose the virtual machine to physical machine mapping more intelligently and improve resource utilization under heavy loads. Our results from this study is currently under submission.

Natural incentives in peer-to-peer streaming

Success of peer-to-peer applications in many cases is attributed to user altruism, where a user contributes some of its own resources to facilitate performance of other users. In this work we demonstrated that there are many scenarios where peer-to-peer resource sharing is a natural behavior that selfish users can use to improve their own performance. In particular we examine such natural incentives that exist in a streaming media application which lead such greedy users to cooperate and share resources with each other in forming an efficient overlay multicast tree. Our work in this area was accepted for publication at NOSSDAV 2005.

Analysis of city wide commercial mesh network

Wireless Mesh Networks (WMNs) represent a new and promising paradigm that uses multi-hop communications to extend WiFi networks. Several WMNs are already deployed and operational. Since the clients pay the same flat rate, the throughput sharing should also be fair. In this project we are working closely with Madcity Networks, that has deployed a city wide mesh network comprising of 750 access points in Madison,WI. We are working closely with the company to identify the performance bottlenecks and other practical implementation challenges. Our analysis from this measurement study appeared at IMC 2008.

Future directions

High-level policy design

The next challenging aspect of my work is in facilitating easy policy design. Prior work in policy design has primarily been in the areas of i) access/security control, and ii) configuration diagnosis and management. This work will examine a more active approach towards policy-driven performance management. Defining performance policies through low-level constructs of traffic components can be fairly tedious and administrators will benefit if the policies can be expressed in a high-level language. Such a design will, however, require an automated mechanism to convert such high level policies into traffic aggregates by inferring behavioral properties of the traffic. How can such mechanisms be facilitated? Work by Karagiannis et. al. (Sigcomm 2005) illustrates an approach to behavioral traffic classification in the Internet, and can be a good starting point for one component of this question. I also believe that work on network policies at IBM Research by Agrawal et. al. at Policy 2005 and IM 2005 has interesting synergies to this component.

Network vs user performance

While most of the above discussion assumes that there are well-defined network-wide objectives which the centralized scheduling system tries to optimize, we will also study implications of user-specified policies. For example, if one of the goals of different users is to achieve conservation of battery power through the centralized scheduling system, can the scheduling algorithm be designed to make the traffic inherently bursty in nature. Such bursty traffic will allow clients the opportunity to power down their wireless NICs when no traffic is targeted towards them. I believe that this work can leverage the advances in Power-Aware Proxy design (for example, that was done at IBM Research by Rosu et. al.) and integrate them with our objectives of a centrally scheduled system.

Summarizing, I believe that the diversity in my research experience will hold me in good stead for working with wide range of networked systems. I have enjoyed working on both wired and wireless systems and the implementation and measurement driven approach of my work has given me a good insight into the working of practical systems. My work will be geared towards building better networked systems with a focus on improving their performance by identifying key bottlenecks and designing elegant yet practical solutions to build a high-performance, deployable system.