# Research Statement

Ashish Patro (http://cs.wisc.edu/~patro)

In recent years, there has been a rapid growth in the deployment of wireless networks such as 3G and 4G cellular networks (e.g. HSPA, EVDO, LTE) and WiFi networks. Simultaneously, there has been a rapid growth in the use of wireless-enabled devices such as laptops, netbooks, smartphones and tablets. As a consequence of these developments, the complex interactions resulting from the diversity of applications, devices, wireless networks and deployments in this ecosystem cause many problems that impact the end user's experience. My graduate research has been in the area of mobile and wireless networking and has been motivated by the following questions: (i) *With the rapid proliferation of wireless networks and devices, how can we understand the impact of network and device related factors on application experience?* and (ii) *How can we leverage various "vantage points" to estimate and manage various RF, network and application level parameters in the wild?*

To address these questions, my research work has primarily focused on:

1. Building tools to measure low layer PHY and MAC statistics, detect and quantify wireless interference (Airshark – IMC 2011, WiFiNet – NSDI 2012 and WiSense – HotMobile 2015).
2. Building frameworks for home wireless Access Points that passively analyze wireless activity to predict performance (WiSe – MobiCom 2013) and mitigate problems (COAP – MobiArch 2014, INFOCOMM 2015).
3. Building frameworks that leverage mobile applications as vantage points to capture diverse network and client device statistics and measure their impact on application usage and revenues (Insight – MobiGames 2012, CoNEXT 2013).

## RESEARCH APPROACH

To identify and tackle important research questions in a new area, my approach is to first understand the recent advances as well as collaborate with others working in the specific area to gather useful insights. This also allows me to benefit from multiple perspectives and identify the important questions. For example, I collaborated with developers of two mobile applications over 3 years for the Insight project that studied impact of device and network characteristics on the "application experience" (e.g., application revenues, battery overhead, user behavior and retention). This effort was beneficial to the developers in gaining deeper understanding about application usage as well as provided valuable insights to the research community. Similarly, I worked with both users and local ISPs for the WiSe project to perform a rigorous study about wireless performance in dense residential environments (e.g., multi-tenant apartment buildings).

I also enjoy building systems and have experienced that coupling academic rigor with expending effort to develop robust frameworks and prototypes yields rich research dividends. When we started working on the problem to diagnose non-WiFi device interference in WiFi deployments, the state of the art consisted of standalone, proprietary and expensive spectrum analyzers. To empower WiFi devices and networks detect and mitigate this interference, we decided to build non-WiFi device detection capabilities over commodity WiFi cards. But, these cards provided much lower resolution data compared to dedicated spectrum analyzers (e.g., 312 kHz resolution bandwidth for WiFi cards vs. 1 kHz for spectrum analyzers). We built a system called Airshark using commodity WiFi cards that provided high non-WiFi device detection accuracy comparable to the standalone spectrum analyzers. It required extensive measurements, modeling and evaluation to build robust non-WiFi device features for Airshark. My emphasis on building robust systems later allowed me to deploy Airshark across 30 residential WiFi Access Points for more than 9 months. This effort provided valuable research data about real-world impact of non-WiFi interference. Then, I built an Android based platform (WiSense[1]) incorporating Airshark. This lead to broader visibility of the project in the research community as well as industry (featured by Google on the Play Store).

---

1 Website: http://www.wisense.io

## RESEARCH RELATED TO DISSERTATION

Following is a brief summary about my work on diagnosing wireless interference and building frameworks that leverage various vantage points to study impact of networks and devices on application experience.

### 1. Systems to detect and quantify wireless interference

**Airshark:** Interference is one of the major performance issues experienced by WiFi networks due to their operation in the unlicensed 2.4 GHz and 5 GHz spectrums. WiFi networks coexist with diverse non-WiFi devices such as cordless phones, microwave ovens, Bluetooth devices and game controllers. We developed Airshark, which is the first system to perform real-time detection of such non-WiFi interferers using off-the-shelf commodity WiFi cards. WiFi cards do not have the capability to decode non-WiFi transmissions making it challenging to detect such interference. We used coarse-grained signal information reported by new-generation WiFi cards and incorporated machine learning techniques to build device specific signatures. Airshark achieved results comparable to specialized spectrum analyzers, with > 90% detection accuracy and low false positives. This work was published in **IMC 2011** and also received **widespread coverage in the press and numerous technical journals** (Slashdot, CRA research highlights). I also built a full prototype and won the **second prize** at the **InterDigital Innovation Challenge 2012**. Currently, we are working with equipment vendors to incorporate Airshark into their products.

**WiFiNet**: We then built WiFiNet, which utilizes multiple Airshark enabled APs to detect, localize and quantify non-WiFi interferers' impact on neighboring WiFi links. Using statistical and machine learning techniques, WiFiNet can distinguish and segregate device transmissions and interference impact of multiple non-WiFi interferers, even of the same type (e.g., two cordless phones). This is much harder than trivial case of segregating WiFi transmitters, since only energy level information is available to commodity WiFi cards about non-WiFi transmissions. WiFiNet can also localize non-WiFi interferers with a median localization error < 4 meters. WiFiNet was published at **NSDI 2012**.

**WiSense**: WiSense, an Android based platform for debugging wireless networks is my most recent work. I augmented the Android kernel to support device drivers required for Airshark, enabling users to leverage mobility of smartphones and tablets (e.g., Nexus phones) to debug both enterprise and home WiFi deployments. WiSense uses an off-the-shelf external USB WiFi card to collect spectrum data for Airshark. A group of WiSense enabled clients act as a distributed sensing platform to build a wireless performance map of a building. I released the kernel patches for driver support and Linux ath9k driver developers accepted them. An initial version of this work will appear at **HotMobile 2015**. Furthermore, **Google has featured the WiSense Android application on the Play Store.**

### 2. Using home Access Points as vantage points to analyze and manage residential WiFi networks

In residential areas, users are mainly dependent on a single WiFi Access Point (AP) as their gateway to the Internet for most of their Internet-enabled devices. Since all wireless traffic at homes pass through these Access Points, it provides an ideal vantage point to monitor as well as manage the traffic of all the wireless links within a home.

**WiSe**: I implemented a framework called WiSe and performed **a 9 month deployment** of OpenWrt based APs (built using off-the-shelf components for the ALIX platform) **in 30 homes**. WiSe uses home APs to collect a diverse set of statistics at the PHY layer (airtime utilization, non-WiFi activity), MAC layer (link quality, local and external WiFi activity) as well as coarse-grained flow level statistics. I incorporated passive techniques into WiSe to measure co-channel WiFi and non-WiFi interference. Using extensive measurements and regression techniques, I also developed a lightweight and passive performance metric called Witt. Witt uses PHY and MAC layer data to predict expected TCP throughput over the wireless link under current conditions. This provides real-time feedback to both administrators and users about the expected throughput without active measurements. An example scenario is — During periods with 60% channel utilization and 30% MAC losses, what is the expected TCP throughput of link X if it starts a bulk download flow right now? Through our deployment, we gained numerous insights about residential wireless experience: (i) While a majority of links performed well, poor WiFi performance occurred during 2.1% of the time, (ii) Due to bursty nature of WiFi and non-WiFi activity, WiFi links experienced short (1 - 10 minutes) periods of interference but upto 80% degradation, and (iii) More than 55% of the observed 305 external APs used a single static WiFi channel, leading to inefficient spectrum utilization. WiSe was published at **MobiCom 2013** and the datasets were made public.

**COAP:** My study motivated the need to mitigate performance issues, reduce interference and improve spectrum utilization in dense residential WiFi deployments. I developed the COAP framework to enable cloud based centralized management of neighboring home WiFi APs. COAP uses the Software Defined Networking (SDN) paradigm to propose software only upgrades and open APIs that can be implemented by Access Point (AP) vendors to enable cooperation and coordination between nearby home APs. The controller provides all necessary management services that can be operated by a third-party (potentially distinct from the individual ISPs). The framework is implemented using OpenFlow extensions and allows COAP APs to share various types of information with the centralized controller — interference (e.g., non-WiFi activity) and traffic phenomenon and various flow contexts, and in turn receive instructions — configuration parameters (e.g., channel) and transmission parameters (through coarse-grained schedules and throttling parameters. Through a deployment of 12 COAP APs in a single apartment building, we observed reduced channel congestion upto 47%. The initial work was published at **MobiArch 2014** and **received the best presentation award**. I also presented demos at **Open Networking Summit 2014**, an industry conference and received positive feedback. The full paper will appear at **INFOCOMM 2015**.

### 3. Understanding mobile application experience in the wild

**Insight:** On mobile devices, the combination of cellular/WiFi network performance, device characteristics, application UI and performance overhead impacts overall application experience. The application itself provides a unique measurement vantage point since it allows us to correlate the aforementioned parameters with the user's activities within the application. I implemented the Insight framework as a library for mobile application developers to collect application analytics as well as perform lightweight active network measurements during application usage. I collaborated with developers to deploy Insight within 2 popular mobile applications — a MMORPG (Parallel Kingdom) and a study application (StudyBlue). Insight was deployed for **more than 3 years** across more than **1,000,000 unique users**. Insight provided interesting observations such as: (i) Poor network performance (latency > 900ms) led to 52% drop in application revenues, (ii) The device form factor (e.g., screen size, slide-out keyboard vs. touch screen) resulted in upto 2x variation in user interactivity. A study about the MMORPG won the **best paper award** at **MobiGames 2012.** Insight was published at **CoNEXT 2013** and was a **best paper nominee (top 4 papers).**

## OTHER RESEARCH

I had the opportunity to collaborate with great mentors and work on interesting projects during my internships. It also enabled me to learn skills required to tackle research problems in areas outside my expertise.

### 1. Indoor localization algorithms for the Android platform

At Google, I worked on building algorithms for WiFi based indoor localization using crowd-sourced data from Android devices. The data comprised of WiFi scans with partial GPS location information (e.g., outside building entrances). The data was very noisy due to the diversity of Android devices and required careful modeling to reduce localization error. The models involved probabilistic techniques such as Maximum Likelihood Estimation to compute devices' location.

### 2. A platform to enable better and smarter driving

At Microsoft Research, I worked on a system to combine vehicular and smartphone sensor data with cloud services to enable better and smarter driving behavior. We used commodity OBD devices to collect sensor data from vehicles (e.g., RPM, engine load, fuel usage), built and deployed a smartphone application to collect driving data from around 20 drivers over 8 months (> 4400 miles). We used mechanical engineering theory and machine learning techniques to develop models that use vehicle and road features to predict fuel consumption before a trip and well as analyze the impact of driving behavior on fuel consumption after a trip. This work is under submission.

## FUTURE DIRECTIONS

My past work on wireless networks and mobile computing led to the development of useful tools to diagnose wireless interference and understand the impact of network and device characteristics on user experience. The major future trends in this space include rapid adoption of diverse network connected devices (e.g., wearables, vehicles, smart routers and home appliances) that support multiple wireless protocols (e.g., NFC, ZigBee, LiFi, DSRC), rapid growth in cellular/WiFi data usage by applications, spectrum allocation problems and smart-sensors (e.g., heart-rate monitors, home energy monitors). These trends throw up a lot of interesting questions. In this broad space, I'm interested in exploring and collaborating on the following research directions.

**Client-based distributed spectrum monitoring Infrastructure:** WiSense used commodity client devices and WiFi cards as a distributed platform for monitoring WiFi network performance. The next step is building a low cost and general-purpose framework to monitor much broader spectrum bands (30 - 7.5 GHz). One way is to attach a frequency translator frontend to external commodity WiFi cards connected to smartphones and tablets. Many important applications can be enabled by such a framework — a distributed sensing platform to manage TV whitespaces, monitor frequency bands to enable policy makers make informed spectrum allocation decisions, and policing the spectrum to identify unauthorized transmitters. Many research challenges need to be tackled to realize this vision, such as dealing with coarse granularity of reported spectrum data, combating limited sensing range with intelligent distributed sensing, energy management and creating policy frameworks to make sensing decisions.

**Building smarter wireless gateways:** With COAP, I presented the idea that home WiFi Access Points or gateways can be centrally managed to improve wireless experience through co-operation and co-ordination. These gateways also have general-purpose compute capabilities in-addition to rich data about wireless activity (WiFi and non-WiFi) which has been not been tapped into yet. I plan to explore new applications that can leverage the wireless gateways as a vantage point. The upcoming generation of WiFi gateways will have additional built-in radios such as LTE, Bluetooth and ZigBee to support Internet of Things (IoT) devices (e.g., temperature sensors, energy monitors). These trends present opportunities to use wireless gateways as a hub for creating novel applications: (a) context-aware applications for home energy management and (b) location-aware applications that leverage information about user and device activity. If wireless gateways can estimate users' real-time locations within homes by tracking their WiFi activity and/or wearable devices, they can assist in optimizing energy consumption in homes. For example, they can make decisions like turning off the lights and controlling heating activity based on user location. Health monitoring is another interesting application. Information about users' movements and network activity available at the gateways can complement other data sources (e.g., wearables and smartphone sensors).

**Improving multimedia experience in wireless networks:** For the WiSe project, I developed the Witt metric to predict estimated TCP throughput using passively observed wireless statistics from home Access Points. New techniques need to be developed that can leverage these statistics to predict application specific performance. For example, the consumption of multimedia services such as over-the-top video (e.g., Netflix) and IPTV services (e.g., AT&T U-verse) is increasing rapidly. I plan to explore how can we measure the impact of home network performance (wired and wireless) on the application related parameters that impact user-experience. In the case of video streaming applications, the main challenge is to estimate different Quality of Experience (QoE) metrics such as the playout-buffer size, buffering periods, start-up times using only traffic properties captured at the AP without any instrumentation at the clients (due to the heterogeneous platforms and devices involved). Using the Access Point or wireless gateways can also help isolate the causes of performance degradation due to problems occurring on the wired ISP path versus the wireless hop. This mechanism can be helpful for ISPs to estimate and improve the QoE experienced by its users across different services (e.g., VOIP, video, gaming).