RESEARCH STATEMENT

Arani Bhattacharya (arbhattachar@cs.stonybrook.edu)

The advent of smaller devices due to the evolution of the Internet of Things (IoT) is making applications of wireless networks more complex. My research focuses on improving adaptivity of wireless networks to this increasing complexity using algorithms and data-driven techniques. I have worked on multiple applications in wireless and mobile applications, encompassing physical, network as well as application layers of the network protocol stack. At the physical layer, my work focuses on monitoring of RF spectrum, which is increasingly becoming a scarce resource due to the increase in amount of data transmitted and the number of devices. I have also worked on making the network layer robust to unpredictable bursts of requests, which has become relevant due to the increasing popularity of smaller IoT devices. At the application layer, many of the smaller IoT devices have low-powered processors and find it challenging to run their applications. I focused on reducing both the running time and energy consumption of applications by enabling them to better use the processors of other available computer systems such as desktops or servers.

Previous Research Experience

My research aims to design algorithms to improve the scalability and performance of emerging wireless applications. Many of these newer applications depend on heuristics, which perform well in the circumstances envisioned by its developers. These heuristics may not work well in a new environment. Thus, my research focuses on ensuring that these applications can adapt well to a wide variety of environment conditions. Most of my techniques have been validated on actual prototypes of systems. Some of these works have been published in reputed venues such as IEEE INFOCOM, IEEE DySPAN, ACM IMC and IEEE Transactions on Cognitive Communications and Networking. There are also a number of works which are currently under peer review.

Spectrum Monitoring: Due to increase in the amount of wireless traffic, RF spectrum is becoming a scarce and thus, expensive resource. However, the high cost of spectrum may lead to increase in data cost for end-users of mobile networks. A key technique of mitigating this is to enable dynamic sharing of spectrum and protecting it from illegal use, by monitoring its usage. This requires scalable techniques of deploying spectrum sensors to monitor spectrum on a large scale. Such large scale spectrum monitoring requires outdoor deployment of sensors. Sensors deployed outdoors are usually powered by a battery, and thus keeping them running costs energy. They also send the data collected over a wireless network, which has a monetary cost. Thus, increasing the number of running sensors, and thus higher accuracy of spectrum monitoring, comes with a cost. We can reduce the cost incurred by limiting the number of running sensors through selection of the sensors that are most relevant. Therefore, limiting the number of running sensors without hurting accuracy of spectrum monitoring requires careful selection of running sensors.

Selecting the optimum subset of sensors of a given size to maximize the accuracy of spectrum monitoring is, however, NP-Hard. This either leads to no performance guarantees or slow execution. To solve this problem, I designed an approximation algorithm to select sensors that provide good performance bounds with the optimal. Designing this algorithm involved proving that the accuracy of spectrum monitoring follows the law of diminishing returns, i.e. improvement in accuracy reduces with an increase in the number of sensors. By proving this, I showed that a greedy algorithm provides an accuracy that is always close to the optimum [**INFOCOM'19a, TCCN'19, INFOCOM'18**]. This work led to an internship offer at IMDEA Networks Institute based in Madrid as well as invitations to present from multiple universities.

I also designed an actual spectrum monitoring system to benchmark existing different hardware and algorithms. This system consists of a server running an IoT framework, and multiple spectrum sensors built using single board computers like Raspberry Pi's connected to software-defined radios. The spectrum sensors used commodity software-defined radios called RTL-SDRs, to sense data and send it to a server over Wi-Fi. Using this setup, I benchmarked spectrum sensors that use FPGA, single-board computers such as Raspberry Pi's and smartphones to characterize their performance in terms of latency, energy consumption and hardware cost. This enabled me to show the best type of spectrum sensors to use depending on the requirement and available monetary budget [**DySPAN'18, PAM'19**]. This is one of the first projects in the wireless networks community that has looked at the type of hardware needed to make spectrum monitoring practical.

Computation Offloading: A key objective of designers of IoT is to enable more complex applications on smaller IoT devices, while reducing their complexity of development. Ideally, developers would not need to worry about the hardware limitations and focus on the features of the applications. One technique of enabling this is for the smaller devices to offload the more compute-intensive components of its application to larger devices with faster processors over the Internet.

Utilizing offloading efficiently requires determining the location of execution of each component of the applications, while ensuring that the latency and energy requirements of execution are satisfied. Previous studies suffered from poor performances in terms of execution time or energy consumption in some applications, as they used heuristics which do not always work well. This is because standard techniques of offloading consider it as an instance of graph partitioning problem, which is NP-Hard. To improve this, I proposed offloading techniques that could reduce both latency and energy consumption of the applications. In contrast to earlier approaches, my technique showed that allowing duplicate execution on different devices of parts of applications can significantly speed up execution. Moreover, this new model of the problem can be solved by a dynamic programming based polynomial algorithm to decide where to execute each application component. This algorithm guarantees the minimum execution time for any given application. Thus, I was able to significantly improve the performance of an emerging technique by rethinking its underlying model assumptions ([**CCNC'17**]). This is expected to lead to much faster execution of IoT applications and thus, make design of more complex applications feasible.

Handling IoT Traffic by Mobile Networks: The increase in the number of IoT devices is gradually changing the pattern of traffic through mobile networks. Unlike traffic from smartphones, this pattern is more unpredictable, and can lead to unexpected rise in the number of network requests and thus, control packets. Such unpredictable traffic needs to be handled by the designers of mobile networks for IoT to evolve further and become more usable. Handling them involves managing information about the incoming connection requests on one or more cloud servers. However, service providers of mobile networks try to minimize the utilization of cloud servers to reduce their cost of usage. Thus, there are two conflicting objectives faced by designers: reducing cloud utilization, and keeping enough cloud resources free to handle unpredictable requests.

I designed and proposed a more lightweight system of handling traffic to manage these conflicting objectives between reducing cloud utilization and handling unpredictable requests. Standard state-of-the-art techniques usually try to balance the incoming requests evenly across a number of cloud servers. I showed that by distributing the requests across the servers in a controlled uneven manner, it is possible to provide better tolerance to sudden bursts of traffic. This ensures that the latency constraints of each request can be handled, while also limiting the number of cloud servers used. This is one of the first studies that allows scaling up of IoT systems by orders of magnitude ([**SOSR'19**]).

360 Video Streaming over Wireless Networks: Recently, 360 degree videos are becoming increasingly available. For 360 degree video streaming to become popular, it is essential to ensure that they can be streamed over mobile networks. This is challenging due to the large amount of data that needs to be transferred over a limited bandwidth.

I handled this challenge by leveraging the fact that a user views only a small portion of the pixels fetched in 360 videos. I first divide the frames of the videos into rectangular spatial units called tiles. By integrating recent advances in computer vision, we predict a few seconds in advance the viewport of the user, i.e. the part of the video that a user is likely to view. This gives us the probabilities of each tile being viewed. I then utilize these probabilities to design an algorithm that maximizes the average bitrate, while limiting the probability of tile misses that the user views ([**INFOCOM'20b, NETWORKING'19**]). This technique is expected to lead to speed up adoption of video streaming over smartphones and other mobile devices. In this way, using co-design of a data-driven and algorithmic approach, it enabled much better user experience of an emerging application.

Impact of my Research

My research is broadly expected to help the evolution of more complex network applications. Next-generation networks are expected to support much larger number of IoT devices and lead to development of more compute and data-intensive applications. My work on spectrum monitoring enables better usage of the limited RF spectrum by enabling its shared usage, which is one of the major design goals of such networks. The work on computation offloading and handling IoT traffic enables more seamless execution of complex IoT applications. Finally, the work on 360 video streaming enables a more data-intensive application on mobile devices. I believe that the next generation of network applications will use a combination of all these techniques – efficient utilization of spectrum, better sharing of compute resources and more selective fetching of data to satisfy the ever-growing demand of more complex applications from users.

Future Work

My broad goal is to identify techniques to make current research prototypes as usable as possible in actual deployment scenarios. This also involves identifying bottlenecks in moving prototypes towards systems by benchmarking and other measurement-based studies. I have identified some key areas where I would like to achieve this goal.

A Better Spectrum Monitoring System: While SpecSense has shown the feasibility of distributed spectrum monitoring, it has a number of limitations. First, it has been shown to be able to detect transmitters run by its designers, but has not been used to detect actual illegal transmitters. Second, its sensors have low compute ability, and cannot run more difficult applications such as detection of type of signals, or actual decoding. It is also challenging to deploy these sensors outdoors due to their high energy consumption. I plan to deal with these challenges:

- **Localizing Illegal Transmitters in the Wild:** Although we currently have a working spectrum monitoring system, multiple challenges remain in improving its utility to stakeholders such as regulators and telecommunication operators. Current spectrum monitoring systems have been tested only by deploying transmitters by its designers. One research question that has not been answered is how effective they are in localizing actual illegal transmitters. To answer this question, I have identified detection and localization of GPS spoofing as a goal to demonstrate its actual use. Anecdotal evidence suggests that GPS spoofing devices, though illegal, are widely used by drivers of trucks and car rental services. However, localizing such spoofing devices using our spectrum sensors is currently challenging since the spoofing devices have a range of less than 100 meters on average. By utilizing the fact that the spoofing devices are installed in moving vehicles and deploying sensors in strategic locations within a city, I plan to accurately identify the vehicles that have installed these devices. I have started this work in collaboration with researchers from IMDEA Networks Institute and Armasuisse.

- **Enabling Complex Spectrum Applications:** While basic prototypes of spectrum monitoring now function well, there is a demand for adding more complex applications. For example, a system like SpecSense can be used by regulatory authorities to dynamically manage the price of spectrum. It is also possible for spectrum sensors to detect the type of signals by using neural networks. However, this requires much better compute capability than what is currently available on the single board computers. Sending the raw data to the servers currently cost too much bandwidth to be feasible. Thus, I plan to utilize computation offloading to enable these more complex applications.

Adaptive Data-driven Streaming Protocol: Although data-driven protocols for video streaming have been proposed, current data-driven techniques require retraining to adjust to unexpected changes in the environment. For example, it is possible that changes in the network conditions might lead to tolerance of lower latency. However, it is difficult to foresee during training the amount of time that would be available, as the quality of network can vary depending on the user's location and other factors. I propose to design an adaptive data-driven streaming protocol, that can automatically adjust its working depending on the network condition. This requires proper understanding of the behavior of data-driven protocols, and study of their response to changes in training and test data.

Performance Guarantees to Handle Unpredictable Traffic: While we have improved the tolerance of handling traffic bursts by mobile networks, our technique depends on some heuristics. So far, it is not clearly known if better algorithms are possible to allocate network requests to cloud servers. Since even small savings of cloud resources lead to substantial reduction in overall cost, this is a vital consideration. I am especially interested in providing performance guarantees for vehicular networks, where such guarantees can go a long way in reducing the number of accidents.

With newer applications such as Industrial Internet of Things, autonomous vehicles and augmented reality becoming commercial, it is becoming increasingly important for the networks to seamless handle different requirements. As in my current research, I plan to utilize techniques drawn from machine learning and algorithms to explore ways of making wireless networks robust and more adaptive to enable these emerging applications.

References

INFOCOM'20a Arani Bhattacharya, Caitao Zhan, Himanshu Gupta, Samir R. Das, Petar M. Djuric, Selection of Sensors for Efficient Transmitter Localization, in *Proceedings of IEEE Infocom 2020*

INFOCOM'20b Mallesham Dasari, Arani Bhattacharya, Santiago Vargas, Pranjal Sahu, Aruna Balasubramanian, Samir R. Das, Streaming 360 Videos using Super-resolution, in *Proceedings of IEEE Infocom 2020*

SOSR'19 Vasudevan Nagendra, Arani Bhattacharya, Anshul Gandhi, Samir R. Das, Scalable and Resource Efficient Control Plane for Next Generation Cellular Packet Core, in *ACM Symposium on SDN Research 2019*

NETWORKING'19 Sohee Kim Park, Arani Bhattacharya, Zhibo Yang, Mallesham Dasari, Dimitris Samaras, Samir R. Das, Advancing User Quality of Experience in 360 degree Video Streaming, *In Proceedings of IFIP Networking 2019*

TCCN'19 Arani Bhattacharya, Ayon Chakraborty, Samir R. Das, Himanshu Gupta, Petar M. Djuric, Spectrum Patrolling with Crowdsourced Spectrum Sensors, accepted for publication in *IEEE Transactions on Cognitive Communications and Networking*

INFOCOM'18 Ayon Chakraborty, Arani Bhattacharya, Snigdha Kamal, Samir R. Das, Himanshu Gupta, Petar M. Djuric, Spectrum Patrolling with Crowdsourced Spectrum Sensors, in *Proceedings of IEEE Infocom 2018*

DySPAN'18 Arani Bhattacharya, Han Chen, Peter Milder, Samir R. Das, Quantifying Energy and Latency Improvements of FPGA-Based Sensors for Low-Cost Spectrum Monitoring, in *Proceedings of IEEE International Symposium on Dynamic Spectrum Access 2018*

PAM'19 Mallesham Dasari, Muhammad Bershgal Atique, Arani Bhattacharya, Samir R. Das, Spectrum Protection from Micro-Transmissions Using Distributed Spectrum Patrolling, in *Proceedings of Passive and Active Measurement 2019*

CCNC'17 Arani Bhattacharya, Ansuman Banerjee, Pradipta De, Scheduling with Task Duplication for Application Offloading, in *Proceedings of IEEE Consumer Communications and Networking Conference 2017*