My research interests are in the area of wireless networks and Internet of Things (IoT) systems. IoT systems extend computing into the physical world around us and promise to transform multiple spheres of human lives: medicine, healthcare, homes, industries, cities, and environment. Imagine networks of micro-robots that can move inside human bodies to deliver targeted medicine, help regulate cancer treatments, and even perform in-body biopsies. Imagine smart homes that organically adapt to their occupants' needs, while also monitoring their health and wellbeing. An array of cheap sensors spread around the city can monitor and regulate traffic in real-time, reduce emissions, and aid navigation of autonomous vehicles.

At a high level, all these applications are about distributed computing machines connected to the Internet. However, these applications have very different constraints (size, power, communication, sensing capabilities). In-home sensors can be large enough to run on batteries or be plugged into wall plugs for power, but in-body robots need to be small and battery-free. Outdoor devices can locate themselves using GPS, but in-body and in-home devices cannot rely on GPS to find their location. While cellular connectivity serves in-city sensors, rural areas like farms do not have cellular connectivity or Wi-Fi to rely on. Building IoT systems under these constraints requires answering multiple questions. How do we power these devices? How do we set up a network in the constrained environments? How do we ensure availability of these distributed sensing systems? How can we make inferences from unreliable data coming in from multiple sensing modalities present in each environment? My research answers these questions to build end-to-end IoT sensing systems in multiple heterogenous settings: in-body, homes, cities, and farms.

My work distinguishes itself along three axes:

- **System Diversity:** I have designed and built IoT systems for a broad range of settings with diverse constraints: from tiny in-body sensors to large scale deployments in farms and cities. IoT systems built by me have spanned the entire scale of constraints along multiple dimensions: power, communication range, throughput, size, deployment scale, sensing modalities, etc.

- **Technological Depth:** In the pursuit of building these IoT systems, I solved core open problems in networking. Existing solutions for communication, localization, and inference do not take into account the constraints for IoT systems. I built new primitives that solved these problems and showed how these primitives have widespread general implications. Specifically, I built the first algorithm and system design that enabled accurate localization using *a single Wi-Fi access point*, unlike past work in localization that requires measurements from multiple vantage points. My contributions also include the first technique and system that allows next-generation cellular networks to communicate using advanced multi-antenna (MIMO) techniques *without any feedback from client devices*, and the first backscatter communication and localization system design *that works even inside human and animal bodies*.

- **Recognition & Impact:** My research has had academic, industrial, and social impact. My research has been published in the premier venues in networking and ubiquitous computing ([SIGCOMM 14; SIGCOMM 15; NSDI 16; SIGCOMM 16; NSDI 17; SIGCOMM 18; CoNEXT 18; UbiComp 18]). I have been awarded the Microsoft Research PhD Fellowship, and the SIGCOMM best paper award. My research has also been prominently featured in the popular press – BBC, Economist, World Economic Forum, MIT News, CBC, IEEE Spectrum, Daily Mail, Science Daily, among others. My work on indoor positioning was deployed at the Venice Architecture Biennale, 2014 – the premier global architecture exhibit lasting six months. My work on data-driven agriculture, FarmBeats, has been deployed by Microsoft at farms in North America, Asia, and Africa and has been operational for a couple of years. Small-scale farmers using FarmBeats have self-reported improved yields and reduced costs. These real world deployments have gained interest from top executives at leading tech and healthcare companies – Satya Nadella (CEO, Microsoft) named one of my projects as one of ten inventions that inspired him in 2017; Bill Gates visited one of our deployments and made a video describing the system.

IoT Systems & Applications

I have built on techniques from signal processing, computer networks, computer logic, computer vision, machine learning, and ubiquitous computing to design and deploy IoT systems in a wide range of settings:

In-body Devices: In-body devices promise to transform modern medicine. Small swallowable capsules can replace 40-inch long tubes that had to be thrust through a person's mouth, throat, and esophagus for endoscopy. Small tumor markers can enable doctors to track tumors and avoid radiation damage to healthy tissues. Micro-robots can perform targeted drug delivery, in-body biopsies, and targeted nerve stimulation.

All of these devices need to communicate the data that they are sensing and we need to identify their location in-body. For instance, endoscopy capsules need to send the images from the G.I. tract and tag the images with the location so that the doctors can isolate the part of the G.I. tract that has abnormalities. However, these tiny devices are constrained by the lack of power sources in-body and cannot carry large batteries to support traditional communication technologies. To solve this problem, I built ReMix [SIGCOMM 18], the first system that can communicate with and localize in-body devices without requiring them to transmit any power of their own. ReMix uses a device outside the body to transmit radio signals. The in-body devices communicate by just reflecting the signal back to the device and hence, do not have to generate a signal of their own. Crucially, ReMix can also use these reflected signals to find the location of the device inside the body (a.k.a. in-body GPS). This allows tracking of in-body devices such as tumor markers using safe, inexpensive radio waves; a capability that has so far relied on expensive and harmful technologies like X-rays & CT.

I collaborated with medical experts at Massachusetts General Hospital (MGH) to test ReMix in human tissue substitutes – chemically synthesized human tissues, and dead animal tissues like whole chicken and pork belly. ReMix can locate devices in-body to centimeter-level accuracy while communicating with them at data rates up to 1 Megabits per second. We recently ran joint experiments with MGH to successfully test the compatibility of ReMix with radiation therapy equipments used for cancer treatments.

Data-driven Agriculture: The global food production needs to double by 2050 to meet the growing demands of the increasing population. Unfortunately, the production is expected to go down, given the challenges of reduced cultivable land, lower water tables, and climate change. Data-driven agriculture is the most promising approach to alleviate this challenge. Data-driven techniques rely on dense measurements of farm's physical characteristics and can enable farmers to reduce input costs, enhance production, and ensure environmental sustainability. Yet, such techniques are limited to few farmers in the developed countries (and almost none in the developing world). This is because of three reasons: (a) dense deployment of sensors is expensive and inhibits the natural workflow of a farmer who needs to drive tractors and bullock carts through this complex maze of sensors; (b) farm sensors are susceptible to harsh weather conditions, making data collection intermittent; (c) farms are located in rural areas with sparse connectivity (cellular or otherwise), forcing labor-extensive manual data collection.

We built FarmBeats [NSDI 17], an end-to-end IoT platform for data-driven agriculture. FarmBeats includes system designs and algorithms that automate data collection, ensure availability in harsh weather conditions, and estimate dense sensor maps of the farm from sparse sensor deployments. The key component of FarmBeats is a new multi-modal sensor fusion algorithm that combines sensor data from sparse on-ground sensors with visual images collected using a drone (an increasingly common farm tool) to create a dense sensor map of the farm. In addition to this sensor fusion, we designed a novel hybrid communication system relying on available TV White Spaces in rural areas and an edge architecture that ensures availability in spite of harsh weather variations.

FarmBeats started as a summer project when I was at Microsoft Research in summer 2015. Over the next 2 years, I designed and developed the multi-modal sensor fusion algorithm, the system for robustness to weather variations, and collaborated with others on the design of the communication framework. I was the lead author on the FarmBeats paper published at NSDI 2017. I have continued to work with teams at Microsoft to deploy FarmBeats in multiple farms in Asia, Africa, and North America. I have worked on technology transfer as well as dissemination through talks, press interactions, and demos. The multi-modal sensor fusion designed by me was one of the core components in the public strategic partnership between DJI, the largest drone manufacturer, and Microsoft. In deployments, FarmBeats has been resilient to long stretches of cloudy days that cripple availability of solar power on the farms and thunderstorms that break down the backbone Internet connection for days. Farmers using FarmBeats have self-reported increased yields and reduced input costs (by 30-45%).

Smart Homes: For a long time, we have been captivated by the vision of smart homes that organically adapt themselves to their occupants' needs. Smart homes can control heating and lighting, alert against intruders, and monitor occupants' daily wellbeing. However, current devices hardly exhibit such smart behavior and need to be controlled by explicit user input either through an app or a voice-based assistant like Google Home.

RF-based indoor localization can play a central role in realizing the vision of smart homes, by enabling the home to be aware of its occupants' locations and react to their presence. It can allow a smart home to track the user as she sits on the TV couch and tune to her favorite channel, alert the parents if, during their absence, the babysitter enters their bedroom instead of the baby's room, and monitor mobility of occupants to check for sign of depression or diseases. However, traditional RF-based localization systems are targeted to enterprise environments, and cannot directly adapt to smart homes. Specifically, they require measurements from multiple vantage points, while homes typically have a single access point. Moreover, such systems rely on users carrying their smartphones *all the time*, which is a false assumption in homes.

I developed a series of systems (Chronos [NSDI 16], and Duet [UbiComp 18]) that enable location-aware smart homes. Chronos is the first system that can enable accurate decimeter-level localization with a single Wi-Fi access point. It allows the single Wi-Fi access point in homes to track user devices. Then, I built Duet, a multi-modal positioning framework for smart homes. Duet ingests two modes of positioning data: location of user devices in the home and identity-free locations of humans obtained by relying on radio reflections. Both these data streams are partially correlated because people intermittently interact with their devices. Furthermore, these data streams are prone to errors because of blockage caused by TV screens, bathroom tiles, concrete pillars, and other metallic objects. Duet builds a first-order logic based reasoning framework that can ingest these intermittent, error-prone data streams and predict the actual location and identity of users in a smart home. I deployed Duet in homes and small offices for several weeks with users using their own devices and requiring no change to their behavior. Across these deployments, Duet estimated the correct semantic location (e.g., bed, TV couch, kitchen) of the user over 95% of the times.

Fundamental Problems in Networked Systems

In designing the above IoT systems, I solved the following important open problems:

Localization using a Single Access Point: Localization of off-the-shelf user devices using radio waves has been a long standing problem in wireless networking. However, all current approaches require presence of multiple vantage points and hence are limited to enterprise environments that have extensive networking infrastructure.

My work in Chronos [NSDI 16] breaks this barrier. It allows a single Wi-Fi access point to accurately localize a client device. The key contribution of Chronos is a new algorithm that enables measurement of time-of-flight between two Wi-Fi devices to sub-nanosecond accuracy. Since RF signals travel at the speed of light, one can use this time-of-flight to calculate the distance between two devices. The main challenge in designing such a system stems from multipath. When a client transmits a signal, it is reflected off all kinds of objects in the environment: walls, computer screens, floor, etc. The signal received at the access point is the sum total of all these reflections. How can one untangle these reflections and identify the signal along the direct path? Chronos makes the observation that modern Wi-Fi systems can operate on multiple frequency bands. At each of these frequency bands, the reflections combine in slightly different ways. Using this observation, I built a non-uniform Fourier Transform variant that relies on the sparsity of the reflections to untangle these multiple reflections. Once these reflections have been untangled, the path with the shortest time of flight can be picked as the direct path and used to locate the client device.

More generally, Chronos changes the fundamental nature of localization systems. While existing systems rely on infrastructure support to deliver accurate localization, Chronos enables one device to locate another without relying on networking infrastructure. We showed that Chronos could use single access points in homes and small businesses to localize smartphones for enabling occupancy detection, restricting access to Wi-Fi based on location, and even to design personal robots that follow you around.

Communication for Next Generation Cellular Networks: The ever-increasing demand for data has forced cellular networks towards advanced multi-antenna (MIMO) techniques. However, use of MIMO techniques in cellular networks requires client devices to measure the wireless channel from the basestation to the client and send it back to the basestation as feedback. This feedback overhead scales linearly with the number of antennas on the base station. Since 5G networks plan to rely on massive antenna arrays on the base stations, this overhead has become a core problem; so much so that the LTE standardization body has recognized this as a challenge for future LTE networks.

I built R2F2 [SIGCOMM 16], a system that can eliminate this feedback requirement, by inferring these wireless channels from natural client transmissions. The key challenge in building R2F2 is that a majority of cellular networks (all networks in United States) use different frequencies for uplink transmissions from the client and downlink transmissions from the basestation. Thus, R2F2 must use transmissions on one frequency (uplink) to determine wireless channels on a different frequency (downlink). R2F2 makes the observation that while the transmission happens on different frequencies, the underlying paths that the signal travels are the same. So, I identified a frequency-invariant representation of the physical paths and designed an optimization problem that could convert wireless channel measurements to physical paths and vice versa. Solving this optimization problem allows the basestation to convert channel measurements on uplink frequency to underlying physical paths, and then use these physical paths to infer the wireless channels on downlink frequency; without relying on user feedback. R2F2 introduced a new bridge between two different areas of wireless networking: communication and localization, and was awarded the Best Paper Award at ACM SIGCOMM 2016.

In-body Backscatter Communication and Localization: Backscatter communication allows devices to communicate without transmitting any power of their own. Research in backscatter technologies has traditionally focussed on in-air devices. However, enabling backscatter across the air-human interface, as required for in-body devices, is very challenging. The human body quite literally stands in the way. The human body reflects 100 million to 1 billion times stronger signal than an in-body device and acts as a strong interference for any backscattered signal. Trying to isolate the target signal in the presence of this interference is equivalent to listening to your breathing in a rock music concert.

To overcome this problem, ReMix [SIGCOMM 18] leverages hardware imperfections. Specifically, most hardware is designed to exhibit linear behavior, i.e., if you transmit signals at frequency f_1 and f_2, you receive signals at frequency f_1 and f_2. However, imperfect design of hardware causes non-linear behavior that mixes these frequencies together and creates new frequencies $2f_1$, $f_1 + f_2$, $2f_1 + f_2$ and so on. ReMix uses a simple non-linear component (diode, millimeter sized) in the in-body devices to introduce non-linearity in the reflected signal. This allows ReMix to filter out the original transmitted frequencies (f_1,f_2) to remove the signal from the human body and focus on the non-linear frequencies ($f_1 + f_2$, $2f_1 + f_2$) from the in-body device.

The use of radio signals from an in-body device for localization poses yet another fundamental challenge. Unlike in-air localization where signal propagation happens in straight-lines, radio signals bend at the interface of various human tissues (air-skin, skin-fat, fat-muscle). This causes traditional geometric models used for localization to break down. This effect is similar to how we tend to erroneously locate objects in water when perceiving them with our eyes. I collaborated with medical experts from Massachusetts General Hospital (MGH) to understand the effect of the human body tissues on radio signals, and used this understanding to build mathematical models that allowed us to accurately localize in-body devices without requiring any information about patient-specific tissue structure.

Multi-modal Inference using UAVs: Traditional inference systems for IoT deployments consisting of UAVs (Unmanned Aerial Vehicles) try to make inference solely based either on UAVs or on deployed sensors. In designing FarmBeats [NSDI 17], I realized that combining the two data streams: visual data from UAVs and ground truth data from deployed sensors offered unique advantages. UAVs provide measurements that are dense in space (continuous visual maps), but sparse in time. On the other hand, ground sensors are sparse in space, but dense in time – the sensors can be frequently polled to provide measurements. We built a new graphical model that combined the two streams to provide dense space-time sensor maps of the farm. The core idea was to train a graphical model that maps visual features to sensor values using the sparse sensor deployments and use the trained model to predict sensor values in parts of the farm which had no sensors. This paradigm of using multi-modal inference mechanism allowed us to reduce sensor deployments by one to two orders of magnitude across farms.

Future Work

My research has focussed on building IoT systems to solve practical real-world problems. I believe, going forward, IoT systems will form the core of three emerging areas in computing. First, questions of sensing, communication, and inference are integral to the promise of digital healthcare. These digital healthcare systems will glean information from wearables, in-body devices, radio signal reflections, and medical records to reduce risk of health emergencies, ensure better treatment procedures, decrease long-term healthcare costs, and generally produce better healthcare outcomes. Second, multi-modal inference mechanisms will form a core part of the next generation wireless and IoT systems, as diverse sensors continue to get deployed in our physical environments. These inference mechanisms will build on sensing using different modes and combine them with inference tools like computer logic, and machine learning. Finally, underlining all of this will be the questions of security, privacy, and integrity of the data being collected using these systems. This becomes even more relevant with sensitive information like health data, sensor data from our homes, etc. I intend to pursue research along each of these three directions.

Digital Healthcare: Digital healthcare entails continuous monitoring of an individual's health, and ensuring timely interventions. I intend to push forward in this direction along two dimensions. First, my work in Duet has shown how we can track people in-home without requiring changes to user behavior. Extracting more insights from this data opens up several possibilities for monitoring the daily wellbeing of individuals, disease progression, hospitalization risk, and ensuring compliance with drug routines. For example, inferring activities of daily living could allow us to prevent hospitalization by enabling early interventions in patients with chronic conditions like CHF (congestive heart failures). Second, I plan to build in-body sensor networks that last for years and provide constant health monitoring. My work on in-body communication and localization has already taken the first step in this direction. Building on this work to design medium access protocols for multiple sensors, zero-power sensing mechanisms, and exploring the interaction of radio waves with human bodies can have widespread impact in healthcare and medicine. For instance, zero-power sensing of pH can provide long-term continuous monitoring for patients with acid reflux and micro-sensors embedded in pills could ensure drug adherence, a cause of 125,000 preventable deaths every year.

Multi-modal Sensing on Smartphones: Over the next few years, commercial devices like our smartphones will become equipped with radio devices capable of transmitting at vastly different frequency ranges. Near field communication (NFC) operates at sub-100 MHz, Bluetooth at 2.4 GHz, Wi-Fi at 5 GHz, and millimeter wave (mmWave) networks at 24 and 80 GHz. These RF frequencies, combined with pre-existing sensors like cameras and microphones, will enable multiple modes of sensing the environment. However, all these modes have very different forms of interacting with the environment, and extracting meaningful insights from these complex interactions is a challenging task. Fortunately, tools in machine learning, particularly deep learning, have recently had success in analyzing complex data. I believe exploring this confluence between multiple sensing modes and machine learning can open up several interesting capabilities for our smartphones, ranging from scanning the environment for Virtual Reality/Augmented Reality applications, to looking inside closed boxes for security, to in-home medical imaging, to material identification, to evaluating food quality.

Security and Privacy in Smart Environments: The advent of IoT and cyber-physical systems have forced us to change our existing notions of security and privacy. Security and privacy risks are not limited to digital information that is stored online but encompass our physical identities and environments. IoT systems have the capability to sense information about us and our surroundings and can change our physical environments by controlling devices and objects. As a result, there is an urgent need to rethink our security and privacy solutions. I believe the capability of IoT systems to sense our physical environment presents us an opportunity to rethink our security solutions by incorporating the physical information itself into our security approach. Specifically, the physical information being sensed and the communication channel will form signatures for verifying next-generation IoT devices. I want to build on this intuition to integrate low-level sensing mechanisms with edge architectures, inference mechanisms, and security protocols to design end-to-end security and privacy mechanisms for IoT systems.

References

[SIGCOMM 18] Deepak Vasisht, Guo Zhang, Omid Abari, Hsiao-Ming Lu, Jacob Flanz, and Dina Katabi. "In-body Backscatter Communication and Localization". ACM SIGCOMM 2018.

[UbiComp 18] Deepak Vasisht, Anubhav Jain, Chen-Yu Hsu, Zachary Kabelac, and Dina Katabi. "Duet: Estimating User Position and Identity in Smart Homes Using Intermittent and Incomplete RF-Data". ACM IMWUT/UbiComp 2018.

[CoNEXT 18] Roshan Ayyalasomayajula, Deepak Vasisht, and Dinesh Bharadia. "BLoc: CSI-based Accurate Localization for BLE Tags". ACM CoNEXT 2018.

[NSDI 17] Deepak Vasisht, Zerina Kapetanovic, Jongho Won, Xinxin Jin, Ranveer Chandra, Sudipta Sinha, Ashish Kapoor, Madhusudhan Sudarshan, and Sean Stratman. "FarmBeats: An IoT Platform for Data-Driven Agriculture". USENIX NSDI 2017.

[SIGCOMM 16] Deepak Vasisht, Swarun Kumar, Hariharan Rahul, and Dina Katabi. "Eliminating Channel Feedback in Next-Generation Cellular Networks". ACM SIGCOMM 2016. Best Paper Award.

[NSDI 16] Deepak Vasisht, Swarun Kumar, and Dina Katabi. "Decimeter-Level Localization with a Single WiFi Access Point". USENIX NSDI 2016.

[SIGCOMM 15] Omid Abari, Deepak Vasisht, Dina Katabi, and Anantha Chandrakasan. "Caraoke: An E-Toll Transponder Network for Smart Cities". ACM SIGCOMM 2015.

[SIGCOMM 14] Jue Wang, Deepak Vasisht, and Dina Katabi. "RF-IDraw: Virtual Touch Screen in the Air Using RF Signals". ACM SIGCOMM 2014.