Research Statement

Bill Tao

My research focuses on advancing wireless communication and mobile systems to new domains and new environments with novel system designs. The future of wireless technologies is set to **connect everything everywhere**, to facilitate the most free and flexible flowing and processing of data and information. That's why the cutting edge of wireless systems is always about expanding the domain of wireless connections in all dimensions. My research focuses on two dimensions of this expansion: **new territories** such as outer space, and **new scenarios** in complicated biomass-based medium.

Low latency network for remote sensing in outer space My work in outer space networking solves the excessive latency of the status quo system caused by the unique patterns of motion and network links of low earth orbit (LEO) satellites. By precisely modeling the movement of satellites, I develop systems that predict satellite network traffic and proactively schedule transmissions on satellites. Also, my research utilizes the generic computation devices such as GPUs onboard the latest generation of LEO satellites to perform edge computing. Through this combined effort, I reduce the image latency by thousands of times and enable quasi-real-time earth imagery service.

Precise wireless signal model in biomass I proposed a precise physics-based model for the propagation of wireless signals in biomass-based environments, such as human bodies or crop canopies. My research has proven the usefulness of such accurate modeling: we design new base stations that can greatly increase coverage for agricultural robots, and we demonstrate precise localization algorithms for medical implants that are robust, generalizable to different setups, and easily deployable.

Impact and contributions My research has made open source contributions to the community. I have released source code for all of my research publications, and they have received multiple artifact badges from the conferences where they were published. My research was also featured on the ACM GetMobile journal, and both RF-Protect [3] and Serval [4] have been covered by the Illinois News website.

My research distinguishes itself by its **theoretical and technical depth**: I design systems by combining the physics rules underlying these systems with learning and optimization tasks, which has proven to provide great benefit to the system designs. For example, by understanding the orbital and link physics for low earth orbit satellites, I built a holistic system that optimizes global network throughput and latency across time outperforming any greedy or heuristic system designs, as demonstrated in Umbra [5] and Serval [4]. By carefully modeling the magnetic field generated by near-field communication (NFC) coils, I built a robust, generalizable and easy-to-deploy in-vivo localization system with high precision. By modeling the propagation of wireless signal on a farm with crops, I designed a smart 5G networked system on the farm with superior signal quality and bandwidth. Besides, my research is also unique in its **diversity**, with research projects and papers spanning the top to the bottom of a networked system, covering scheduling, image processing, cyber-physical design, signal processing and privacy problems.

1 Wireless systems in new territories: the outer space

As we expand our footprints in the universe, we need new technologies to help us stay connected with each other. Recently, we have seen a 10X growth in LEO satellite launching activities in the past decade. This new type of satellite is equipped with advanced sensors like high-definition cameras, as well as onboard computation and storage devices. Such developments necessitate a shift from satellites as specialized devices to satellites as computers in the contemporary global distributed computer system. For example, the Starlink constellation is incorporated into the Internet infrastructure to provide standard-grade network access service globally. On the other hand, earth observation satellites such as Planet Dove are being wired into the existing cloud computing architecture to provide continuous insights from earth imagery.

This push of using satellite as communicating and computing devices posts new challenges to the design and implementation of the underlying system: on one hand, they are equipped with better sensors that generate a large amount of multimodal sensory data that needs to be transmitted and processed. On the other hand, satellites have unique dynamics. They are constantly and quickly moving and changing their connection topology with the ground; they have strict limit on the weight and power that they can consume; and they are hundreds of kilometers away from the earth and each other, which makes the networking connections weak and limited.

My research has focused on solving the **large latency** (up to days) of sensory data delivery caused by these unique characteristics of satellites. My research has discovered inefficiencies across the current networking stack when applied directly to satellites. I take a holistic approach to optimize the satellite sensing pipeline, and to use the limited network capacity to deliver as much data as possible, as fast as possible. By this holistic approach, our system can reduce the image latency **from days down to minutes**.

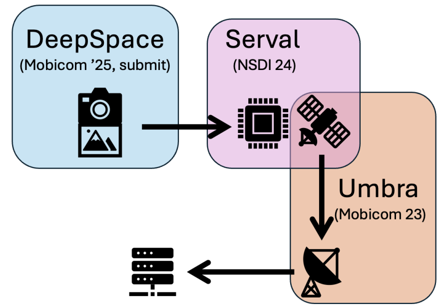

Figure 1: My research solves the network latency holistically at all network layers and all steps in the image transmission pipeline

Optimize traffic routing on the network layer In Umbra [5], we analyzed the long latency for LEO satellite data transmission and discovered that the intuitive "greedy" way that tries to use as much bandwidth as possible each time will cause **Uneven Queueing Effect (UQE)**, a load imbalance unique to LEO satellite systems due to the unique patterns of their motion and network link to the ground. This effect causes some ground stations to be oversubscribed, with others are underutilized if not scheduled smartly, wasting the backhaul bandwidth of the ground stations. To solve this problem, we propose "withhold scheduling" which tries to calculate the optimal ground station to connect to as well as the optimal bandwidth to use for each satellite at each point in time, instead of greedily using all available bandwidth. We do so by predicting the satellites' positions and ground connections in the future and build a time-expanded network structure to represent the traffic scheduling problem. We then form the network optimization problem as a maximum flow problem using a special network structure called "time expanded network" (TEN). We were able to reduce the network latency by 4 times using this method.

Optimize image compression on the presentation layer In DeepSpace [2], we analyzed imagery data sent to the ground and found great redundancy in it. We discovered that satellite images over the same region are usually repetitive and contain marginal new information. Given the tight link budget of satellites, we must reduce the redundancy in the satellite image transmission. To do this, we designed an adaptive image compression algorithm tailored to the satellite environment. It trains a mixture of experts (MoE) model using an ensemble of image super-resolution neural networks to reconstruct high-quality images from the optimally compressed data sent down from the satellites.

Smart edge computing on satellites on the application layer In Serval [4], we further tackled the task of reducing latency of image delivery. Upon analysis, we found that **not every image is equally desired**: most (99%) images are "normal" images with no events going on in it, only less than 1% of the images contain emerging events such as natural disasters or war. These images with events are usually of much higher interest to the users. To further improve the user QoE on the earth imagery service, we recognized that the new generation of LEO satellites were equipped with advanced computing devices, including GPUs. We thus proposed to perform edge computing on the LEO satellites to proactively sort out highly demanded images and send them in high priority. However, edge computing is challenging on LEO satellites: the energy consumption and size are strictly constrained for the compute devices, while the satellite camera generates around 1TB of image per day that cannot all be processed by the on-board computing. In addition, there are over 200 satellites moving more than 10 km per second in a single constellation, which makes edge computing more challenging. To address this challenge, we designed a smart edge computing framework that uses the technique of **query bifurcation** to proactively offload the compute to the ground using the predictability of satellite movements, and we were able to achieve minute-level latency in earth image delivery.

2 Wireless signals in new scenarios: signal through biomass

Traditional wireless research mostly uses radiofrequency (RF) signals and assumes that they travel in free spaces like the air. The technologies built from this research, such as WiFi, Bluetooth, or 5G, have been used to connect a variety of devices like smartphones, laptops, cars, planes, ships, and satellites. However, there are also a variety of devices that don't live in the free space but also require network connection. For example, medical implants are situated deeply in human tissues, and agricultural robots often operate under crops. All these devices need connection to the Internet to send their sensory data and receive operating instructions. However, connecting these devices to the network is hard: the biomass they are in is rich in water that absorbs wireless signals. For example, WiFi signal cannot penetrate 10cm layer of water. The problems get even more complicated when we need to localize these devices, as the structures of the human body are complicated and standard localization algorithms like triangulation don't readily apply here.

My research has aimed to extend wireless connectivity to devices in such challenging environments. My approach is to incorporate a precise and detailed modeling of the electromagnetic properties of biomass, and build the signal models on top, which has been proven to provide a strong theoretical guidance on system designs that successfully achieve high-quality network connection and sensing.

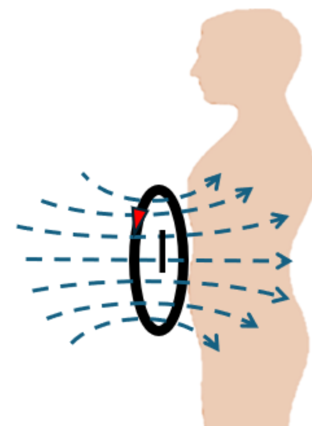

Figure 2: InnerCompass uses a physical model for signal propagation in human body

Wireless localization in human body In the future of medicine, we envision that smart implantable medical devices will be able to provide continuous, precise, and detailed health monitoring to patients. However, this must be built upon precise and efficient system designs, for example, the algorithms to locate these devices in human bodies. However, state-of-the-art algorithms use machine learning models trained on thousands of "fingerprinting" data collected from a specific calibration device that does not generalize to diverse environments such as patients with different body shapes. In InnerCompass [6], instead of collecting large "signal fingerprint" data to train the localization models, we designed the first precise physics-based model on near field communication signal strength (RSSI). By combining the theory on magnetic fields an circuits, we were able to accurately calculate the signal strength at any given point, and thus locate the medical devices at the same precision as other baselines while requiring a **thousands of times less data collection** and generate to any new environment or setup that the localization system was never trained on.

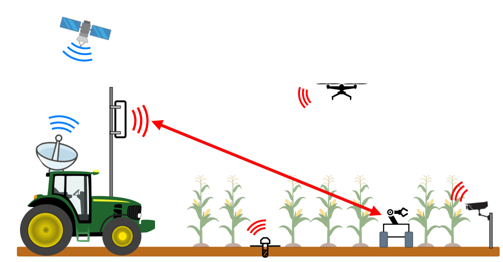

Figure 3: Byon uses an optimally placed CBRS to provide network on the farm

Wireless connectivity through plantation In BYON [1], we solved connection problems in another challenging environment: the farm. Farms are generally not fully covered by standard 4G base stations due to their size and remote location. However, modern farms also require agricultural robots that operate under the canopy of the crops to perform tasks such as monitoring growth or applying fertilizer. Due to their location, these robots cannot be easily covered by standard cellular signals. To solve this challenge, we modeled the path attenuation of cellular signals in the crops and found that the reason behind the inefficiency in standard cellular service is the placement of the base stations: they are not placed at the optimal height and suffer too much from attenuation in crops to provide good coverage for robots. We proposed the system "BYON" that features a CBRS base station that moves on the farm with the robots and designed an algorithm to find the optimal placement of the base station to achieve optimal signal quality based on the signal strength model we developed. As a result, our system was able to achieve a 5dBm average increase in signal strength which corresponds to almost 2X improvements in network throughput. We verified that our system makes it possible for a real-time teleoperation of robots with high-definition video streaming.

Privacy against adversarial wireless human sensing In RF-Protect [3], we focus on the privacy impact of sensing technologies in these new environments. Past work has enabled sensing in indoor non-line-of-sight conditions using a special type of radar called FMCW, which enables the sensors can "see through walls" and locate humans inside homes. However, this technique poses privacy threats to users as adversaries can now "sense" their location and movements in their homes without their knowledge or concern. To contain this privacy threat, we built a smart backscattering device called "RF-protect" that is tailored to the physical nature of FMCW signals: it modulates the signal in a way to fabricate "fake" reflections, that can be similar to legitimate reflection from any given location, and reflect it back to the sensor. To make the generated fake human trajectory look realistic, we employed generative adversarial network (GAN) to generate reflections indistinguishable from realistic ones, thus preserving user privacy against adversarial sensors.

3 Future work

My research has focused on expanding the territory of wireless networks, both into new spaces and into new scenarios. However, it was but a scrape on the surface of the full potential that the ever-expanding wireless network can achieve. I believe that there are exciting open questions to explore in both dimensions of my previous research, both in the short- and long-term. I am eager to further push the boundaries of wireless systems in my future work.

3.1 New territories for wireless

Short term: provide service directly to off-the-shelf devices While LEO satellite-based network access has seen promising success, the drawback is that it requires specialized user end (UE) terminals to operate. This is not suitable for small devices such as cell phones or IoT devices, as they always move and cannot be attached to a special big antenna while moving. Both industry and academia are trying to have off-the-shelf devices connected to satellites by directly providing LTE/5G service to cellphones in rural areas, and there have been publications that try to migrate protocols such as 5G to space. However, it is worth noting that 5G is not just a protocol but also contains system designs such as for radio access network (RAN) management, etc. that also need to be adapted to satellite systems. We will not have an efficient satellite internet infrastructure unless we make a holistic system design.

Long term: support wireless networking in deeper space In the long term, we dream of migrating to the Moon, the Mars, or even outside the Solar system. This will need wireless systems deployed to these remote locations as well. For example, right now all space rovers sent to the Moon and Mars have to send all their sensory data to earth for processing. Is it possible to deploy an edge computing system, or even better, build a datacenter on the Mars to process the sensory data on the "edge"? Also, when the Earth and Mars are moving, their distance could change anywhere between 3 light minutes to 22 light minutes, which greatly impacts both the latency and SNR of the wireless links. What do these dynamics mean for the wireless system? Can we deploy "space routers" smartly to ensure a continuous and reliable link? These new systems require a deep understanding of gravity dynamics in the Solar system, and I plan to once again precisely model the dynamics of deep space such as solar system, and use the model to guide us to solve for an optimal deployment of network systems in deep space.

3.2 New environment for wireless

Short term: physical security in emerging wireless systems With wireless systems becoming more powerful and widely applied, it's also necessary to vet their security: can we trust that our wireless systems are robust against adversarial attacks? Can we trust that our wireless systems are safe and won't attack us themselves? These are all important questions to explore. While the majority of research on privacy and security focuses on the network layer and up, I believe it is equally important to secure the physical signals. For example, as magnetism-based near field communication (NFC) is being widely applied in touchless payment systems, it is necessary to obtain strong theoretical guarantees on the security of the physical signals, for example, that they are unexploitable in a relay attack, which cannot be defended by any other protocol designs beyond the physical layer. One approach to this is to precisely model the statistical characteristics of the NFC signal, so we can quantify both the difference between the relayed signal and the genuine signal, as well as the theoretical limit of our ability to distinguish these attacks.

Long term: wireless sensing of neural activities In InnerCompass [6], I was able to precisely model the propagation of the magnetic signal within the human body. This type of signal is easy to detect and can be precisely modeled. Localization is just one application of this, and there are many more to explore: apart from artificial devices, any biological process that generates an electric current also generates a magnetic field that can be captured and analyzed. This includes, particularly, neural excitations. If we can use magnetics to detect nerve activities, they can be used to detect brain activity. We will be able to build brain-computer interfaces that will allow a much deeper conversation between humans and computers. With nerve activity detection, we can also build medical devices for handicapped patients, such as artificial hands and legs that precisely move according to owners' will just like real ones, by detecting current down the nerves, and doing so in a non-intrusive and less painful way.

Selected Publications

[1] anonymous. Byon: Bring your own networks for digital agriculture applications. In *NSDI 2025, submitted*, 2025.

[2] anonymous. Deepspace: Reliable super resolution powered satellite imagery based earth observation system. In *ACM Mobicom 2025, submitted*, 2025.

[3] Jayanth Shenoy, Zikun Liu, Bill Tao, Zachary Kabelac, and Deepak Vasisht. Rf-protect: privacy against device-free human tracking. In *Proceedings of the ACM SIGCOMM 2022 Conference*, pages 588–600, 2022.

[4] Bill Tao, Om Chabra, Ishani Janveja, Indranil Gupta, and Deepak Vasisht. Known knowns and unknowns: Near-realtime earth observation via query bifurcation in serval. In *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)*, pages 809–824, 2024.

[5] Bill Tao, Maleeha Masood, Indranil Gupta, and Deepak Vasisht. Transmitting, fast and slow: Scheduling satellite traffic through space and time. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, pages 1–15, 2023.

[6] Bill Tao, Emerson Sie, Jay Shenoy, and Deepak Vasisht. Magnetic backscatter for in-body communication and localization. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, ACM MobiCom '23, New York, NY, USA, 2023. Association for Computing Machinery.