I strive to bridge theory and practice in modern Cyber-Physical Systems. Theoreticians often abstract away relevant aspects of problems to achieve elegant results and prove important properties about their analysis and design. While some insight can still be drawn from these advancements, they usually leave a wide gap that the system designer needs to fill in order to have a working implementation. Unreliability occurs as a consequence of unrealistic assumptions, and inefficiencies arise as a consequence of non-trivial constraints imposed by the real-world. Finally, cost is rarely taken into account. On the other hand, a complete bottom-up approach, based usually on experience and domain knowledge, is not scalable. This is because intuition tends to fade in the face of scale and complexity. I strive to introduce systematic design principles into networked embedded sensing systems with the goal of designing inference algorithms that can achieve high reliability, cost-effectiveness, and high performance in Cyber-Physical Systems (CPS). As opposed to pure theoreticians, I work close to the physical layer, and I understand its practical constraints, which I include in my models. As a consequence, my systematic approach takes advantage of the particular structure of the system and provides solutions that not only are near optimal according to the chosen design metric, but also readily implementable. Knowledge of the system allows me to perform opportunistic design, by extracting useful information from sensors and devices deployed for other purposes.

My thesis demonstrates this approach applied to several indoor localization applications, ranging from precise mobile phone localization that can be used for navigation and mobile augmented reality, to extreme systems designed for firefighting scenarios wherein multiple incomplete sources of data need to be opportunistically fused together. In all of this, not only did I develop reliable algorithms, but I also successfully demonstrated the feasibility of my design in several real-world deployments. As a result, my research has been published at such premier highly selective conferences in CPS as IPSN '14 **(cited 130+ times)**, IPSN '18, SenSys '17, RTAS '17, RTAS '15, RTSS '13, ICCPS '13, IPIN '16, and VLCS '14. My work has been demonstrated live at four conferences, been deployed in more than two dozen environments, received a patent, **won the international Microsoft Indoor Localization competition twice**, received a **best demo** award, spawned a **startup**, and led to funding from NSF, SRC, NIST, and industry.

1 Indoor Localization Research

GPS has revolutionized the way we interact with the outdoor environment, and yet, we don't have pervasive and accurate systems for indoor localization. Due to the wide variation in application requirements and indoor spaces, it is unlikely that a single technology will effectively address all indoor localization problems. At the same time, it is impractical, due to cost, to deploy several different localization systems in all buildings at a global scale. Working between theory and systems allows me to approach this issue in a unique manner. I build new hardware platforms by opportunistically reusing existing infrastructure and designing new sensing and signaling schemes by exploiting low-level properties of sensors. For this, I draw from my theoretical knowledge to build systems. I then design opportunistic estimation methods that can accurately localize even in unreliable conditions, where traditional estimation techniques would fail. This naturally requires an understanding of practical systems to scope the design of algorithms. Finally, I bridge theory and practice by building tools that allow system installers to deploy infrastructure with predictable behavior and high confidence about its performance and reliability.

New Hardware Platforms. I designed one of the **earliest visible light communication (VLC) systems** to send data from overhead LED lights to phones. The phones use proximity to lights in order to determine their location. The challenge in building this system is that in order for the lights to be flicker-free, they have to operate at a much higher frequency than what can be sampled by the camera frame rate. To overcome this fundamental limitation, my insight was to exploit the low-level rolling-shutter effect of camera sensors to capture a time-varying light signal as a spatially varying image. Further, I used the exposure control as a temporal filter and the focus control as a spatial filter to improve the signal-to-noise ratio [1, 2]. I extended this to design a novel hybrid communication system where a single light simultaneously sent two independent data streams at different rates to cameras and photodiodes [3] by leveraging the fact that they have different filtering properties. In this work, I applied signal processing, communication, and estimation theory to the design of these networked wireless embedded systems.

I also contributed to building an **ultrasonic time-of-flight (ToF) platform** that localizes unmodified mobile devices [4, 5, 6]. This method is applicable to emerging smart speakers for localization. I experimented with and characterized several **emerging ToF technologies from industry** including WiFi, BLE, and UWB. I realized that ToF ranging is promising and will eventually find its way into future smart devices. However, ToF systems face challenges in scaling beyond single-rooms, since the practical environmental conditions are far removed from the models assumed by theoretical algorithms. Hence, we require new location estimation algorithms and tools that support practical system setup based on theoretical limits in localization accuracy.

Novel Estimation Algorithms. We have solvers for location estimation in ideal scenarios with good measurements from multiple beacons, and have solvers for location tracking where we update the location based on prior estimate and new measurements. However, the realistic case where traditional location solvers fail occurs when we have to acquire the initial location without a prior and we may not have sufficient line-of-sight measurements from beacons. In addition, measurements may be incorrect due to lack of line-of-sight between transmitting beacons and the receiver. In order to cope with this issue, past approaches simply over-deploy beacons, incurring higher deployment and maintenance cost, or do not initialize until the users walks around, causing a significant delay in providing a position estimate.

In contrast to existing approaches, I **opportunistically estimate location** based on the available measurements and the floor plan knowledge [7]. I use the insight that in indoor spaces, we can use the absence of measurements from beacons as useful information, and we can check for consistencies between the measurements and the floor plan. I use models that are simple enough to generalize across indoor environments and ranging technologies, yet practical enough. For instance, I model the beacon coverage using optical ray-tracing and the only assumption I make on an indirect signal path is that it is longer than the direct path. This approach is effective since indoor spaces are not random in geometry. I have shown how this approach significantly reduces the amount of infrastructure, and is robust to disturbances and uncertainties. This location estimation solver implemented on our ultrasonic platform **won the Microsoft Indoor Localization Competition in 2015**. In this work, I used my insights gained from practical system deployments to design estimation algorithms.

Another challenging practical problem on the estimation side is acquiring an initial orientation. Unfortunately, a compass is not sufficiently reliable indoors due to the large amount of metal in buildings. To solve this problem, I designed a novel approach where I crowd-source a dense magnetic field map from pedestrians with mobile phones by fusing data from beacons, on-board cameras, and inertial sensors. Subsequently, I leverage this map as a reference to **instantly acquire orientation** upon startup. Using this concept, my collaborators and I built an end-to-end **multi-user persistent augmented reality (AR)** system [8]. This work won the **best demo** award at IPSN 2018 [9]. I then extended this work to support **continuous location updates**. This system **won the Microsoft Indoor Localization competition in 2018** with ultra-wideband beacons.

The aforementioned techniques have been implemented on an ultrasound-based localization system that spawned a startup. A pilot has recently been deployed for AR-based product finding in a retail store.

Tools for Theory to Practice. To the best of my knowledge, a systematic approach to the efficient deployment of time-of-flight indoor localization beacons is missing. Beacon deployment is often seen as an installer's problem. However, by jointly designing estimation algorithms and the beacon placement approaches, and by quantifying the beacon placement in terms of the localization accuracy it can provide, one can at the same time place fewer beacons and achieve greater robustness while estimating location.

I decided to take on this challenge and designed systematic **beacon placement** algorithms [10]. My first insight was to use the floor plan geometry in a clever way to reduce the number of beacons compared to conventional placement. My second insight was to quantify the quality of a beacon placement by adopting the Cramer-Rao lower bound on the location estimate, which is expressed analytically by the geometry of beacons. I built on these concepts and implemented beacon placement algorithms in a toolchain where system installers can specify accuracy and coverage requirements for a floor plan and obtain a suitable placement which aims at minimizing the number of employed beacons. After some research into existing methodologies, I found that tools from computation geometry theory can be extended and then used to solve these problems in a rigorous manner. To this goal, I struck up a collaboration with Prof. Jie Gao from Stony Brook University and her research group. We mathematically formulated the beacon placement problem and proposed algorithms with provable guarantees about minimality of beacons under performance constraints [11]. In these works, I used my understanding from practical systems to design theoretical tools based on estimation, optimization, and geometry, which we then use for practical deployments.

To address another issue associated with deployment, I built an algorithm for **automatic beacon mapping**, applicable in scenarios where we deploy an ad-hoc, on-the-fly, localization systems. This algorithm is also applicable to future internet-of-things applications where infrastructure can be time-varying, and where devices appear, disappear, and move about the environment. For this goal, I leveraged existing algorithms in robotics and designed a pedestrian-based simultaneous localization and mapping (SLAM) algorithm [8]. I implemented the algorithm to map ultra-wideband beacons in several deployments. The algorithm is also currently used for asset tracking applications.

Other Areas in CPS. In related works, I have also contributed to time-synchronization [12, 13]. Time and location are both key primitives for CPS. Earlier in my studies, I worked on non-intrusive load monitoring [14, 15] using a wireless sensor network of electromagnetic field sensors.

2 Future Directions

CPS are becoming pervasive and growing in scale with applications ranging from smart cities, smart transport, to personalized applications like body area networks. In the future, I want to support paradigms where physical and computational resources are used by multiple applications and shared by multiple tenants. For instance, tracking of an asset will span smart buildings, smart transport, and smart cities. A second example is a shared network of drones serving multiple applications, such as imaging for smart agriculture, and transportation of critical resources to a hospital for an emergency. Multi-tenant problems are already beginning to plague smart building infrastructures that should be capable of being reconfigured on-the-fly based on the context of the building occupants.

In the near-term, I want to build theoretical and practical frameworks for CPS in the built environment for applications ranging from disaster resilience or smart buildings, to smart manufacturing or smart medical facilities. These systems should enable any user to scan the environment, discover services, and view and interact with both objects and digital content in real-time. In addition to application-specific sensing, this requires advancements in several areas including system integration, rapid deployment methods, integration with next-generation communication systems, secure sensing, and creating mixed reality systems for interaction. Some preliminary directions are:

Rapidly Deployable Systems. Future CPS will require rapid infrastructure deployment in an ad-hoc manner. As a concrete step towards this lofty goal, I would like to explore robust systems for localizing first responders and building occupants in serious fires without relying on existing infrastructure. I led the proposal for a firefighter localization project which is funded by National Institute of Standards and Technologies (NIST) that has started to look at the role of opportunistically sharing information between clusters of first responders to refine location estimates. The proposed method uses beacons on firetrucks, wearable devices on firefighters, a long-range low-power wireless network, and mobile network localization and mapping algorithms. My next steps are to evaluate the impact on various technologies under smoke, develop methods to train low-accuracy inertial sensors using vision-based high accuracy sensors, and integrate infrared sensors and drones for increasing resilience. I am working with NIST and industry on emerging beaconing technologies that could potentially be part of all buildings in the exit signs, primarily for emergency (e911), but could be also useful for bootstrapping other systems.

Secure Sensing. Current sensing systems can easily be attacked since the system models do not account for a possible attack. It is therefore critical to secure these systems. As a first step, I am tackling the problem of secure location sensing. Secure location is a key feature to enable device discovery and interaction in the internet-of-things with heterogeneous untrusted devices. My early conversations with security research groups suggest that security needs to be included at system design rather than added as an afterthought. This requires a complete overhaul of all the layers of the localization stack, starting from the physical layer design, to the location estimation, and infrastructure setup algorithms. I will leverage the growing body of research on CPS resilience and the experience gained at Carnegie Mellon interacting with my advisors and peer PhD students at CyLab and adapt it to the CPS domain I am targeting, taking advantage of its nuances and peculiarities. More broadly, I will once again work across theory and systems, for re-designing sensing systems to be designed-in secure.

Mixed Reality. Mixed reality is an ultimate realization of CPS with tight coupling between the digital and physical worlds in real-time, with a human-in-the-loop. This includes augmenting the human's perception of reality through various senses and enabling the human to realize actuation in the physical world digitally. Among many challenges, mixed reality requires accurate representation of the digital display, as well as smart devices and physical objects in the same reference frame, for which current vision-based methods are insufficient. As a start, I have shown how we can improve state-of-the-art mobile Augmented Reality (AR) using beacons and magnetic fields [8]. In future, I would develop new sensing techniques and design principles for fusing multiple sensors, e.g. vision, emerging localization technologies, and VLC, at the low-level to create robust mixed reality systems. More broadly, I am interested in building architectural frameworks for integration of various sensing, actuation, and interaction technologies.

3 Building my Research Group

I plan to have a large, interdisciplinary research group. Irrespective of how students balance theory and practice in their work, my goal is for them to have a systematic and analytical approach to design, as well as an appreciation for practical real-world implementation. While working on indoor localization, I interned at Apple's location group, worked with Texas Instruments' emerging localization technology, and met with users of localization systems ranging from firefighters to museum staff who wanted to deploy localization systems. These interactions helped me think about the problem more holistically and motivated me to keep my work grounded in real-world problems. I am looking forward to diversifying into new domains with collaborations.

I would develop a collaborative culture in my research group. I am co-advised, and have had the opportunity to be part of two research groups during my PhD. One of my external collaborations started off with a casual discussion with a faculty at a conference, and another collaboration started off when I met a faculty who visited Carnegie Mellon to give talks. In addition, I have been part of two SRC multi-university research centers, Terraswarm and CONIX, throughout my PhD. Being part of these research centers has given me the opportunity to experience firsthand the impact of collaborations across groups and universities, and to build a wide academic network. I will create opportunities for my students to grow through such collaborations and interactions within and outside the university.

During the course of my PhD, I have significantly contributed to the writing of proposals by bringing new ideas and fresh perspectives on the proposed work. I am therefore confident in my ability to secure funding from industry and government agencies to build my group and sustain my proposed research agenda.

References

[1] Niranjini Rajagopal, Patrick Lazik, and Anthony Rowe. Visual light landmarks for mobile devices. In *Proceedings of the 13th international symposium on Information processing in sensor networks (IPSN)*, pages 249–260. IEEE Press, 2014.

[2] Niranjini Rajagopal, Patrick Lazik, and Anthony Rowe. Demonstration abstract: how many lights do you see? In *Proceedings of the 13th international symposium on Information processing in sensor networks*, pages 347–348. IEEE Press, 2014.

[3] Niranjini Rajagopal, Patrick Lazik, and Anthony Rowe. Hybrid visible light communication for cameras and low-power embedded devices. In *Proceedings of the 1st ACM MobiCom workshop on Visible light communication systems*, pages 33–38. ACM, 2014.

[4] Patrick Lazik, Niranjini Rajagopal, Bruno Sinopoli, and Anthony Rowe. Ultrasonic time synchronization and ranging on smartphones. In *Proceedings of the 21st Real-Time and Embedded Technology and Applications Symposium (RTAS)*. IEEE Press, 2015.

[5] Patrick Lazik, Niranjini Rajagopal, Oliver Shih, Bruno Sinopoli, and Anthony Rowe. Alps: A bluetooth and ultrasound platform for mapping and localization. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*, pages 73–84. ACM, 2015.

[6] Patrick Lazik, Niranjini Rajagopal, Oliver Shih, Bruno Sinopoli, and Anthony Rowe. Demo: Alps– the acoustic location processing system. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*, pages 491–492. ACM, 2015.

[7] Niranjini Rajagopal, Patrick Lazik, Nuno Pereira, Sindhura Chayapathy, Bruno Sinopoli, and Anthony Rowe. Enhancing indoor smartphone location acquisition using floor plans. In *Proceedings of the 17th ACM/IEEE International Conference on Information Processing in Sensor Networks*, pages 278–289. IEEE Press, 2018.

[8] Niranjini Rajagopal, John Miller, Krishna Kumar Reghu Kumar, Anh Luong, and Anthony Rowe. Improving mobile augmented reality re-localization using beacons and magnetic-field maps. Under Submission.

[9] Niranjini Rajagopal, John Miller, Krishna Kumar Reghu Kumar, Anh Luong, and Anthony Rowe. Welcome to my world: demystifying multi-user ar with the cloud: demo abstract. In *Proceedings of the 17th ACM/IEEE International Conference on Information Processing in Sensor Networks*, pages 146–147. IEEE Press, 2018.

[10] Niranjini Rajagopal, Sindhura Chayapathy, Bruno Sinopoli, and Anthony Rowe. Beacon placement for range-based indoor localization. In *Indoor Positioning and Indoor Navigation (IPIN), 2016 International Conference on*, pages 1–8. IEEE, 2016.

[11] Haotian Wang, Niranjini Rajagopal, Anthony Rowe, Bruno Sinopoli, and Gao Jie. Efficient beacon placement algorithms for time-of-flight indoor localization. Under Submission.

[12] Maxim Buevich, Niranjini Rajagopal, and Anthony Rowe. Hardware assisted clock synchronization for real-time sensor networks. In *Real-Time Systems Symposium (RTSS), 2013 IEEE 34th*, pages 268–277. IEEE, 2013.

[13] Adwait Dongare, Patrick Lazik, Niranjini Rajagopal, and Anthony Rowe. Pulsar: A wireless propagation-aware clock synchronization platform. In *Real-Time and Embedded Technology and Applications Symposium (RTAS), 2017 IEEE*, pages 283–292. IEEE, 2017.

[14] Niranjini Rajagopal, Suman Giri, Mario Berges, and Anthony Rowe. A magnetic field-based appliance metering system. In *Proceedings of the ACM/IEEE 4th International Conference on Cyber-Physical Systems (ICCPS)*, pages 229–238. ACM, 2013.

[15] Niranjini Rajagopal, Suman Giri, Mario Berges, and Anthony Rowe. Demo abstract: a magnetic field-based appliance metering system. In *Proceedings of the 12th international conference on Information processing in sensor networks*, pages 307–308. ACM, 2013.