Aditya Arun

(510) 771 4618
aarun@ucsd.edu
https://wcsng.ucsd.edu/aarun/

Research Statement

A device's location are central to providing physical intelligence. A phone's displayed content can be tailored based on location, a robot's actions are governed by its location, or buildings may reconfigure their HVAC systems to serve their users better. More broadly, with physical intelligence, we can adapt a device's generalized intelligence and provide situational awareness. General intelligence, primarily provided by LLM's and foundational ML models, requires troves of data and novel network architectures to capture the similarities. Similarly, I envision imbuing our devices and systems with *physical intelligence* will require us to "see" the environment more holistically, develop and deploy systems to capture this data, and most importantly, do it in a secure and privacy-preserving manner. My broader research vision is to provide this physical intelligence to robotics, mobile devices, and smart building systems. However, I find our current sensing systems are poorly positioned towards providing this physical intelligence. GPS primarily aids in determining this location for many outdoor locations. But, gleaning this context in many indoor situations and GPS canyons becomes hard. Humans usually discern this context by gathering visual clues from our environment. Along similar lines, visual sensors like cameras or LiDARs are primarily used to understand a device's position with respect to its environment. However, these visual sensing modalities have numerous gaps when used for localization and mapping. Failure under occluded settings, dynamic lighting scenarios, or monotonous and featureless environments are just a few. Consequently, my research has primarily been guided by the vision that we need better and more robust ways to sense the environment around us. My research in developing wireless sensing systems, leveraging WiFi, UWB and other RF-signals is a step towards fulfilling this vision. My work has

1. taken advantage of ambient WiFi signals to improve a robot's ability to localize and map itself in challenging indoor environments (P2SLAM [ICRA+RAL '22], ViWiD [Under submission]): I developed SLAM algorithms that operate independently of visual sensors for localization and navigation, thereby improving the robustness and compute-efficiency of these systems.
2. built deployable and cm-accurate ultra-wideband localization systems to serve the next generation of extended reality applications (XRLoc [Sensys '23], ULoc [IMWUT '21']): I developed systems to provide fine-grained localization and tracking of UWB tags for safety monitoring, AR/VR world building or logistics management.
3. architected state-of-the-art indoor localization systems (Dloc [Mobicom '20]) and automated the deployment of these systems (LocAP [NSDI '20]) in large environments: I have developed end-to-end pipelines to simplify the deployment of indoor localization to provide sub-meter scale positioning accuracy for mobile and IoT devices.

These core contributions have paved the way for devices and robots to "see" via their RF radios to augment their spatial understanding in challenging scenarios. They have also positioned me to leverage other spectrums, including mm-wave, ultrasound, and sub-Ghz, to provide a more holistic understanding of our spaces. Providing physical intelligence through these multi-faceted sensing schemes will allow autonomous systems to be deployed more safely, improve AR/VR user experiences, and enhance the context under which our devices operate.

Broader Impact: I have driven additional impact through three core ideologies during my PhD career.

Developing open-source systems: Along the way in developing the above applications, I have built foundational sensing systems leveraging WiFi (WiROS [demo, IPSN '23]) and UWB (VRTrac [demo, Mobisys '22]) for sensing and localization. These systems are being used at a growing list of universities, including the University of Washington, UIUC, Chinese University of Hong Kong, and Johns Hopkins University. I have additionally open-sourced first-of-their-kind wireless sensing datasets (WILDv2 [Kaggle '22], P2SLAM), downloaded over 150 times across 50 universities worldwide. Developing widely applicable open-source systems and datasets is an essential priority as it amplifies the impact of my research. I will continue to prioritize this as a faculty as well.

Solving present-day problems: During the onset of COVID-19, I leveraged my expertise in RF sensing to develop contact tracing with low false-positive rates, one of the key hindrances in deploying wide-scale contact tracing (BluBLE [COVID Special Session, Sensys '20]). This work was **prototyped by Aarogya Setu, India's contact tracing agency**. On the flip side, I have also created systems that can be potentially misused to gather users' location information. To protect users' privacy, I have built novel signal transmission schemes that users can use to disguise their location while maintaining strong communication links (Mirage [Hotmobile '23]).

Driving collaborations: A primary hindrance to future communication systems, multi-modal sensor networks, and edge-distributed computing will be inadequate time synchronization between devices in these networks. Additionally, widespread industry backing and input will be essential to ensure swift technology adoption. Consequently, I have facilitated broad industry collaboration with NICT, Japan, and Meta under the Open Compute Project, where we leverage physical layer wireless signal responses to characterize and correct clock frequency and time drift.

Many of these works have garnered comprehensive media coverage from Wall Street Journal [13], TechXplore [17], EurekaAlert [12], Cosmos [14], The San Diego-Union Tribune [18] and UCSD News [11, 15].

Helping robots see the world differently

One of my research work's cornerstones has been incorporating WiFi-based sensing into a robot's simultaneous localization and mapping (SLAM) stack. Introducing this sensing modality enables **more reliable SLAM operation in dark and monotonous indoor environments** by providing a new way for a robot to "see" the world around it. Consequently, this work is rooted in my primary vision for robots to have access to as many diverse sources of sensing for operation as possible. A first step towards this is to reduce SLAM's reliance on visual-based sensing for operation. SLAM is one of the most fundamental operations performed by a mobile robot, and failures here echo across the robot stack, impeding its autonomy, higher-level operation, and safety. Visual sensors have their inherent drawbacks, so in ping-pong SLAM (abbreviated as P2SLAM [2]), my colleagues and I devised a novel way to fuse WiFi CSI information instead to bring the worlds of WiFi localization and SLAM together. However, we observed two fundamental problems existing in each world. First, a large amount of drift was accumulated during quick yaw rotation in SLAM systems. Second, WiFi localization systems needed pre-determined WiFi access point locations prior to deployment. Our key insight to tackle these problems was to utilize angle-of-arrival (AoA) measurements and jointly determine the robot's and access point's 2D pose (location and orientation). This provided additional resilience to yaw rotation and allowed out-of-box operation in unknown environments. Consequently, we improved **odometry trajectory estimates by** $6\times$. Additionally, we open-sourced our datasets, which have been used and downloaded over 20 times across five universities and three countries in the past year.

As SLAM operations scale in size, "loop closures" need to be applied to the predicted robot trajectory to correct accumulated errors. SLAM systems can create constraints between different parts of the trajectory by applying correlations between what is currently seen and what was seen in the past. However, this correlation requires additional compute and memory storage. Loop closure hence become a necessary evil in SLAM systems. Instead, ViWiD [5] (under submission) showcased, for the first time, **the operation of SLAM systems without the aid of loop closures**. This required real-time integration of the WiFi-sensing systems built in P2SLAM. We devised signal averaging schemes to attenuate reflections in the environment and measure direct-path AoA $200\times$ more efficiently. We developed smart variable initialization schemes to ensure our graph optimizations converge in real-time. We developed and open-sourced the first ROS-compatible WiFi sensing framework to stream real-time, synchronized WiFi measurements for sensor fusion. Consequently, through 1500 m of cumulative travel, we found an end-to-end $4.3\times$ reduction in compute compared to camera-based SLAM [16] and a $4\times$ improvement in memory consumption over lidar-based SLAM [10]. The additional robustness provided by WiFi measurements also improved the 90^{th} percentile translation errors by $\sim 40\%$ and orientation errors by $\sim 60\%$ compared with purely camera-based systems.

Developing localization for next-gen applications

Physical intelligence is an obvious necessity for robots; however, it is equally imperative for mobile devices, AR/VR systems, and smart building automation. Through my localization research, **I have paved the way to bring physical intelligence to these systems as well**. Specifically, my work has helped automate systems that provide more operational context for smart buildings (DLoc [Mobicom '20], LocAP [NSDI '20]), make our factories, warehouses, and construction sites safer and more efficient (ULoc [IMWUT '21]), and simplify building digital twins of our physical world more seamlessly (XRLoc [Sensys '23]). The foundational systems built in these works will create a launchpad for the future development of physically intelligent systems.

Building large scale localization systems

A large part of developing physical intelligence is collecting troves of data to train models and understand trends. A core data point that needs to be collected is the indoor location of devices for better contextual understanding. However, most accurate localization schemes poorly scale to large spaces due to lack of automation. Consequently, my colleagues and I developed two key works to enable the deployment of large-scale wireless indoor localization technologies to ensure mobile devices have better contextual understanding. First, DLoc [9] leveraged deep learning techniques to better model the wireless channel in an environment and predict the user's location. Through this work, we primarily delivered best-in-class median and 90^{th} percentile localization accuracy, outperforming state-of-art systems by $\sim 80\%$. Additionally, we developed efficient automation schemes to collect large-scale datasets for model training. This open-sourced, first-of-its-kind dataset was downloaded over 130 times across 50 universities across five countries in the past three years. Second, my colleagues and I helped develop an essential pre-requisite for indoor localization – mapping and localization of the access point in a building. We discovered that measuring antenna separation within 5 mm of error in the global map was the key driver for decimeter-level localization. LocAP [8] measured this autonomously by relying solely on WiFi signals exchanged between a client and an AP. Our key idea was the relationship we discovered between the relative phase change measured between two antennas on the access point and their relative coordinates. We achieved 3 mm of median error in predicting the antenna separation through this observation and experiments using off-the-shelf WiFi access points.

Pushing the limits on localization

The most impactful way to leverage physical intelligence is to improve safety systems in dangerous working conditions like smart factories, construction sites, and warehouses. These safety monitoring systems demand accurate localization with sub-10 cm of errors. However, I found most systems in the market failed to provide the required accuracy, poorly scaled to large spaces, and were power-hungry at the tag, limiting their operation times. To fill this gap, my colleagues and I developed **the first-of-its-kind UWB localization module to measure the 3D Angle-of-arrival**. In ULoc [19], we prototyped a board with eight clock-synchronized UWB modules arranged in an L-shape to measure the azimuth and elevation angle of arrival from a single UWB transmission at the tag. It localized tags with a median accuracy of ~ 10 cm in dynamic settings with $9\times$ more power efficiency at the tag compared to time-of-travel-based systems. We additionally demoed this system [3] and open-sourced the hardware files for the benefit of the community. We are now actively supporting users of our hardware at both UCSD and Johns Hopkins University.

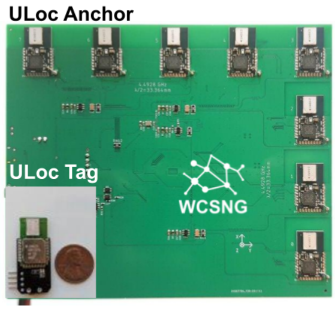

ULoc anchor and tag hardware

Can we do it better and more easily?

Few-centimeter accurate localization to track people and assets plays a key role in bringing the physical world into the virtual world. However, for these localization systems to be consumer-friendly, they need to be easily deployable. ULoc and most other accurate localization systems often require the cumbersome deployment of multiple anchors for reliable location estimates. In XRLoc [6], my colleagues and I developed a localization module (sized $\sim 1m$ in length) with the capability to **measure a UWB tag's location to within a few centimeters**. This simplified deployment to setting up a single module, but single-point localization severely degraded localization accuracy by over $10\times$ to many tens-of-centimeters. Through a novel fusion of time and phase measurements of UWB signal within a particle filter, we improved UWB localization to ~ 2 cm of median accuracy in dynamic scenarios. Consequently, this work has positioned UWB-based localization as a strong contender to camera-based tracking in AR/VR while also promising more robustness and privacy.

Wireless sensing to drive immediate impact

Through my research, I have driven impact across three additional applications, primarily driven by my desire to solve socially important and industry-relevant problems. Furthermore, solving these problems has utilized similar fundamentals to delivering the above applications, indicating the broad scope of wireless sensing.

BluBLE: developing effective contact-tracing

Contact tracing was employed to monitor the spread of the virus early in the pandemic, which often involved interviewing sick patients to understand their potential contacts with others. This manual and laborious process was the gold standard in contact tracing as it curtailed false positives. Instead, many BLE-based contact tracing applications were deployed, which relied on BLE beacon discovery to indicate contact with a person to automate the process. However, this binary measure (presence or absence of a beacon) of contact leads to many false positives, as BLE signals travel through wooden walls or sheetrock. Realizing this key problem with existing systems, in BluBLE [4], we devised ML algorithms to reliably classify peer-to-peer contacts (with a median accuracy of 96%) and create a social contact graph (with an accuracy of at least 88%). This reliably automated the contact tracing process and was later **prototyped with Aarogya Sethu, India's contact tracing platform**.

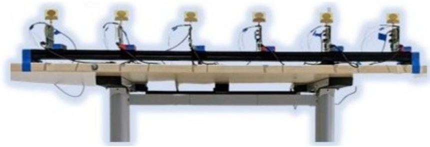
XRLoc localization module

Preserving location privacy

There has been recent concern about the use of WiFi localization to track people in public spaces to build better customer profiles. Furthermore, my past research in improving localization accuracy deepens this problem's impact. Hence, I found it imperative to develop systems that can allow users to protect their privacy from being localized unknowingly by these systems. In Mirage [7], my colleagues and I prototyped and conducted preliminary testing to leverage multipath in the environment to confuse an access point's ability to measure the accurate angle of arrival. Mirage degrades[1] the AoA performance by $\sim 2\times$ to 46° and localization performance by $\sim 5\times$ to over 10 m, while the spatial multiplexing gains and data throughput at the AP remain consistent. In effect, Mirage helps to undo over two decades of research in improving localization performance.

[1] likely the only time I use this word in a positive context

Wireless time sync

Time and frequency synchronization across distributed computing and sensing systems is critical in load-balancing, job scheduling, and jointly sensing the environment. In my ongoing work, which is a **collaboration between NICT, Japan, and Meta** under the Open Compute Project, we seek to provide this sync across WiFi communication channels using off-the-shelf WiFi IoT chipsets. The key idea we are exploring is using WiFi CSI measurements to estimate the carrier frequency offset and the sampling time offset between a pair of nodes to apply closed-loop corrections to the onboard base clock. Through this, we hope to provide a nano-second level of time synchronization and a milli-hertz level of frequency syntonization via a packet handshake between the devices. The preliminary results of this work were presented at OCP's Global Summit [1].

Future work: Pushing the bounds on wireless sensing

To further the world of automation beyond large language models and RL-based robot control, it is imperative to get a broader understanding of the environment. Our automated systems need to have better physical intelligence about their surroundings for improved human-computer interactions and better decision-making. Ironically, current visual-based sensing techniques create a blind spot for these systems. I would like to continue exploring how we can "see" the world differently to plug gaps in the current visual-based methods. Along these lines, I will **expand the use of RF-based sensing** to perceive the world better and **explore using ultrasound sensing and communications** to build more collaborative marine robot systems.

Wireless throughout the robot stack: Many robot systems require sensing beyond the limited field of view provided by visual-based sensors to be more effective in their tasks. We can surmount these challenges by continuing to incorporate RF-based sensing throughout the robot stack. A few fundamental problems that I envision we can solve are:

Exploration: An essential task of autonomous robots is to explore their environment to completion quickly. However, limited visual sensing and lack of GPS indoors can prevent a robot from planning its path optimally. Instead, I hypothesize WiFi and BLE devices can provide beyond line-of-sight understanding to improve a robot's exploration efficiency.

Collaboration: Limited visual sensing can lead to a limited spatial understanding and, hence, ineffective use of the robot fleets. I will explore extensions to P2SLAM and ViWiD to build collaborative robot systems that leverage WiFi-based sensing to provide more accurate and real-time estimates of the robots' positions and the global map.

Control: Developing robot control systems based on wireless sensing would allow another layer of certainty for robots operating in new environments. I will explore using ultrawideband mm-wave radars to develop end-end sensing and control algorithms to take further advantage of the robustness of RF-based sensing.

Wireless under water: Through my advisory role at YonderDeep, an undergraduate student org, we are developing and testing O-RAN[2] systems to provide better connectivity for marine robots. There is a clear and urgent need to develop real-time communication and fine-grained sensing systems for teams of marine robots to improve mapping and reconnaissance for defense, environmental, and exploratory purposes. A few fundamental problems I will explore are:

Acoustic-mesh networks: Teams of underwater marine robots will require peer-to-peer connections for both sensing and communications. I will explore the development of these mesh systems, which in turn can communicate with on-shore communication infrastructure, to improve robot collaboration.

Underwater 3D SLAM: These teams will additionally require accurate positioning and mapping systems. Lack of GPS, unclear visual landmarks, and full 6-DoF sensing make this problem challenging. I would like to explore and build novel acoustic systems to enable localization in these environments.

Wireless for AR/VR: Through my research, I have built better localization and tracking systems to transfer the physical world into the digital world. By providing few-cm scale localization accuracy and mm-level tracking accuracy, it has provided a viable way to transport everyday objects into the virtual world. However, many challenges continue to exist to make this transition seamless. A few challenges I will explore are:

Full body tracking: Current IMU and camera-based full-body pose systems are underconstrained or highly privacy-invasive. Instead, I would like to develop deterministic and easily deployable UWB localization systems to provide full-body real-time and accurate full-body pose estimates.

Immersive spatial audio: Besides crafting visual experiences, creating immersive audio experiences is equally crucial. However, delivering realistic, binaural audio is challenging without a millimeter-accurate location of the listener's ears. I would like to explore building more accurate localization schemes to extend my current research in delivering spatial audio applications.

[2] O-RAN: Open Radio Access Networks

Papers

[2] **Arun, Aditya**, Ayyalasomayajula, R., Hunter, W., & Bharadia, D. (2022). P2slam: Bearing based wifi slam for indoor robots. (Vol. 7, pp. 3326–3333). IEEE Robotics and Automation Letters.

[3] **Arun, Aditya**, Chang, T., Yu, Y., Ayyalasomayajula, R., & Bharadia, D. (2022). Real-time low-latency tracking for uwb tags. In *Proceedings of the 20th annual international conference on mobile systems, applications and services* (pp. 611–612).

[4] **Arun, Aditya**, Gupta, A., Bhatka, S., Komatineni, S., & Bharadia, D. (2020). Bluble, space-time social distancing to monitor the spread of covid-19. (pp. 750–751). Proceedings of the 18th Conference on Embedded Networked Sensor Systems.

[5] **Arun, Aditya**, Hunter, W., Ayyalasomayajula, R., & Bharadia, D. (2022). Viwid: Leveraging wifi for robust and resource-efficient slam, arXiv preprint arXiv:2209.08091.

[6] **Arun, Aditya**, Saruwatari, S., Shah, S., & Bharadia, D. (2023). Xrloc: Accurate uwb localization to realize xr deployments, Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems.

[7] Ayyalasomayajula, R., **Arun, Aditya**, Sun, W., & Bharadia, D. (2023). Users are closer than they appear: Protecting user location from wifi aps. (pp. 124–130). Proceedings of the 24th International Workshop on Mobile Computing Systems and Applications.

[8] Ayyalasomayajula, R., **Arun, Aditya**, Wu, C., Rajagopalan, S., Ganesaraman, S., Seetharaman, A., ... Bharadia, D. (2020, February). LocAP: Autonomous millimeter accurate mapping of WiFi infrastructure. In *17th usenix symposium on networked systems design and implementation (nsdi 20)* (pp. 1115–1129). Santa Clara, CA: USENIX Association. Retrieved from https://www.usenix.org/conference/nsdi20/presentation/ayyalasomayajula

[9] Ayyalasomayajula, R., **Arun, Aditya**, Wu, C., Sharma, S., Sethi, A. R., Vasisht, D., & Bharadia, D. (2020). Deep learning based wireless localization for indoor navigation. (pp. 1–14). Proceedings of the 26th Annual International Conference on Mobile Computing and Networking.

[19] Zhao, M., Chang, T., **Arun, Aditya**, Ayyalasomayajula, R., Zhang, C., & Bharadia, D. (2021). Uloc: Low-power, scalable and cm-accurate uwb-tag localization and tracking for indoor applications. (Vol. 5, pp. 1–31). Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies.

News Articles

[11] Labios, L. (2022a). Using everyday wifi to help robots see and navigate better indoors. *UCSD News.* Retrieved December 8, 2023, from https://today.ucsd.edu/index_redis_default_site.php?/story/using-everyday-wifi-to-help-robots-see-and-navigate-better-indoors&/pressrelease/using-everyday-wifi-to-help-robots-see-and-navigate-better-indoors

[12] Labios, L. (2022b). Using everyday wifi to help robots see and navigate better indoors. *EurekaAlert.* Retrieved December 8, 2023, from https://www.eurekalert.org/news-releases/953372

[13] Mims, C. (2022). The end of car keys, passwords and fumbling with your phone at checkout. *Cosmos.* Retrieved December 8, 2023, from https://www.wsj.com/articles/the-end-of-car-keys-passwords-and-fumbling-with-your-phone-at-checkout-11641618007

[14] Nasrullah, Q. (2022). Using wifi to help robots navigate more accurately indoors. *Cosmos.* Retrieved December 8, 2023, from https://cosmosmagazine.com/technology/wifi-robots-navigate-more-accurately/

[15] Rojas-Rocha, X. (2020). Bluble: Estimating your covid-19 risk with accurate contact tracing. *UCSD News.* Retrieved December 8, 2023, from https://today.ucsd.edu/story/bluble-estimating-your-covid-19-risk-with-accurate-contact-tracing

[17] University of California, S. D. (2021). New upgrades to old wireless tech could enable real-time 3d motion capture. *TechXplore.* Retrieved December 8, 2023, from https://techxplore.com/news/2021-09-wireless-tech-enable-real-time-3d.html

[18] Winkley, L. (2020). Your phone could be telling you if you've been exposed to covid-19. here's why it's not. *The San Diego Union-Tribune.* Retrieved December 8, 2023, from https://www.sandiegouniontribune.com/news/health/story/2020-08-04/your-phone-could-be-telling-you-if-youve-been-exposed-to-covid-19-heres-why-its-not

References

[1] (2023, September). Retrieved from https://www.opencompute.org/events/past-events/2023-ocp-global-summit#time_appliances_project

[10] Hess, W., Kohler, D., Rapp, H., & Andor, D. (2016). Real-time loop closure in 2d lidar slam. IEEE.

[16] Rosinol, A., Violette, A., Abate, M., Hughes, N., Chang, Y., Shi, J., ... Carlone, L. (2021). Kimera: From slam to spatial perception with 3d dynamic scene graphs. SAGE Publications Sage UK: London, England.