Lecture 8: Indoor Localization

Chenshu Wu

Department of Computer Science

2025 Spring

香 港 大 學
THE UNIVERSITY OF HONG KONG

HKU AIoT LAB
香港大學人工智能物联网實驗室

Contents

A Real Event

- Your car was stolen and parked somewhere.
- The only information that you can still have is the car App
  - for remote engine start & lock
  - that offers a relative distance to the owner's phone

- *What would you do to find the car?*

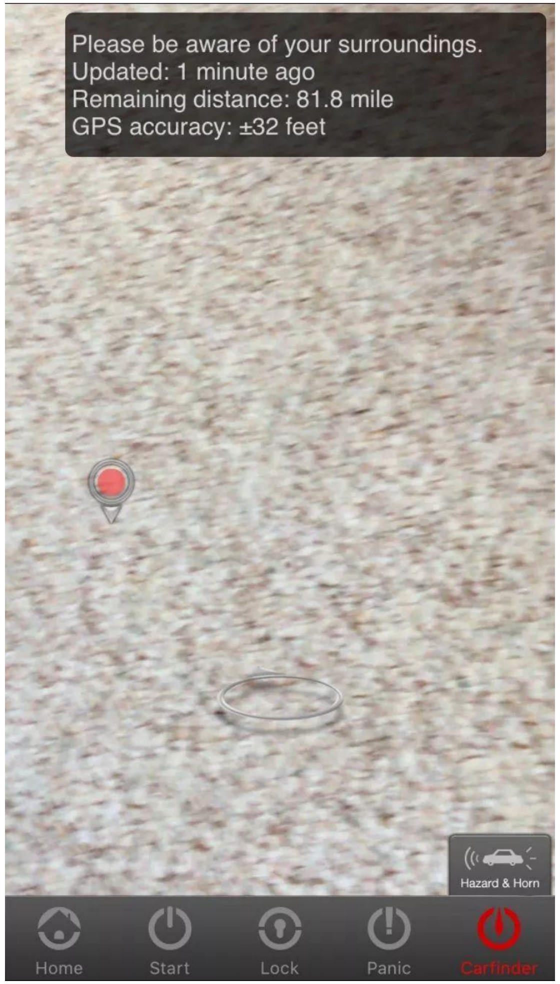

Mazda Mobile Start

WHERE Am I?

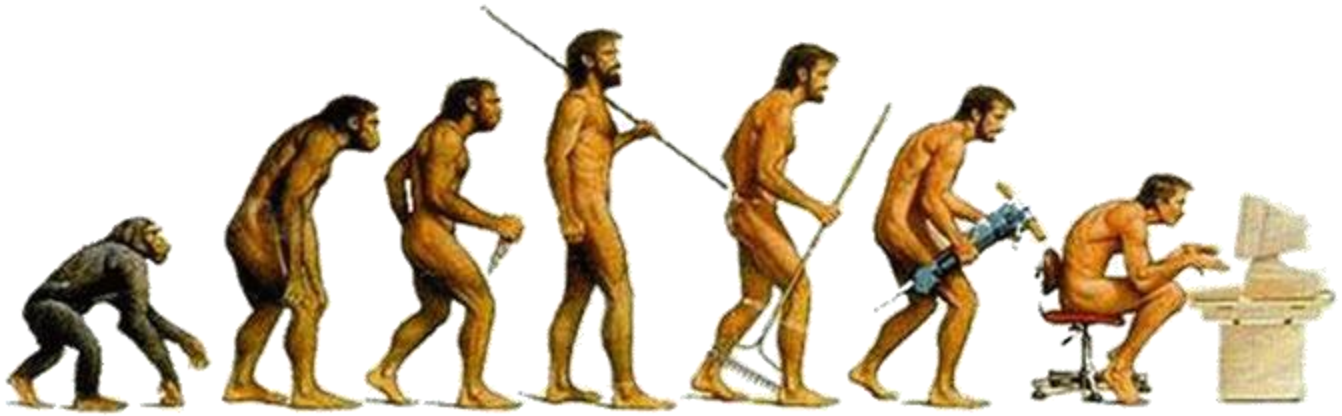

Who Are You,

Where Are You Going,

Where Have You Been?

Localization: A long history…





"Localization" ≈ GPS

1957
Sputnik

Transit 1959

NAVSTAR 1974

1978-85
Block I Satellite

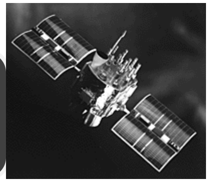

Civilian use 1983

1989-94
Block II satellite
Hand-held RX

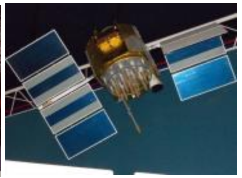

Worldwide use FREE 1994

1999
Benefo
n Esc

Final deactivation of selective availability 2000

2001
In-car navigation

A-GPS 2004

31 satellites / 3m 2014

2015
GPS III

Indoor Positioning

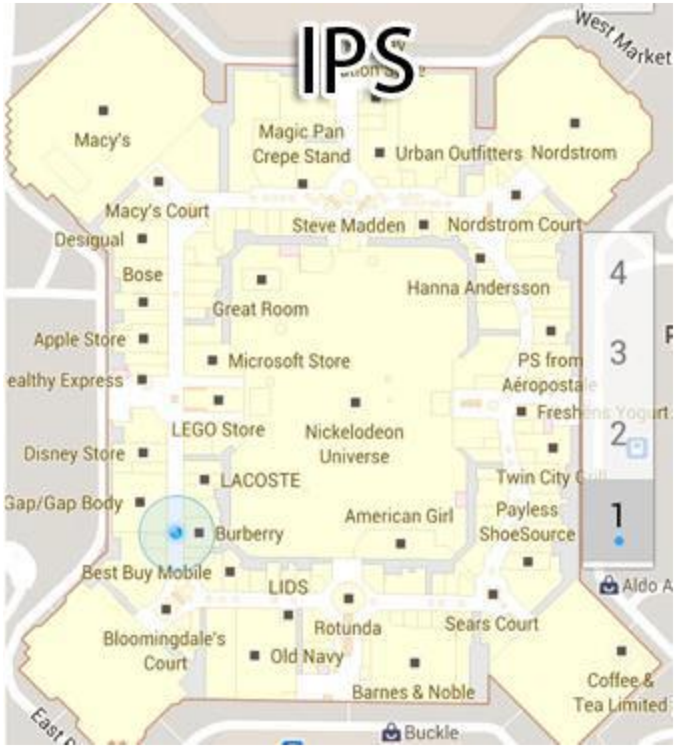

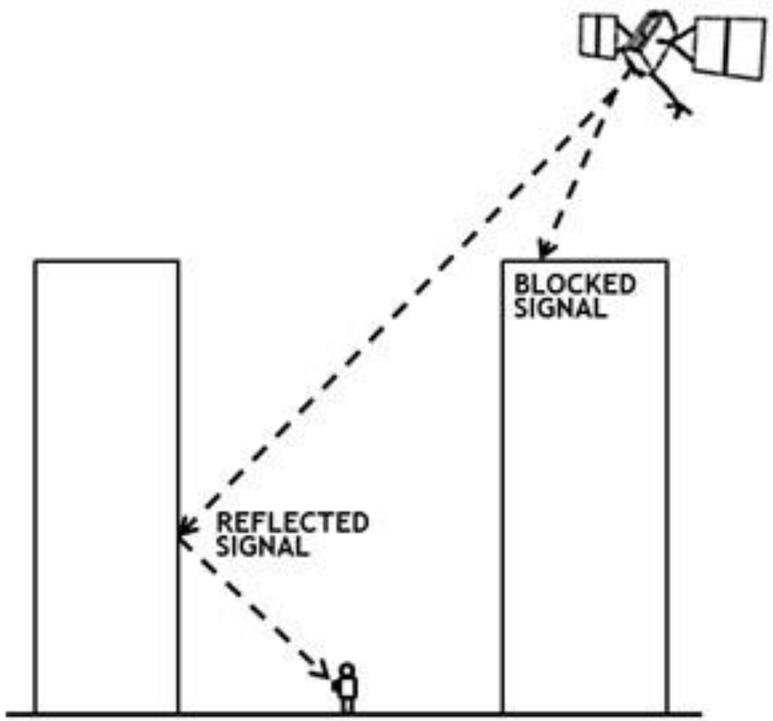

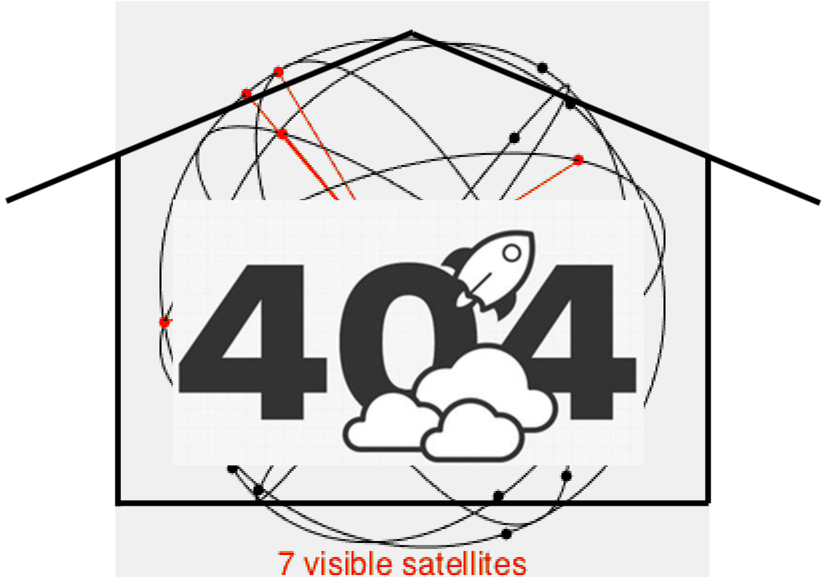

Sorry
GPS signals NOT found!

Indoor Positioning

Robot Navigation

VR Gaming

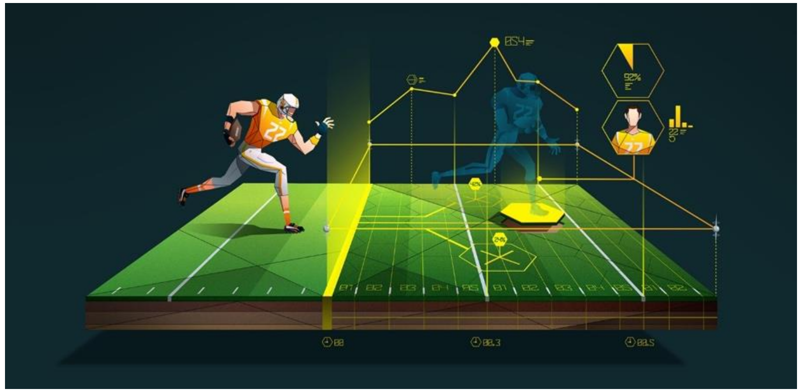

Sports Analytics

Mobiles & Wearables

Robots

Drones

Device-based vs. Device-free

- Two different contexts
  - Device-based: A user carries a certain device in order to be located
  - Device-free: A user can be located without carrying/wearing any devices
- Our focus: Device-based approaches
  - GPS
  - Smartphone localization
  - Robot/asset tracking
  - …
- Device-free approaches are more related to contactless wireless sensing (last topic).

Indoor Positioning: 30+ years

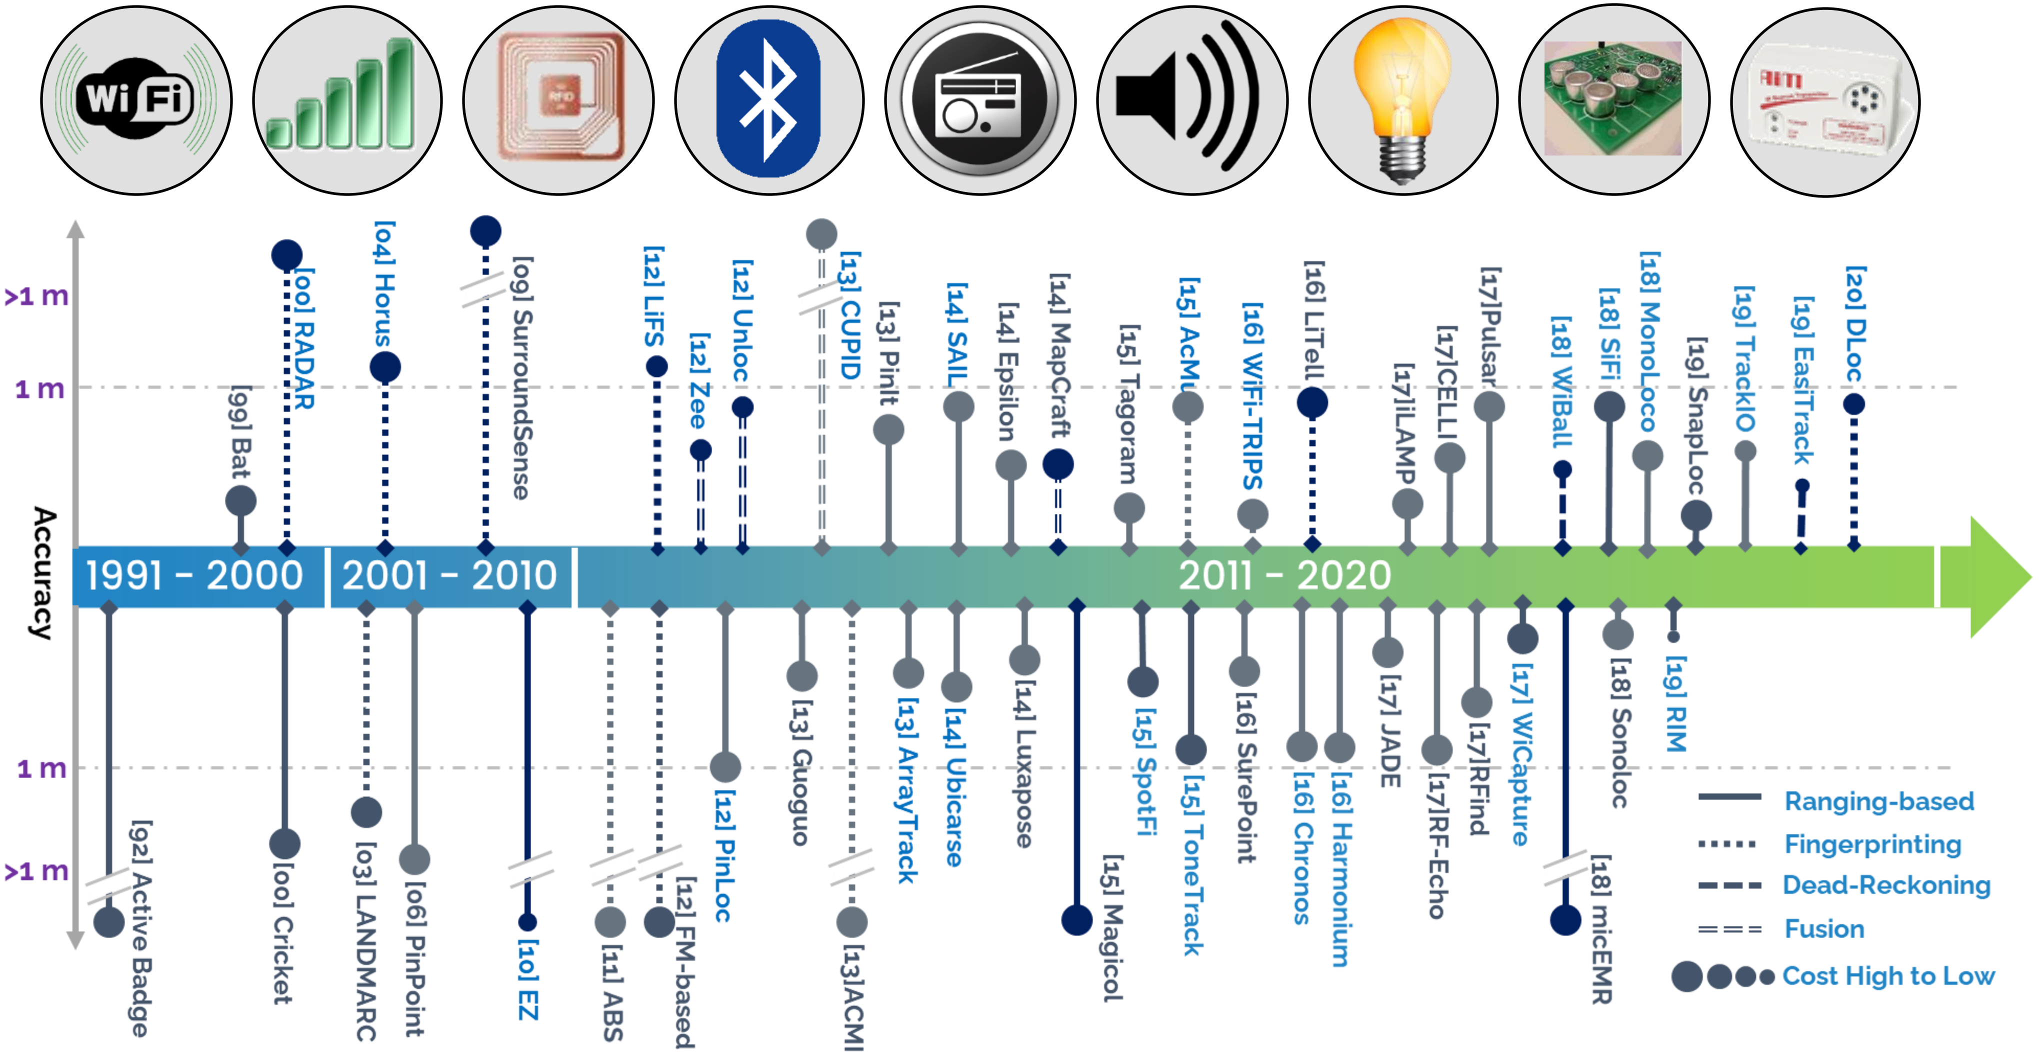

An Idea Indoor “GPS”

- Despite 30+ years of worldwide efforts, we still do not have a practical solution today that scales to the world.
- Why?

- Among many reasons
  - No worldwide “GPS” – infrastructure
  - Complex indoor environments
  - Higher accuracy requirement than outdoors
    - ~1 m needed to differentiate neighboring rooms, aisles in supermarkets…

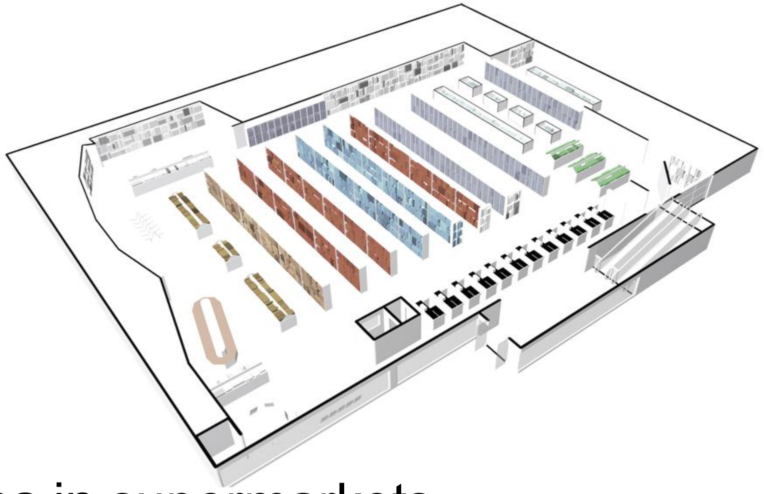

An Idea Indoor “GPS”

- Accurate
  - ~1 m
- Robust
  - Environmental changes/dynamics
- Scalable
  - Worldwide buildings and global users
- Easy-to-install
  - Infrastructure-free
- Coverage
- Sustainable

Example Architecture

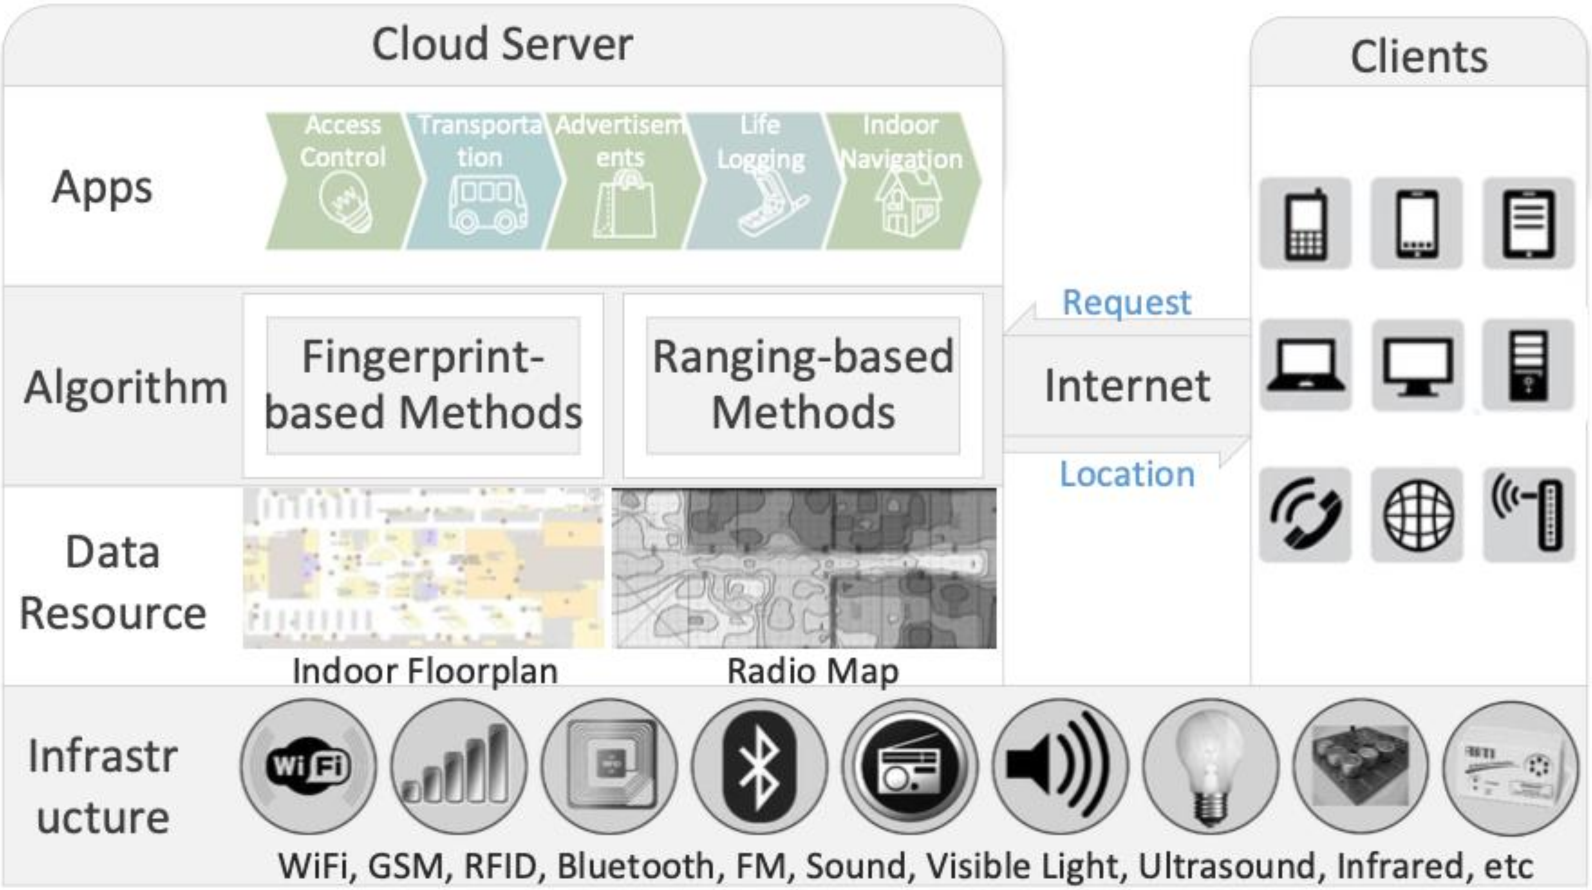

Three Mainstream Approaches

Triangulation
Angulation / Lateration

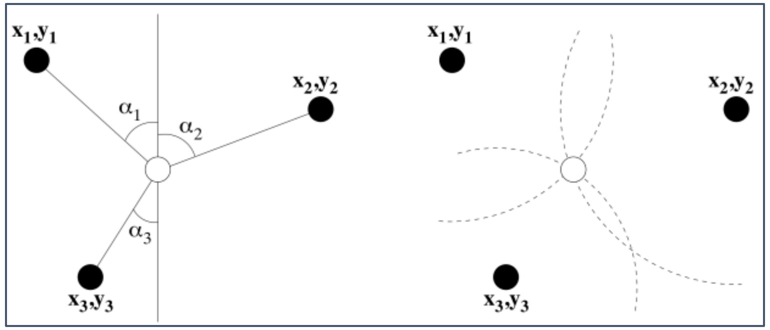

Fingerprinting
Location matching

Dead-Reckoning
Inertial Tracking

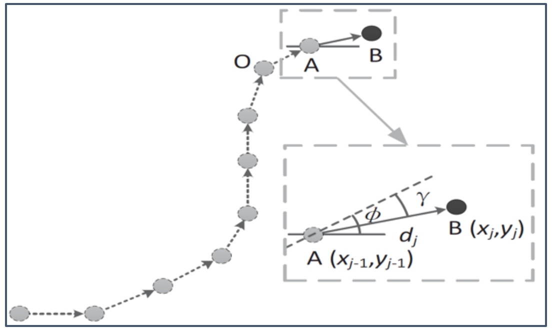

- Metrics
  - Accuracy: ~1 m
  - Cost: hardware, installation, deployment, maintenance, etc
  - Coverage: How large space can be supported?
  - Scalability: How many users/buildings to support?

Early Systems (1)

Active Badge

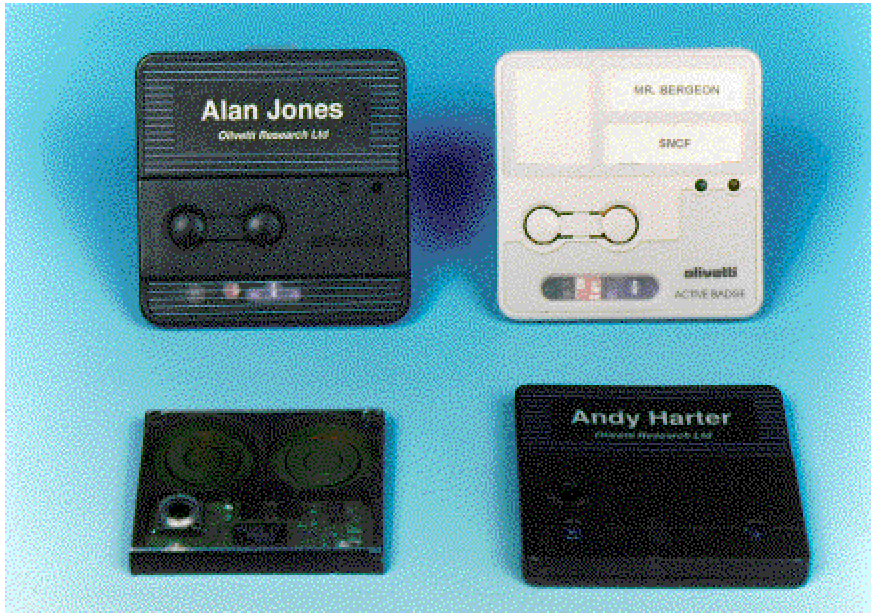

Designed and prototyped between 1989 and 1992
By Andy Hopper etc, Olivetti Research Lab (ORL)

Signals: infra-red signals
Beacons: Pre-deployed networked infra-red receivers
Tags: small active badge

Technology: landmarks
Accuracy: room scale

Early Systems (2)

Bat Ultrasonic Location System

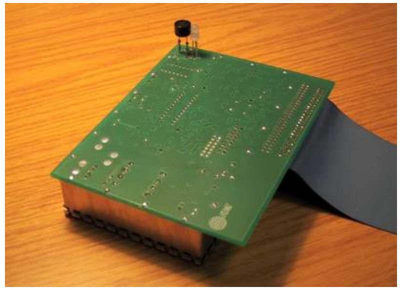

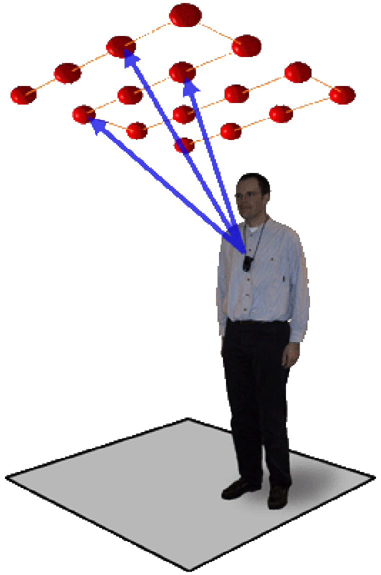

Designed and prototyped between 1997 and 2001
By Andy Hopper etc, AT&T Cambridge Lab

Signals: short pulse of ultrasonic
Beacons: pre-deployed networked ultrasonic sensors
Tags: ultrasonic transmitter (a *Bat*)

Technology: triangulation
Accuracy: centimeter

Early Systems (3)

Cricket

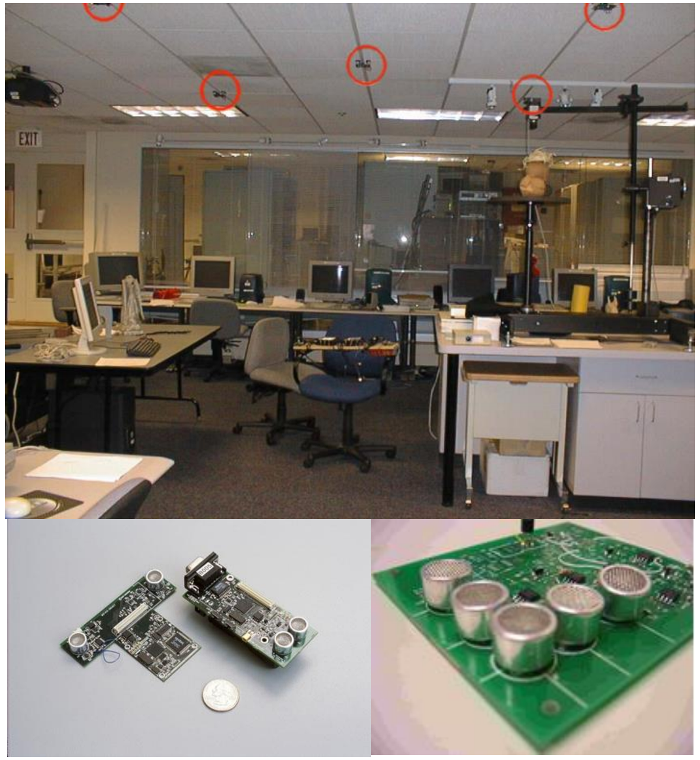

Designed and prototyped between 2000 and 2006
By Hari Balakrishnan etc, CSAIL MIT

Signals: ultrasonic signals & RF signals
Beacons: Ceiling-mounted, transmitted concurrent RF and ultrasonic signals
Listeners: small active badge

Technology: landmarks
Accuracy: centimeter / room-level granularity

Trilateration/Triangulation

- The approach of GPS

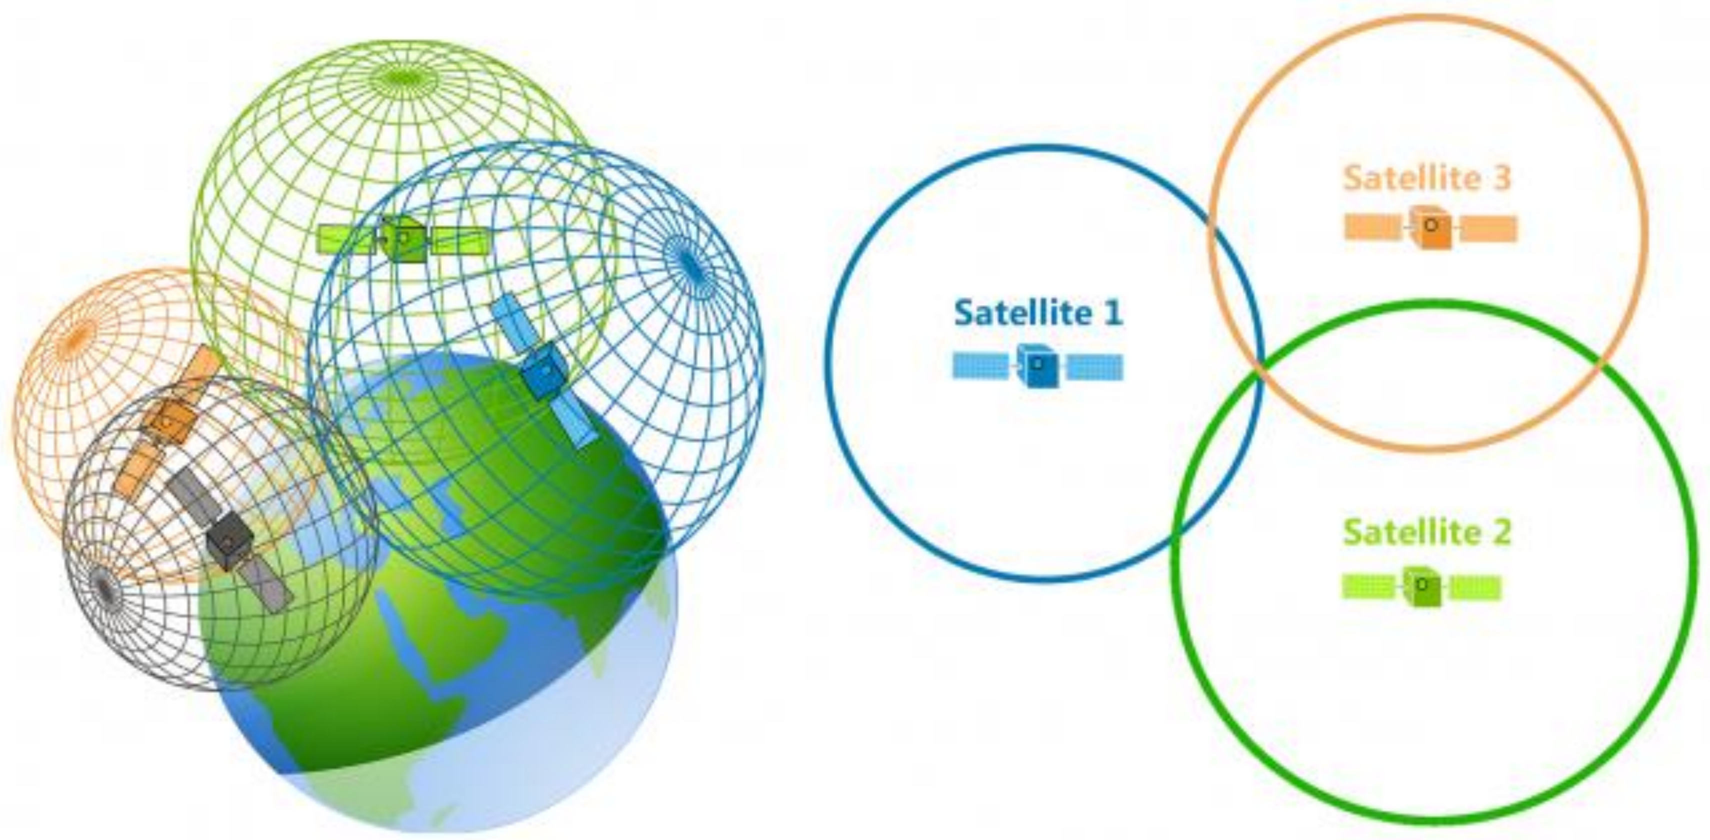

Trilateration/Multilateration

- Given
  - $R_i = (x_i, y_i), i = 1,2,3, \cdots N$
  - d_i: distance from X to R_i
- Solve $X = (x, y)$: $\hat{X} = \underset{X}{\operatorname{argmin}} \sum_{i=1}^{N} ||X - R_i||$
- A Solution using Least Squares Method

$$\begin{cases} (x - x_1)^2 + (y - y_1)^2 = d_1^2 \\ \vdots \\ (x - x_N)^2 + (y - y_N)^2 = d_N^2 \end{cases} \Rightarrow AX = b \Rightarrow \hat{X} = (A^T A)^{-1} A^T b$$

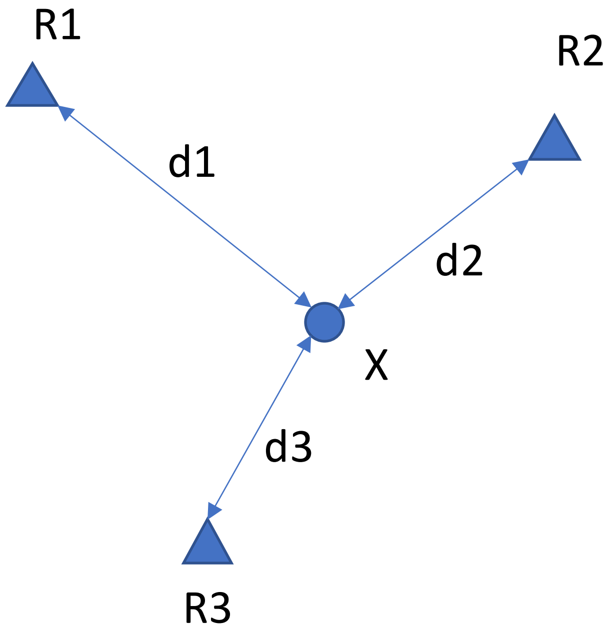

Trilateration/Multilateration

- Ranging is the key
  - What infrastructure/technology to use?
  - What ranging approach to use?

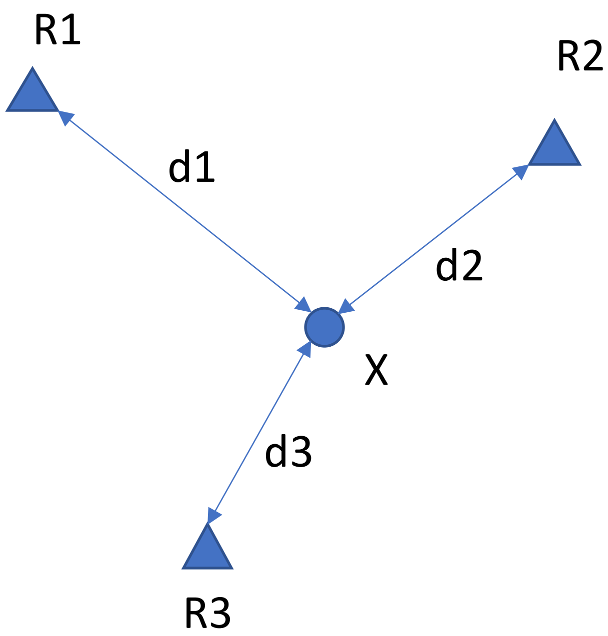

- Recall ranging

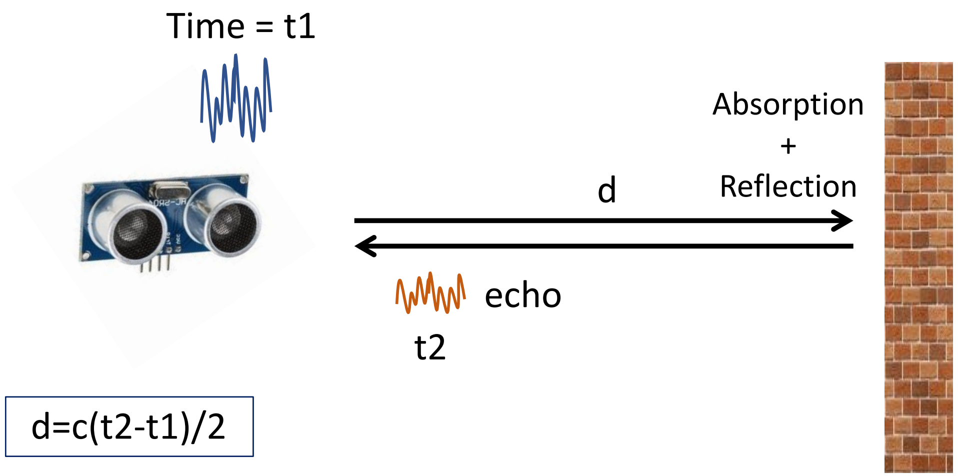

Range Resolution depends on the bandwidth

$$d_{res} = \frac{c}{B}$$

Acoustic Ranging

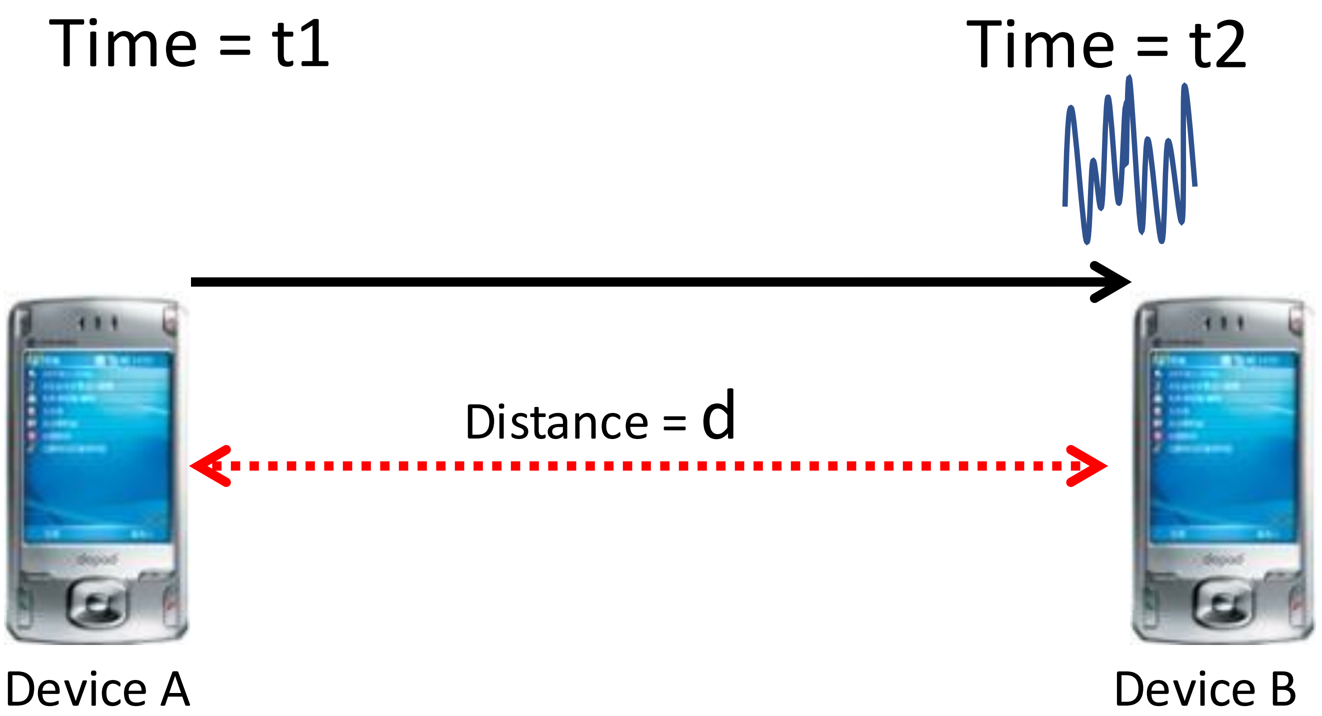

BeepBeep: a high accuracy acoustic ranging system using COTS mobile devices, SenSys'07

Root causes of inaccuracy

- Clock synchronization errors
- Sending/receiving uncertainties
  - 1 ms error in time = 34 cm error in distance
  - 1 cm ranging accuracy requires 30us timing accuracy

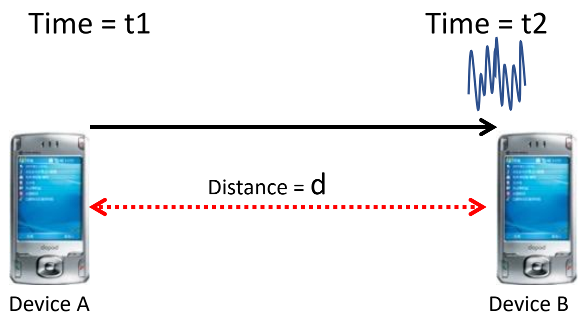

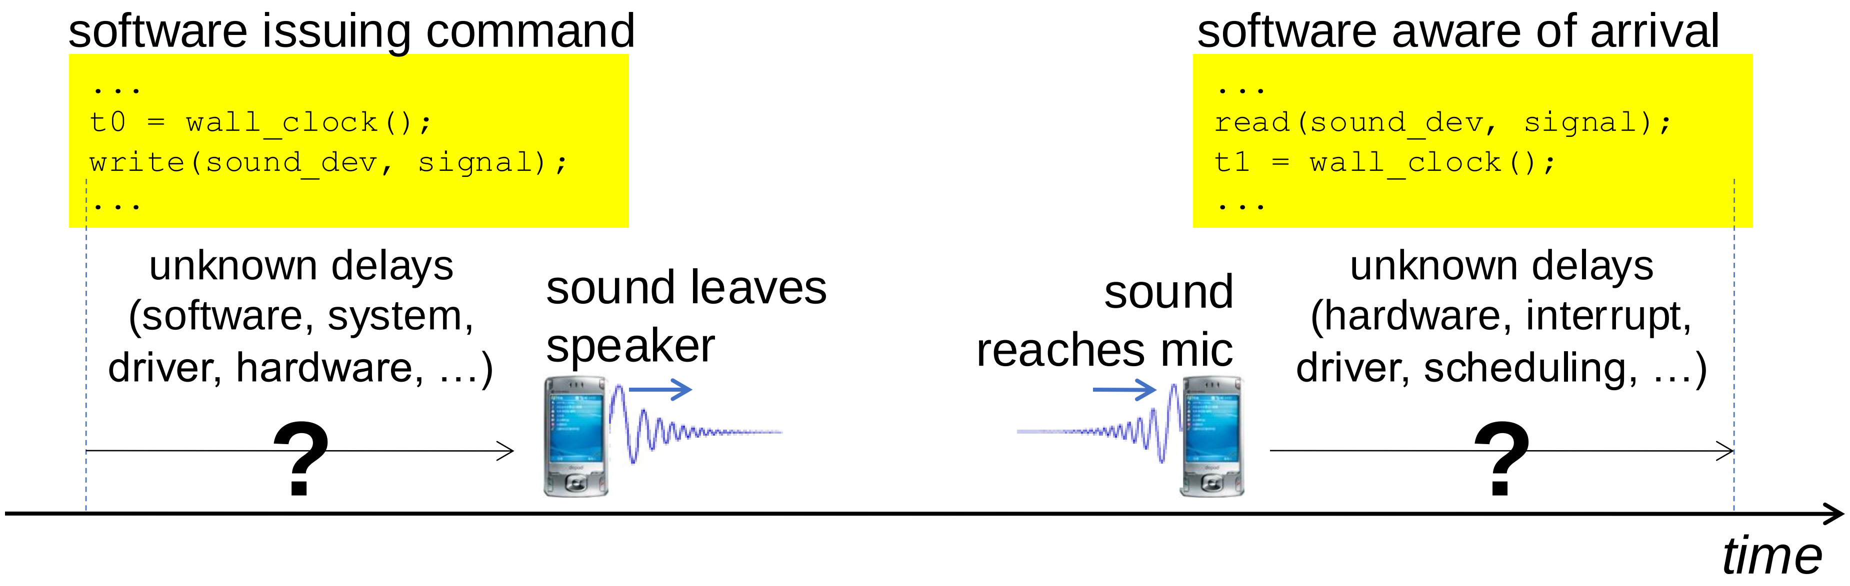

Acoustic Ranging: BeepBeep

1. Device A emits a beep while both recording
2. Device B emits another beep while both continue recording
3. Both devices detect TOA of the two beeps and obtain respective ETOAs
4. Exchange ETOAs and calculate the distance

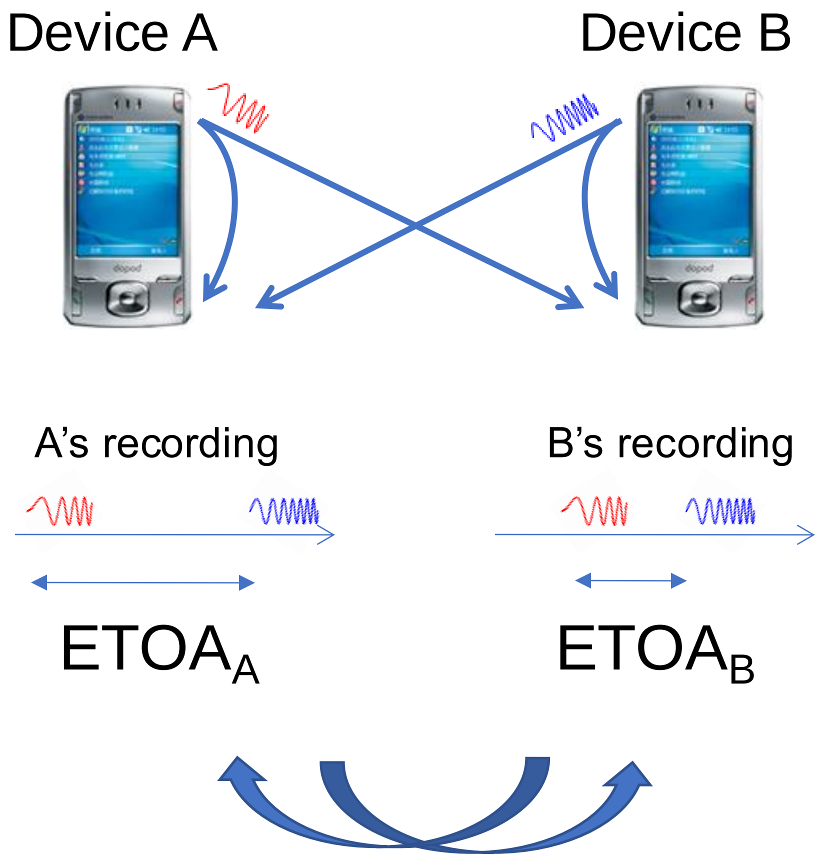

$D_{AB}=|ETOA_A-ETOA_B|/2$

Timeline

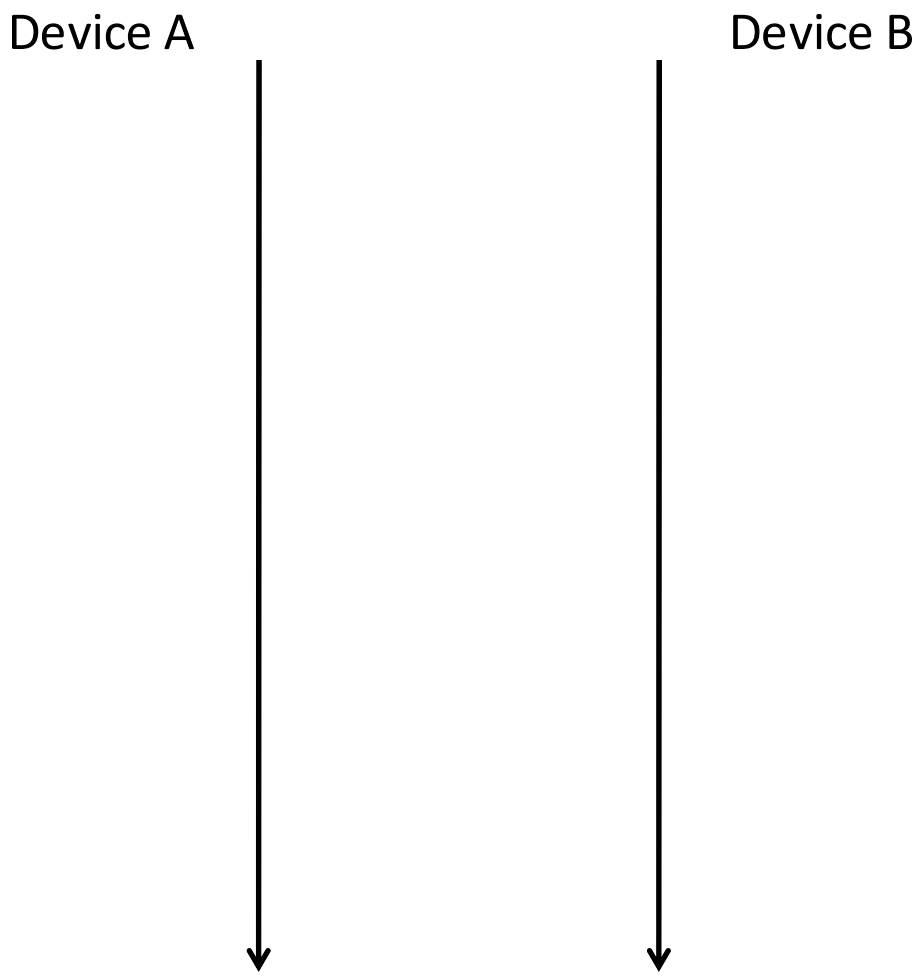

Timeline

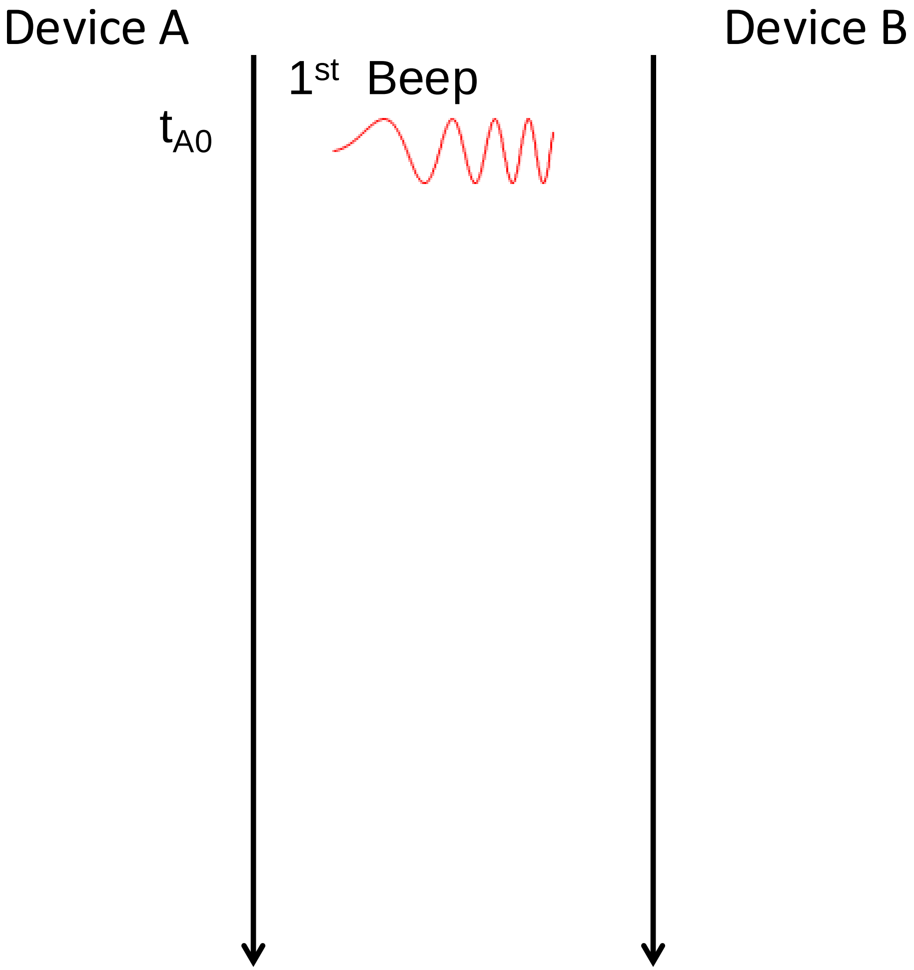

Timeline

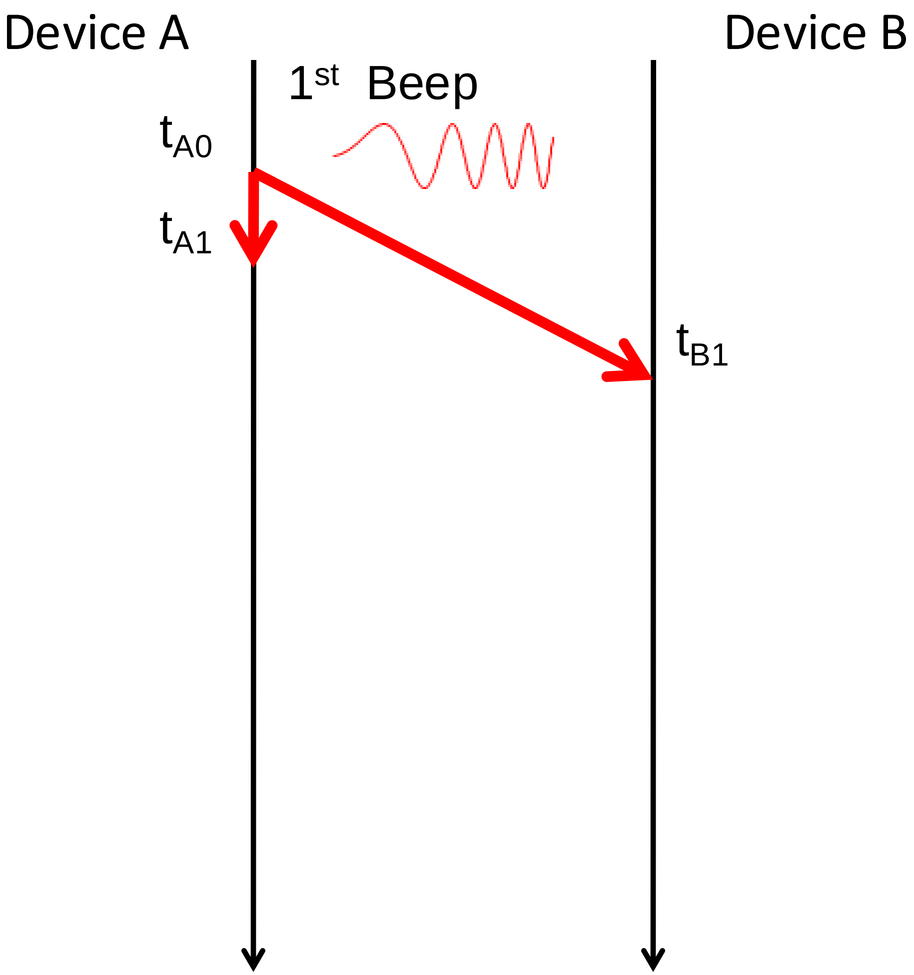

Timeline

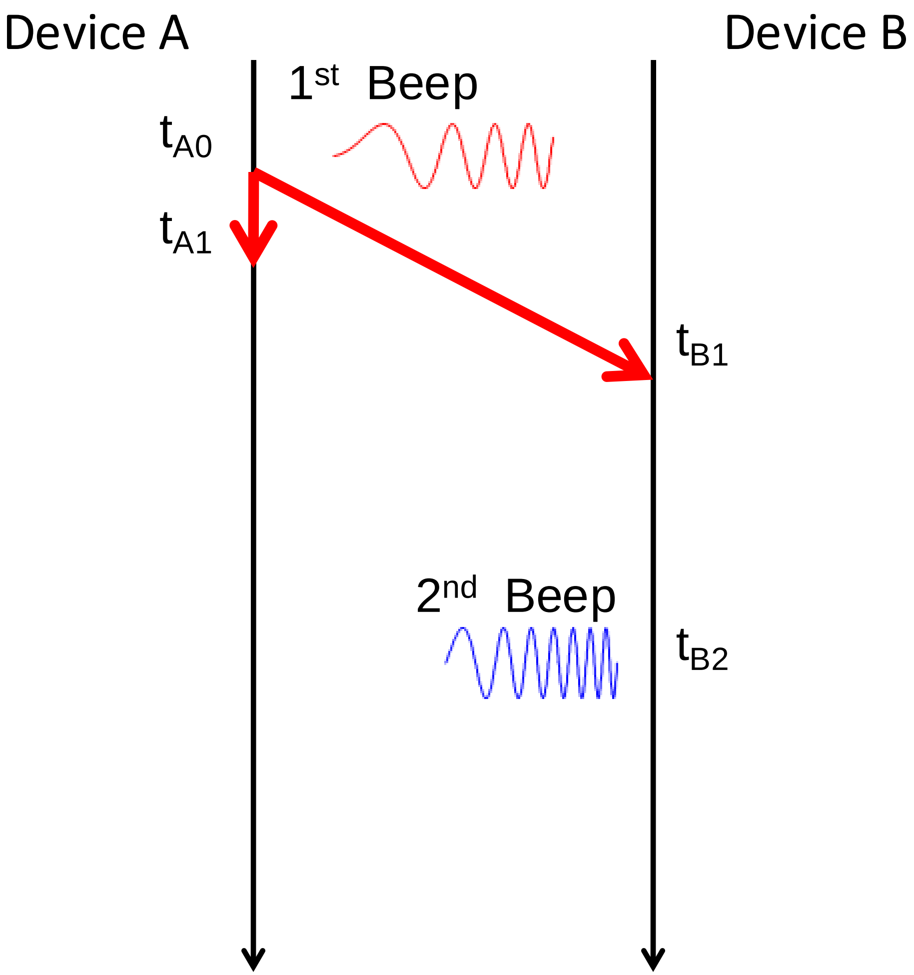

Timeline

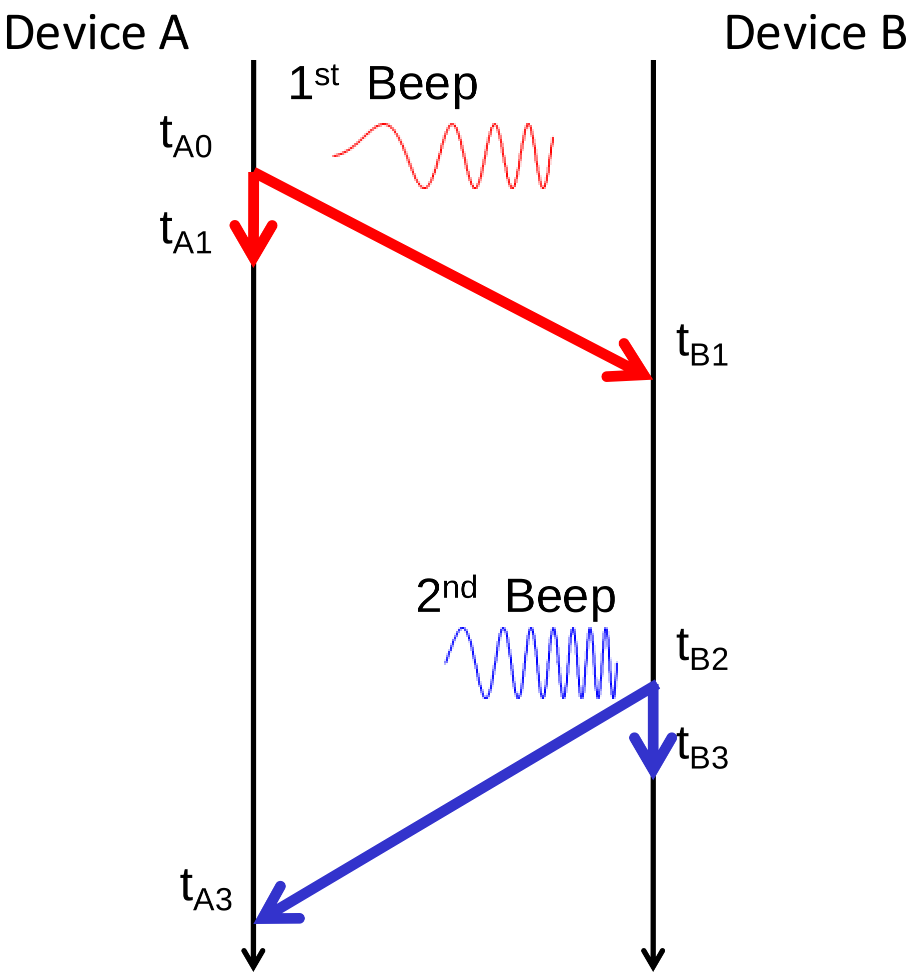

Timeline

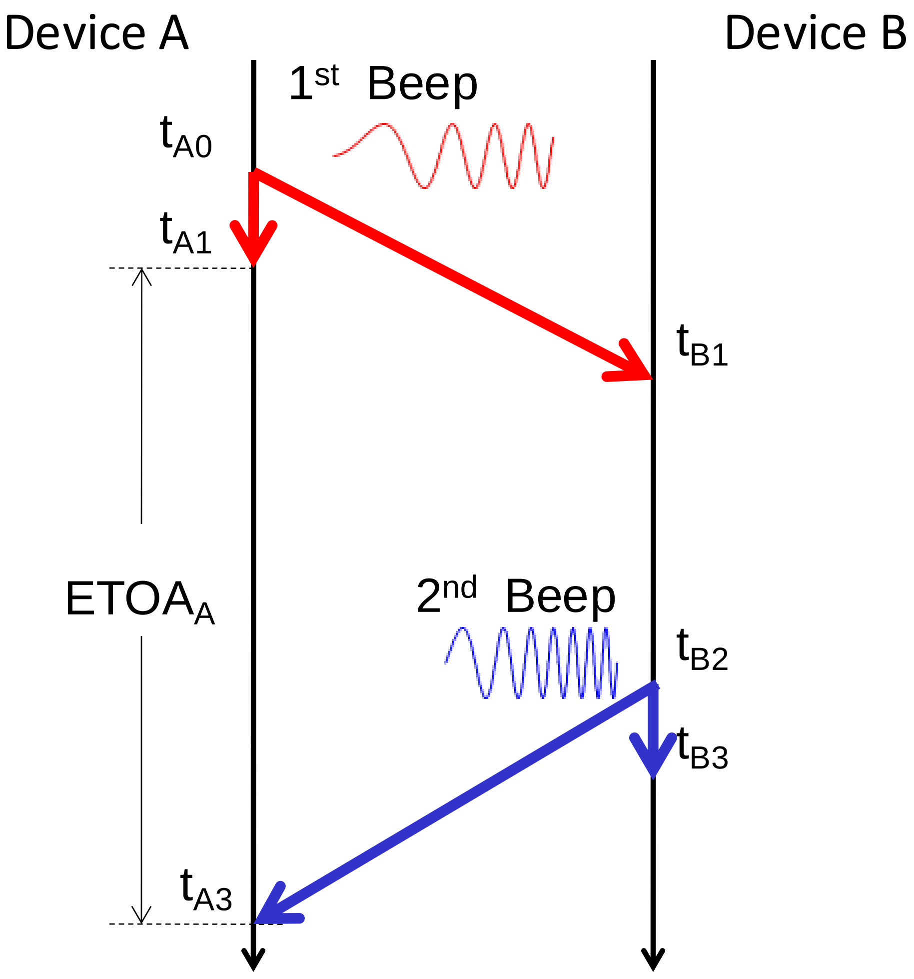
Device A
Device B
1st Beep
t_{A0}
t_{A1}
t_{B1}
$ETOA_A$
2nd Beep
t_{B2}
t_{B3}
t_{A3}

Timeline

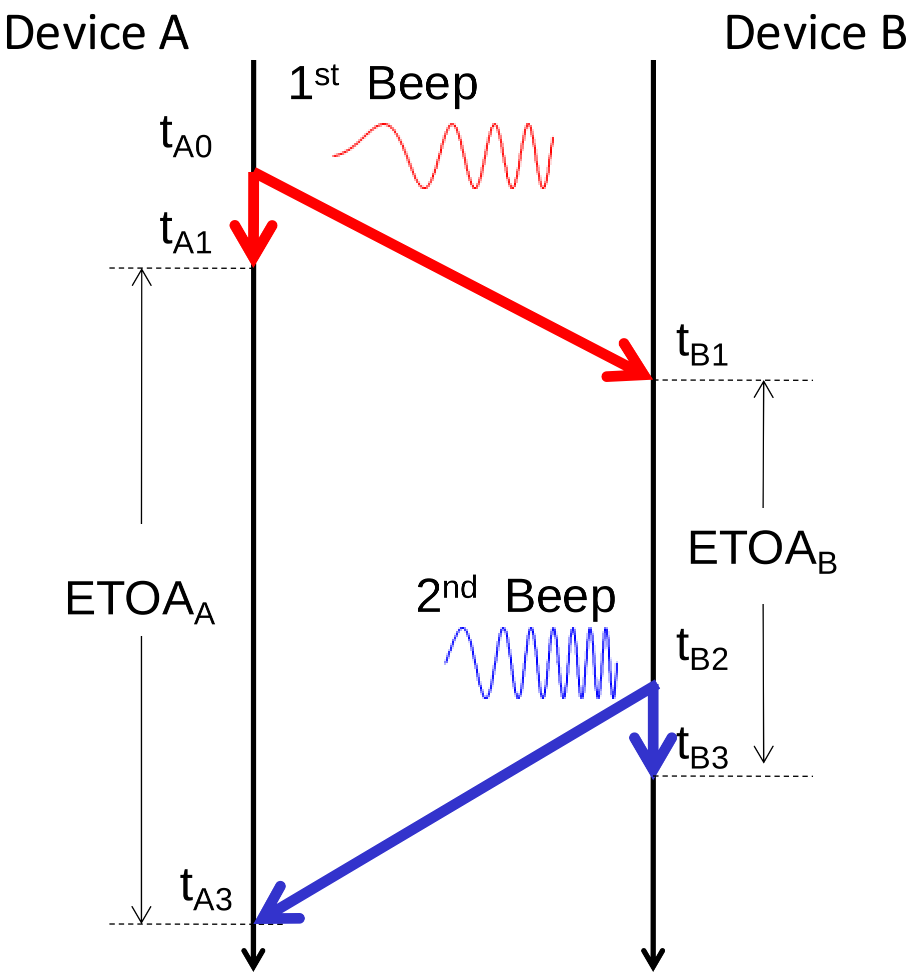

Device A
Device B
1st Beep
t_{A0}
t_{A1}
t_{B1}
$ETOA_B$
$ETOA_A$
2nd Beep
t_{B2}
t_{B3}
t_{A3}

Timeline

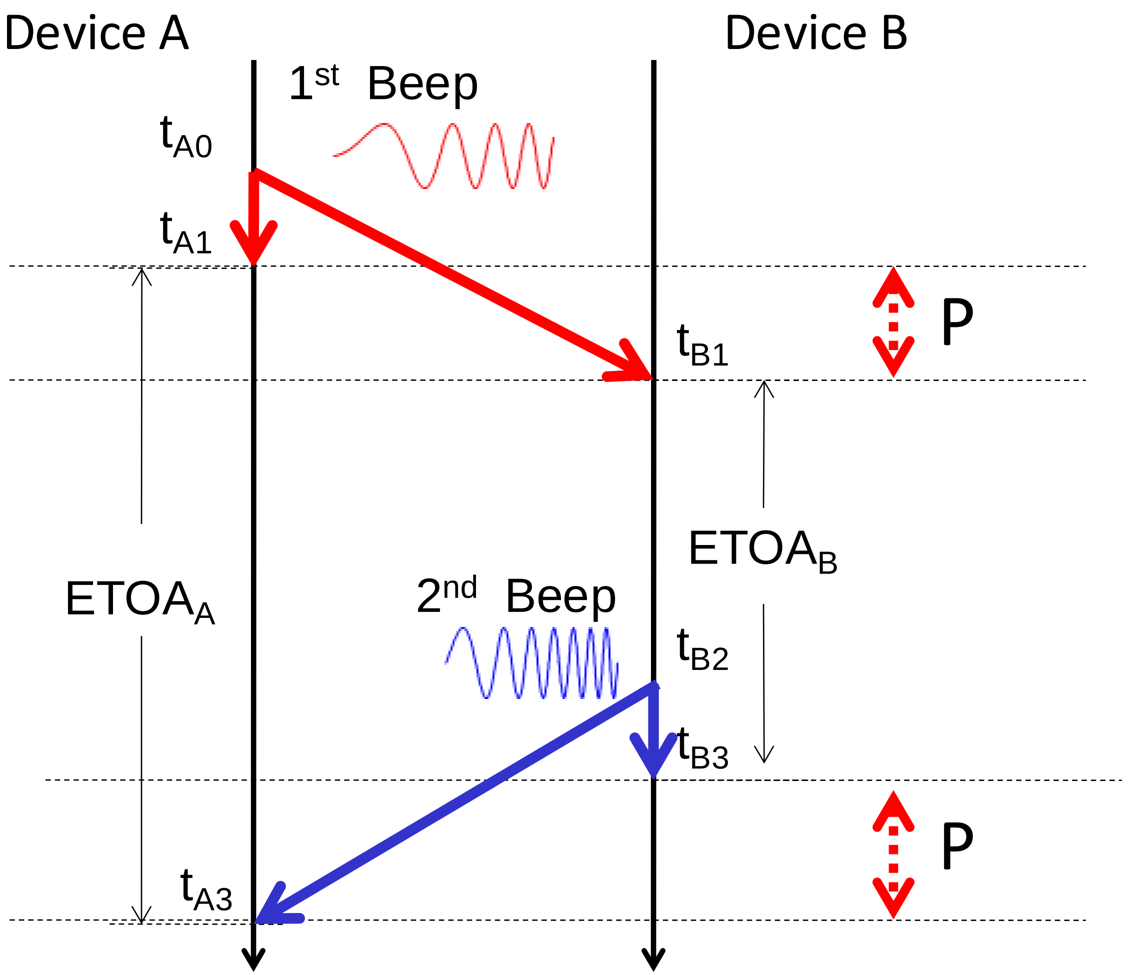

Device A
Device B
1st Beep
t_{A0}
t_{A1}
t_{B1}
P
ETOA$_B$
ETOA$_A$
2nd Beep
t_{B2}
t_{B3}
P
t_{A3}

Timeline

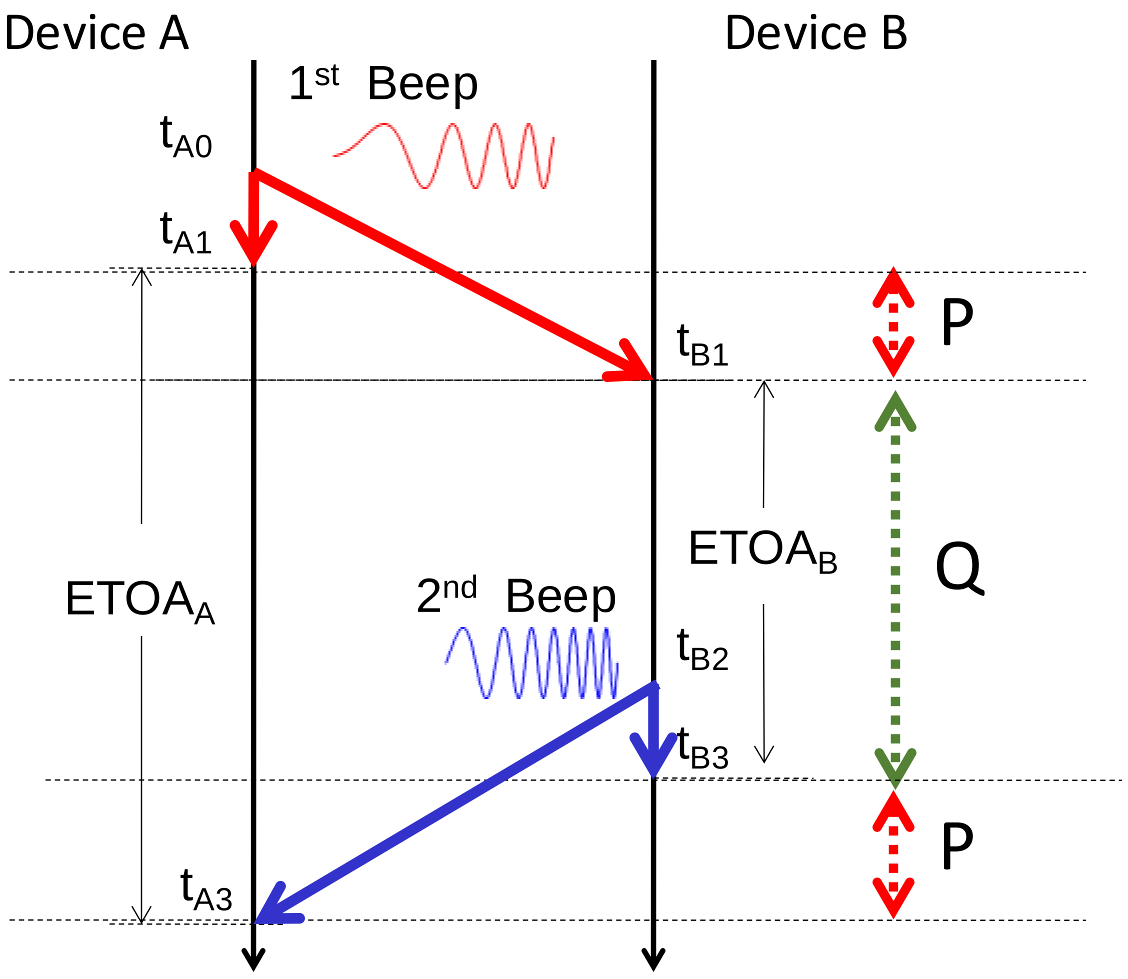

Device A
Device B
1st Beep
tA0
tA1
tB1
P
ETOAB
Q
ETOAA
2nd Beep
tB2
tB3
P
tA3

Timeline

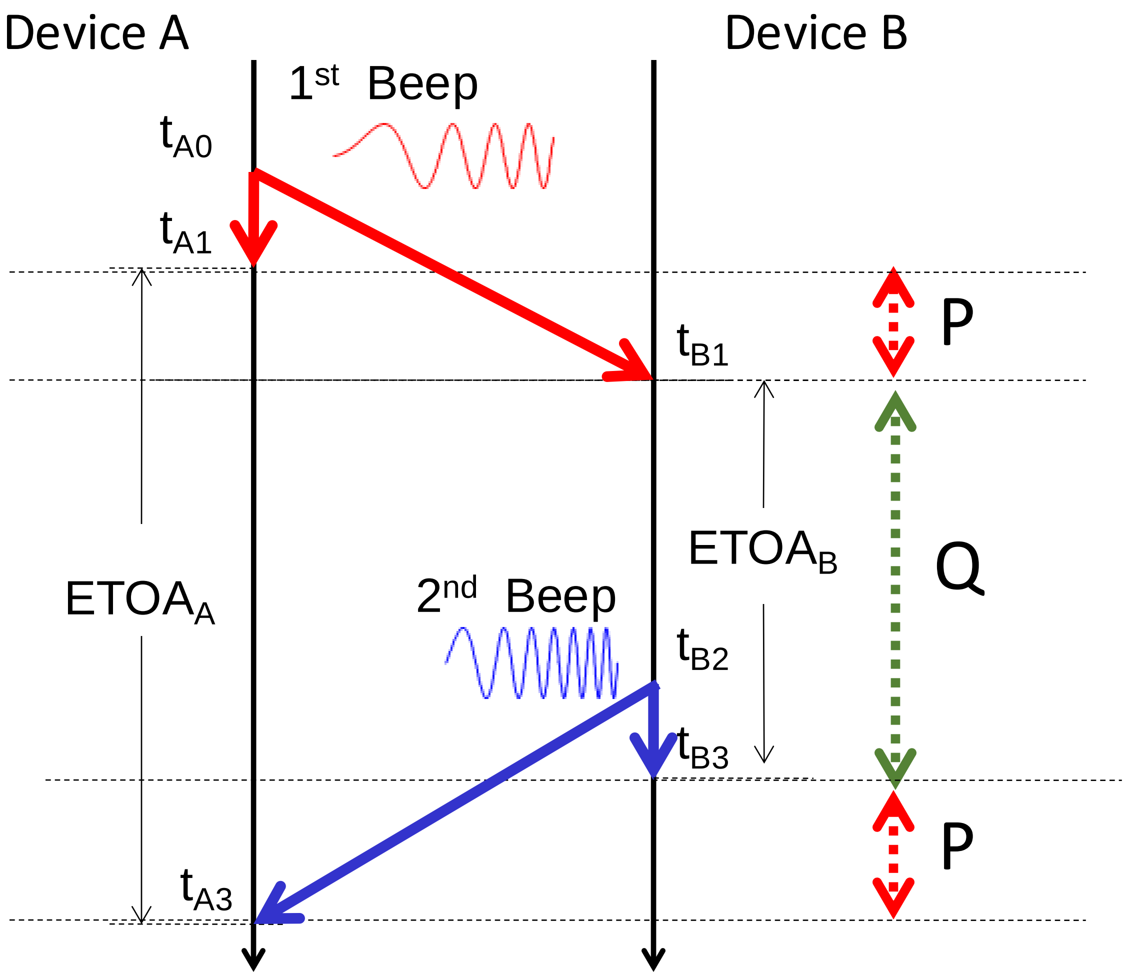

$$|ETOA_A - ETOA_B|$$

$$= (P+Q+P) – (Q)$$

$$= 2P$$

$$d_{B,A}+d_{A,B}= c\cdot[(t_{A3}-t_{A1})-(t_{B3}-t_{B1})] +d_{A,A}+d_{B,B}$$
$$= c\cdot(ETOA_A-ETOA_B)+d_{A,A} +d_{B,B}$$

$$\begin{aligned}
D &= \frac{1}{2}\cdot(d_{A,B}+d_{B,A}) \\
&= \frac{c}{2}\cdot\left((t_{B1}-t_{A0})+(t_{A3}-t_{B2})\right) \\
&= \frac{c}{2}\cdot(t_{B1}-t_{B2}+t_{B3}-t_{B3}+t_{A3}-t_{A0}+t_{A1}-t_{A1}) \\
&= \frac{c}{2}\cdot\left((t_{A3}-t_{A1})-(t_{B3}-t_{B1})+(t_{B3}-t_{B2})+(t_{A1}-t_{A0})\right) \\
&= \frac{c}{2}\cdot\left((t_{A3}-t_{A1})-(t_{B3}-t_{B1})\right)+\frac{1}{2}(d_{B,B}+d_{A,A})
\end{aligned}$$

WiFi RSSI-based Ranging

- The spread in RSS for a given distance is huge, making inversion to estimate the distance from RSS ill posed
- No path-loss model, no matter how complex, can overcome this problem.
- Using CSI helps, but does not solve the problem.

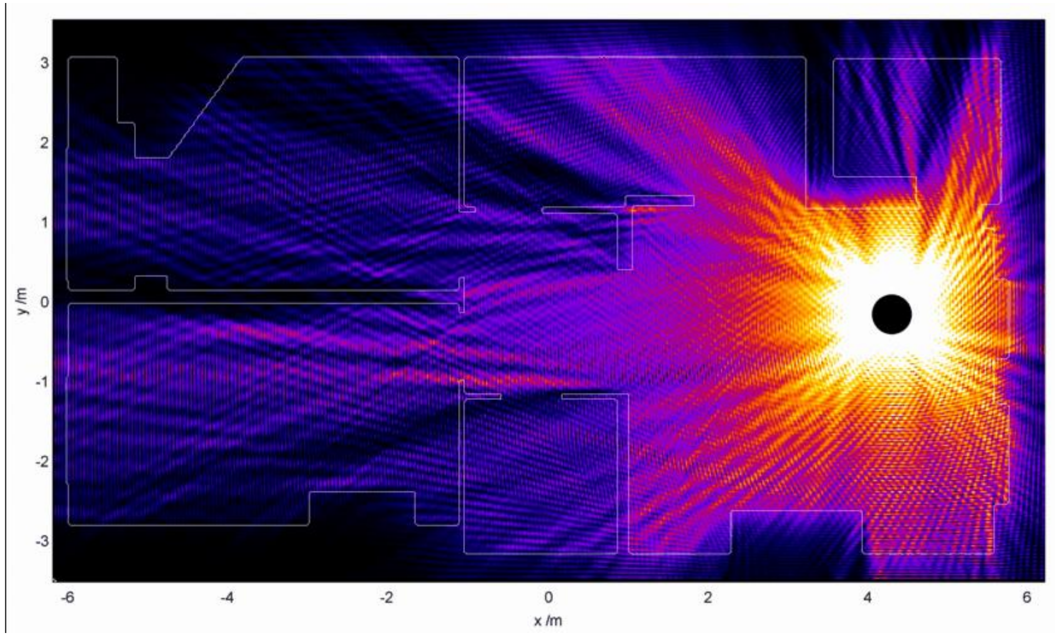

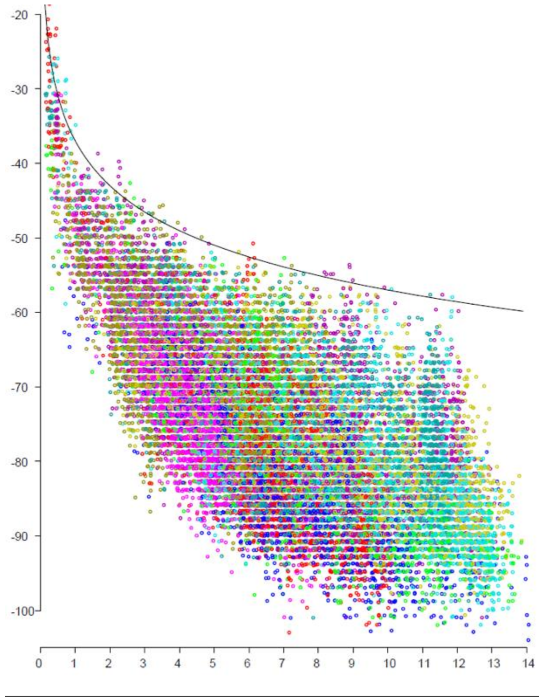

RSSI-based Ranging

- There are many path loss models!
- Log-Distance Path Loss

$$P_d = P_{d_0} - 10\gamma \lg(\frac{d}{d_0})$$

  - P_d: RSS in decibel measured at a distance of *d* (in meters)
  - P_{d_0}: The received power (RSS) at a distance d_0 (usually takes the value of 1 meter), assumed as a constant empirical value (e.g., -40 dB) given Tx power.
  - γ: path loss exponent

Building type	Frequency of transmission	γ
Vacuum, infinite space		2.0
Retail store	914 MHz	2.2
Grocery store	914 MHz	1.8
Office with hard partition	1.5 GHz	3.0
Office with soft partition	900 MHz	2.4
Office with soft partition	1.9 GHz	2.6
Textile or chemical	1.3 GHz	2.0
Textile or chemical	4 GHz	2.1
Office	60 GHz	2.2
Commercial	60 GHz	1.7

WiFi RSSI-based Ranging

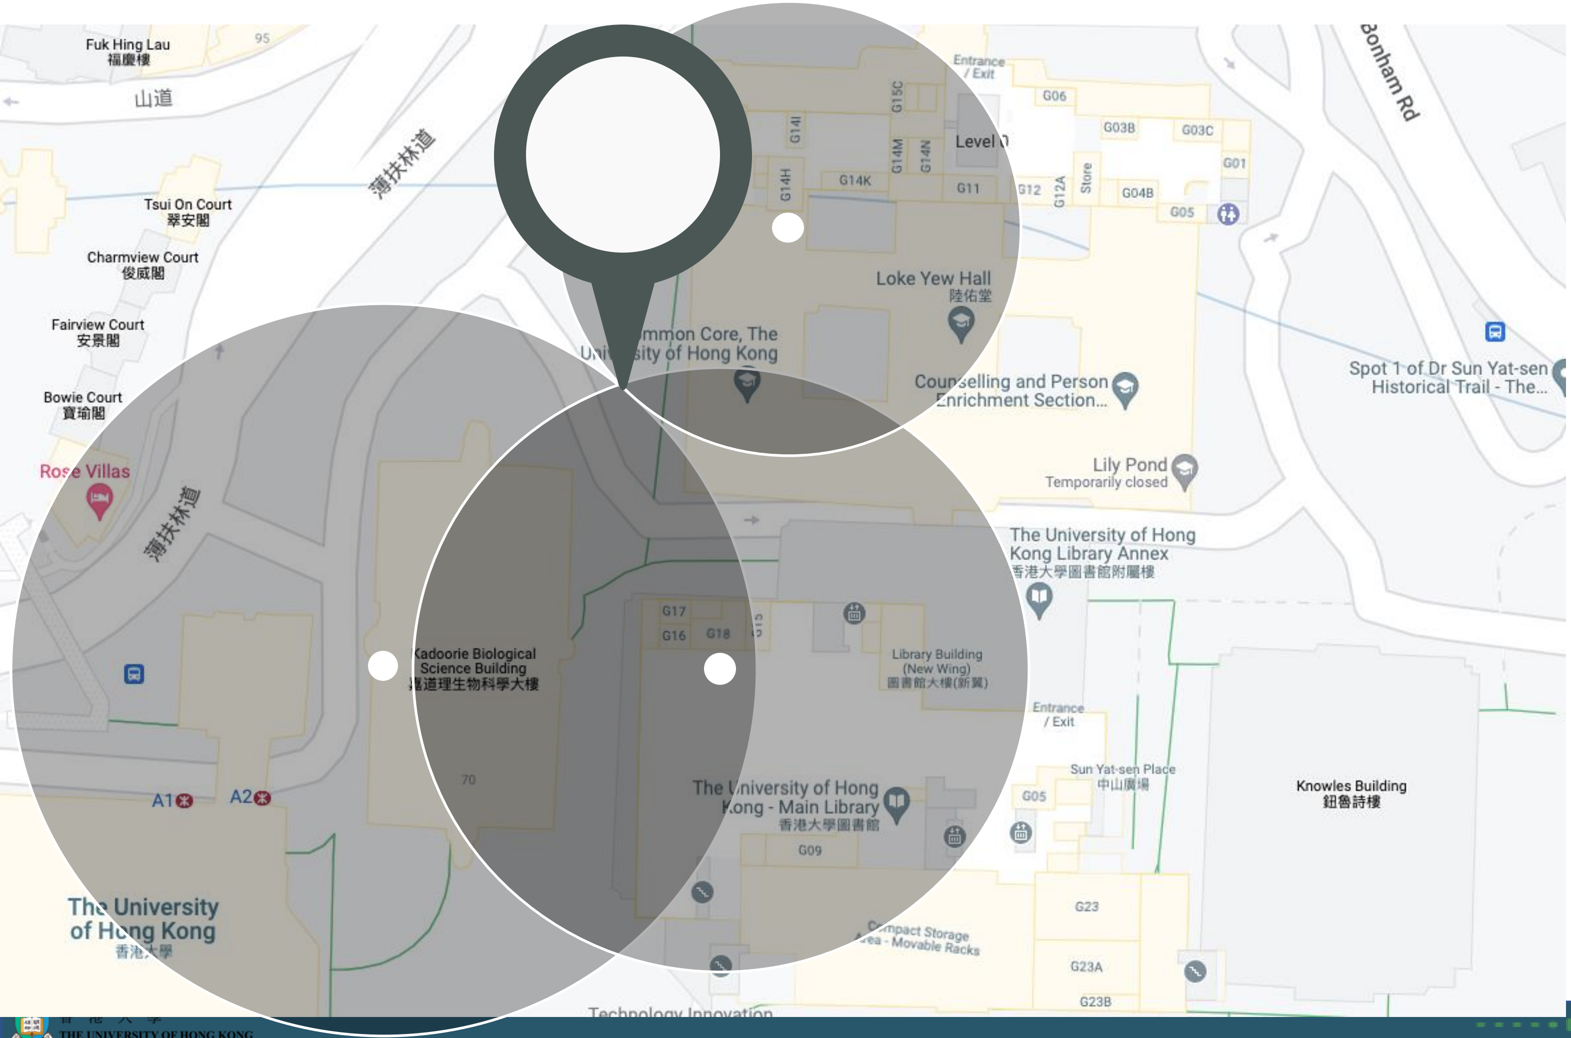

RANGING

Estimate distance from channel measurements

RSSI: Signal strengths decays logarithmically over distance

ToF: Time of Flight
AoA: Angle of Arrival

WiFi RSSI-based Ranging

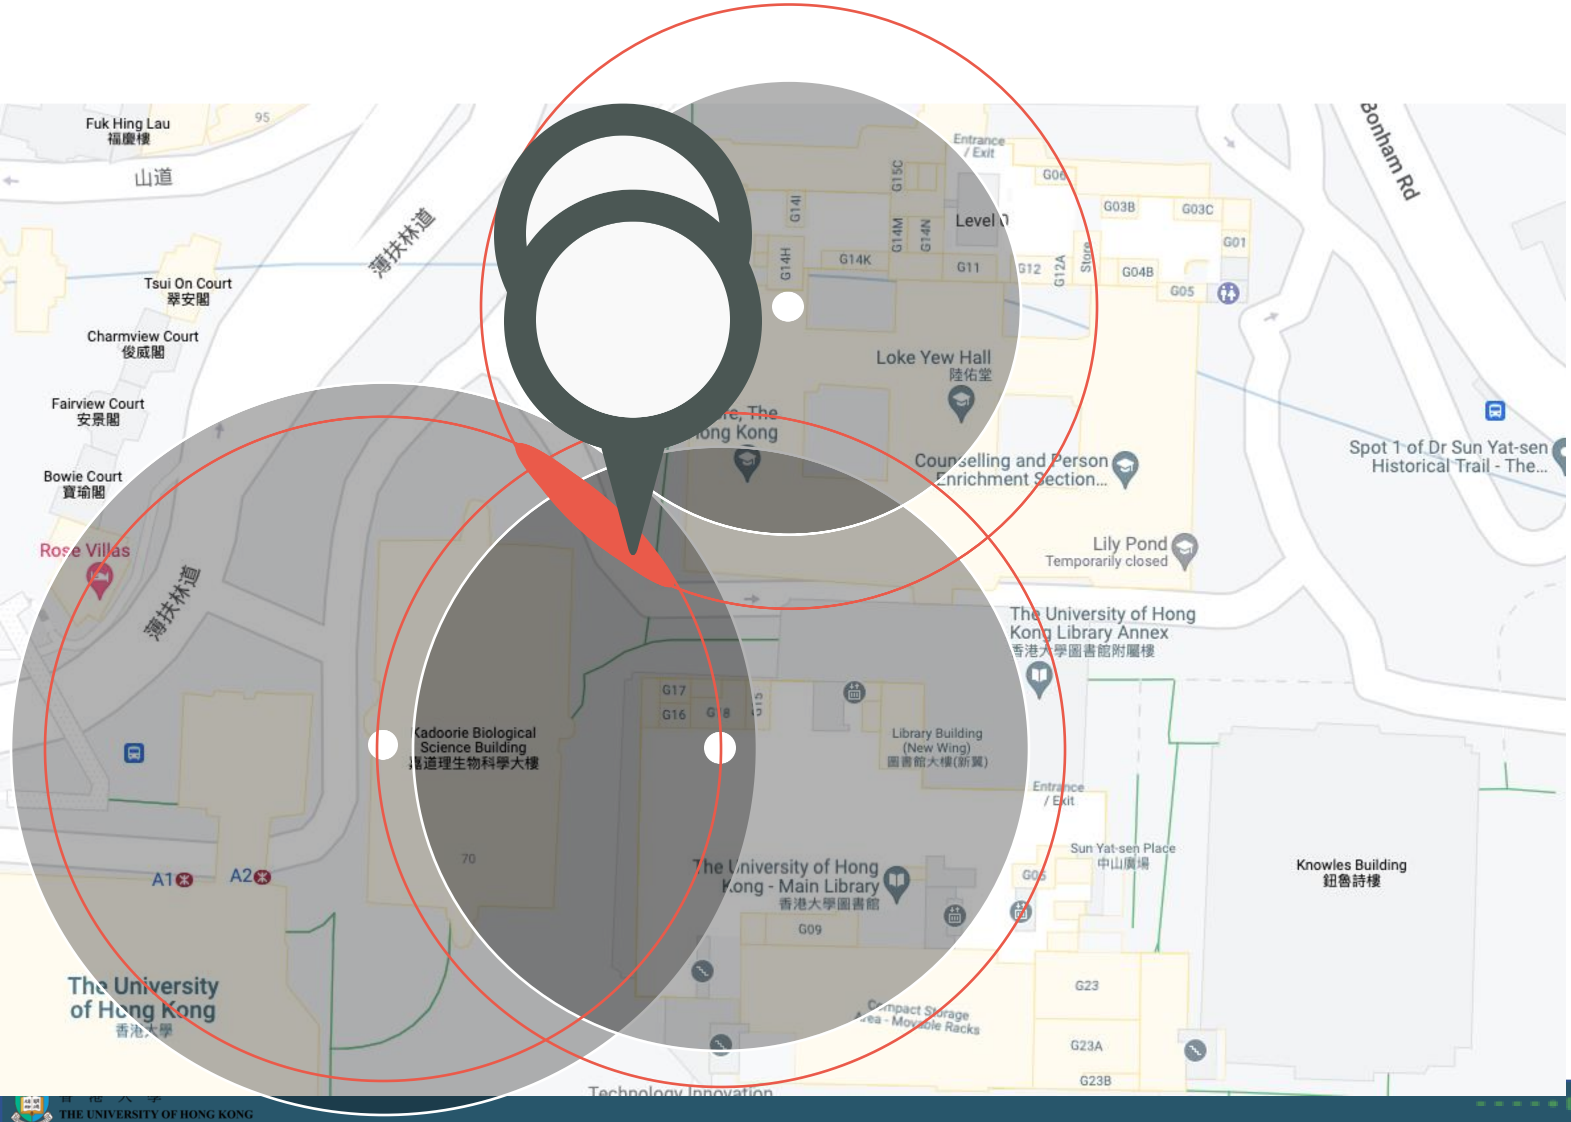

Ranging

Estimate distance from channel measurements

RSSI: Signal strengths decays logarithmically over distance

ToF: Time of Flight
AoA: Angle of Arrival

WiFi Ranging: FTM

- Wi-Fi Fine Timing Measurement (FTM)
  - IEEE 802.11mc FTM RTT

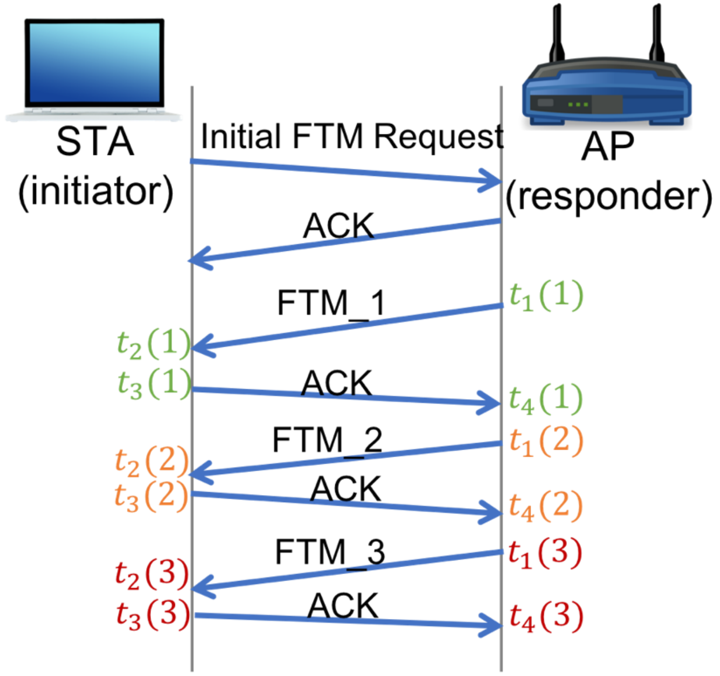

Single Burst with 3 FTMs per Burst

The RTT is calculated for *n* FTM messages:

$$RTT = \frac{1}{n}\left(\sum_{k=1}^{n} t_4(k) - \sum_{k=1}^{n} t_1(k)\right) - \frac{1}{n}\left(\sum_{k=1}^{n} t_3(k) - \sum_{k=1}^{n} t_2(k)\right)$$

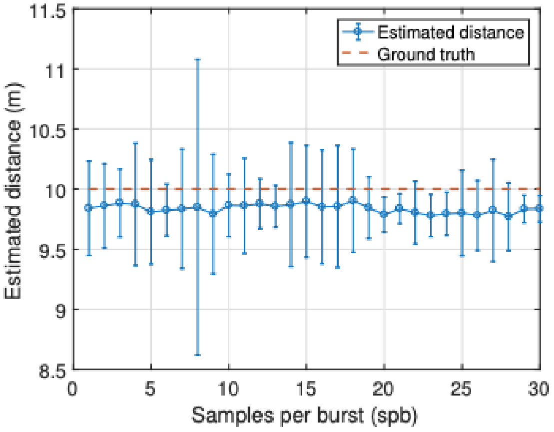

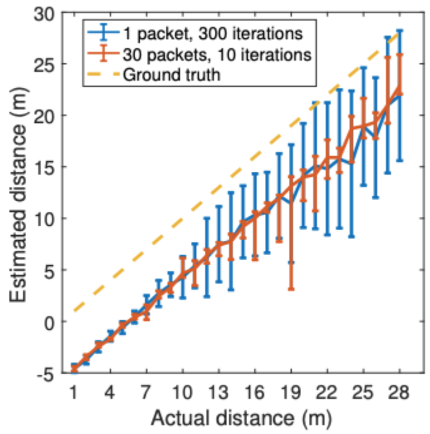

More problems about Trileteration

- Ranging accuracy is only one concern

- Hardware cost
- High cost for installation and maintenance
- Prior knowledge of the anchors

- *Where are the ”indoor satellites”?*

Fingerprinting

RADAR

The first fingerprint-based system
Leading a new epoch / 2000

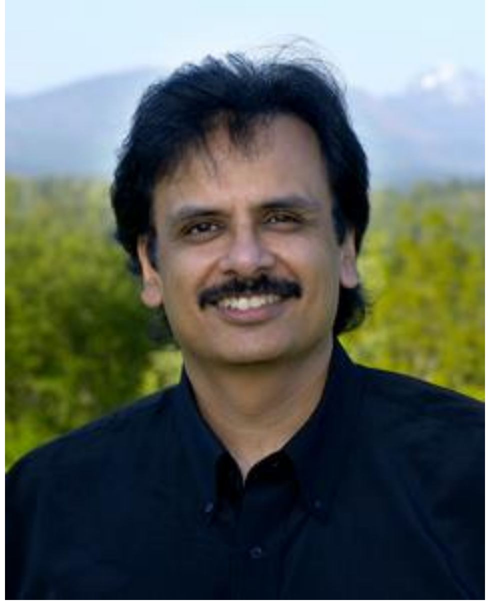

Paramvir / Victor Bahl

HORUS

Improved upon RADAR
/ 2004

Moustafa Youssef

LANDMARC

First RFID Fingerprinting System
/ 2004

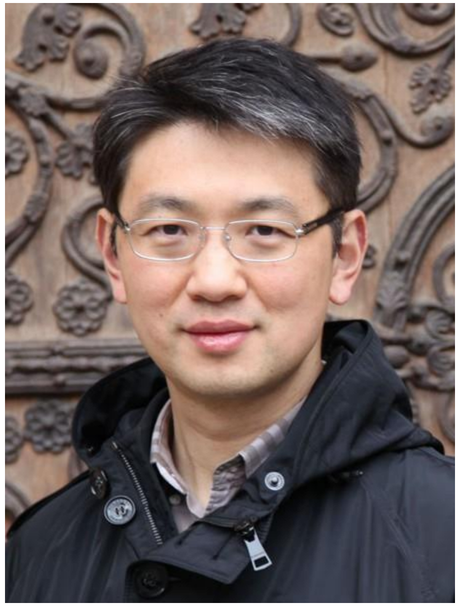

Yunhao Liu

What fingerprints?

- WiFi: One of the most ubiquitous signatures
- RFID
- Acoustics/sound
- Geomagnetism
- FM signals
- Light

Spatial Distinction

Temporal Invariance

Geo-Magnetism Fingerprinting

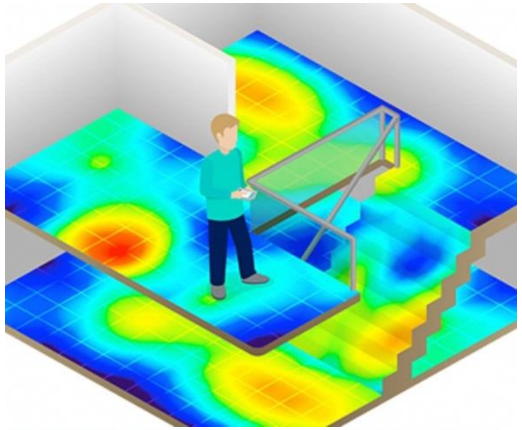

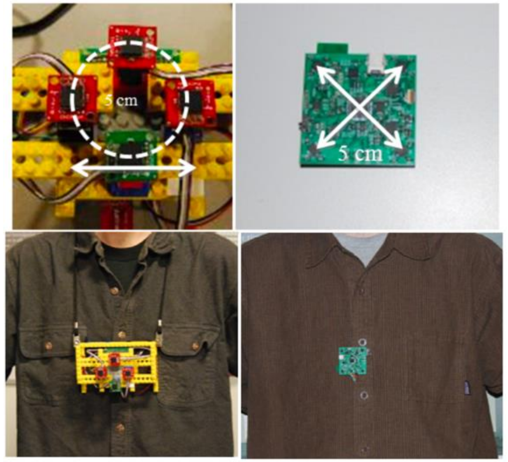

Database: <mag_x_i, mag_y_i, mag_z_i, Location_i>

Observation: <mag_x, mag_y, mag_z>

Find the ‘i’ (or a sequence) for which RMS difference between the observation and the stored magnetic value is minimum.

Indoor Location Sensing Using Geo-Magnetism, MobiSys’11

WiFi Fingerprinting

- Existing WiFi ≈ Infrastructure free

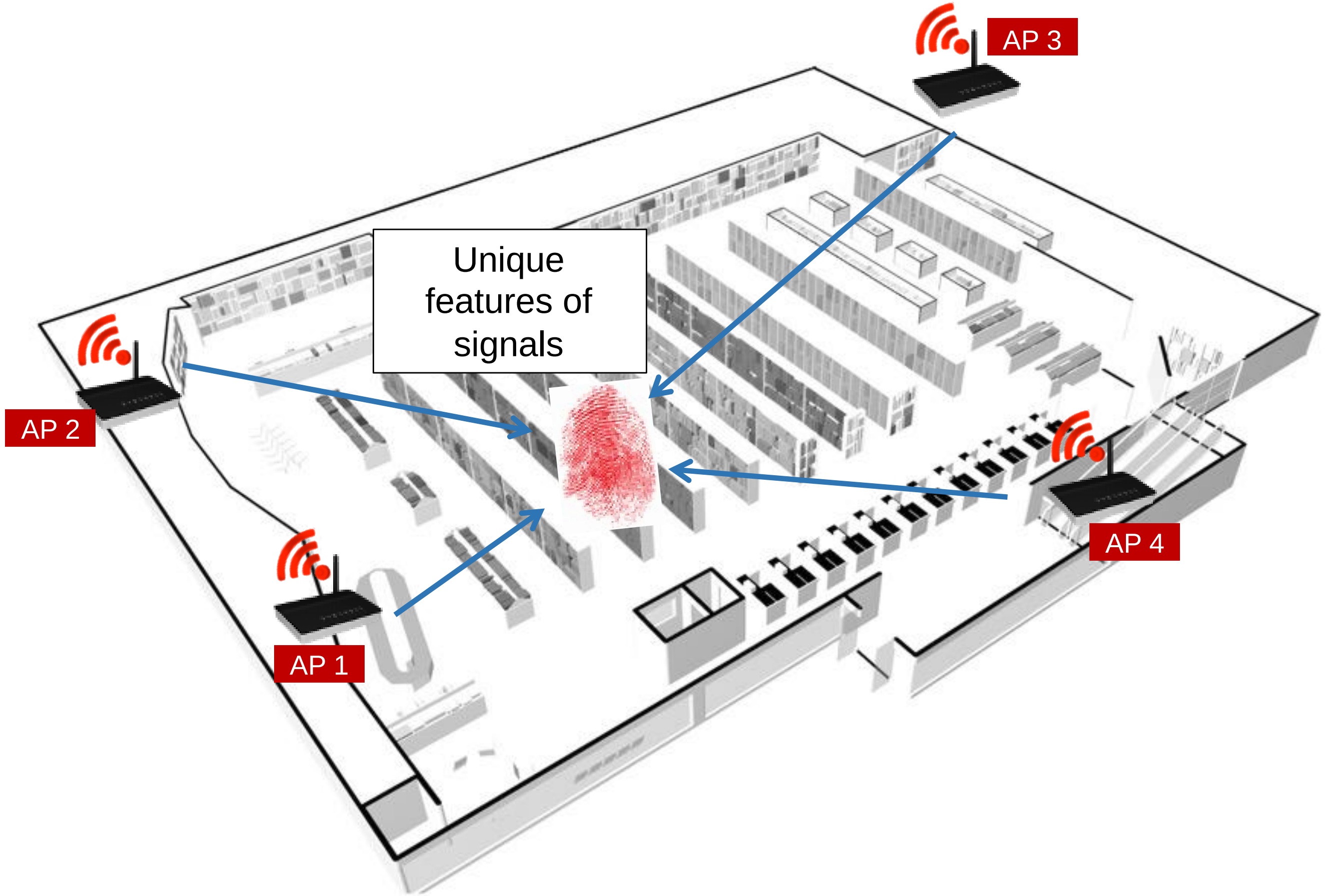

WiFi Fingerprinting

- Existing WiFi ≈ Infrastructure free

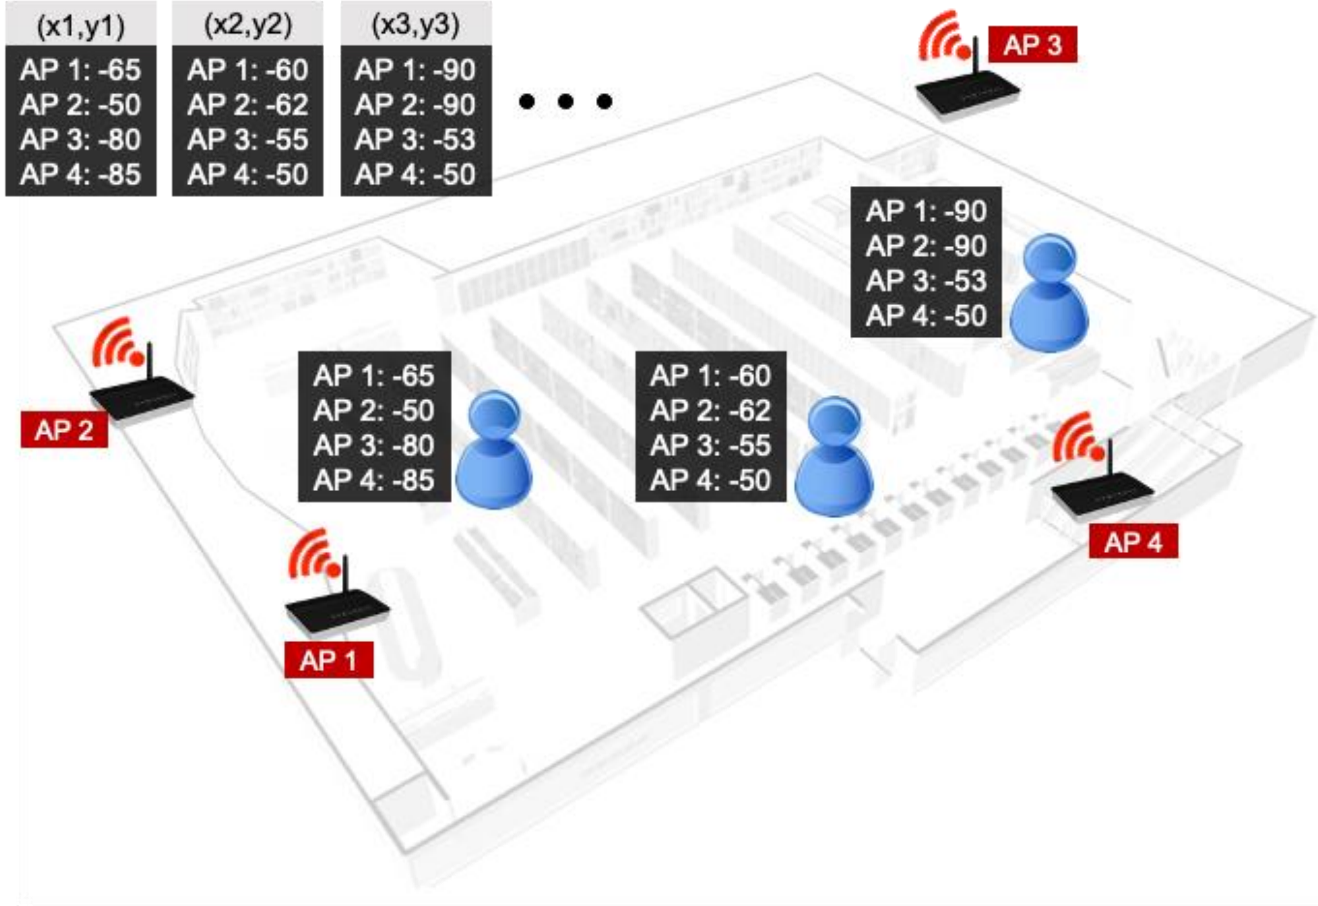

WiFi Fingerprinting

- Offline phase: Building the fingerprint database
- Online phase: Handle location query and find the best match

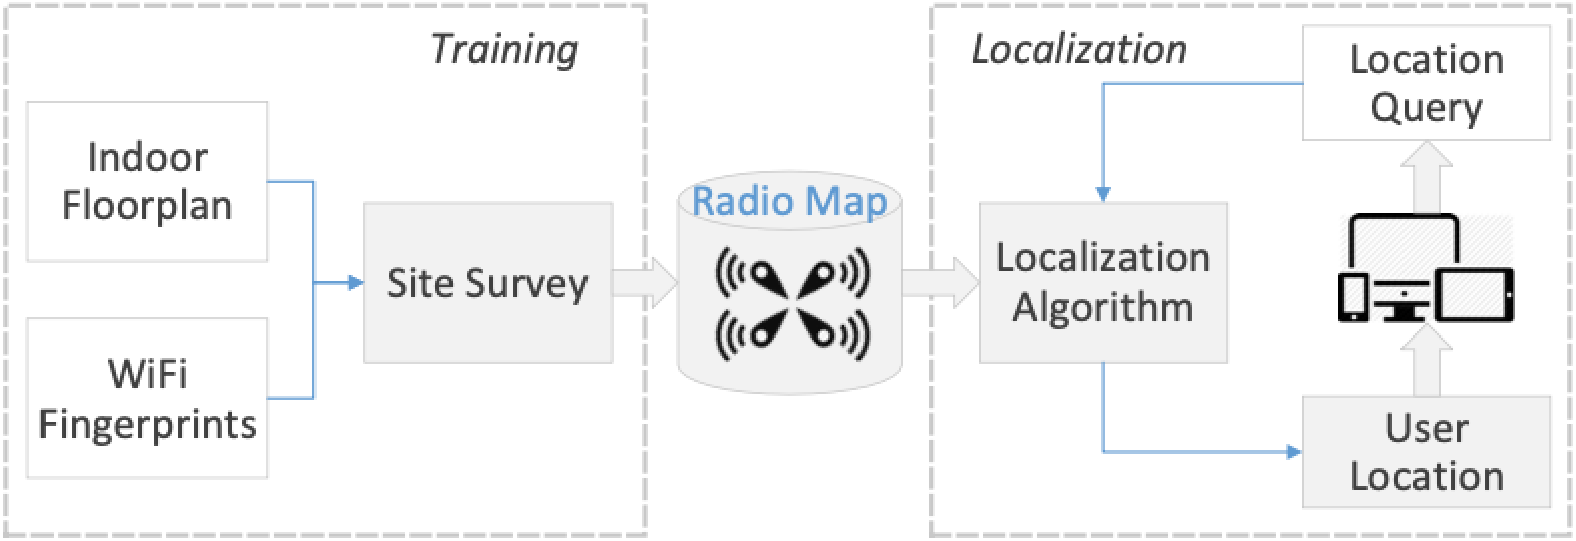

How to build fingerprint database?

- RSS as unique feature of a physical location
- Site Survey: Build fingerprint database of RSS-location records
- Estimate location by finding best-matched item

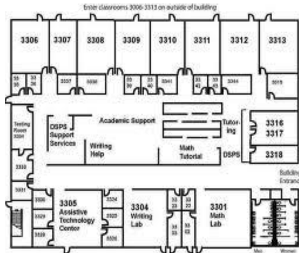

Floor plan

Site survey

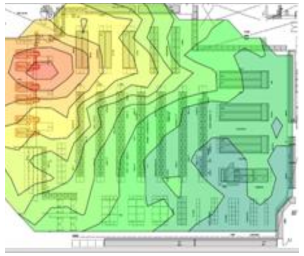

Radio map

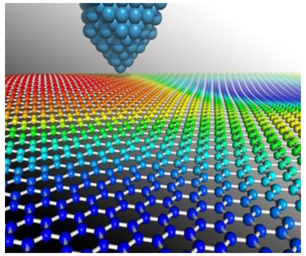

Localization

Problems of Site Survey

- Time-consuming and labor-intensive
  - Leverage mobile crowdsourcing
- Environmental changes (Recall RSS-based human detection?)
- Need to recalibrate periodically

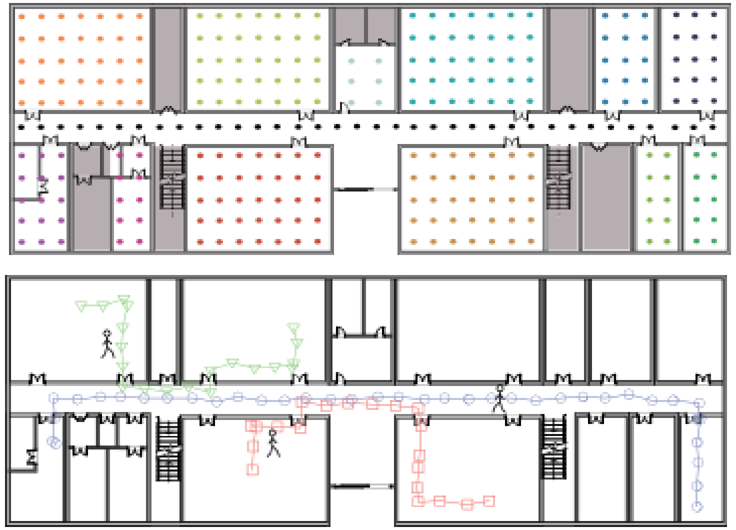

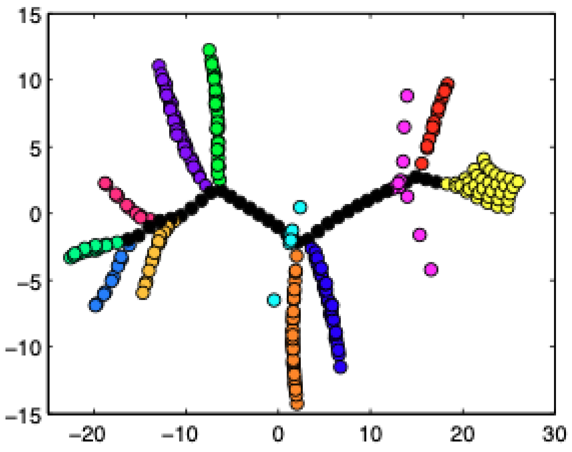

(a) 2D stress-free floor plan

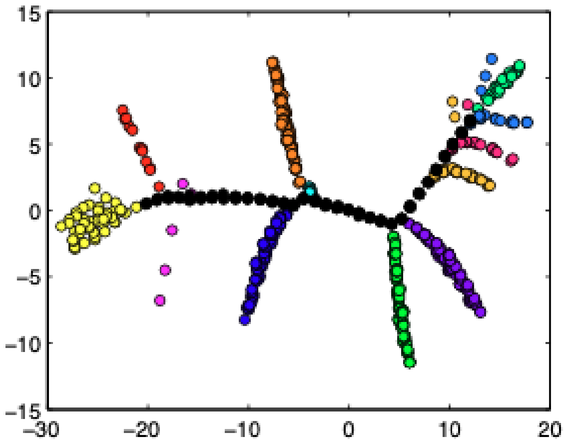

(a) 2D fingerprint space

Smartphones based Crowdsourcing for Indoor Localization, ACM MobiCom'12/IEEE TMC'15

Limitations of Fingerprinting

- Limited Accuracy
  - Spatial ambiguity: RSS doesn't provide enough resolution
  - Temporal variability: RSS varies significantly over time
- Low hardware cost but still high deployment cost
  - Time-consuming and labor-intensive
  - Relieved by crowdsourcing
- Still one of the most practical approaches, used partially in Google/Apple maps
  - Made (more) practical and usable nowadays with big data and AI models

CSI Fingerprinting

- Achieving 1cm accuracy & robust to environment changes!
- TRRS as distance measure
  - Time-Reversal Resonating Strength
  - Cosine similarity?

$$\kappa(H_1, H_2) = \frac{|H_1^{\mathrm{H}} H_2|^2}{\langle H_1, H_1 \rangle \langle H_2, H_2 \rangle}$$

H_1: CSI at location 1
H_1: CSI at location 2

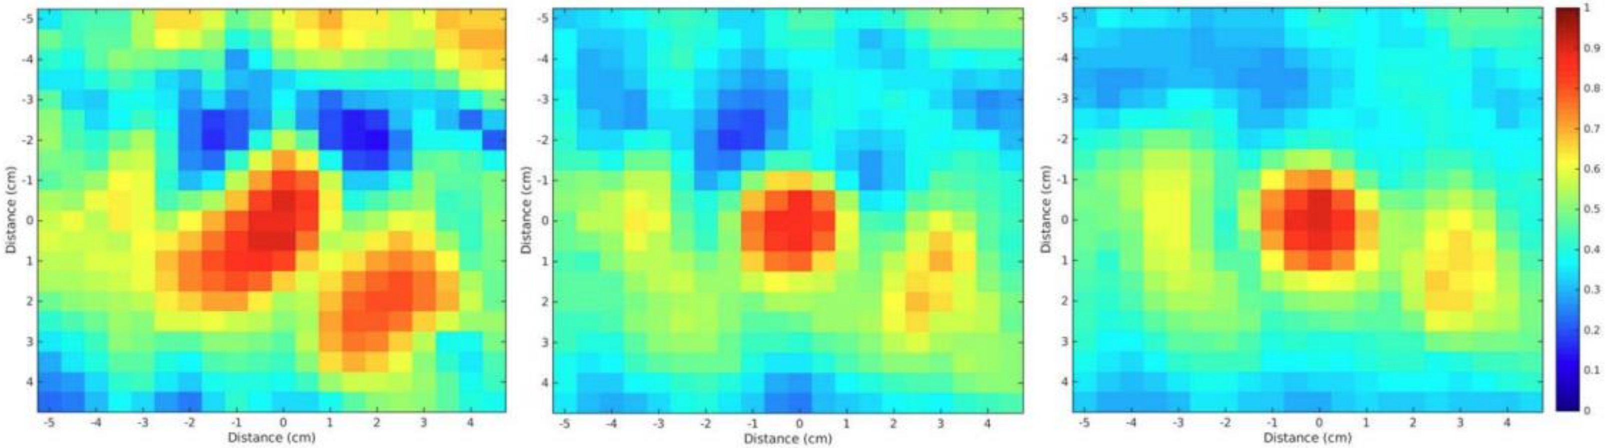

- *Problem?*

Inertial Tracking

- Basic tasks
  - Distance/displacement estimation
  - Direction estimation
  - Integrate distance and direction over time to track locations

- Pros
  - Infrastructure-free
  - Scalable
- Cons
  - Accumulative errors
  - Difficult to infer a user’s heading direction (different from device orientation)
  - Unconstrained user behavior for pedestrian tracking

Inertial Tracking

- a.k.a Dead-reckoning
- PDR: Pedestrian Dead-Reckoning
- Truly infrastructure-free

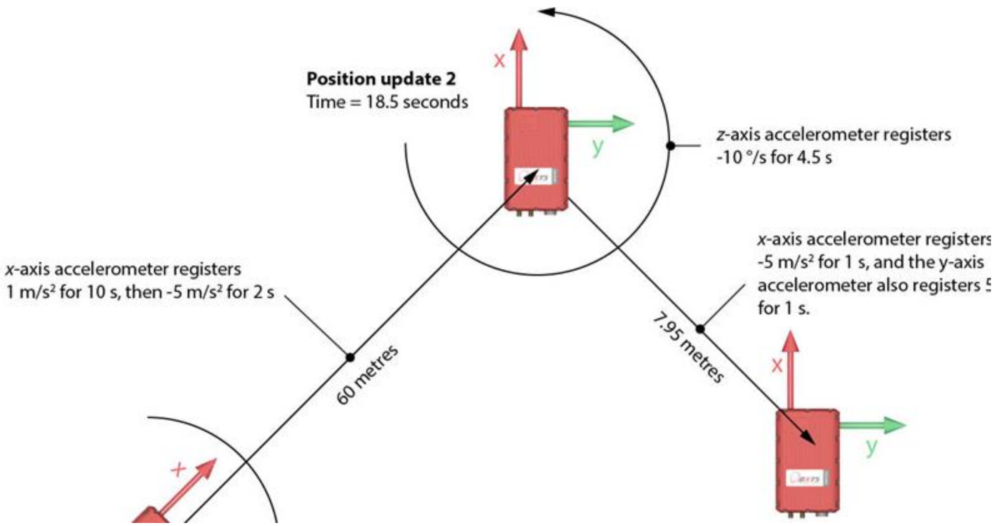

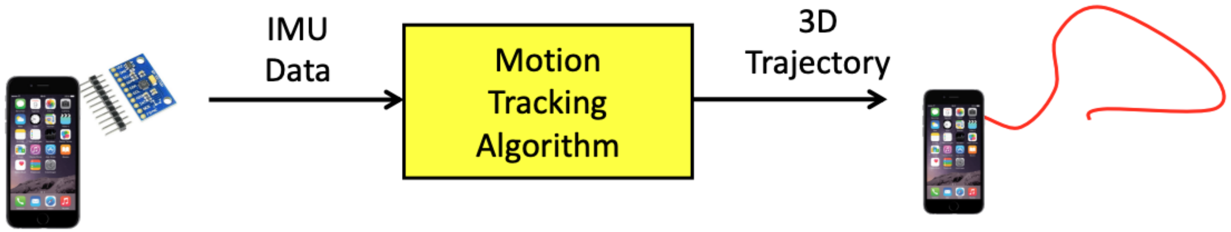

Open Problem in mobile computing
“No one has the solution…But people making progress”

Inertial Tracking

- Can we solve tracking with these inputs?

Initial Location → Algorithm

IMU = Accel + Gyro + Mag → Algorithm

Algorithm $\xrightarrow{\text{3D Trajectory}}$ Map

$[x_0 \quad y_0 \quad z_0]^T$

Initial Location → Algorithm

IMU = Accel + Gyro + Mag → Algorithm

Algorithm $\xrightarrow{?}$ $\begin{bmatrix} x_t \\ y_t \\ z_t \end{bmatrix}$

Zoom into IMU data:

$\text{Accel.} = \begin{bmatrix} a_x \\ a_y \\ a_z \end{bmatrix}$ (Linear Acceleration)

$\text{Gyro.} = \begin{bmatrix} g_x \\ g_y \\ g_z \end{bmatrix}$ (Angular Velocity)

Delta rotation in unit time

$\text{Mag.} = \begin{bmatrix} m_x \\ m_y \\ m_z \end{bmatrix}$ (Earth's Magnetic North)

Inertial Tracking

- One possible solution: Direct integration

$$\text{Initial Location} \xrightarrow{[x_0 \quad y_0 \quad z_0]^T} \boxed{\text{Algorithm}} \xrightarrow{?} \begin{bmatrix} x_t \\ y_t \\ z_t \end{bmatrix}$$

IMU = Accel + Gyro + Mag $\longrightarrow$ Algorithm

$$\boxed{\begin{bmatrix} x_0 \\ y_0 \\ z_0 \end{bmatrix} + \iint_0^t (Accel.)\, dt^2}$$

Accel. = $\begin{bmatrix} a_x \\ a_y \\ a_z \end{bmatrix}$ is measured in local reference frame, and needs to convert into the global reference frame.

- Big Problem: Acc drifts, gyro drifts, significantly
  - Huge (!!!) accumulative errors of time

Pedestrian Dead-Reckoning

- Any good idea to get a better/reasonable estimate of distance?

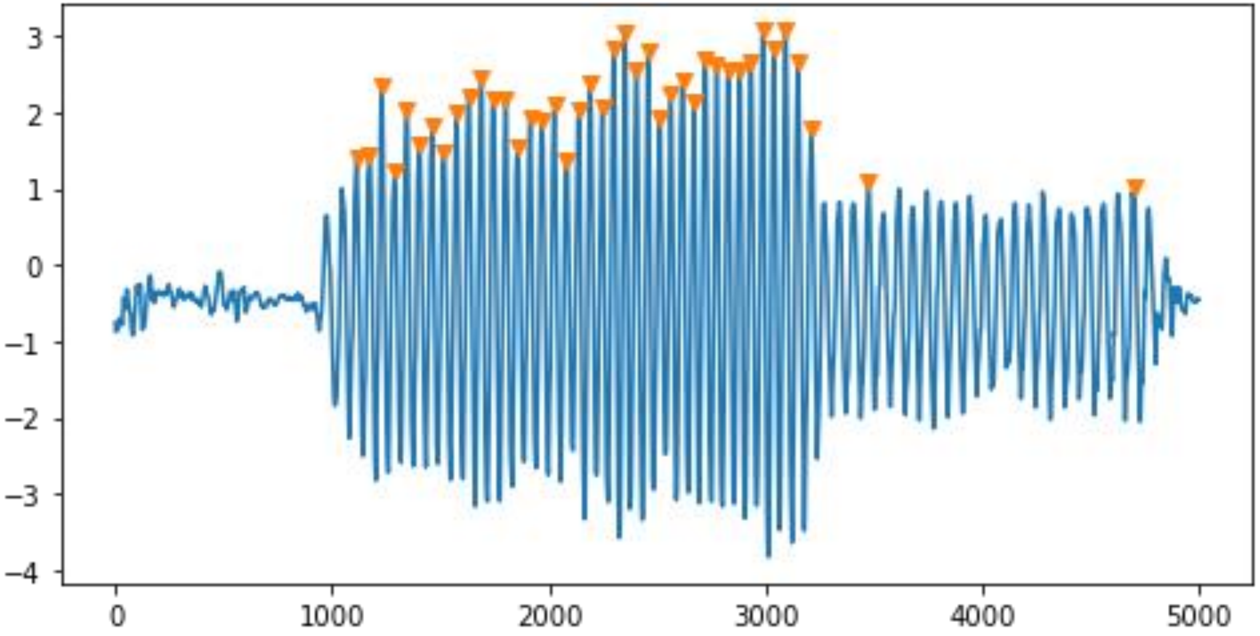

Step Count: 63 steps

Double Integration : -551m
(using magnitude)
?????????

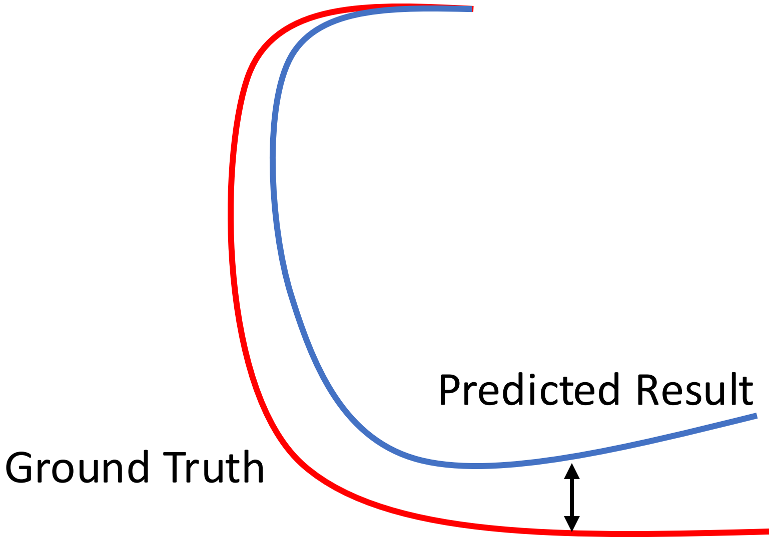

Direction also drifts significantly over time

3D Orientation

- The 3D rotation needed for coordinate transformation
  - [Frontwards, rightwards, upwards] → [Northwards, eastwards, vertical]

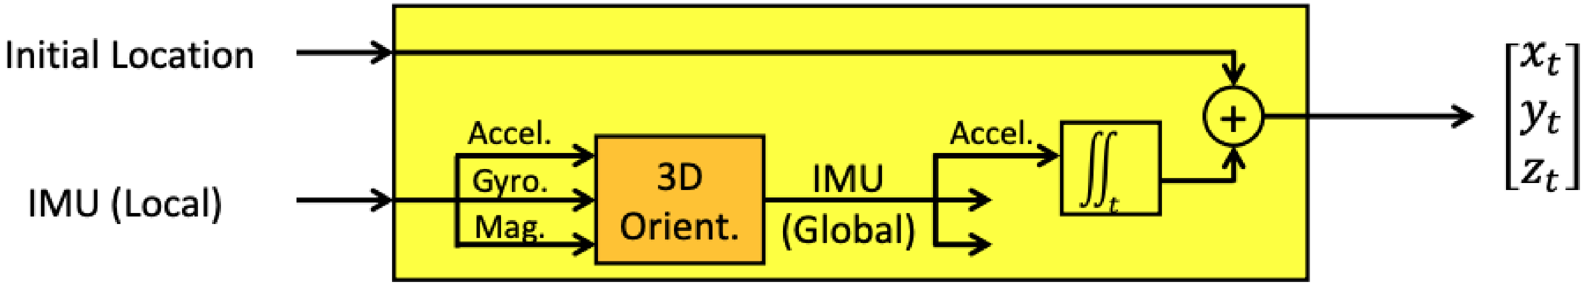

3D Orientation

- Main opportunities
  - Constant gravity
  - Magnetic north

- Key idea: What rotation is needed such that
  - **Gravity** is exactly in the **downward** direction
  - **North** is exactly in the **frontward** direction

3D Orientation

- For static objects, can rely mostly on gravity + North
  - Does not work well for moving objects
  - Any motion will affect the reported acceleration and pollute the gravity estimate

- Another idea: Integrate angular velocity from gyroscope for continuous estimation
  - Gyro also drifts, only useful in short time scales

$$\boxed{\text{Initial Orientation}} + \boxed{\int_0^t (Gyro.)\, dt} = \boxed{\text{New Orientation} \quad \text{(at time } t\text{)}}$$

3D Orientation: Sensor Fusion

(More) Accurate 3D orientation all the time!

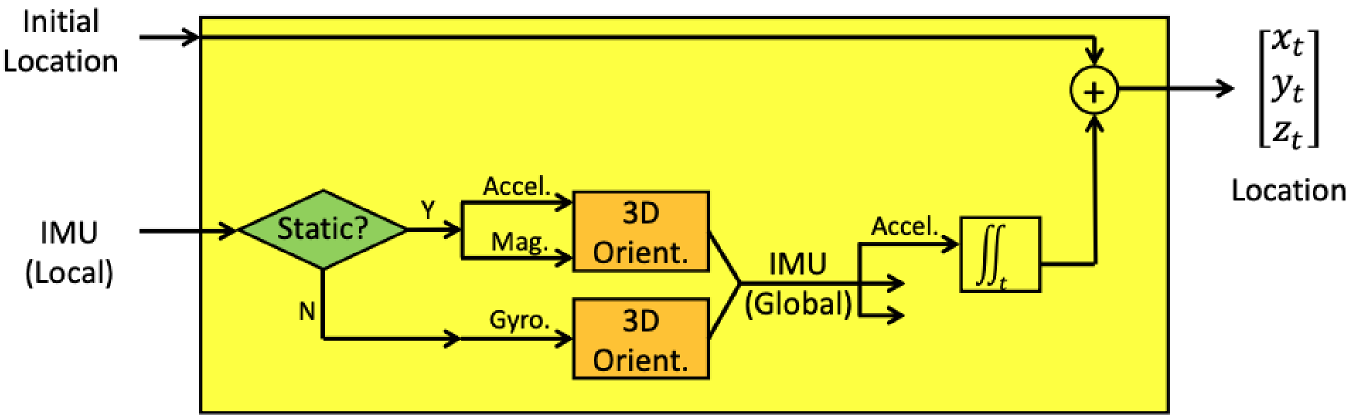

- If static: Rely mostly on gravity + North
- If moving: Rely mostly on gyro integration
- Gravity as the main reference anchor

3D Orientation

- What if the object is not often static?
- Many different sensor fusion algorithms
  - No good solution today…
  - Count on you to solve the problem...

Pedestrian Dead-Reckoning

- Any good idea to get a better/reasonable estimate of distance?

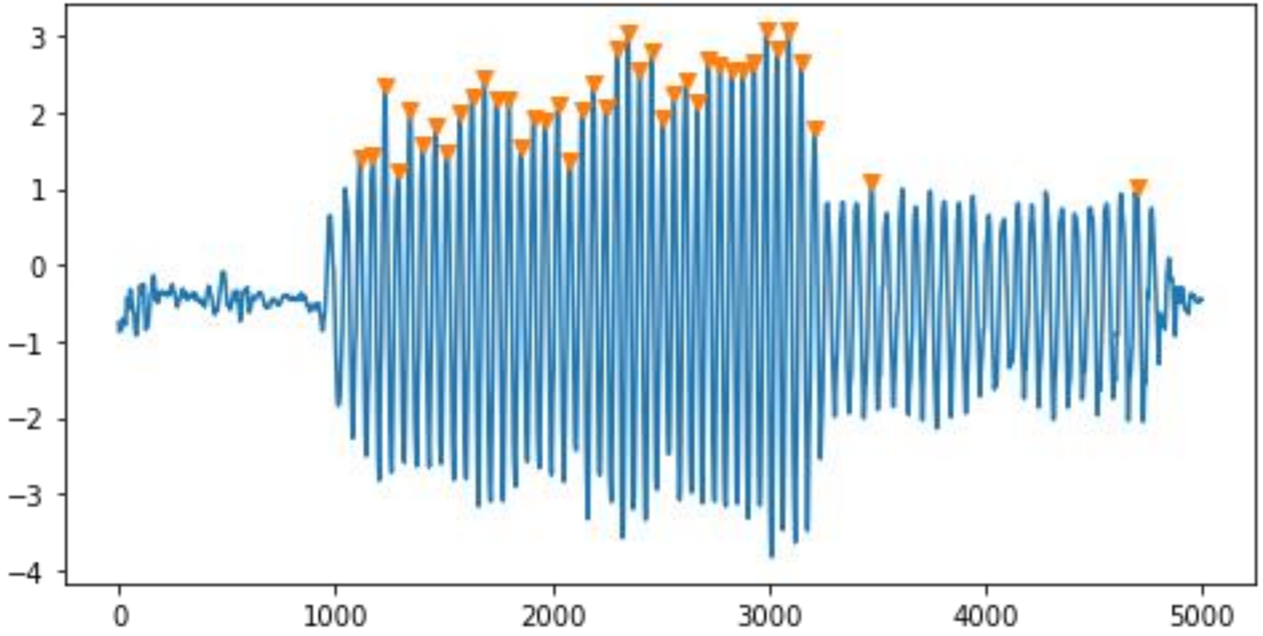

Step Count: 63 steps

Double Integration : -551m
(using magnitude)
?????????

Pedestrian Dead-Reckoning

- Any good idea to get a better/reasonable estimate of distance?

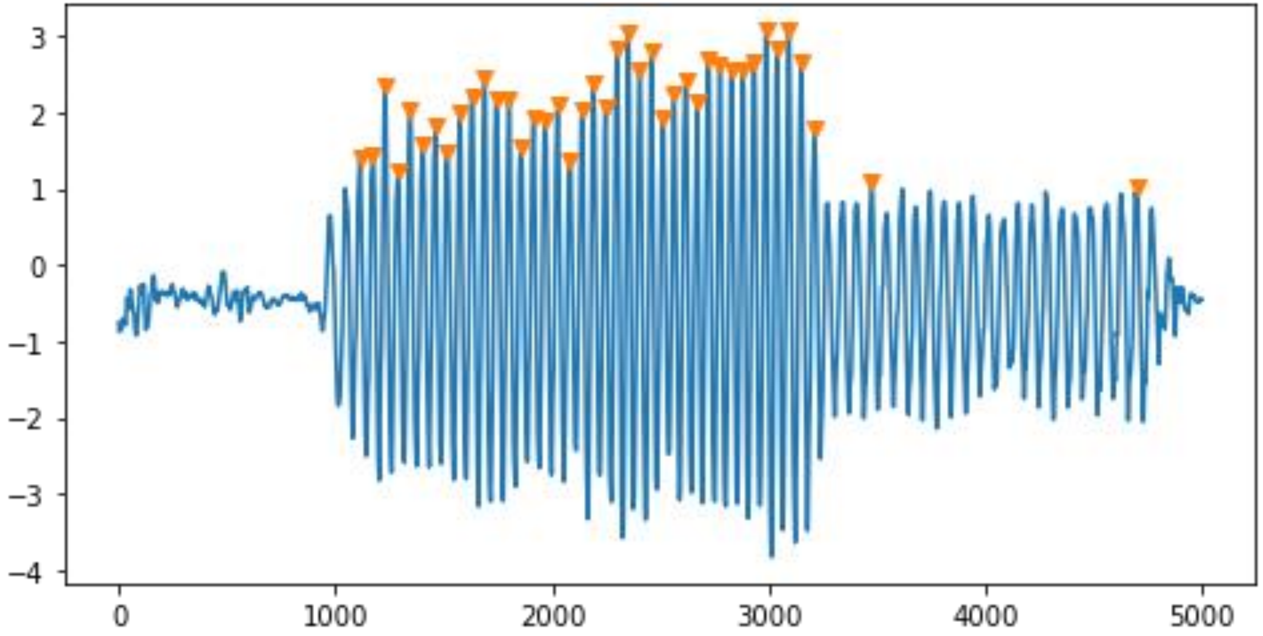

Step Count: 63 steps

Double Integration : -551m
(using magnitude)
?????????

Step counting instead of double integration

$\boxed{\text{(\# of steps) x (stride length)}}$

How to get stride length?

- Fixed value
- Estimation given height
- Dynamically estimated

Visual Inertial Tracking

- Visual-Inertial Odometry (VIO)

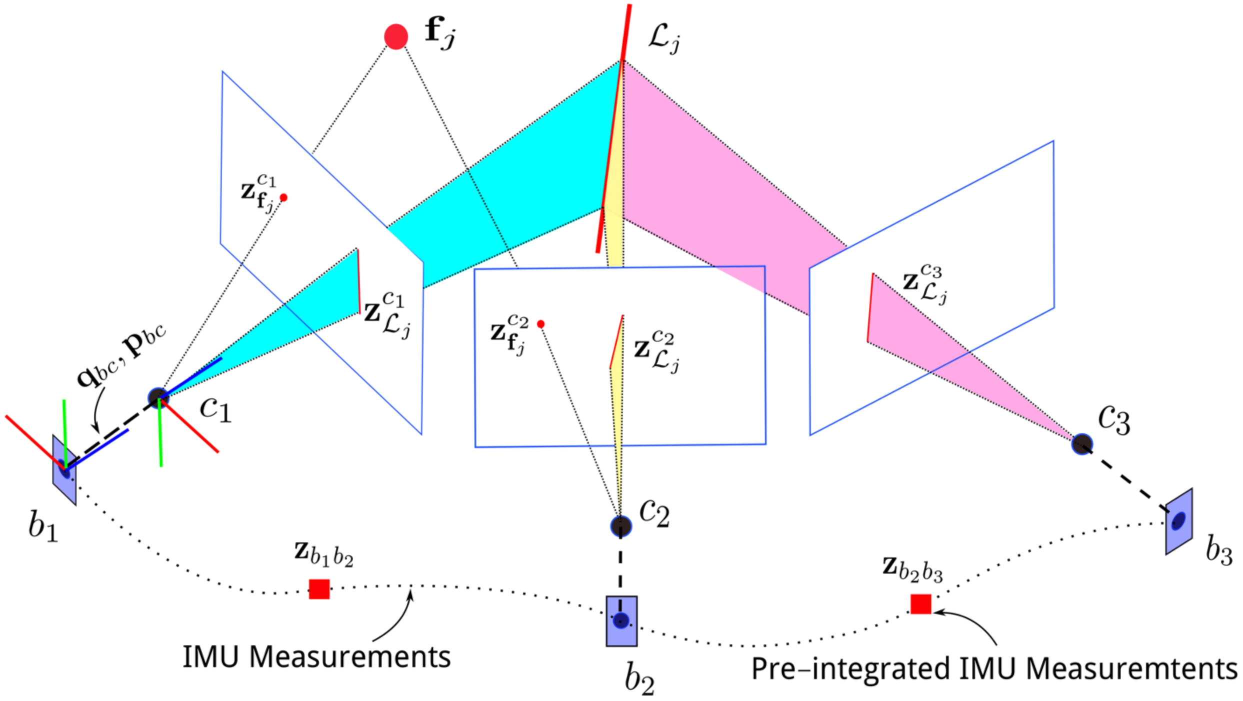

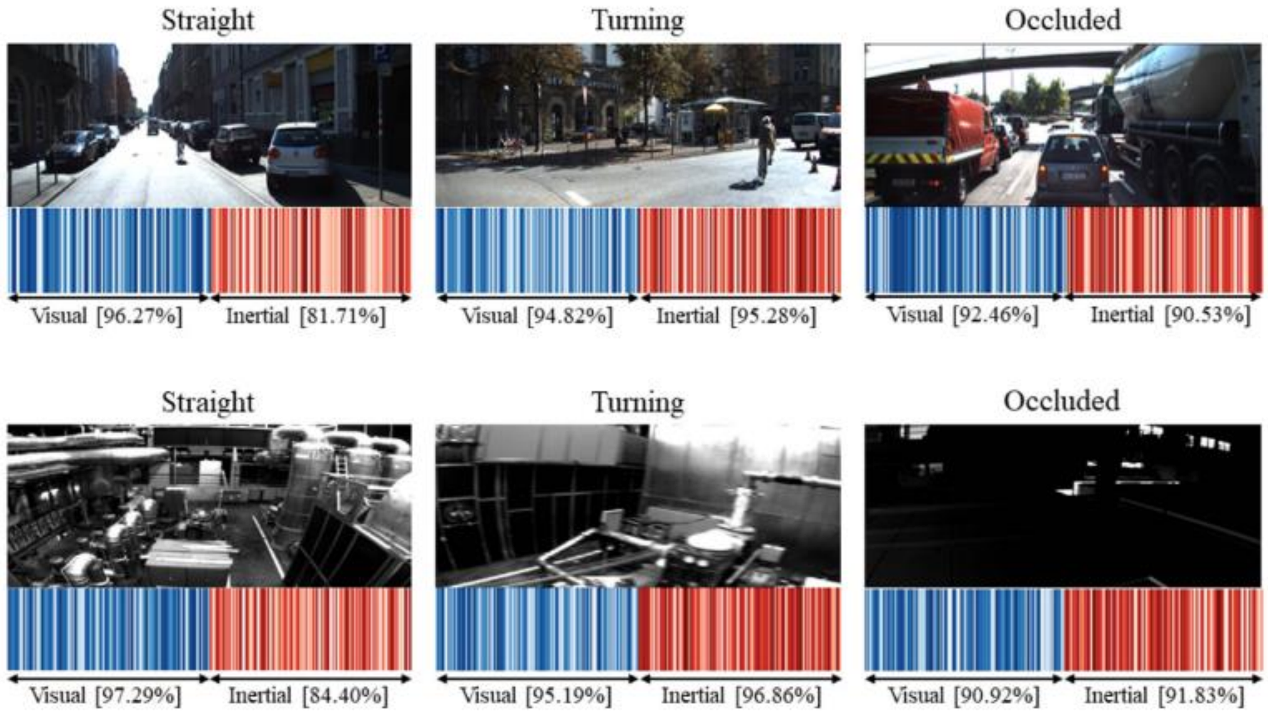

Neural Inertial Tracking

- Using deep neural networks to learn
  - The distance, velocity, and/or positions
  - And thus predict the moving trajectories

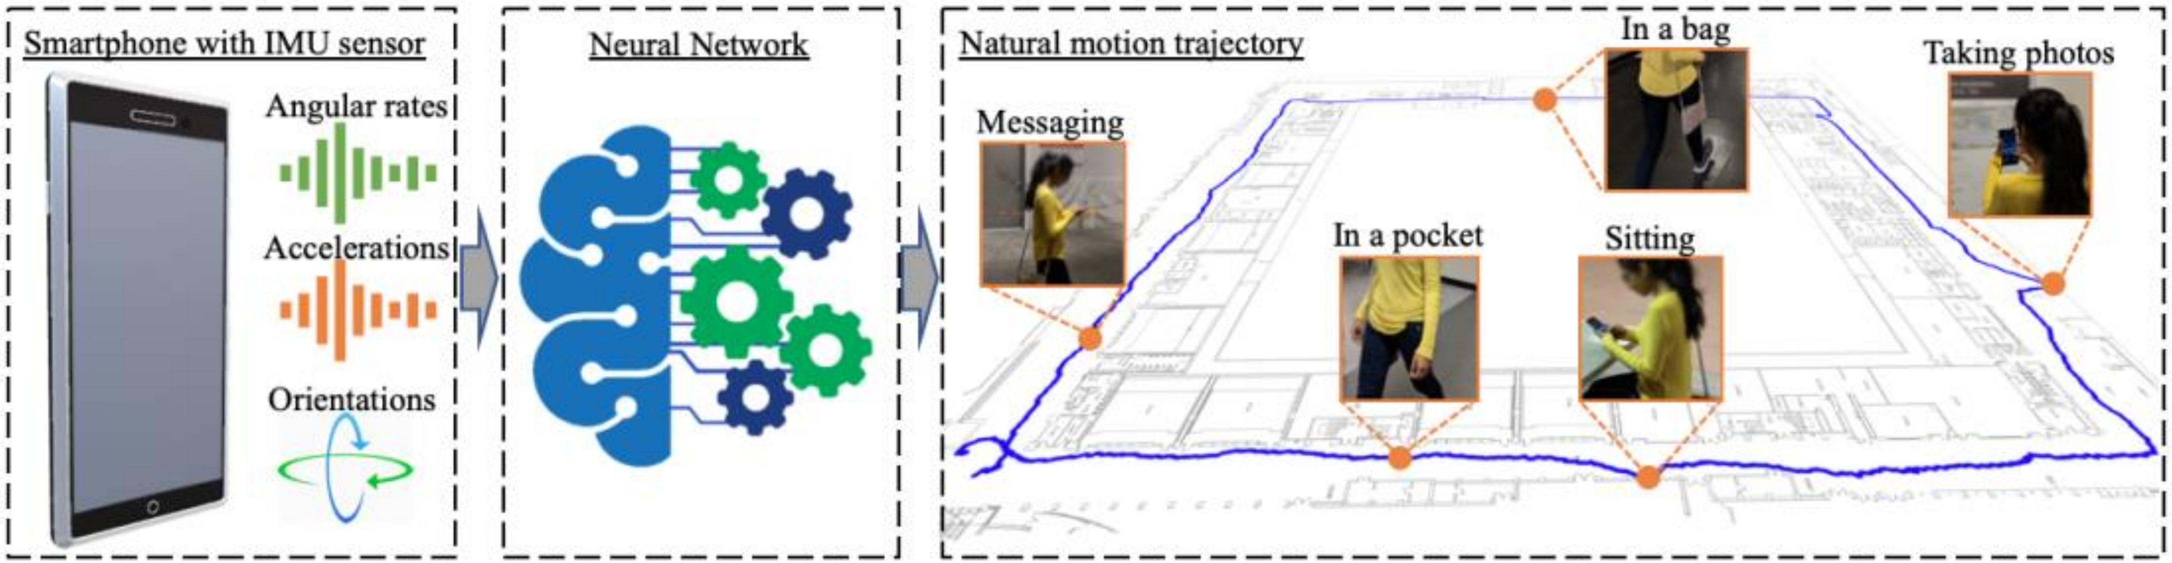

Tracking Results

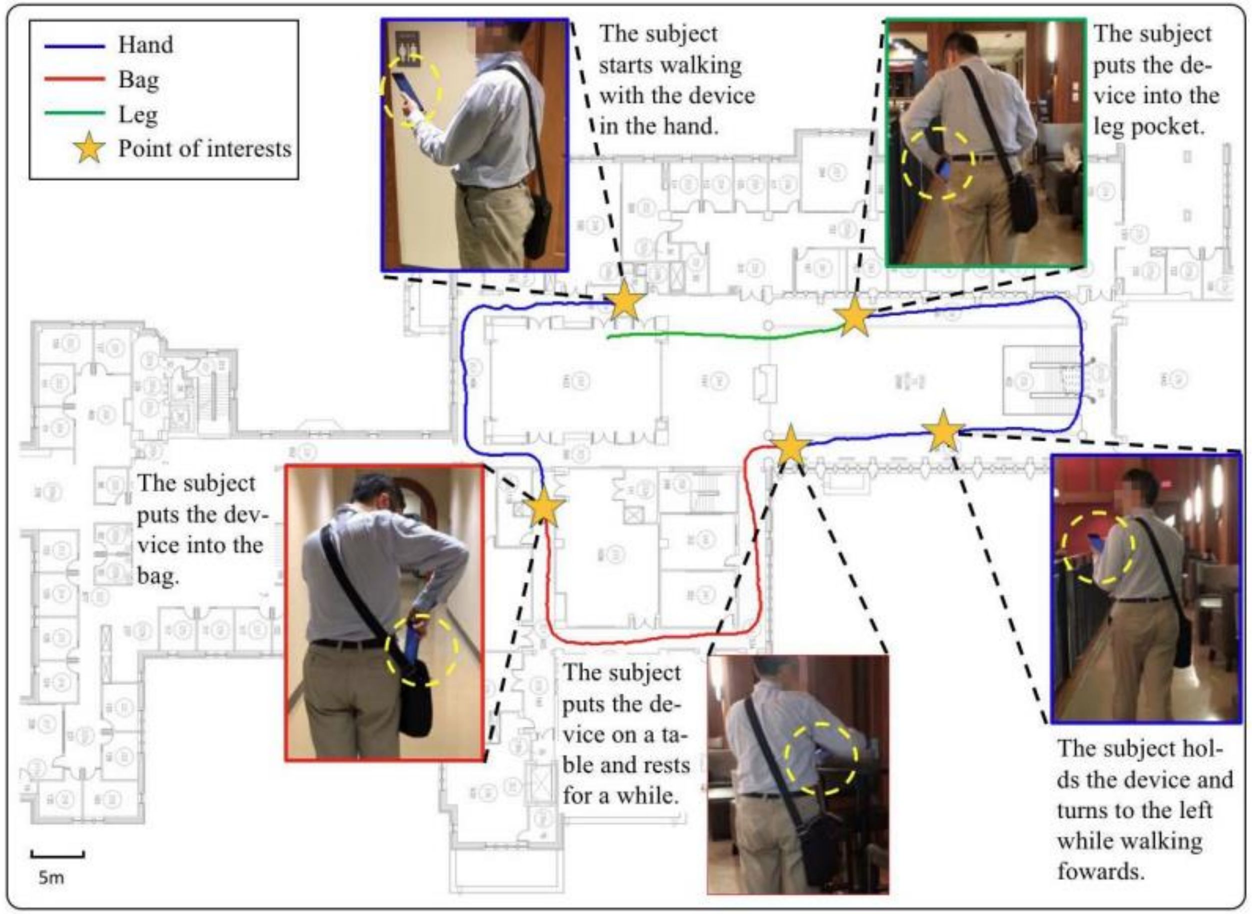

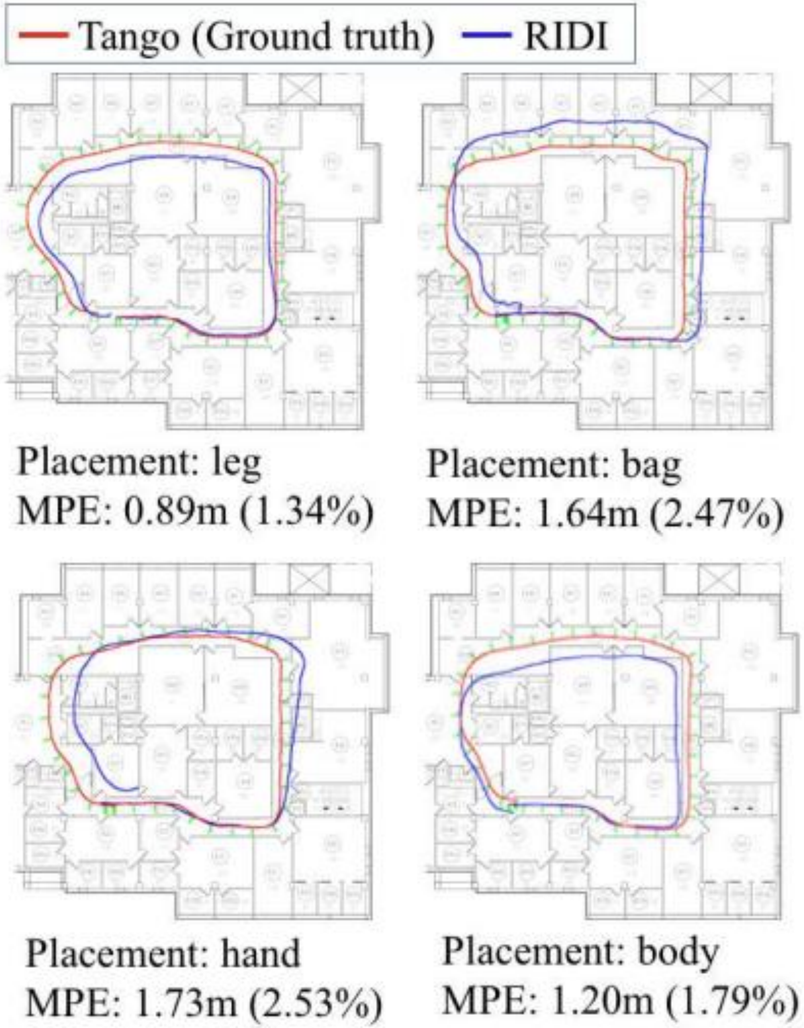

NeurIT: Time-Frequency Block-recurrent Transformer

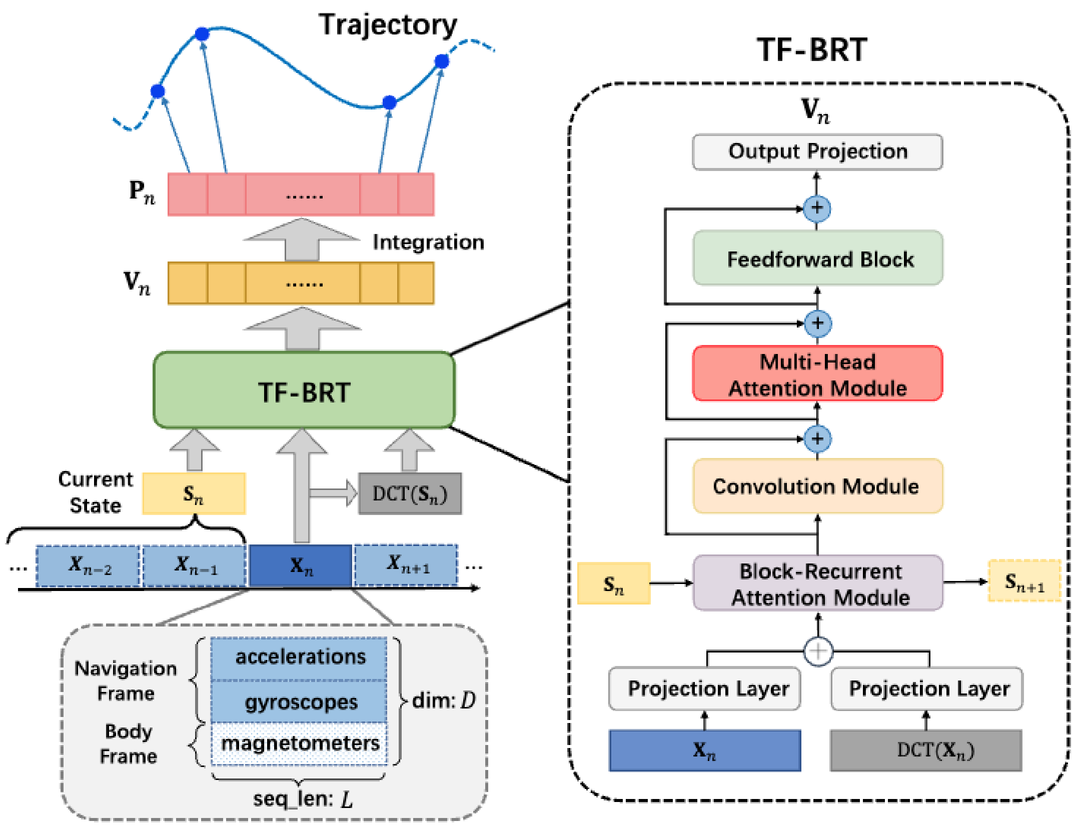

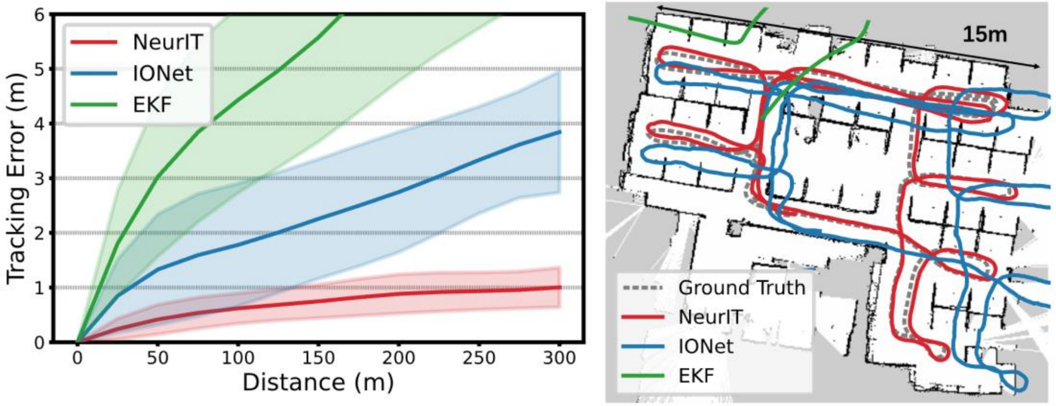

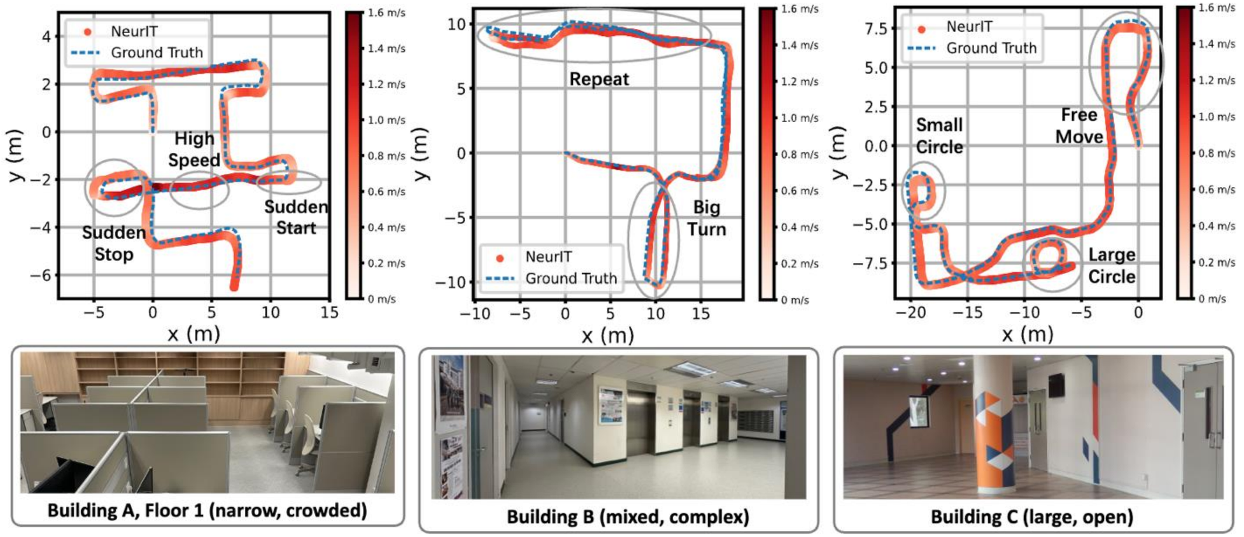

Xinzhe Zheng, Sijie Ji, Yipeng Pan, Kaiwen Zhang, Chenshu Wu, “NeurIT: Pushing the Limit of Neural Inertial Tracking for Indoor Robotic IoT”, 2024.

Inertial Measurement Unit Recap

Motion Parameters

- Moving Distance
- Heading Direction
- Rotating Angle

IMU

- Accelerometer: Measuring the linear acceleration
- Magnetometer: Reporting the absolute orientation
- Gyroscope: Calculating the angular velocity

Significant limitations in precise and robust motion estimation:

- **Accelerometer**: Noisy readings, step counting for distance
- **Gyroscope**: Accumulative errors due to integration
- **Magnetometer**: Environment interference, cannot infer heading direction

RIM: RF-based Inertial Measurement

- Turns COTS WiFi radio into precise IMU that measures motion parameters at centimeter accuracy:
  - Moving distance, Heading direction, Rotating angle

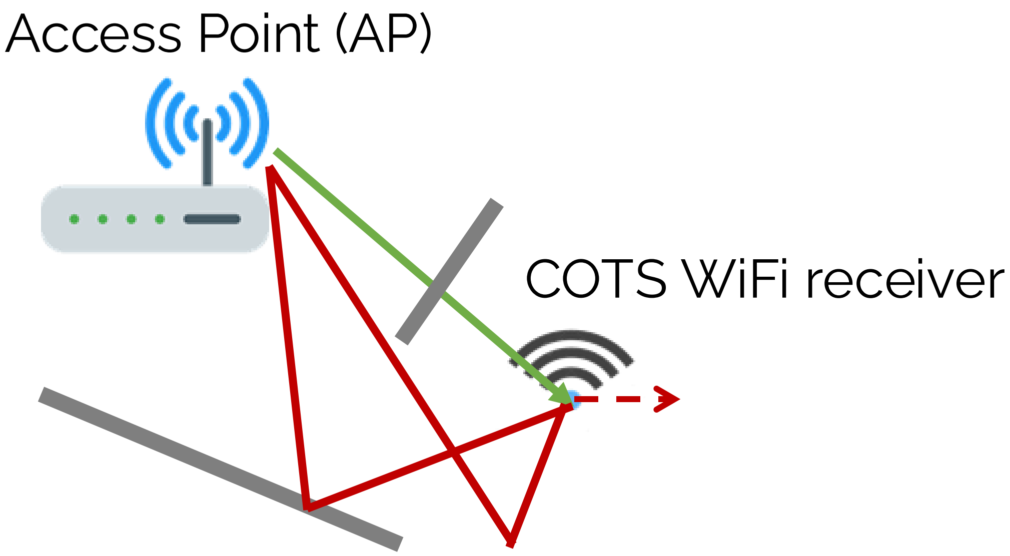

- One single arbitrarily placed AP
- No additional infrastructure
- Not require large bandwidth or many phased antennas
- No need of a priori calibration
- Works for LOS & NLOS

Wu, C., Zhang, F., Fan, Y., & Liu, K. R.. RF-based inertial measurement. ACM SIGCOMM 2019.

Virtual Antenna Alignment

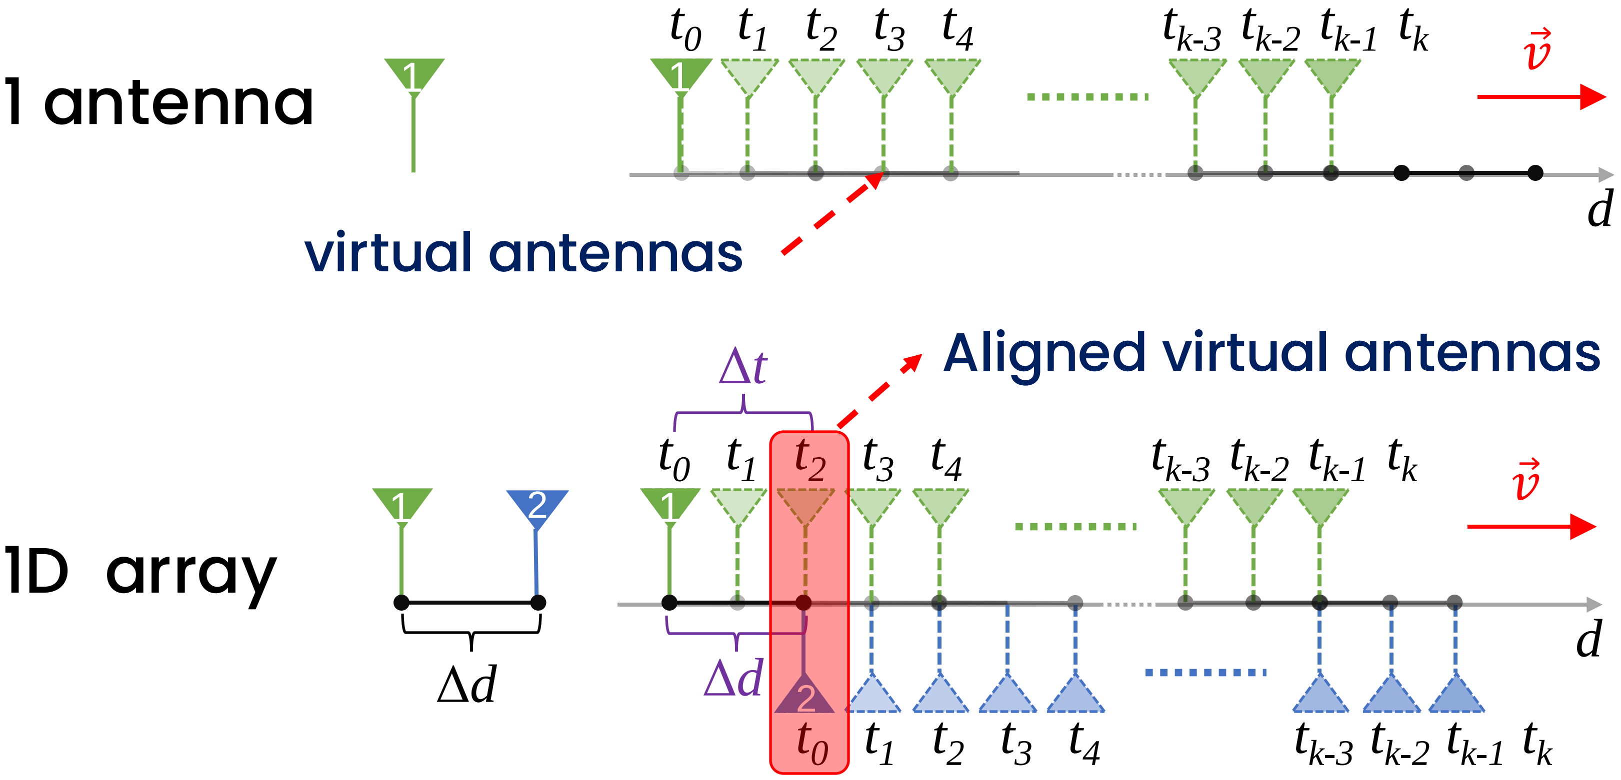

Multipath Profiles as Virtual Antennas!

$$\hat{v} = \frac{\Delta d}{\Delta t}$$

$$\hat{d} = \int_{t_0}^{t_k} \hat{v} dt$$

Moving distance ✔

Super-Resolution Virtual Antenna Alignment

How to accurately pinpoint the space-time point that two virtual antennas are aligned with each other, at sub-centimeter resolution?

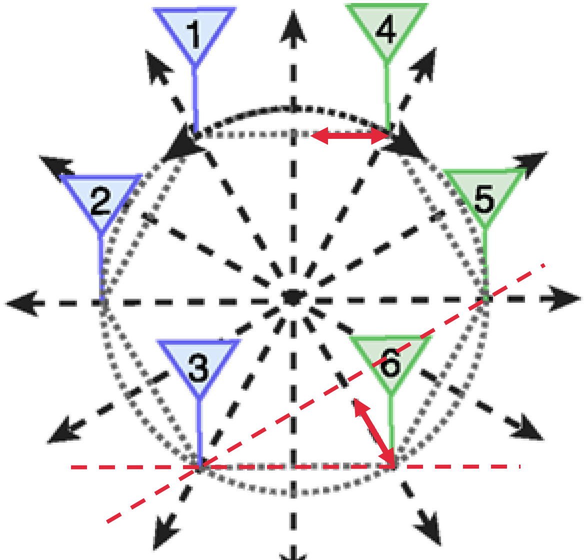

$\Delta d = \lambda/2$

e.g., 1cm error = ~ 50% error in speed = 30◦ heading error = 22◦ rotation error

Time-Reversal Principle

Time Reversal Transmission

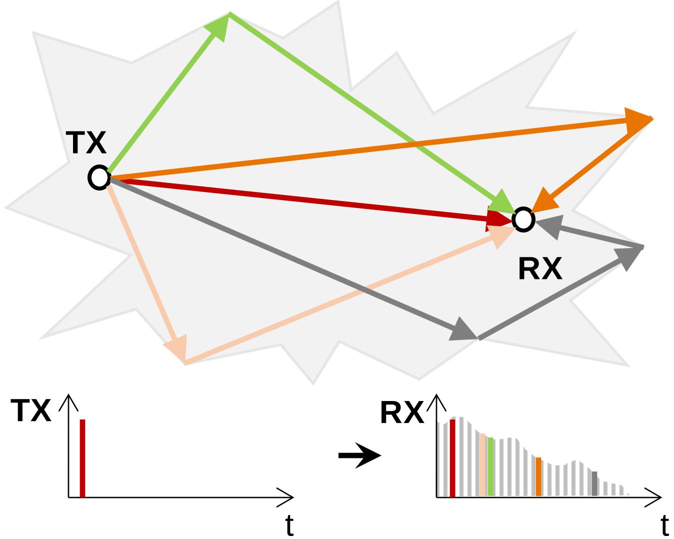

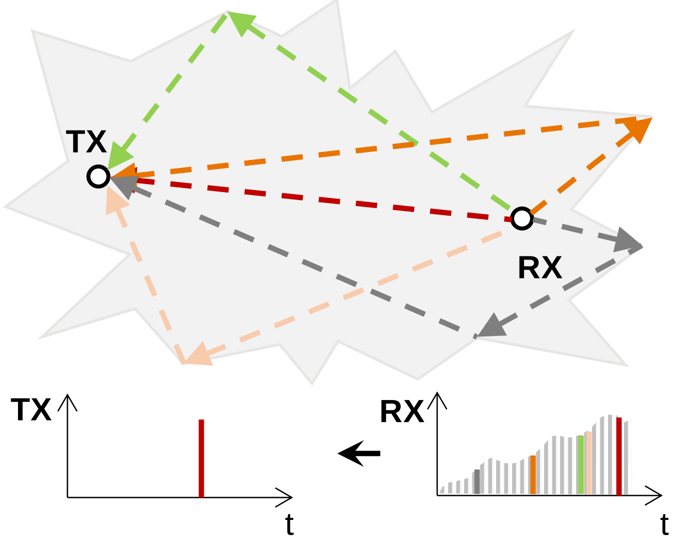

Time Reversal Resonating Effect

TR Resonating Strength (TRRS)

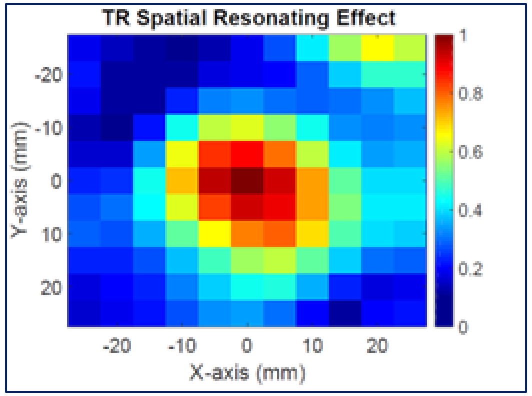

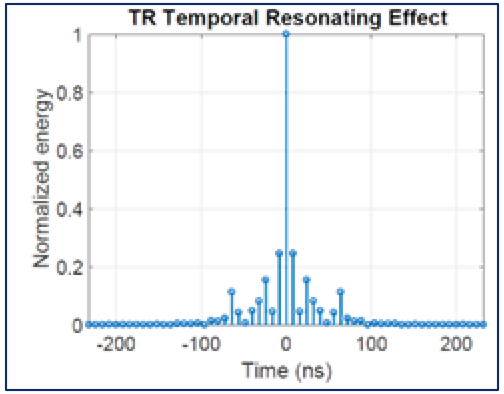

Wu, Z. H., Han, Y., Chen, Y., & Liu, K. J. R.. A time-reversal paradigm for indoor positioning system. IEEE TVT 2015.
Zhang, F., Chen, C., Wang, B., Lai, H. Q., Han, Y., & Liu, K. J. R.. WiBall: A time-reversal focusing ball method for decimeter-accuracy indoor tracking. IEEE IOTJ, 2018

Time Reversal Resonating Strength (TRRS)

- **Time-Reversal Focusing Effect**: The received CSI, when combined with its time-reversed and conjugated counterpart, will add coherently at the intended location but incoherently at any unintended location, creating a spatial focusing effect

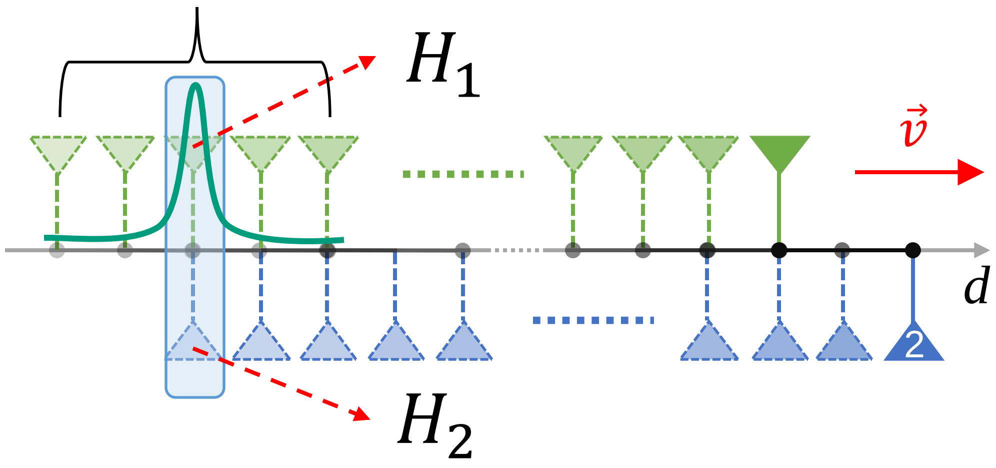

TRRS

$$\kappa(\mathbf{h}_1, \mathbf{h}_2) = \frac{\left(\max_i \left|(\mathbf{h}_1 * \mathbf{g}_2)[i]\right|\right)^2}{\langle \mathbf{h}_1, \mathbf{h}_1 \rangle \langle \mathbf{g}_2, \mathbf{g}_2 \rangle} \quad \text{(CIR)}$$

$$\kappa(H_1, H_2) = \frac{|H_1^{\mathrm{H}} H_2|^2}{\langle H_1, H_1 \rangle \langle H_2, H_2 \rangle} \quad \text{(CFR)}$$

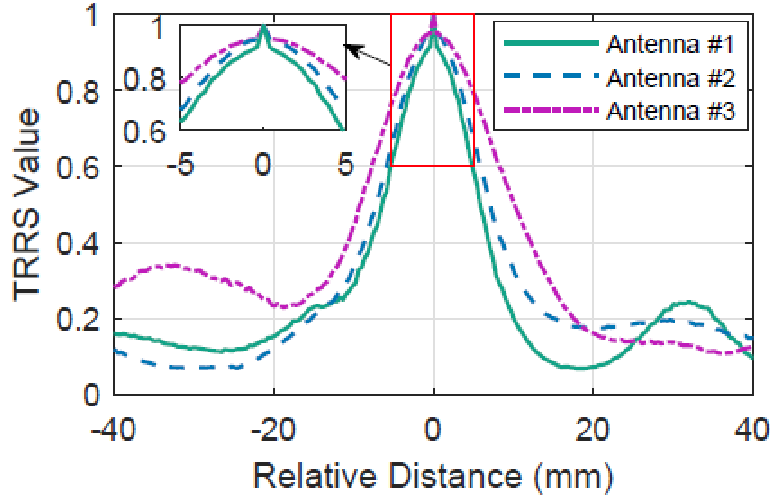

TRRS Resolution

- The peak value as high as possible
- The peak width as narrow as possible
- The above two properties as robust as possible

Virtual Massive Antennas

- Overcome distortions in TRRS: Leveraging consecutive multipath profiles as massive virtual antennas

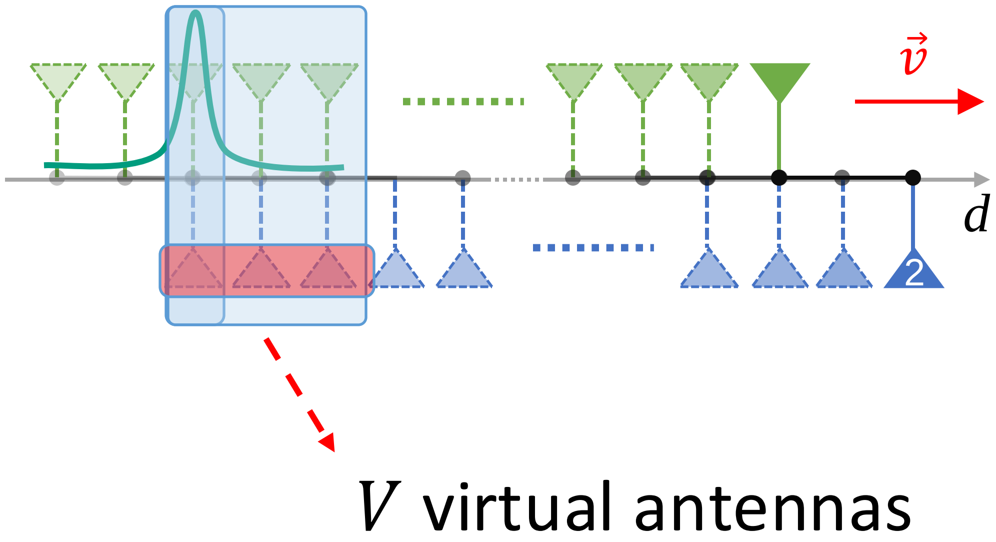

$$\kappa\left(P_i(t_i), P_j(t_j)\right) = \frac{1}{V} \sum_{k=-V/2}^{V/2} \bar{\kappa}\left(H_i(t_i + k), H_j(t_j + k)\right)$$

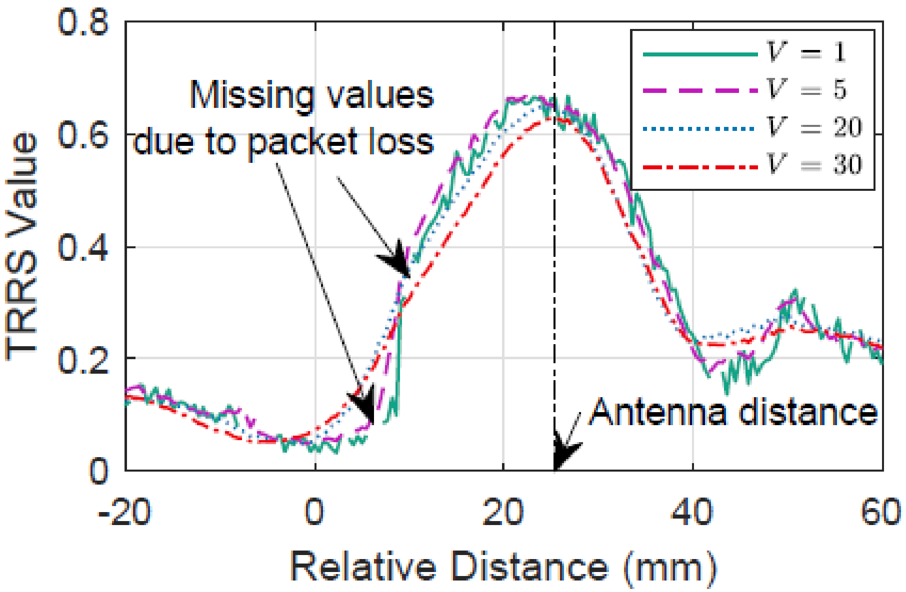

Tracking Alignment Delay

- Continuously track alignment delay via Dynamic Programming

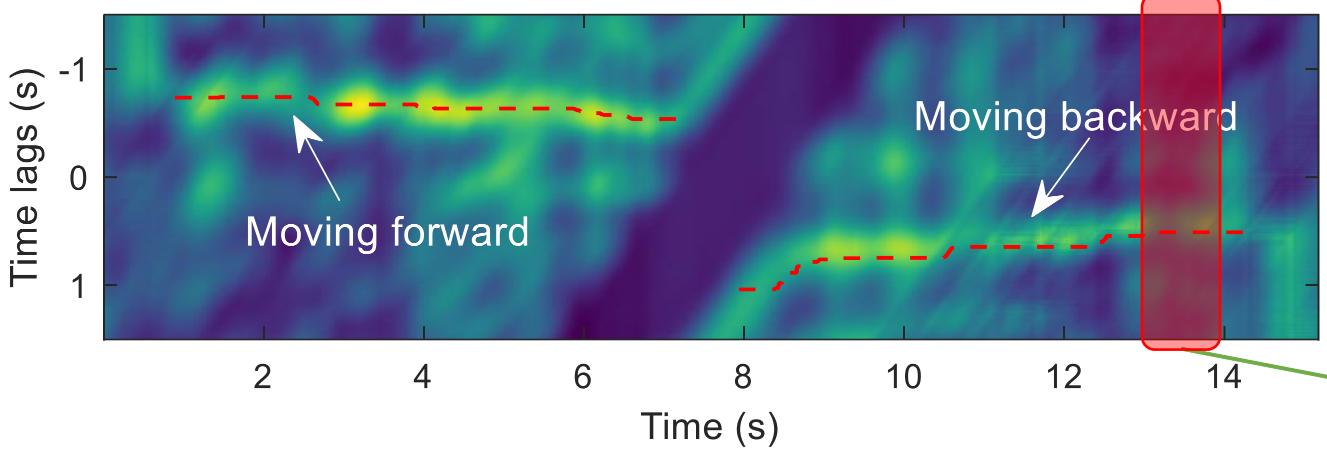

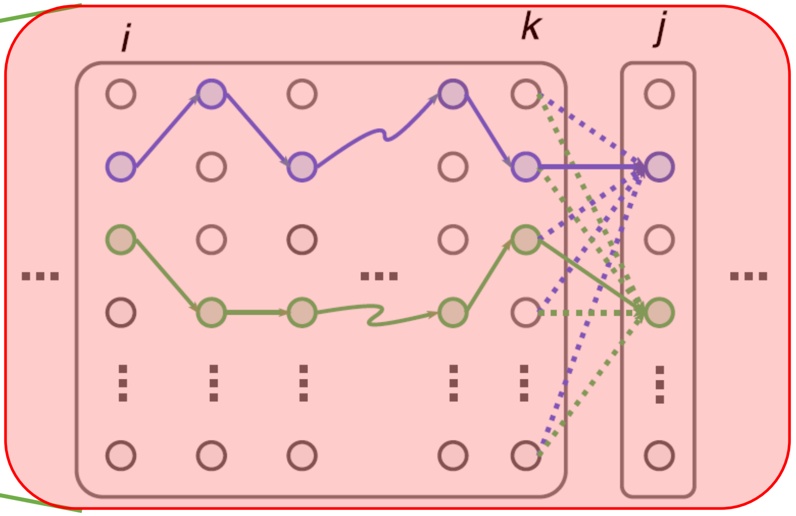

$$S(\boldsymbol{q}_i \rightsquigarrow q_{jn}) = \max_{l \in [-W, W]} \{S(\boldsymbol{q}_i \rightsquigarrow q_{kl}) + S(q_{kl} \rightsquigarrow q_{jn})\}$$

$$S(q_{kl} \rightsquigarrow q_{jn}) = e_{kl} + \boxed{e_{jn}} + \boxed{\omega C(q_{kl}, q_{jn})}$$

Peak value (negative) Cost function

$$\boxed{n^\star = \arg\max_{n \in [-W, W]} \{S(\boldsymbol{q}_i \rightsquigarrow q_{jn})\}}$$

A WiFi Ruler with RIM

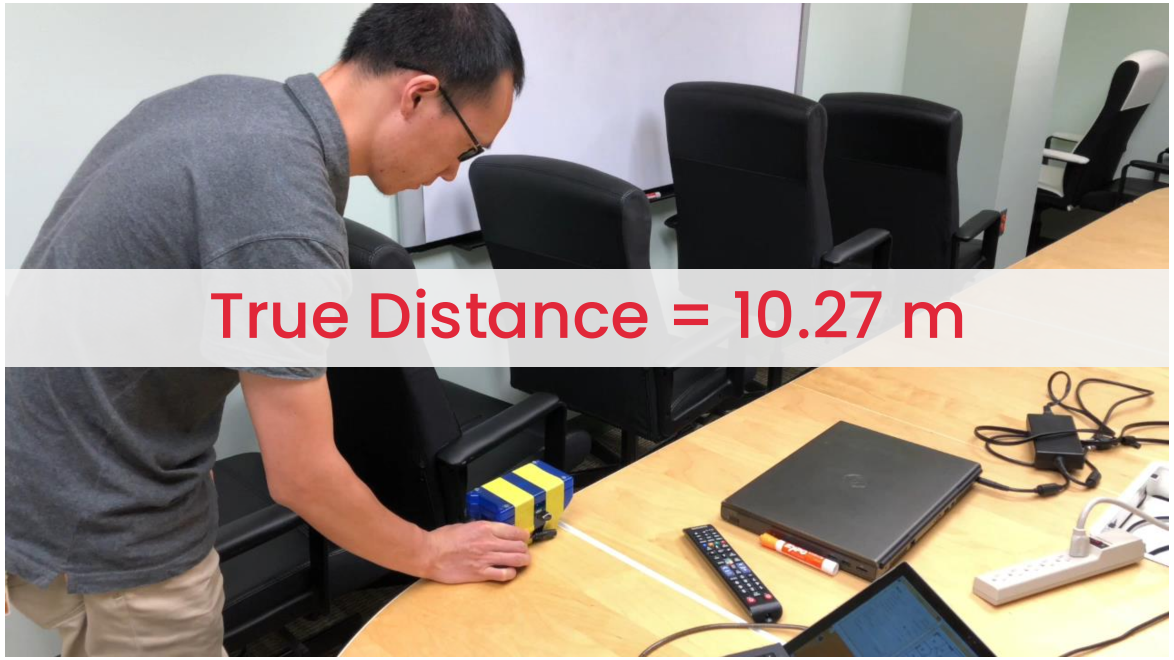

Measure the perimeter of a big round table

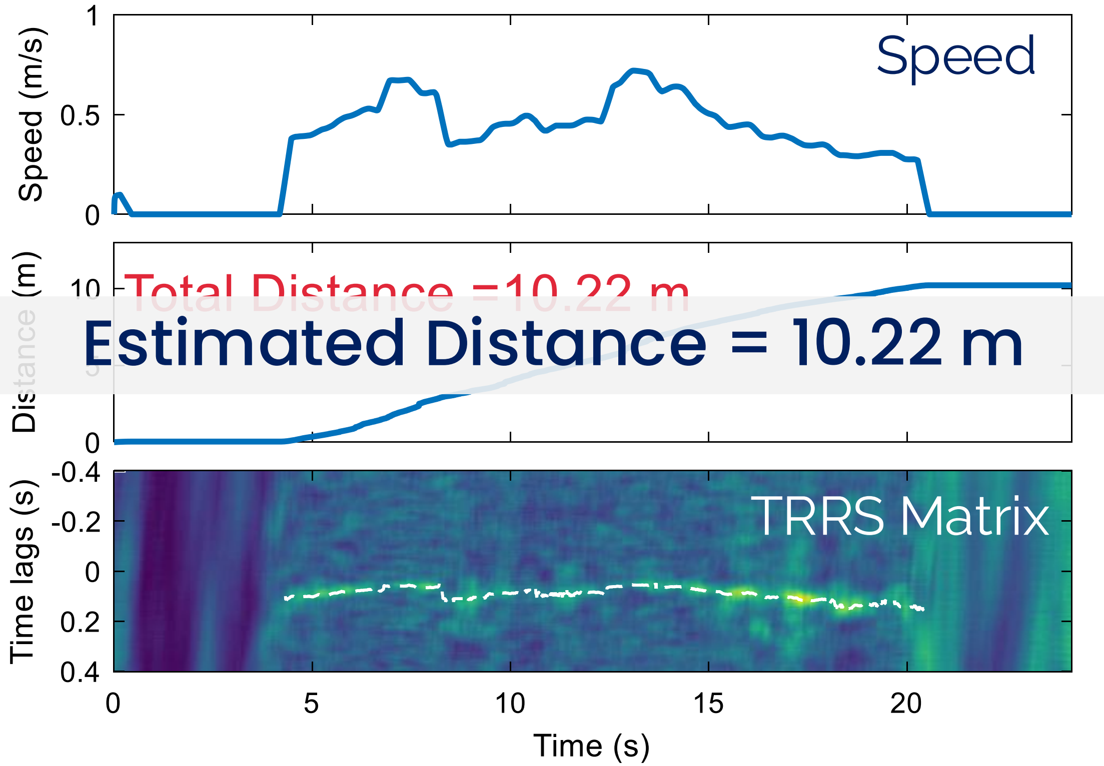

Results

- **1** single AP, **7** different locations
- Both LOS and NLOS (40m away through multiple walls)
- 200Hz sampling rate on a 40MHz channel in the 5GHz band

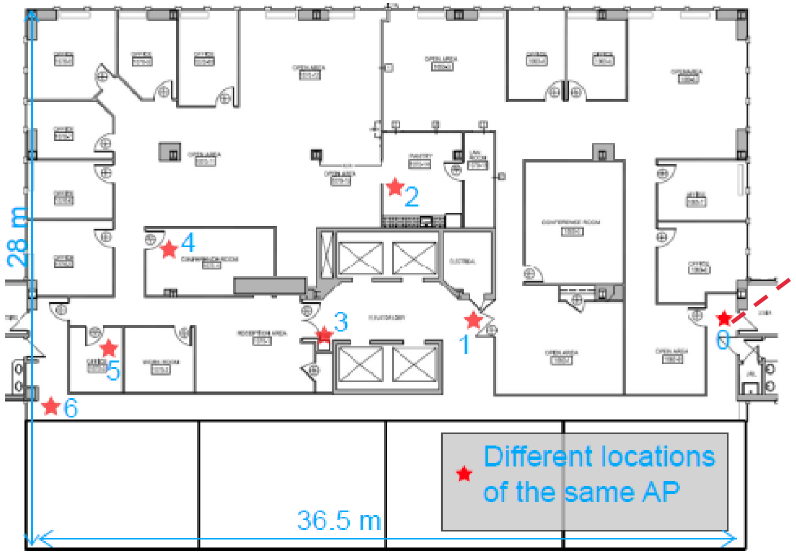

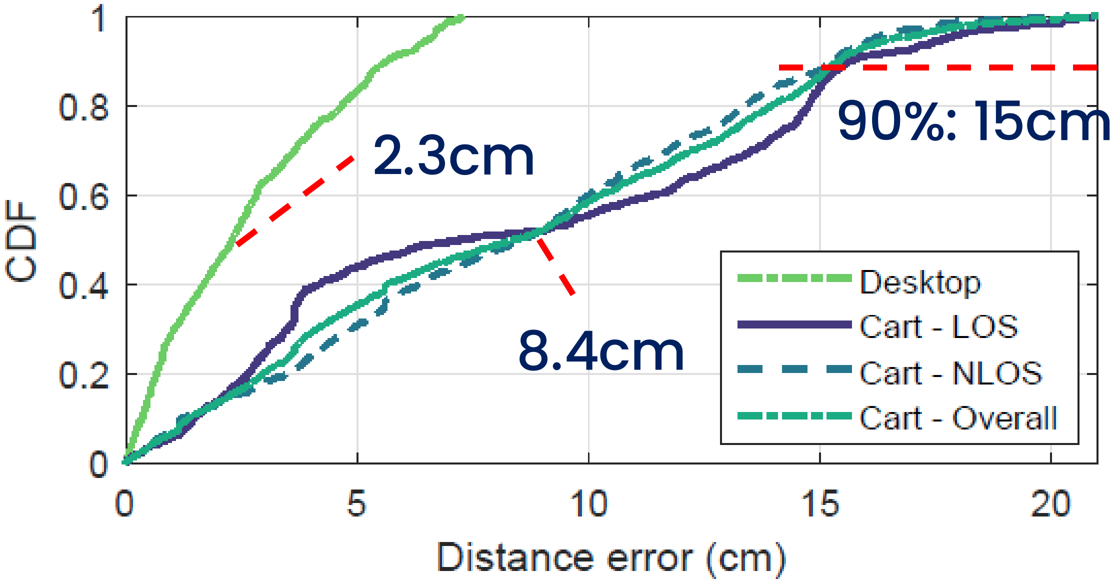

How is RIM useful in practice?

- It tolerates certain deviation.
- Good for robot/cart/asset tracking, not ideal for human tracking.

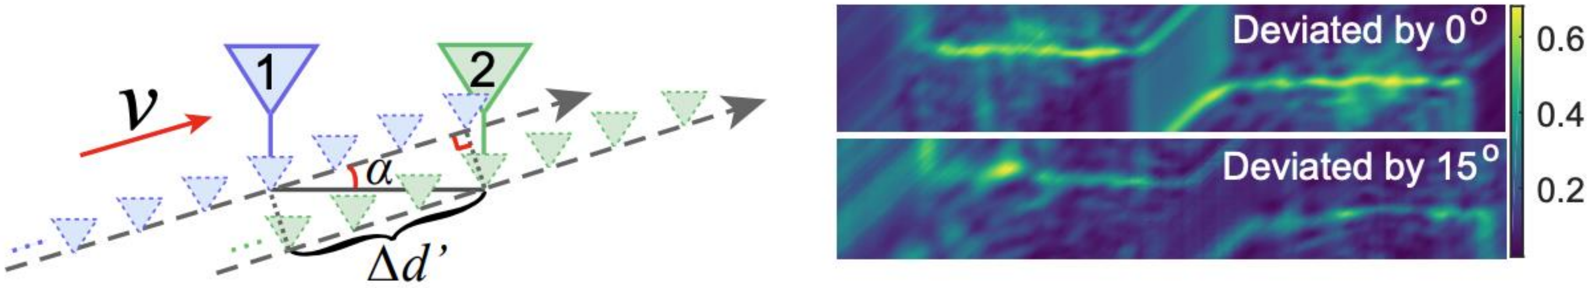

(a) (b)

Problems

- **However accurate it predicts, the errors always accumulate**
- Useful for short-term tracking
- Fusion with other modalities
  - Augment GPS (GPS alone may not be accurate)
  - Visual-inertial odometry
  - WiFi SLAM (Simultaneous Localization and Mapping)
  - Mapping

How to overcome drifts?

- Find global/absolute references to overcome local/relative errors

- External information
  - WiFi, GPS, Bluetooth, Vision…

- Internal information
  - Use IMUs differently, e.g., to find landmarks with unique motion patterns

EasiTrack: RIM + Indoor Maps

Large-Scale Decimeter-Level Indoor Tracking with a Single AP

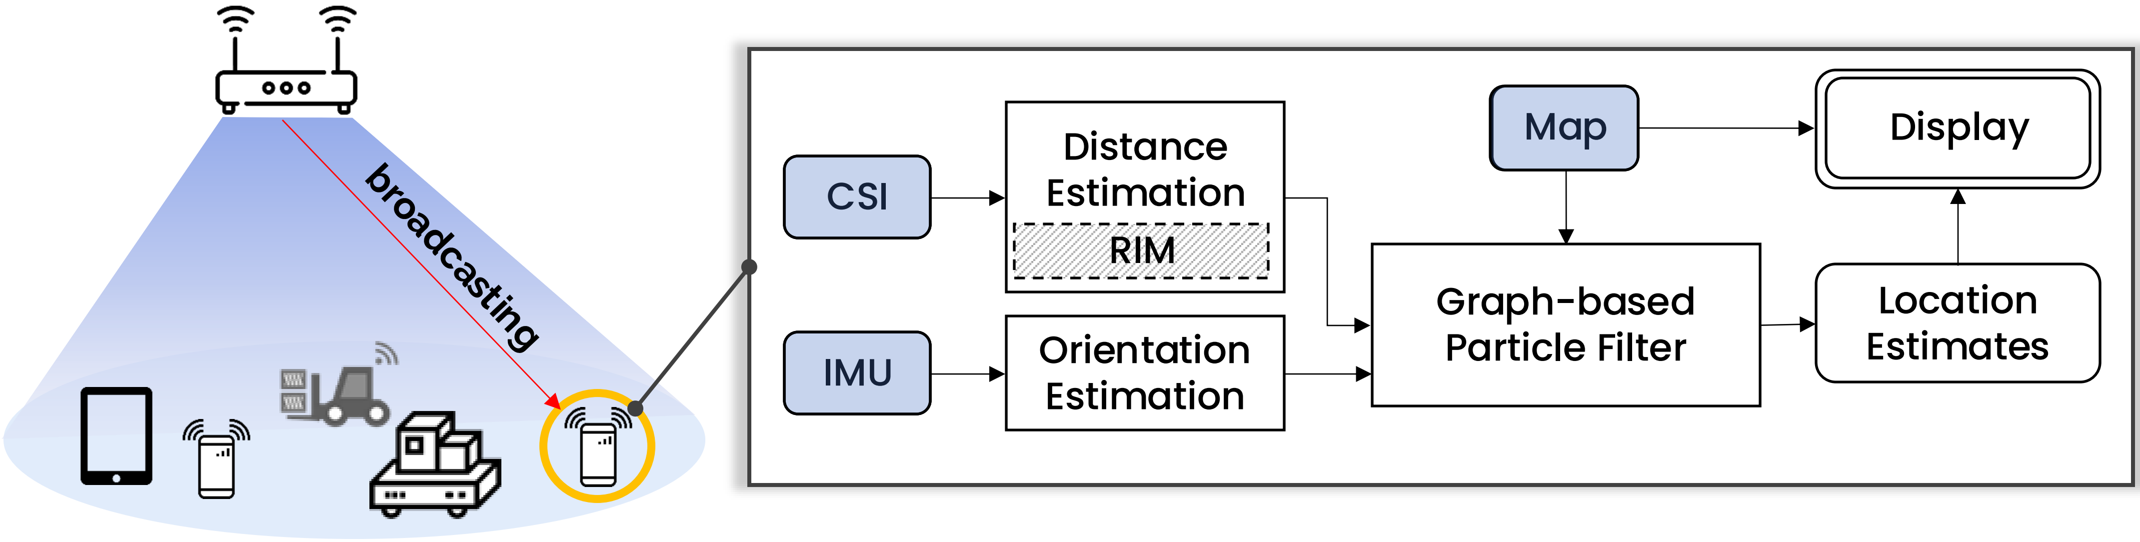

EasiTrack: Easy, Accurate, Scalable Indoor Tracking

Wu, C., Zhang, F., Wang, B., & Liu, K. J. R. EasiTrack: Decimeter-Level Indoor Tracking With Graph-Based Particle Filtering. IEEE Internet of Things Journal, 2019.

Bring Maps to Indoor Tracking

- Maps impose constrains of movements
  - E.g., people do not penetrate walls

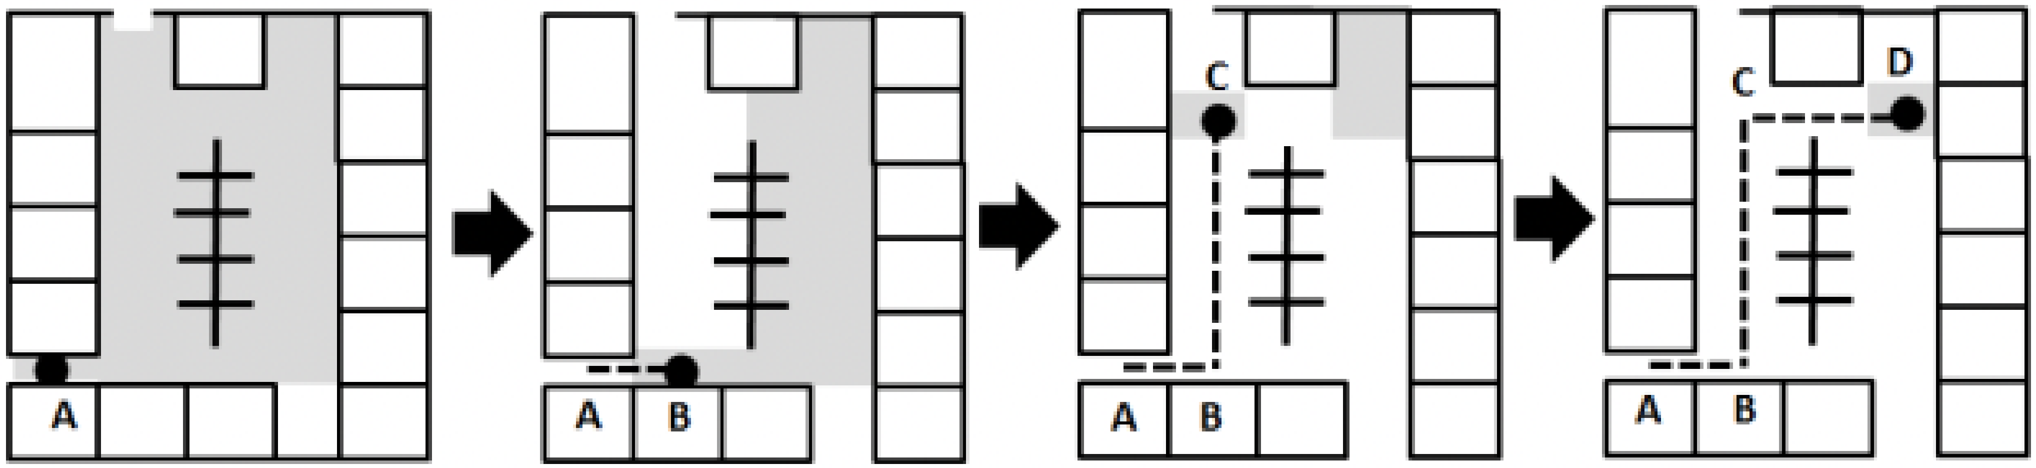

Particle Filter

- Sequential Monte Carlo methods
  - Represent the posterior distribution of some stochastic process given noisy and/or partial observations with a set of samples (i.e., *particles*)
- (1) Prediction
  - Move to the next position with a **motion model**
- (2) Updating
  - Update the likelihood weight of each particle using **measurements**
- (3) Resampling
  - Sequential Importance Resampling (SIR)
  - Overcome the *degeneracy problem*: most of the weights are close to zero
- (4) Estimation
  - Determine the *target* by the particles

The power of randomness

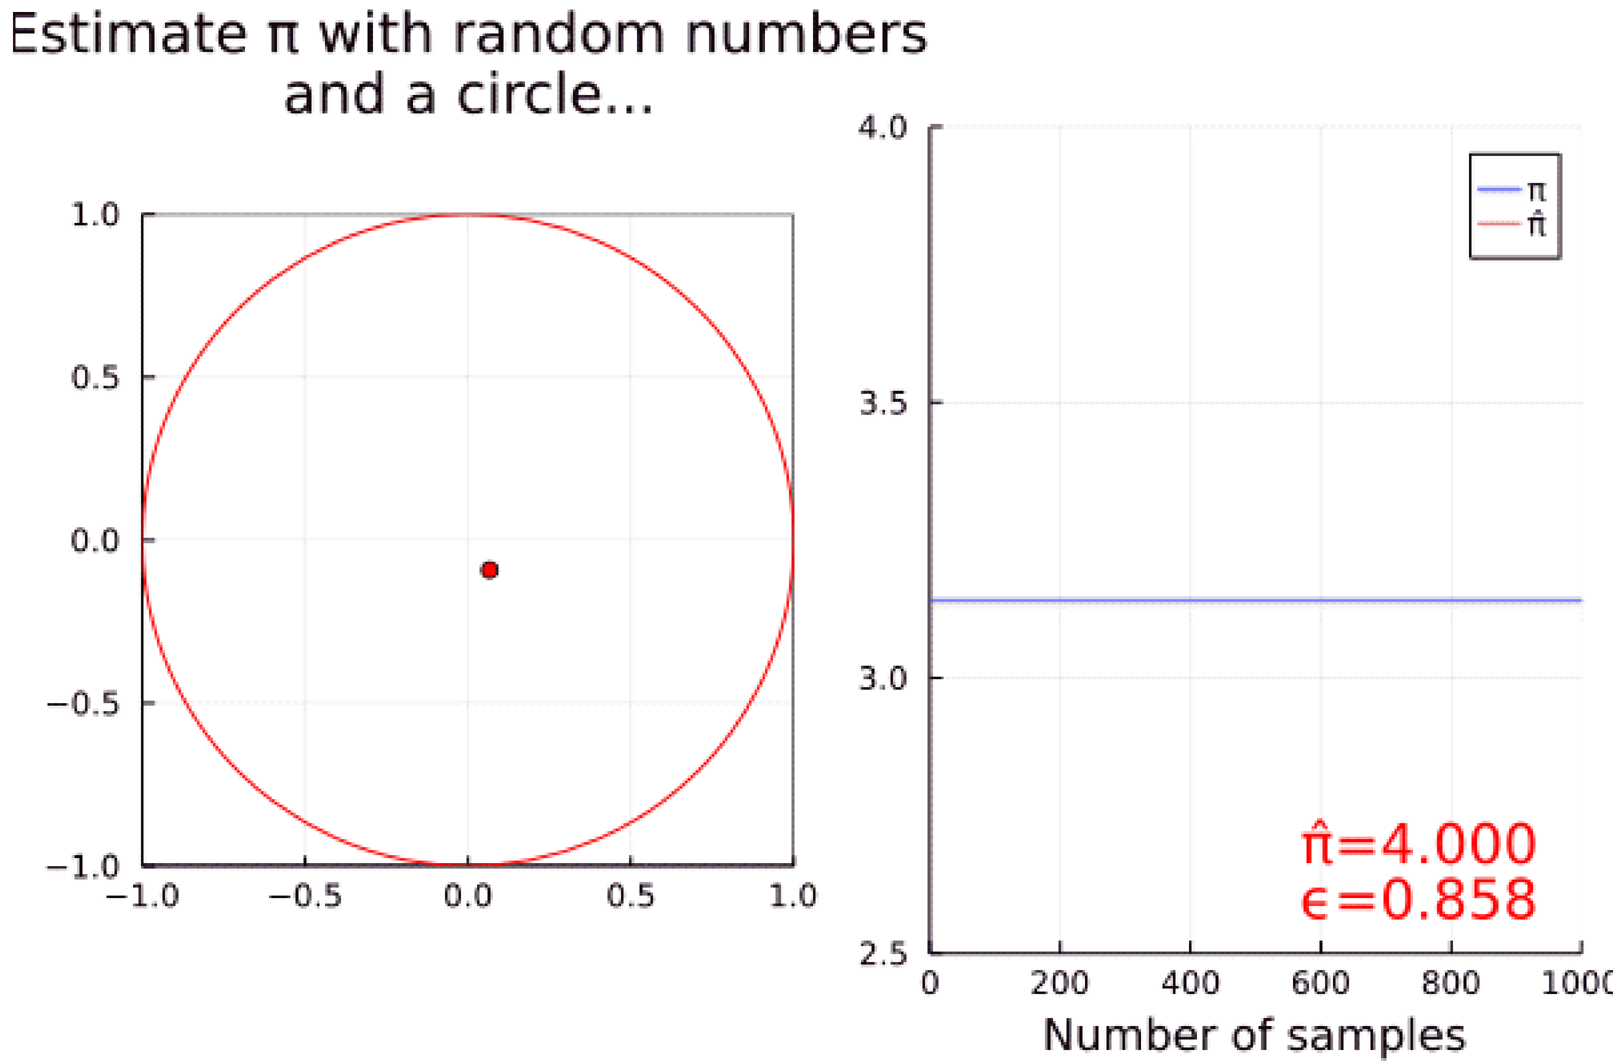

An Illustrative Example

Importance Sampling

Weighting

Sequential Importance Resampling

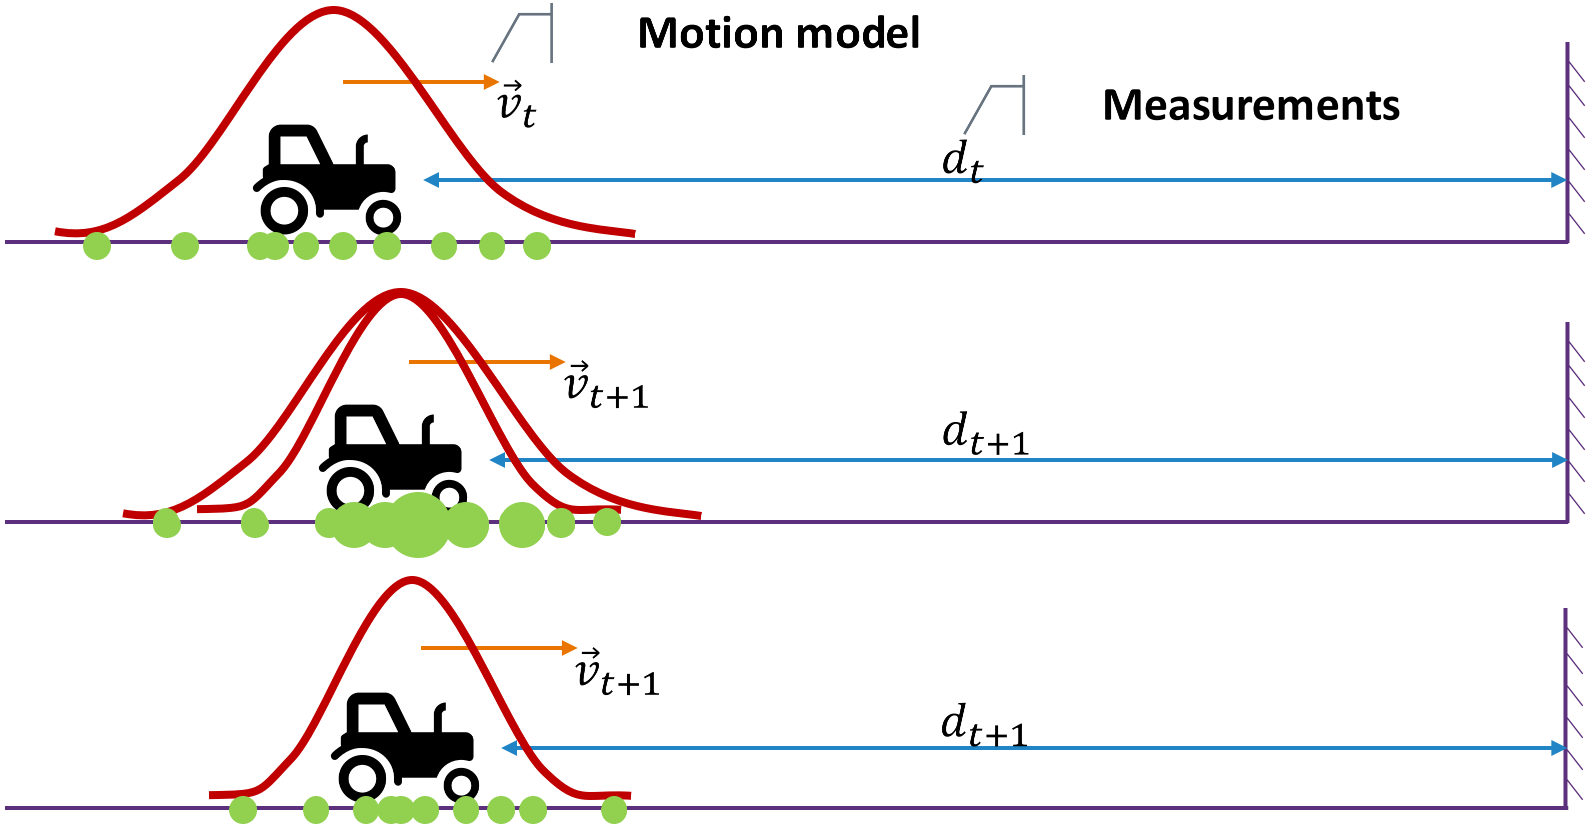

Particle Filter based Map Correction

- (1) Prediction
  - Move to the next position with $(\boldsymbol{\theta}, \boldsymbol{\Delta d})$ **by position engine**

    Motion model
- (2) Updating
  - Update the likelihood weight of each particle using **Map Info**

    Measurements
- (3) Resampling
  - Sequential Importance Resampling (SIR)
- (4) Estimation
  - Determine the *target* by the particles

Key Challenges

- Without secondary measurements, how to specify the importance and determine the weights of particles?
    - Typical systems have some additional measurements (e.g., laser ranging, WiFi-based estimations) for this purpose
- Without global ranging, how to overcome accumulative errors?
    - Errors, in particular direction errors from IMU sensors accumulate significantly over time

Particle Weighting (1): Hit and Die

- Initially, each particle gets an equal weight of 1/N
- Any particles that hit the inaccessible areas (e.g., a wall) during a move (prediction) will die; others survive
- Set the likelihood weights of “dead particles” to be 0.
- For any living particle, weight it by its “***Distance-to-live***”

Particle Weighting (2): Distance-To-Live (DTL)

- Derived by particle status (position, direction) and map constraints
- DTL: the max accessible distance from current position along the current moving direction
- Max-DTL: force overlarge DTLs to be the same

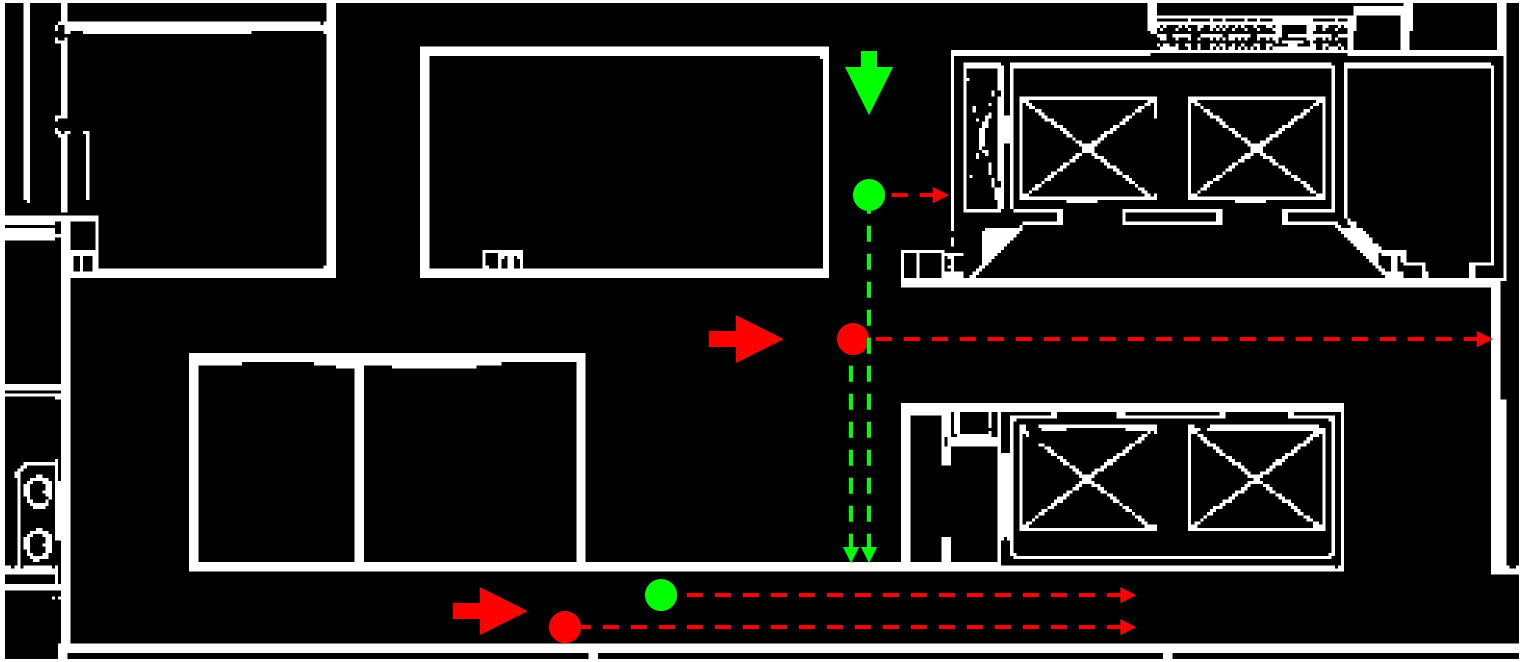

Particle Filter Tracking

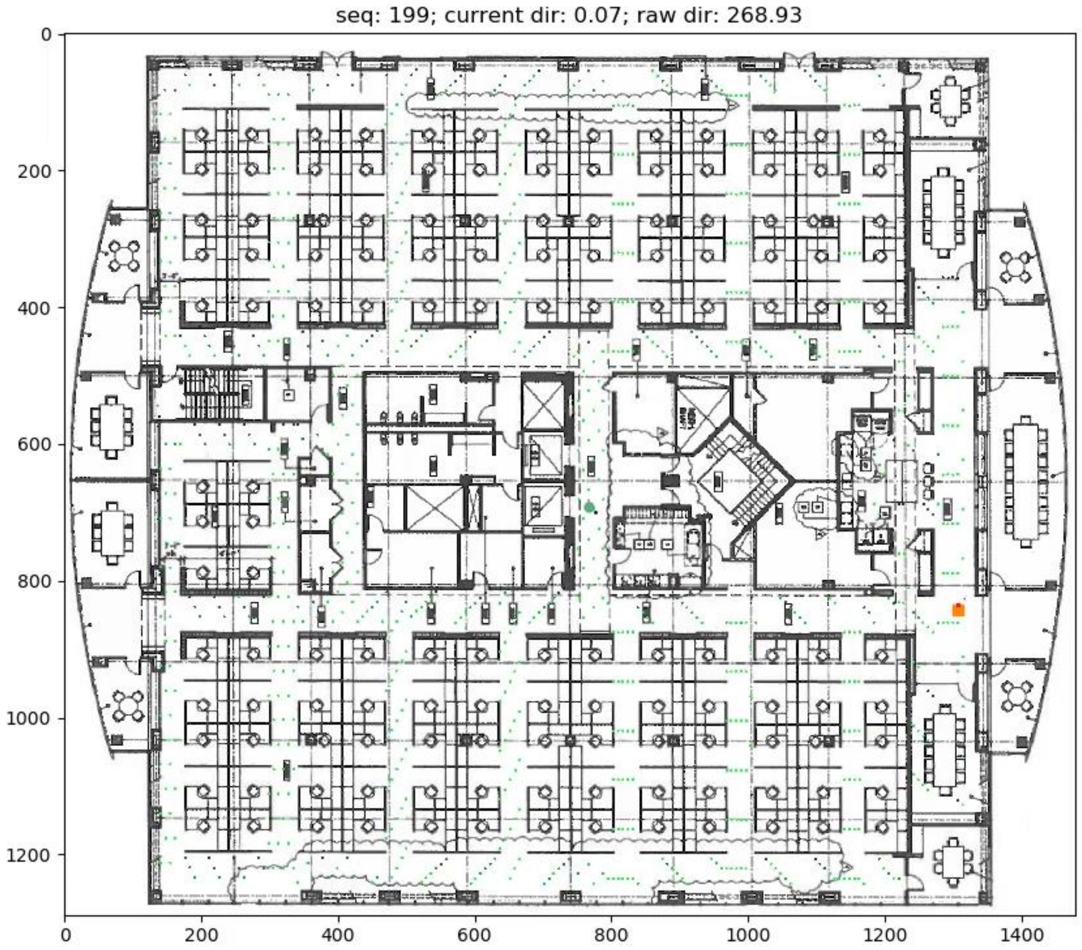

Inertial Tracking with Maps

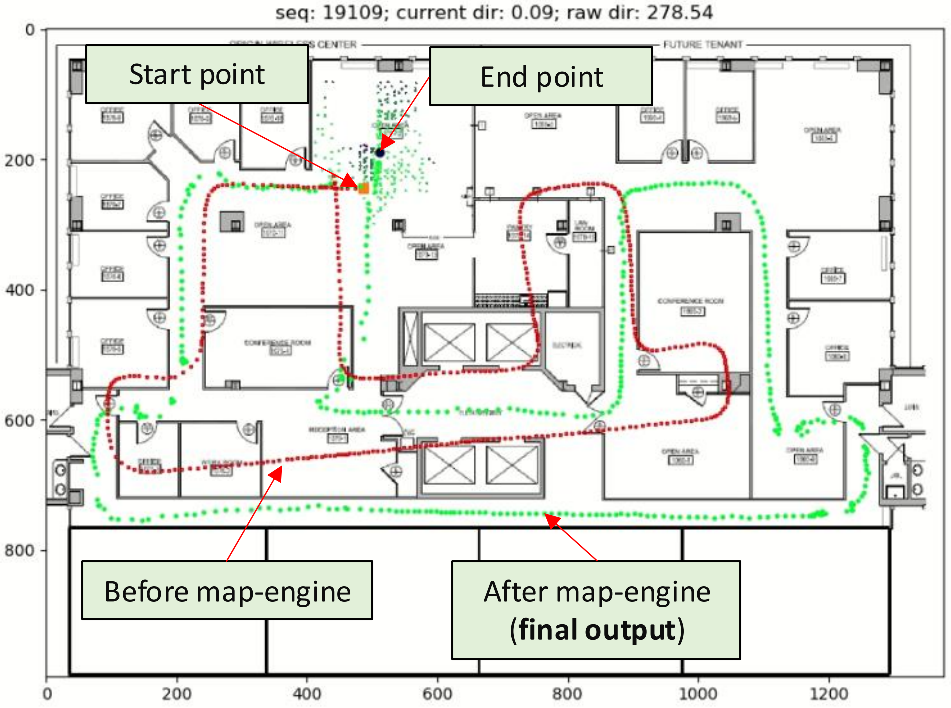

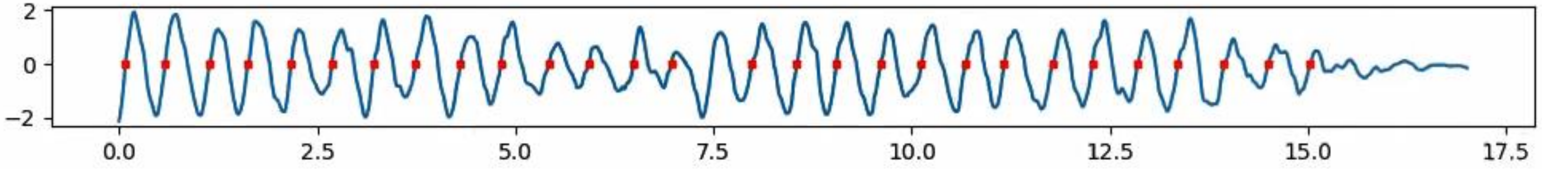

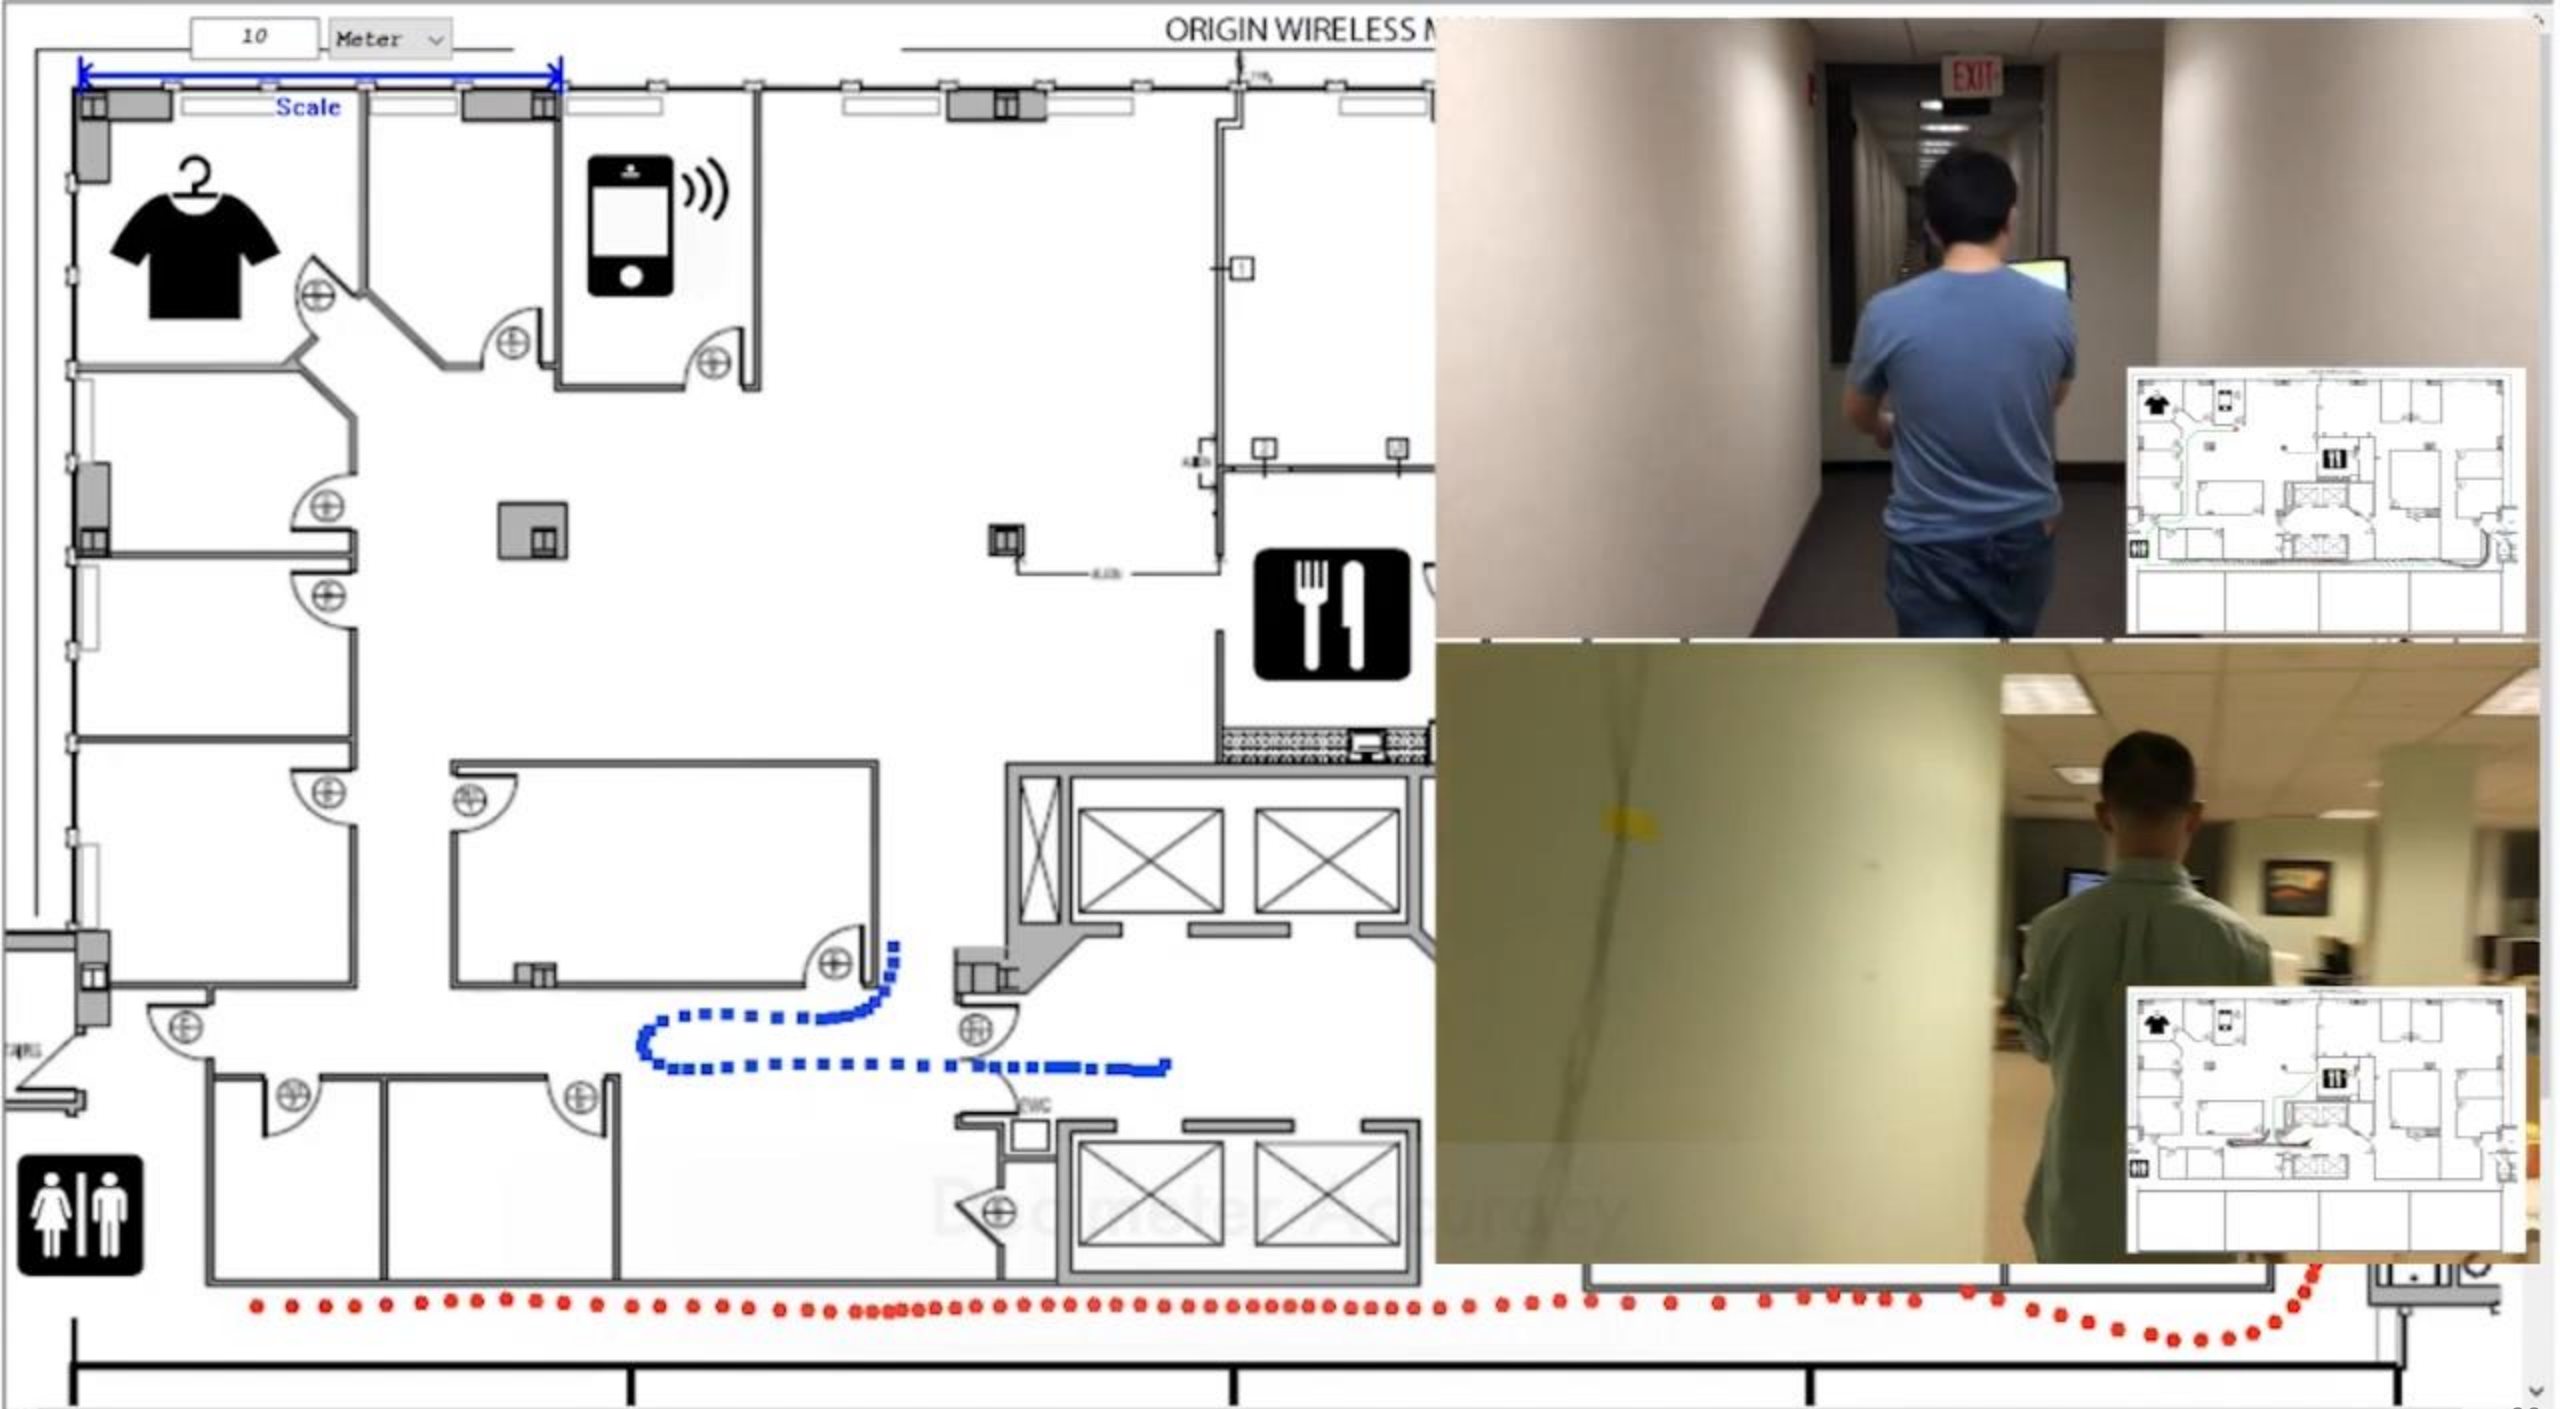
10
Meter
ORIGIN WIRELESS M
Scale
EXIT

Demos

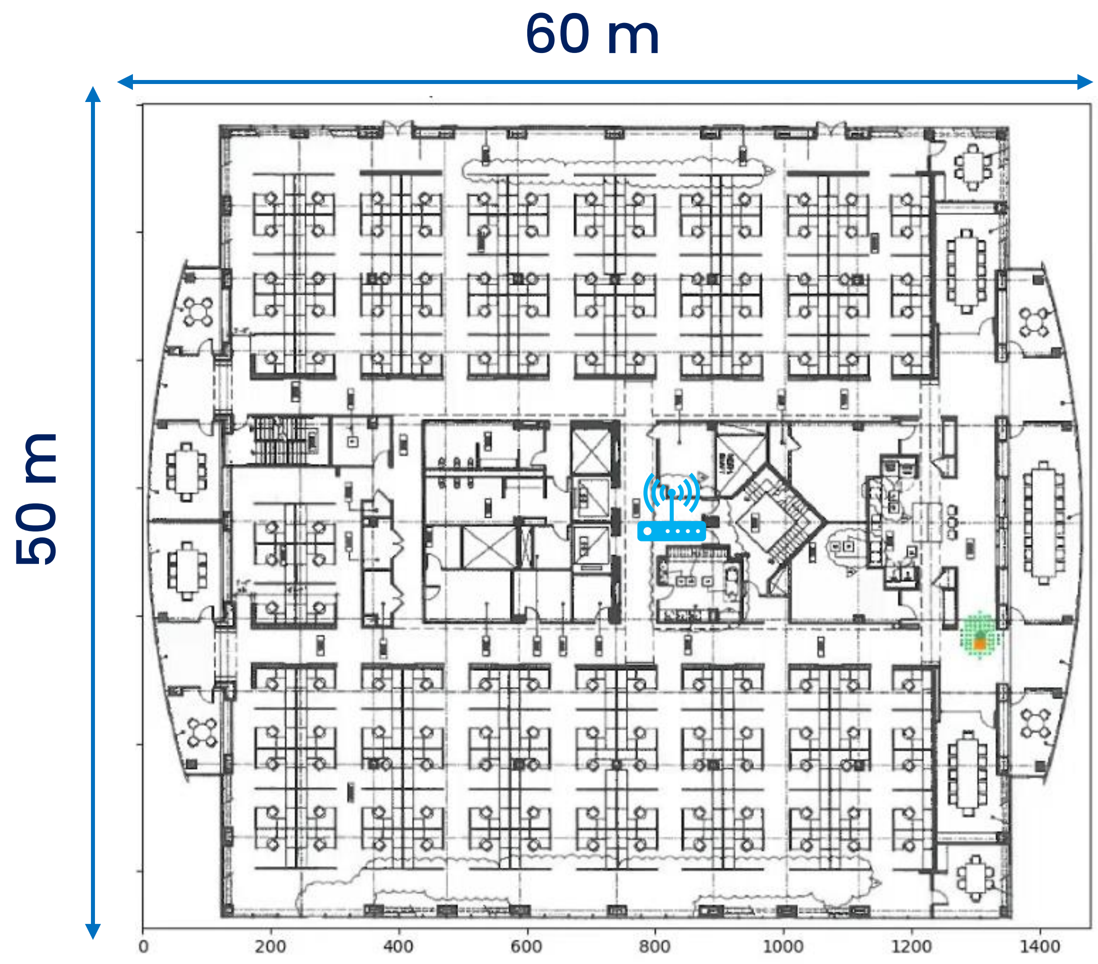

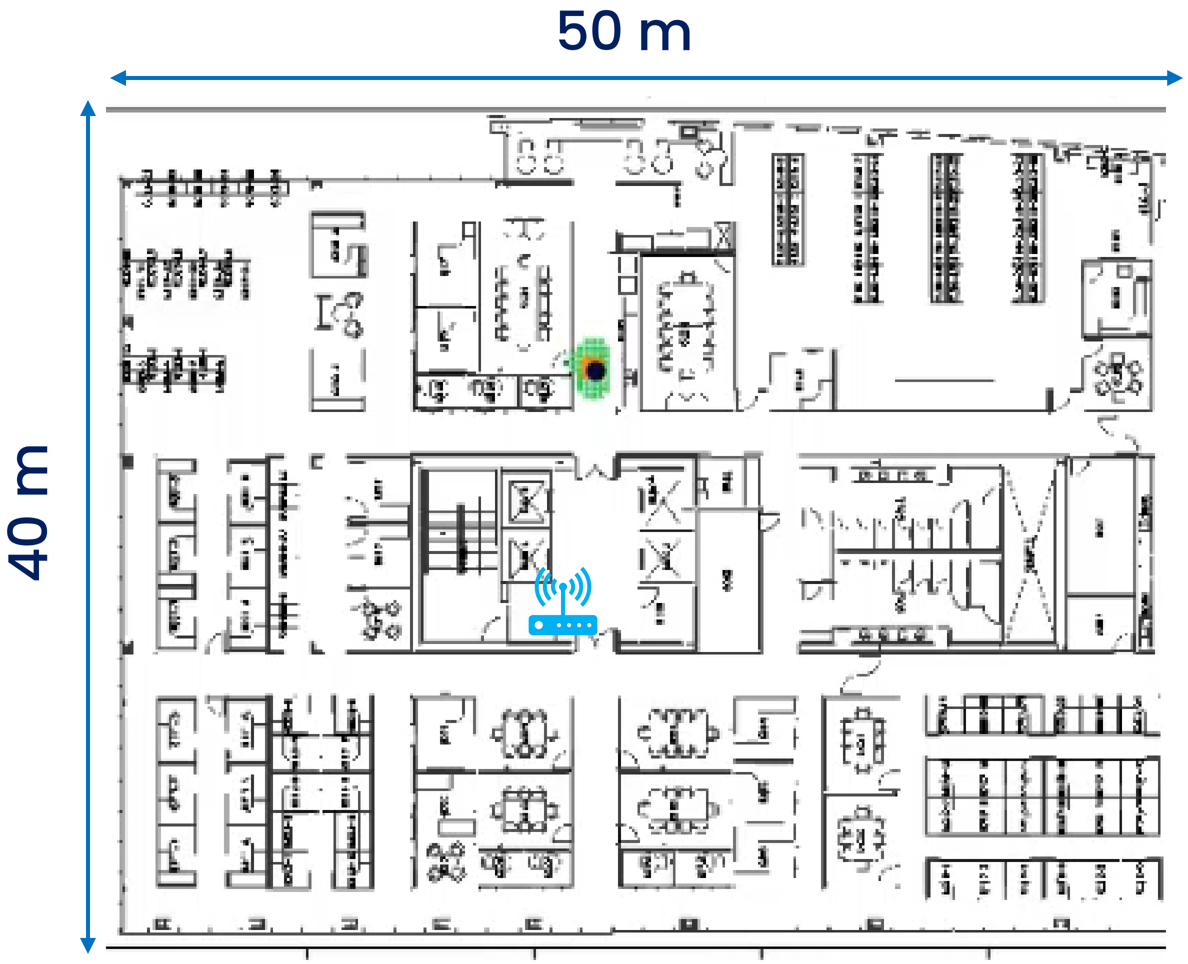

Raw trace w/o map

Final trace w/ map

Map-based Correction

- No other measurements (WiFi RSS, BLE, etc) needed
- Only a plain image of indoor floorplan
  - Represented as a binary image indicating accessible and inaccessible locations
- PF-based design is easy to implement and efficient to calculate

A Step Towards "Indoor GPS"

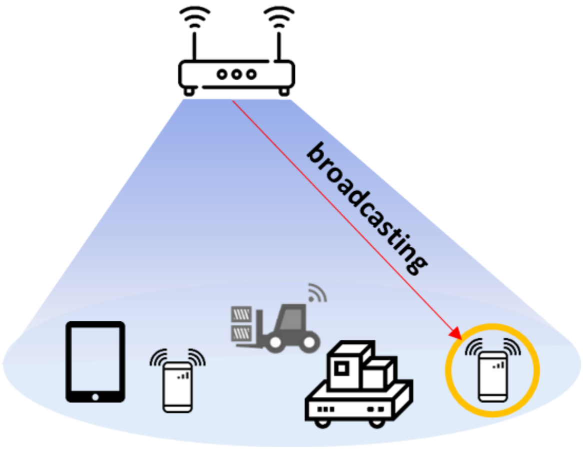

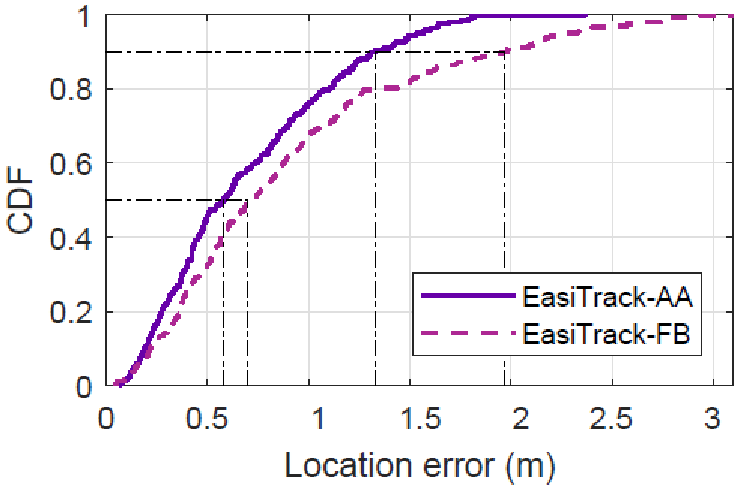

Accuracy

Decimeter-level (even in NLOS)

Robust

To environmental changes / people

Installation

A single unknown AP, easy to install

Coverage

One AP for 3,000 m², including NLOS

Scalability

Massive clients (like GPS) and buildings

Location: A long way to go…

Questions?

- Thank you!