Handling blockage and mobility

Grand Challenges for mmWave Networking

- Shorter wavelengths, higher attenuation
  - ~1000x higher attenuation than WiFi or LTE
- Use highly directional, electronically steerable phased-arrays to overcome propagation loss
  - Introduces new challenges: blockage, mobility

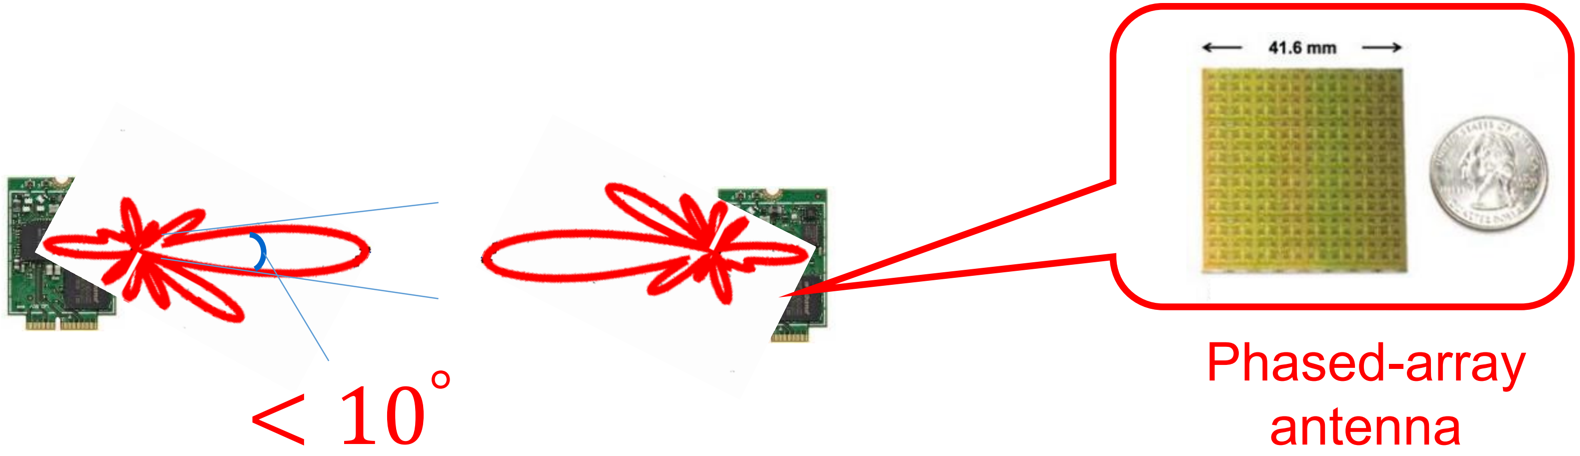

Phased-array antenna

Grand Challenges for mmWave Networking

- Mobility

Tx and Rx beams must keep aligned

- Blockage

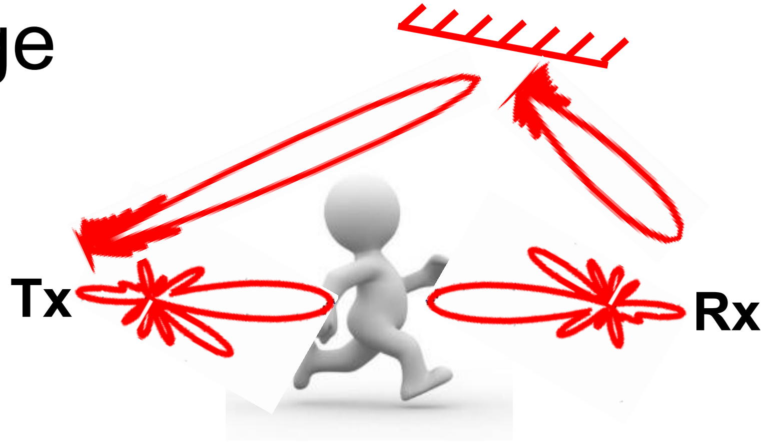

Needs environment reflection to overcome blockage

How severe is the blockage/mobility problem?

- Signal attenuation of a directional link
  - The body absorbs majority of the energy from a directional transmitter

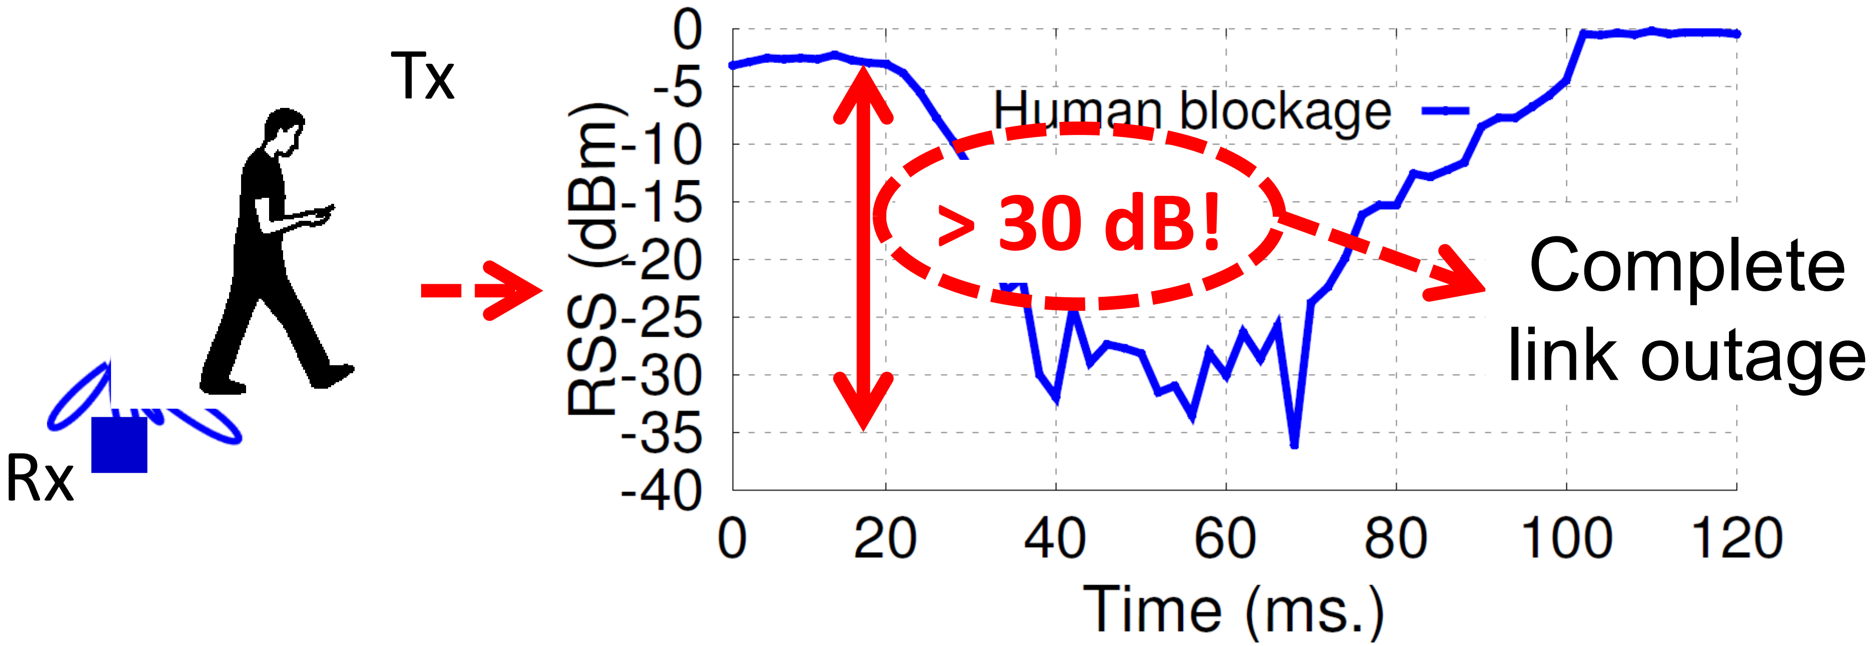

How severe is the blockage/mobility problem?

➢ Throughput drop due to signal attenuation and blockage

- Experimental setup

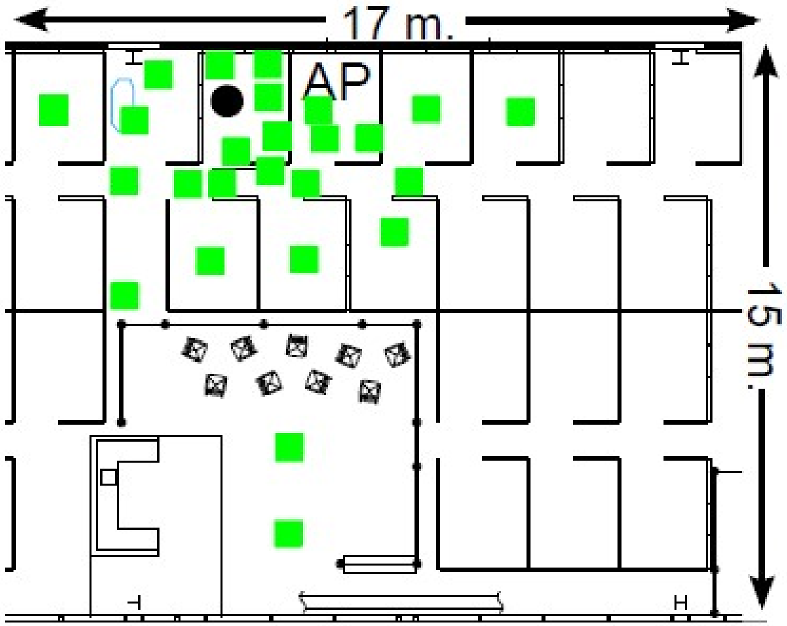

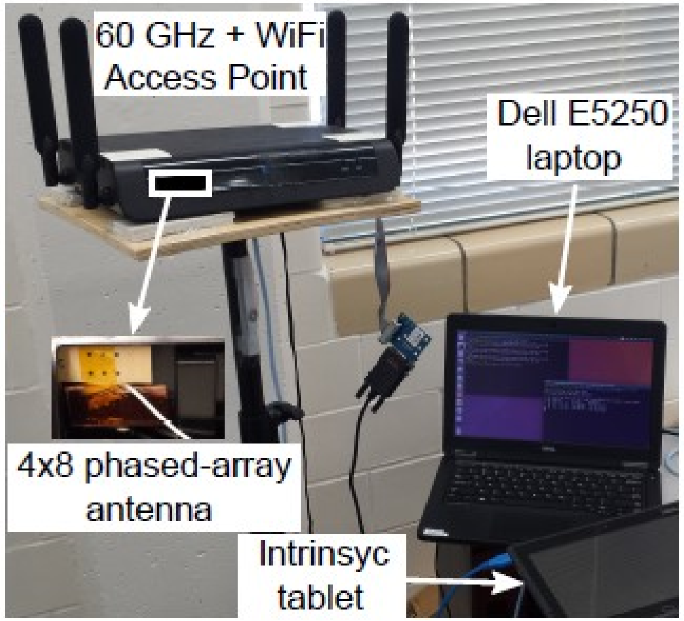

* S. Sur et. al., ACM MobiCom'17

How severe is the blockage/mobility problem?

- Throughput drop due to signal attenuation and blockage
  - Results

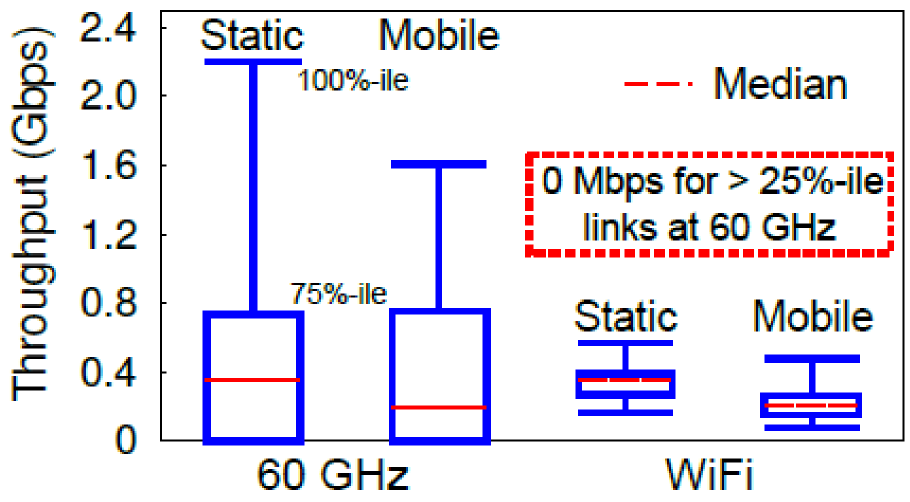

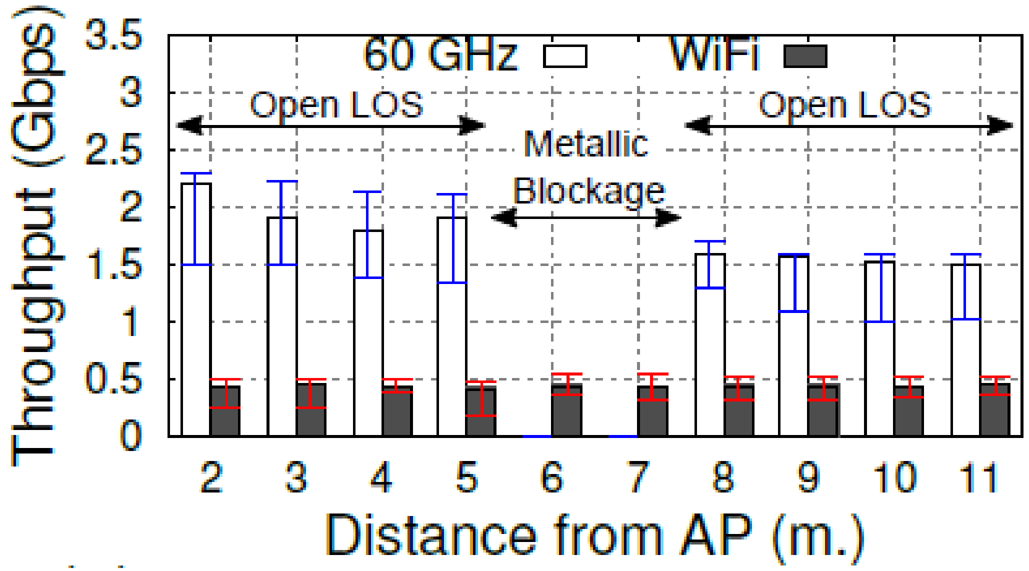

How severe is the blockage/mobility problem?

➢ Theoretical recovery time (from triggering to completion)

Phased-array size	1 client	4 clients
8	0.51 ms	1.27 ms
16	1.01 ms	2.53 ms
64	4.04 ms	304.04 ms
128	106.07 ms	706.07 ms
256	310.11 ms	1501.11 ms

* Hassanieh et. al., arXiv 1706.069335v1

➢ Non-trivial protocol level operations and decision making

- Beam searching overhead grows with the number of beams
- When to trigger the beam searching? (Tradeoff: overhead vs. responsiveness)
- There is no guarantee that beam searching can result in a usable pair of TX-RX beams

How severe is the blockage/mobility problem?

- Measurement of recovery time
  - Qualcomm 802.11ad radios, 32 element phased-array, 128 beams
  - Measure time to converge to best beam after blockage
  - CDF over 50 trials
  - Link outage effect is amplified at higher layer (TCP results later)

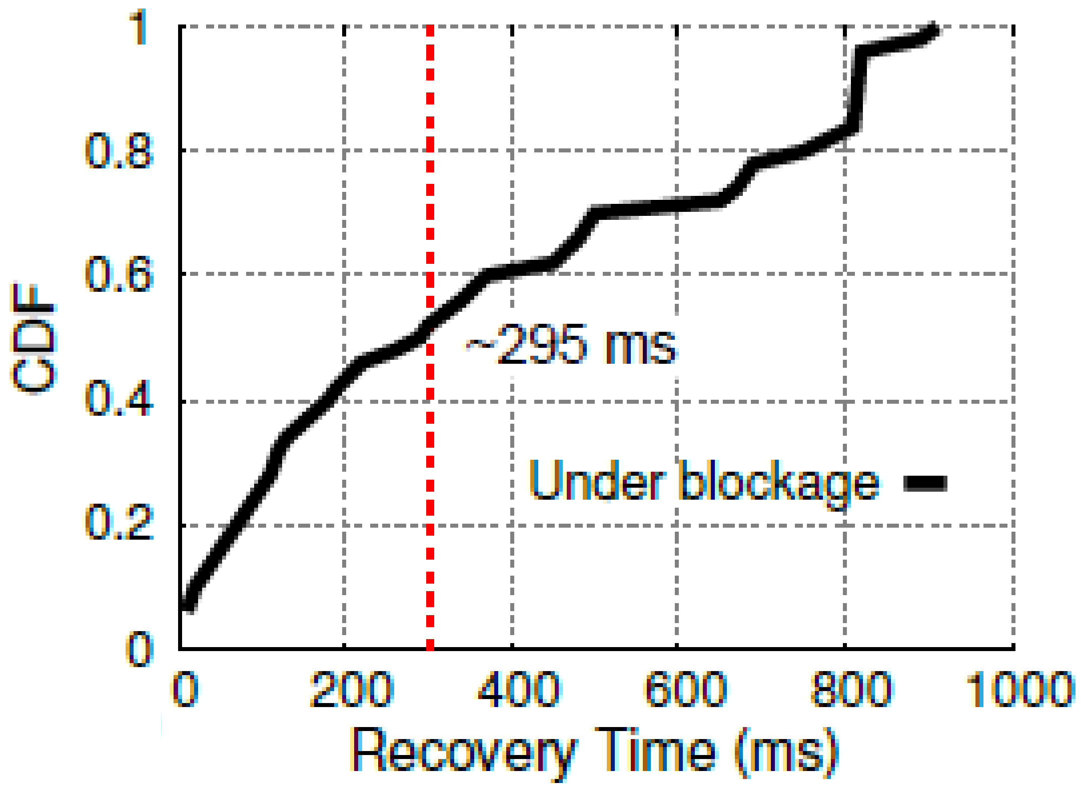

* S. Sur et. al., ACM MobiCom'17

Design principles to handle mobility/blockage

- Fast beam realignment protocols
  - Predictive and proactive beam switching
    Example: BeamSpy (S. Sur et al., NSDI'17)
  - Sensor assisted beam searching
    Example: Pia (T. Wei et al., MobiCom'17)

- New network architectures
  - Multi-node coordination
    Example: Pia (T. Wei et al., MobiCom'17)
  - Multi-band cooperation
    Example: MUST (S. Sur et al., MobiCom'17)

BeamSpy: predictive link recovery under blockage

- Working conditions
  - Quasi-stationary TX and RX

- Working principles
  - Measure the channel of current TX/RX beams
  - Predict the channel of other beams, without beam scanning overhead!

* "*BeamSpy: Enabling Robust 60 GHz Links Under Blockage*",
Sanjib Sur, Xinyu Zhang, Parameswaran Ramanathan, Ranveer Chandra, USENIX NSDI'16

Key insights: correlation between beams

- Blockage in a beam *drops performance* of other beams

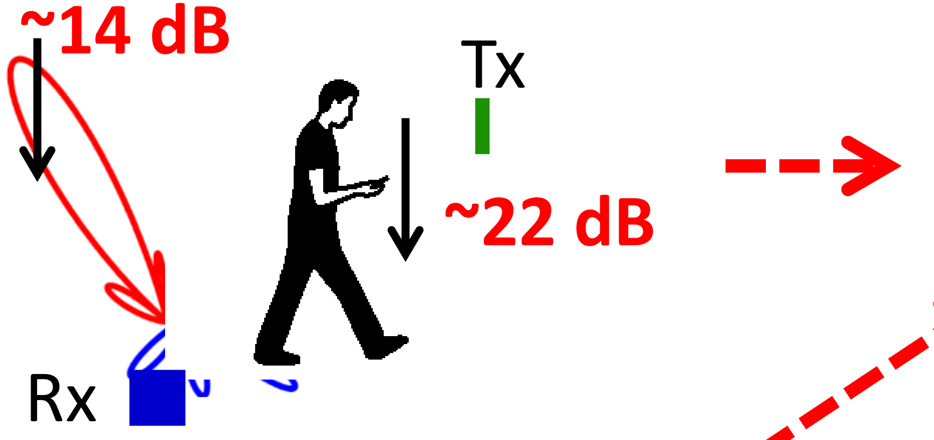

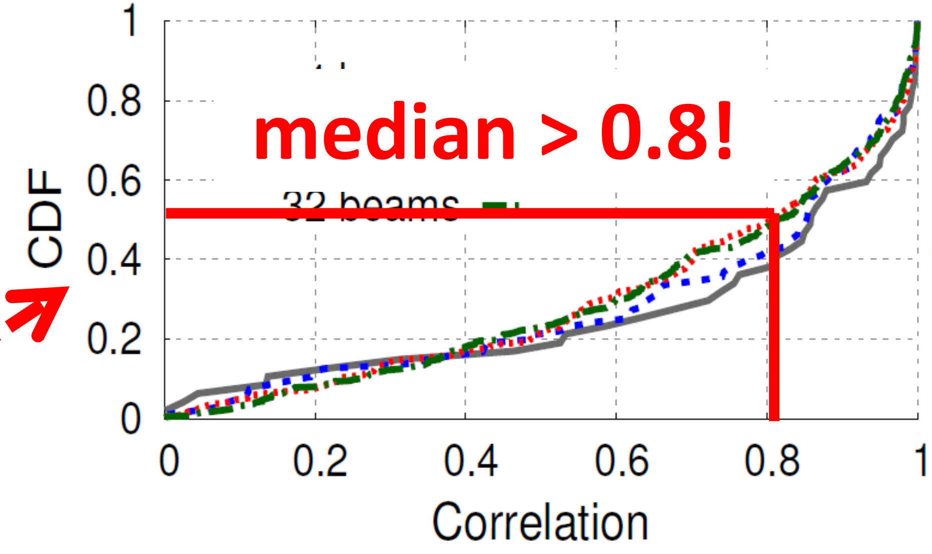

RSS drop correlation of other beams w.r.t. strongest beam

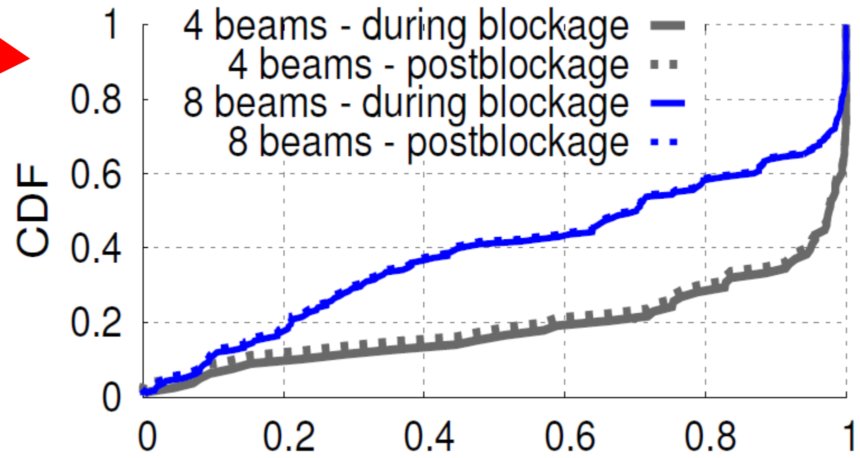

- Correlation remain *unchanged* irrespective of blockage!

Key insights: correlation between beams

- Why does correlation exist?

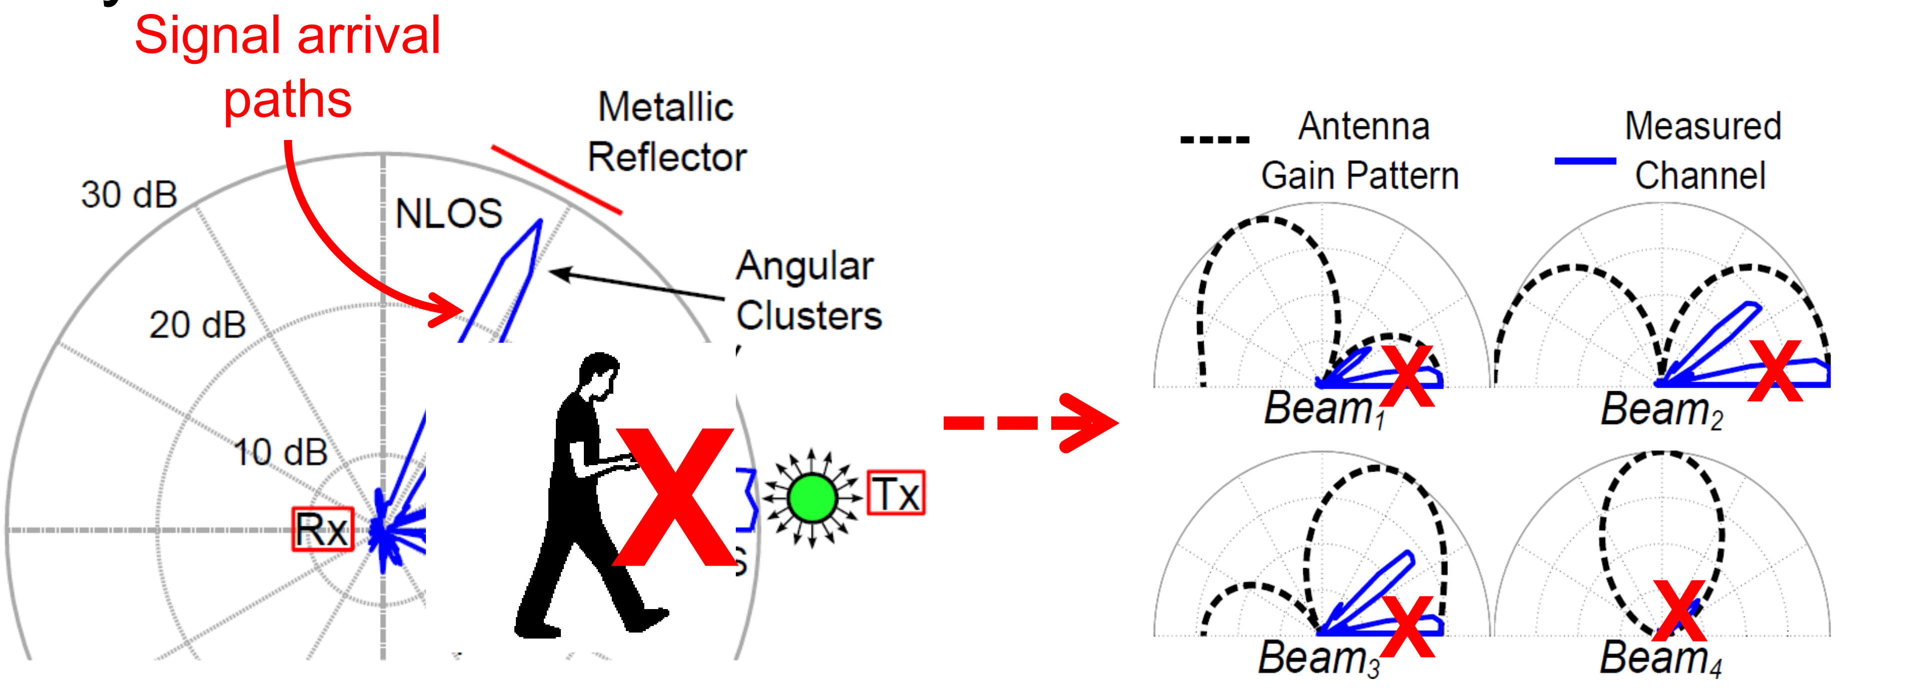

Sparse signal arrival paths are *shared* between beams, thus blockage causes correlated RSS drop in all beams!

Modeling the correlation through a sparse channel model

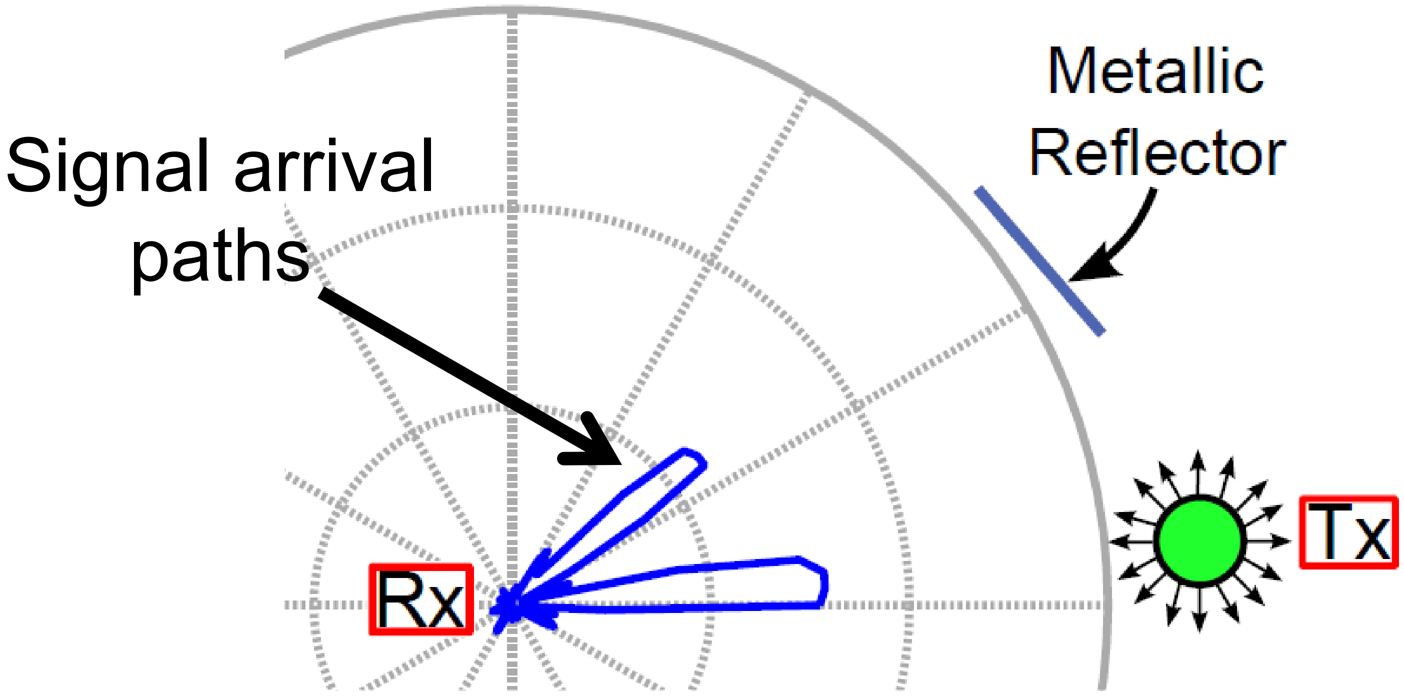

Modeling the correlation through a sparse channel model

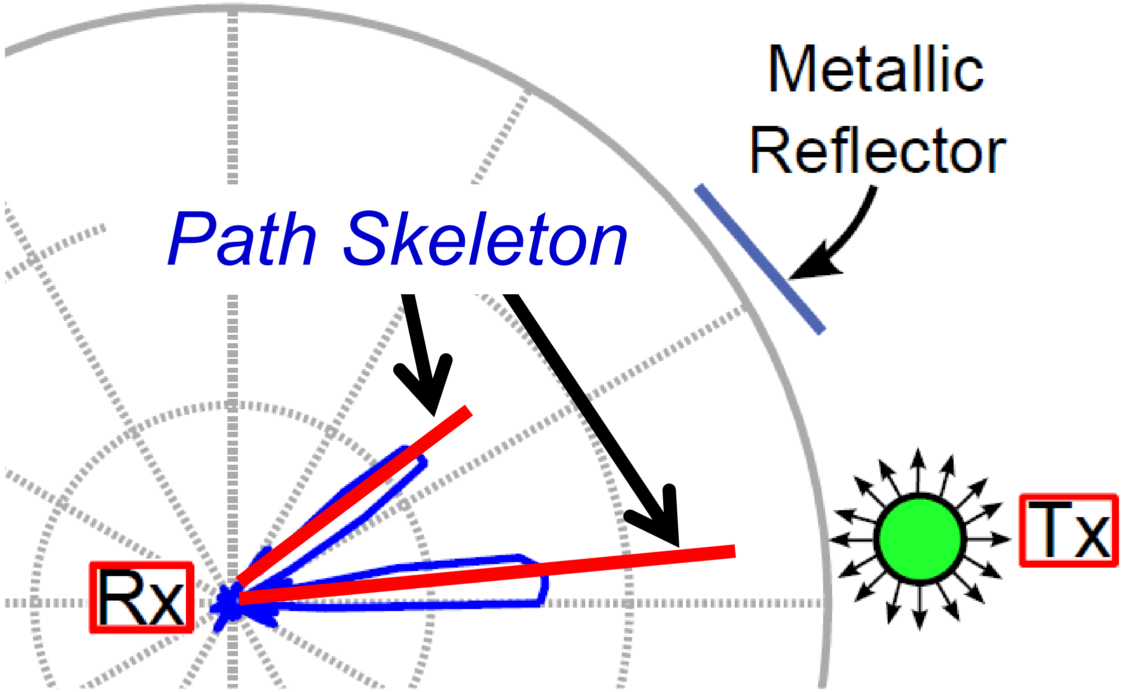

Modeling the correlation through a sparse channel model

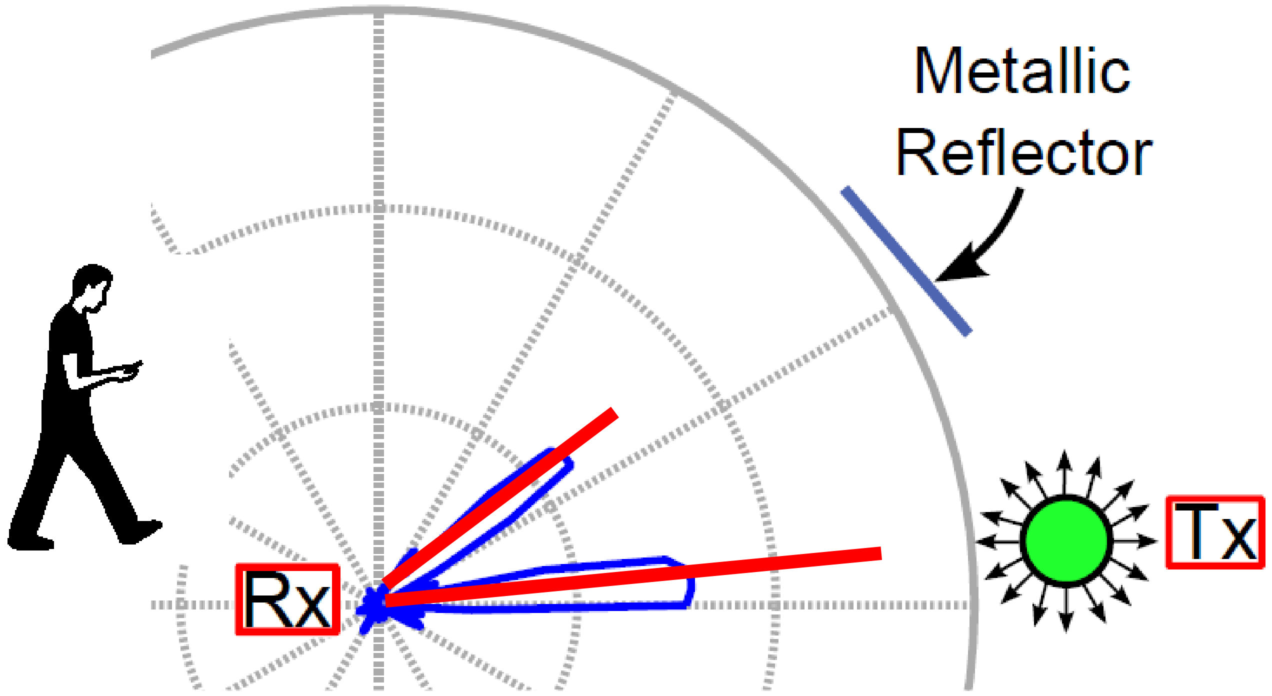

Modeling the correlation through a sparse channel model

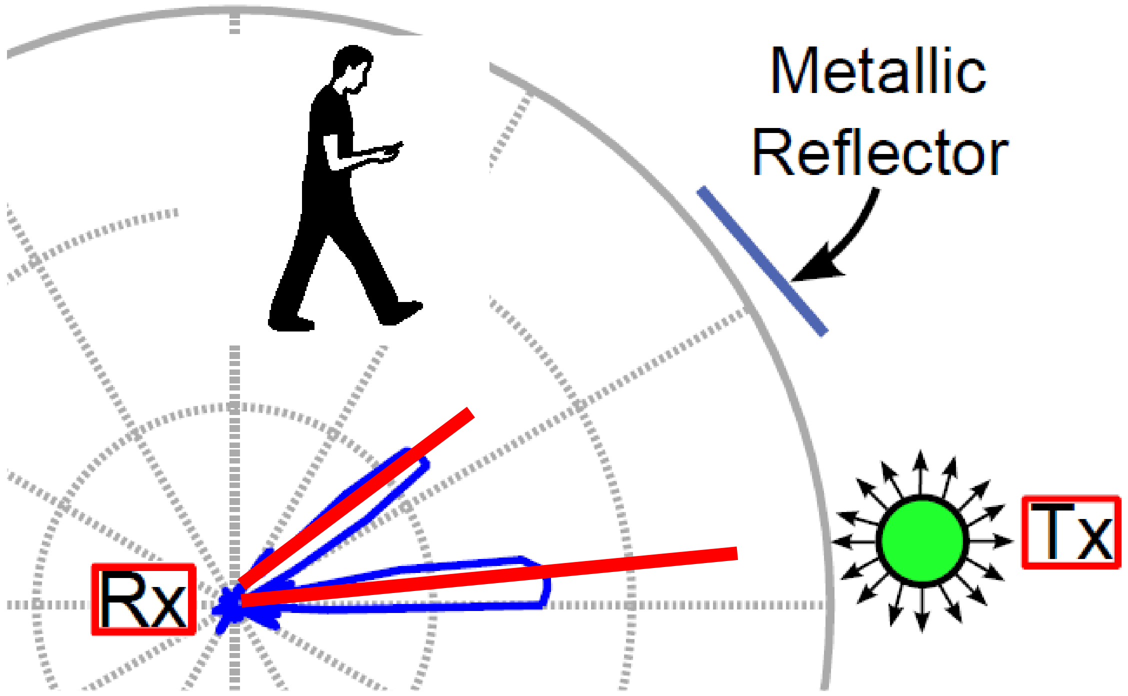

Modeling the correlation through a sparse channel model

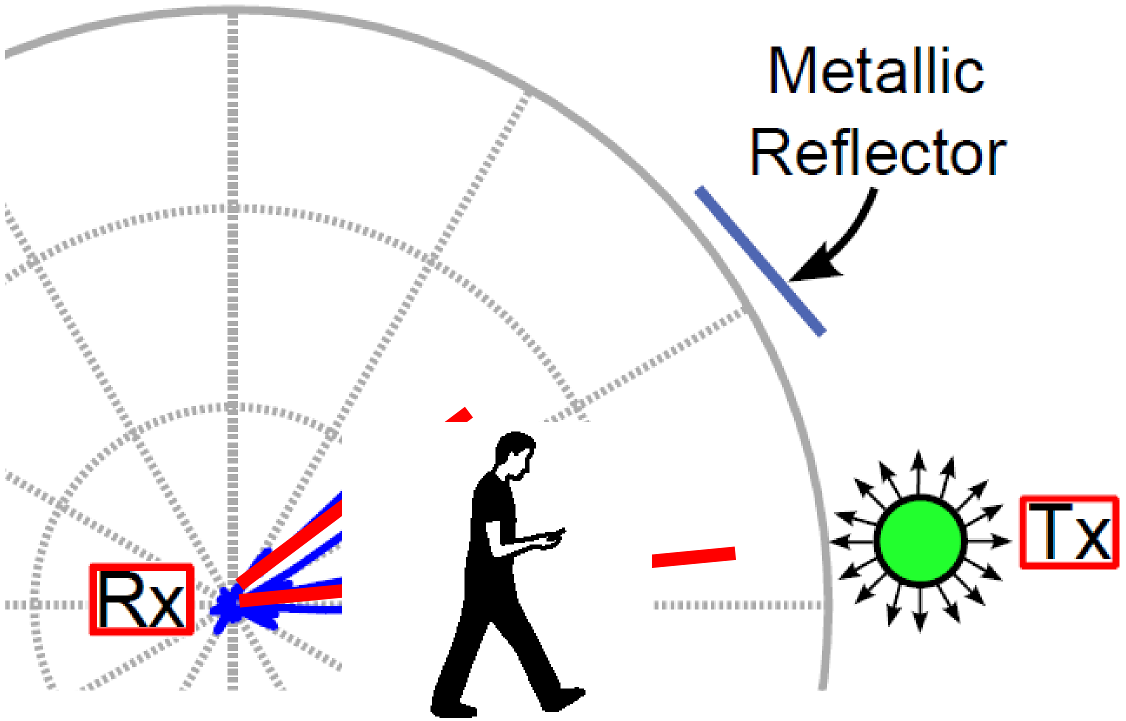

BeamSpy workflow

At deployment time

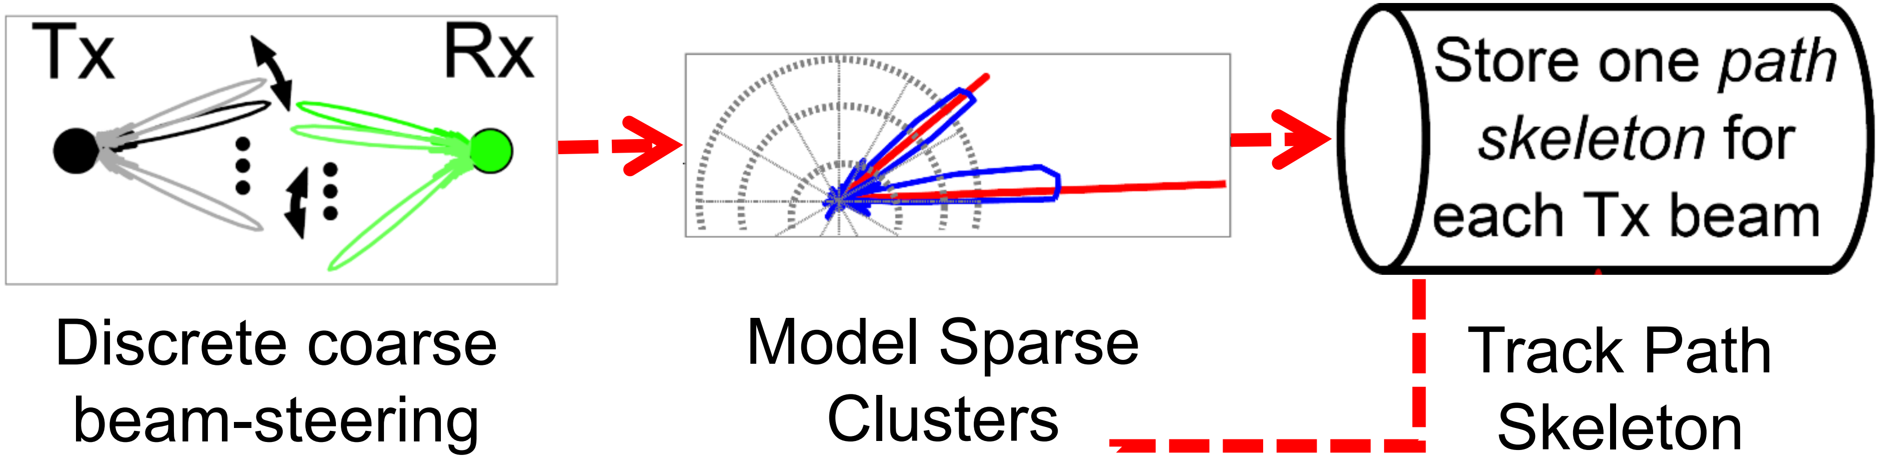

Discrete coarse beam-steering

Model Sparse Clusters

Track Path Skeleton

At run time

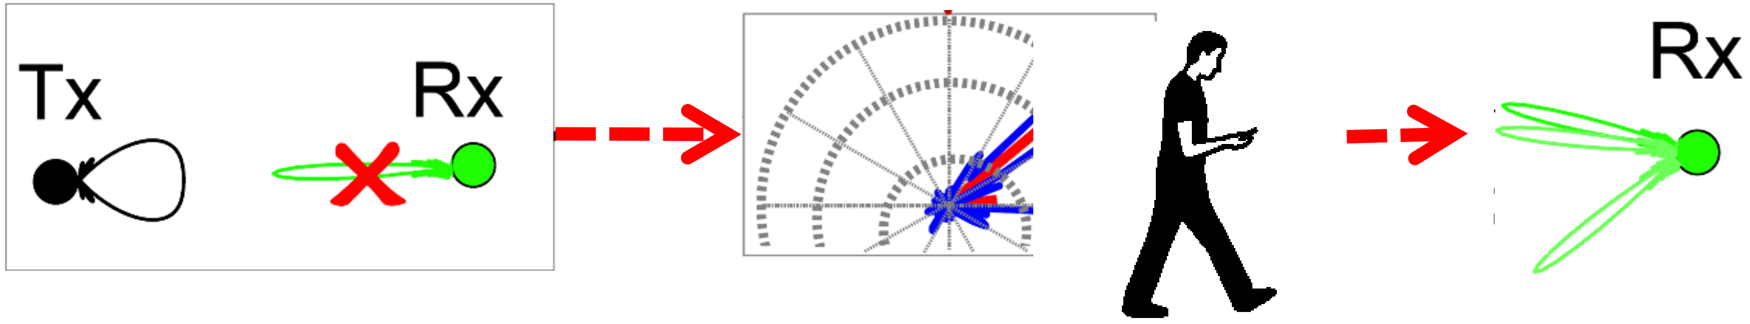

Beam is blocked

Identify state of Path Skeleton

Predict RSS of other beams from new state

Modeling the correlation through a sparse channel model

- How does the prediction work

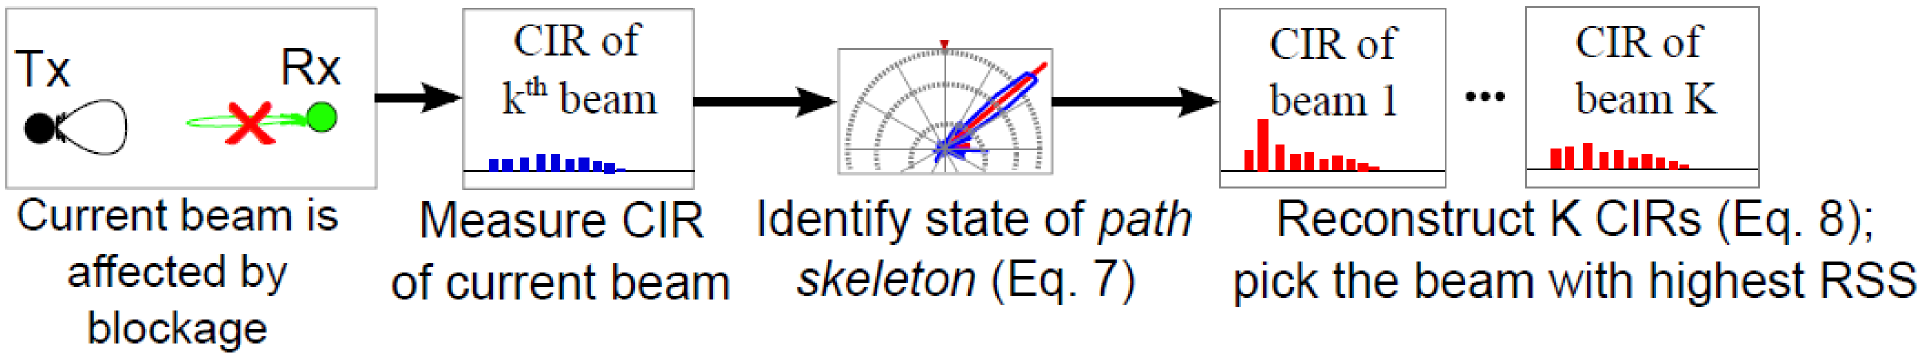

- Look at *current beams condition* under blockage → *Identify the state* of sparse cluster → Virtually *reconstruct* performance of rest of the beams and pick the best one.

BeamSpy performance

- Accuracy of best beam direction prediction under blockage

Close to 70% even with 32 beams!

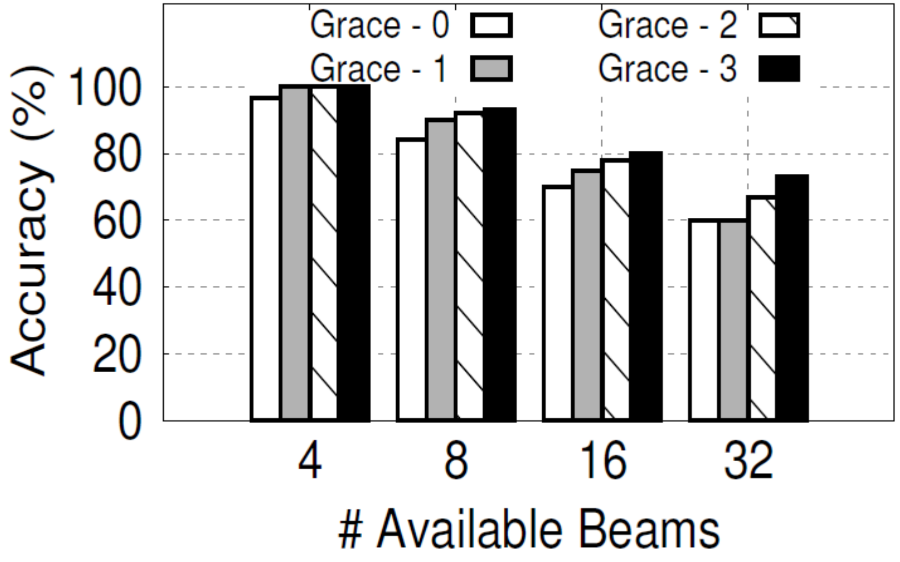

- Predicting RSS of the best beam under blockage

Prediction error (90%-ile) is within ±3 dB for 32 beam

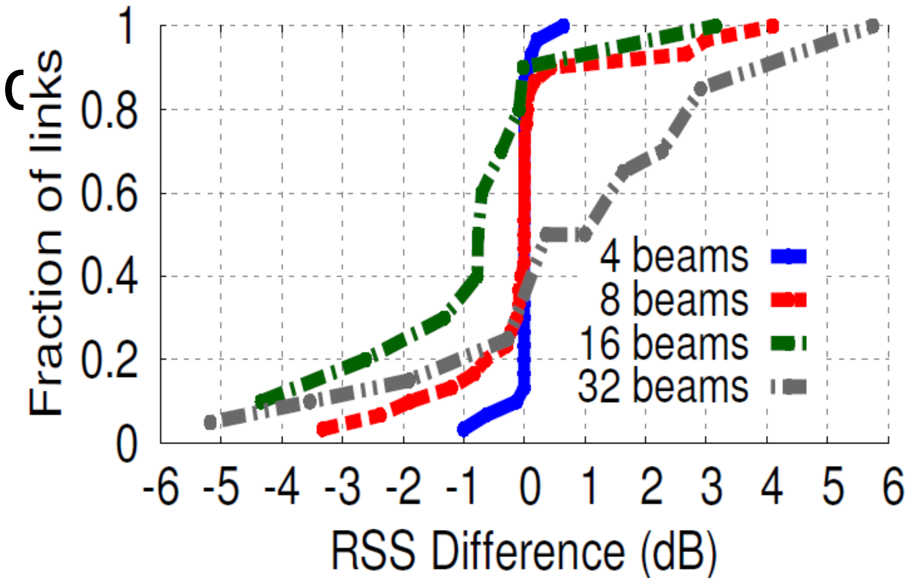

BeamSpy performance

➢ Link performance gain under blockage

Throughput performance close to oracle.

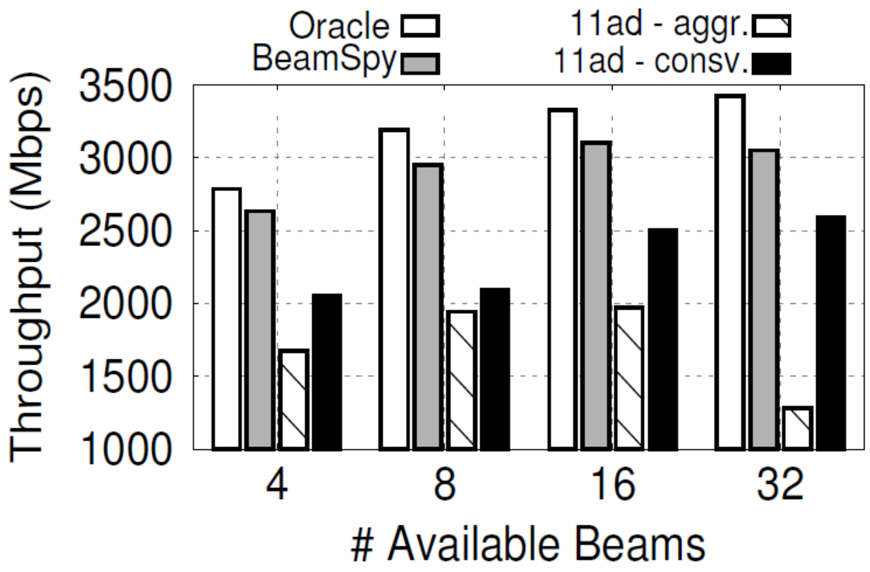

Towards seamless coverage and mobility support

- BeamSpy works for quasi-stationary TX/RX
- Can we make mmWave networks as mobile and ubiquitous as WiFi?
- Non-trivial! Even for room-level mobility/coverage
  - Limited TX/RX coverage due to directionality and lack of multipath
  - Blockage, mobility, and even minor orientation change can cause beam misalignment

Pia: Pose information assisted 60 GHz networks

- Design principles
  - Cooperation between APs to ensure coverage
  - Leverage mobile client's pose information (x,y,z coordinate and elevation/azimuth angle) to select the best AP
  - Leverage pose information to select the best beams to maximize spatial reuse

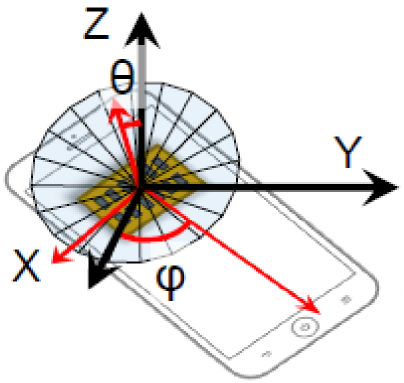

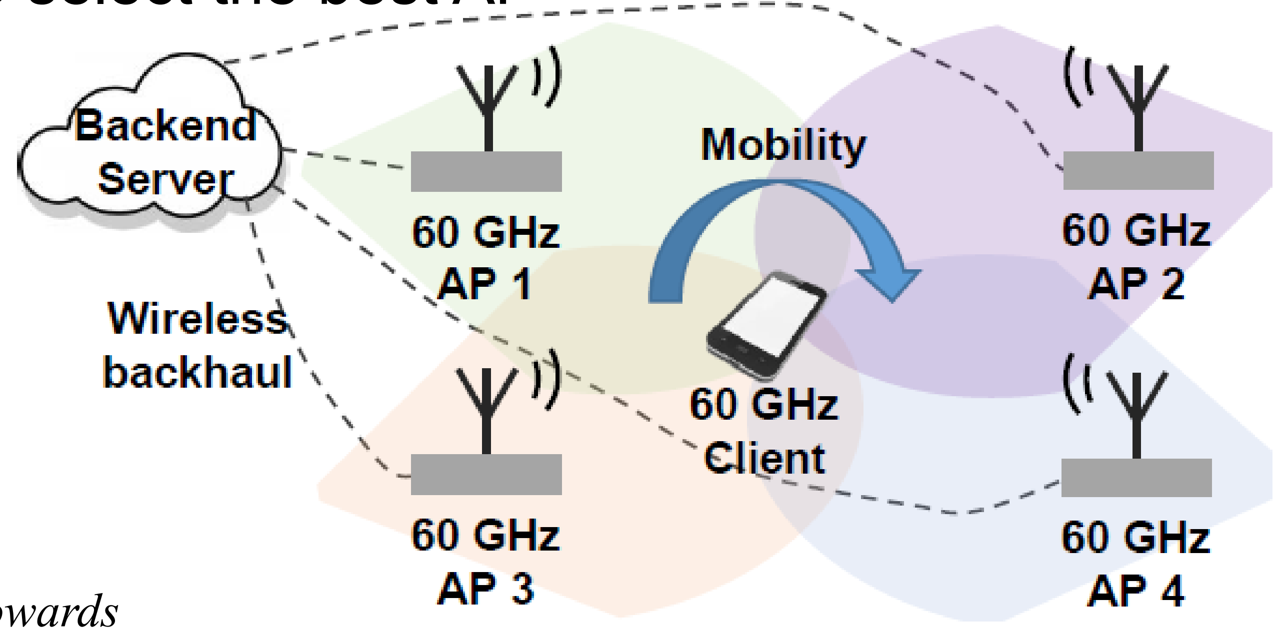

* "*Pose Information Assisted 60 GHz Networks: Towards Seamless Mobility and Coverage*", Teng Wei, Xinyu Zhang, ACM MobiCom'17

How does pose change affect link performance

- Vary relative angle between TX and RX
  - Throughput almost constant with an 160 degree field-of-view (FoV)
  - Throughput drops dramatically when out of FoV

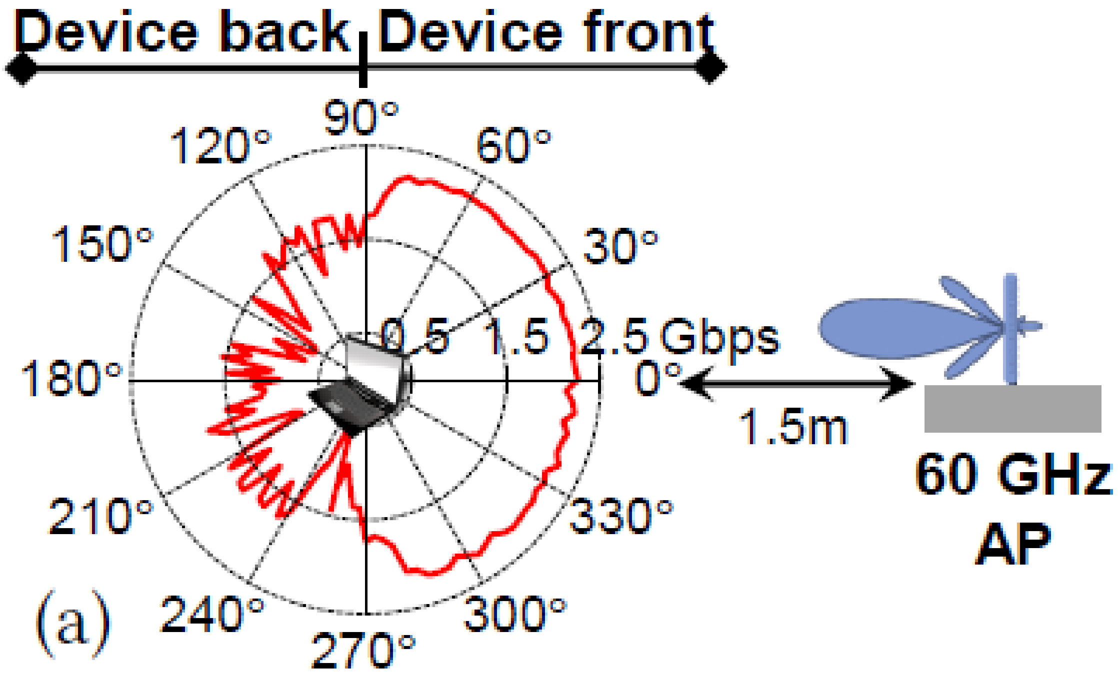

How does pose change affect link performance

- Vary relative angle between TX and RX
  - For room level coverage, in/out of FoV matters more than distance

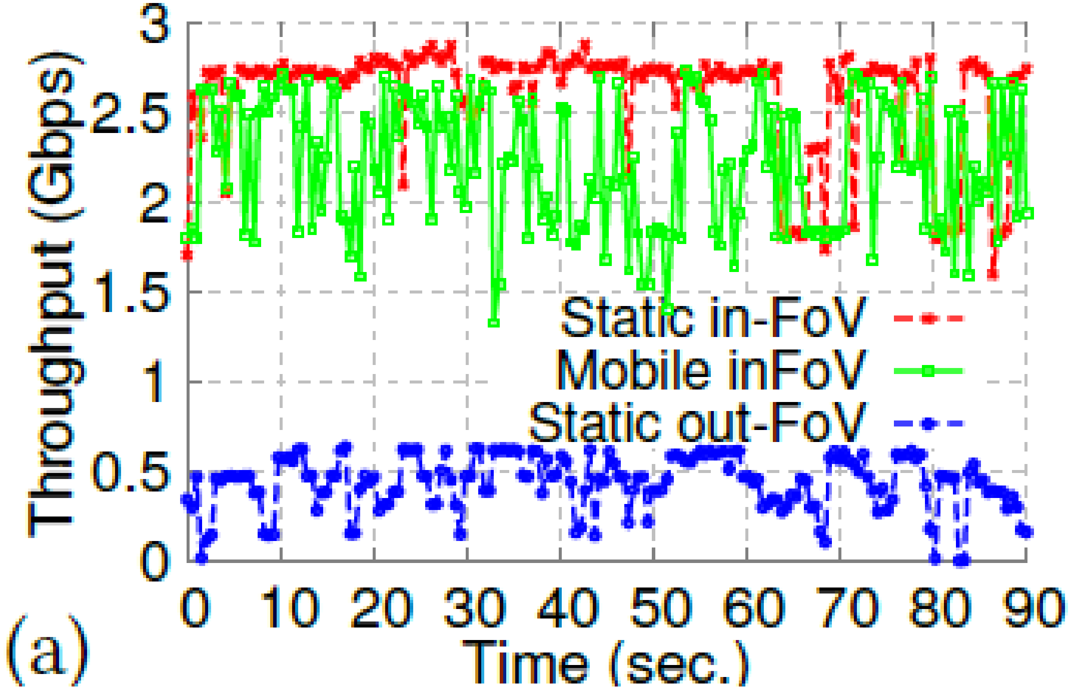

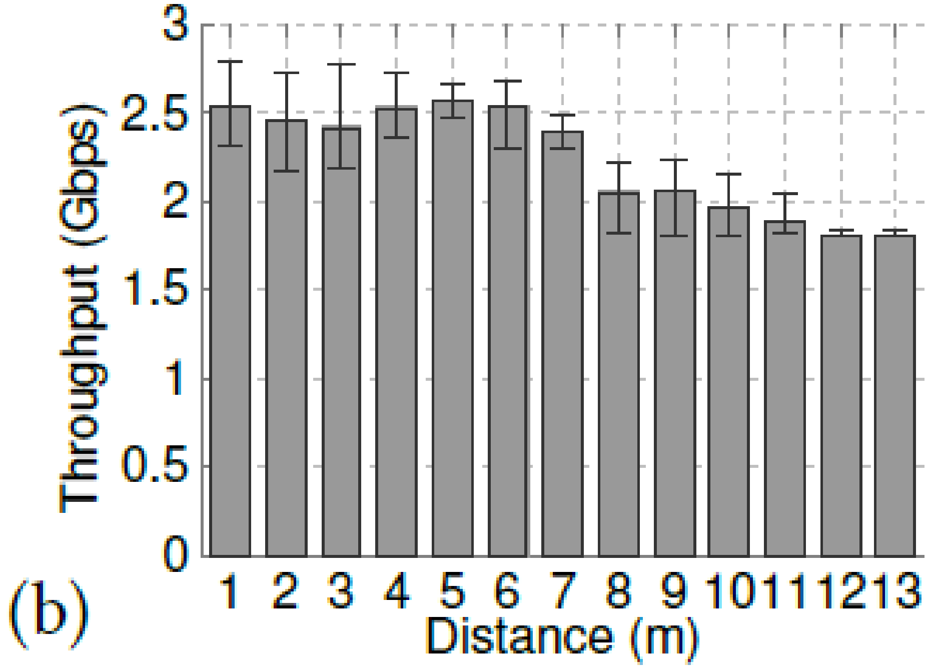

Pia work flow

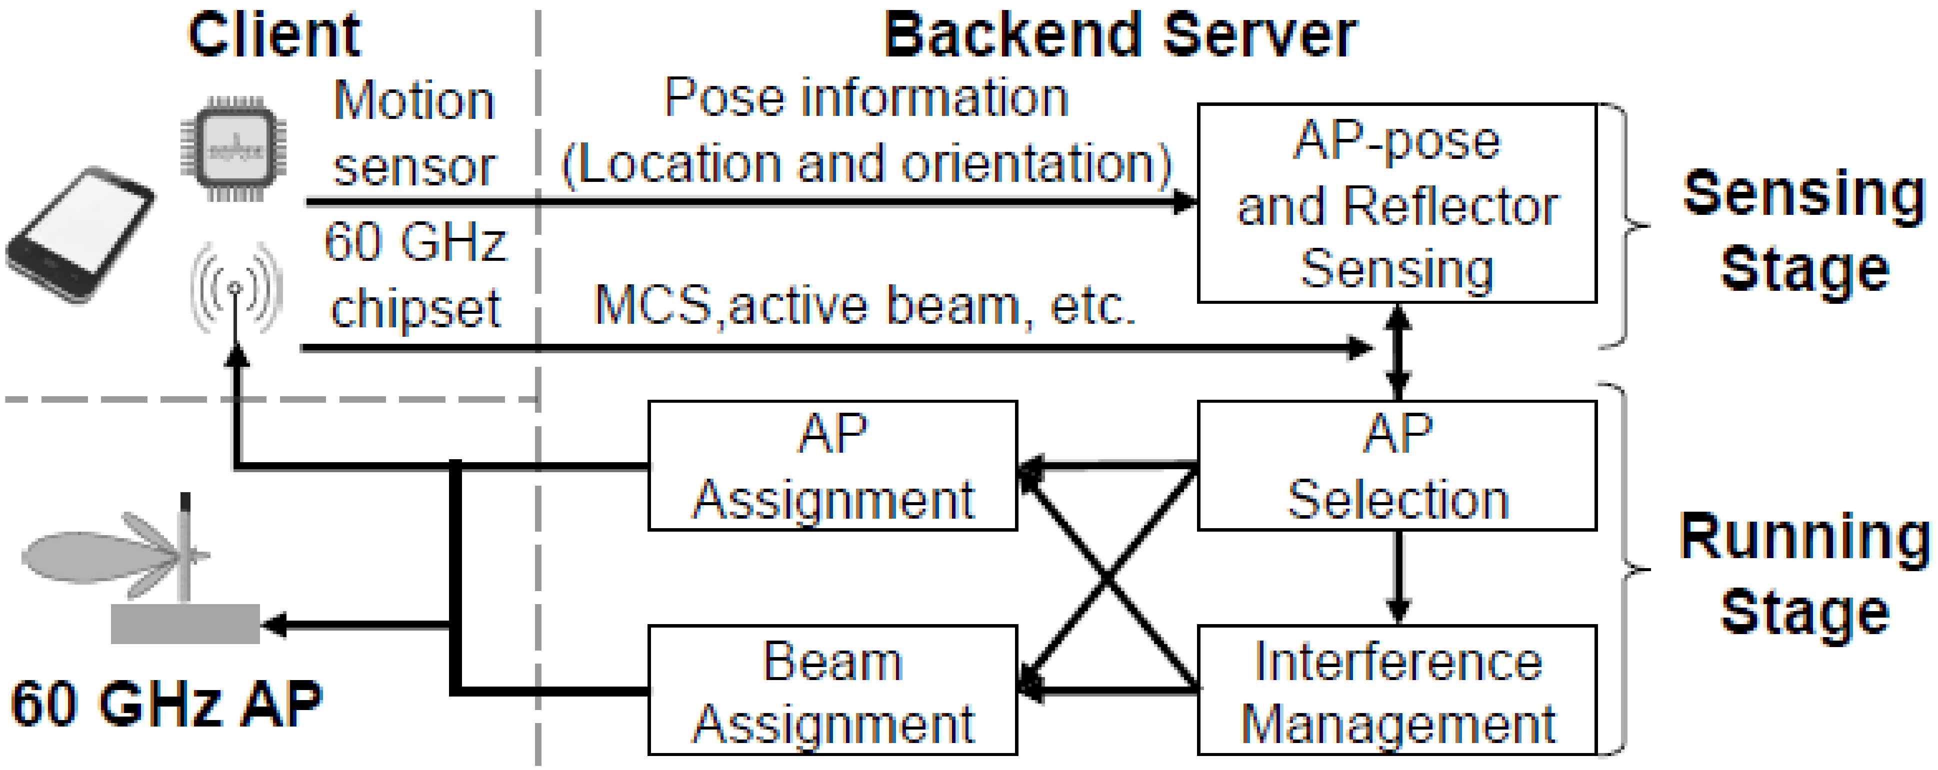

Pia: AP selection

- Proactive AP switching instead of reacting to link outage
  - Predict pose: simple kinematic model

$$\hat{P}_c(t+1) = P_c(t) + \Delta\hat{P}_c(t),$$
$$\Delta\hat{P}_c(t+1) = \Delta P_c(t),$$

  - Predict in/out of FoV based on relative pose between client and AP
  - Switching before outage

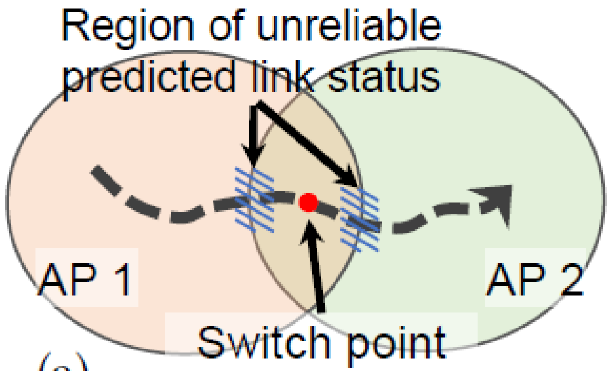

(a)

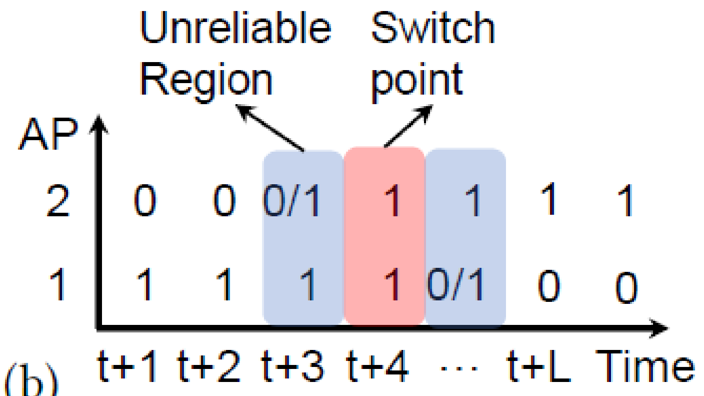

(b)

Pia: AP selection

➢ How does a client know the APs‘ pose?

- One-time initial training, to obtain APs’ global pose info
- Statistical estimation

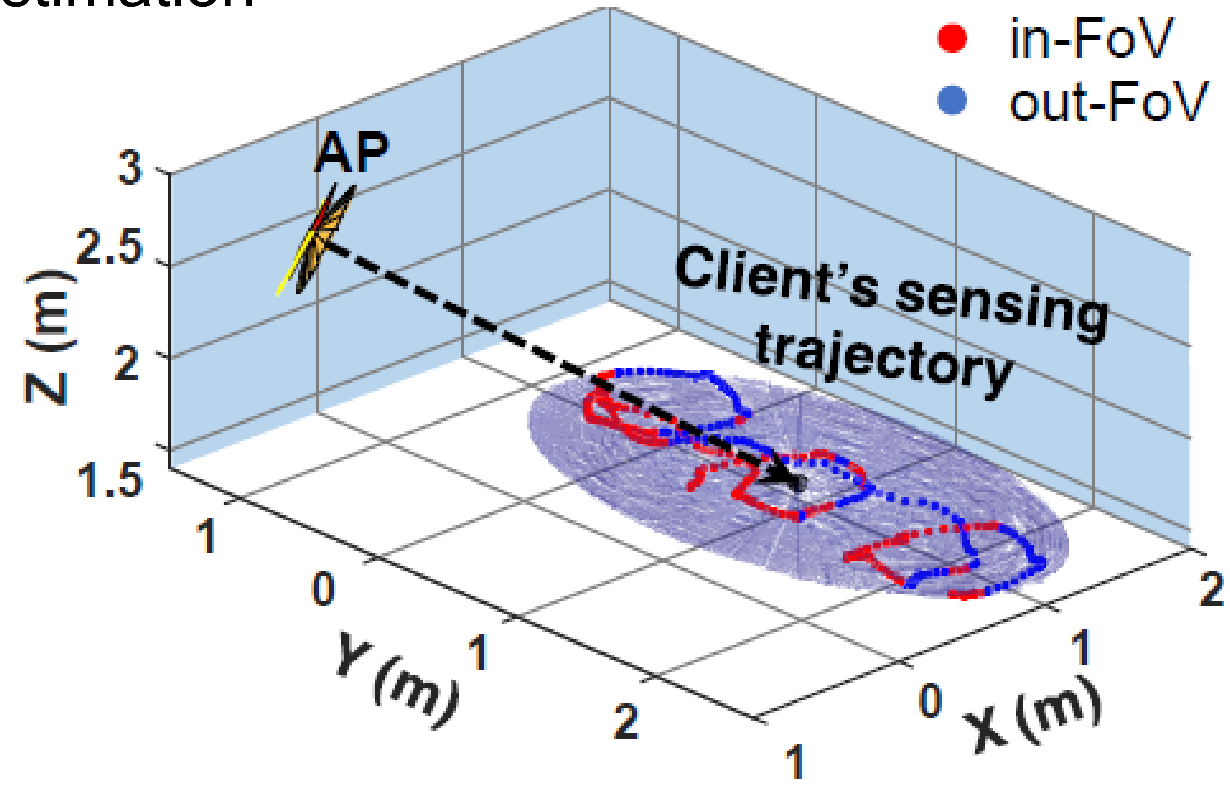

Pia: beam selection for spatial sharing

- Non-trivial due to imperfect directionality of phased-arrays
  - Strongest beam is not necessarily the throughput-optimal one

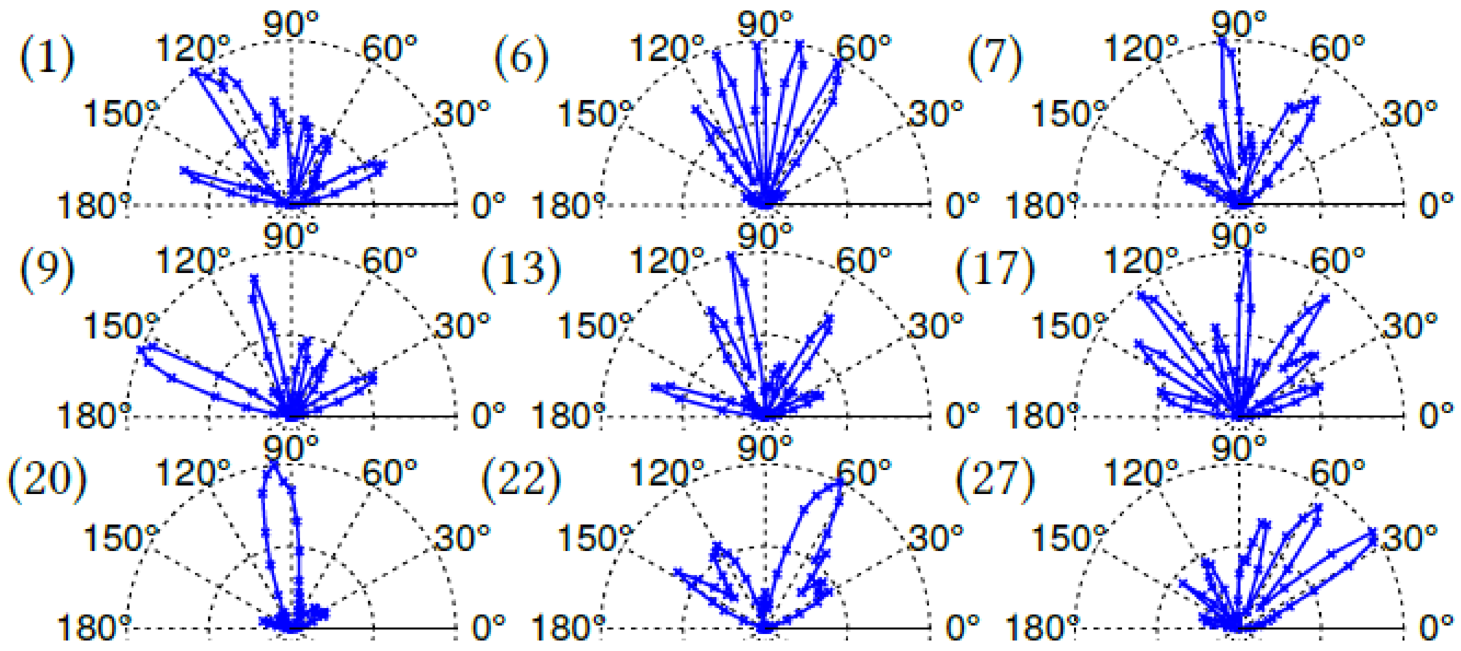

Measured beam patterns from a commercial 802.11ad device.

Pia: beam selection for spatial sharing

- Joint beam and AP selection problem.
  - Beam strength map (BSM) as a basic data structure
  - Objective: maximize SIR

$$\max_{A,B} \frac{1}{N_c} \sum_{i=1}^{N_c} \frac{BSM[A(i), i, B(i)]}{\sum_{j=1}^{N_s} INF_{max}(j, i)},$$

A(i): AP assignment for client i;
B(i): beam assignment for client i;
INFmax(j,i): max interference from AP i to client j

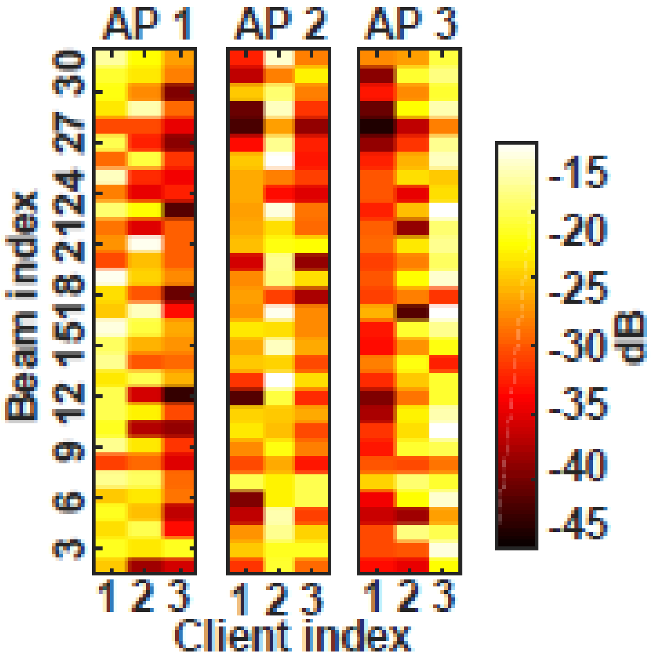

Example BSM between 3 APs and 3 clients.

  - Computational cost too high. Approximate using signal to leakage ratio (SIR).

Pia: testbed verification

➢ Experimental setup

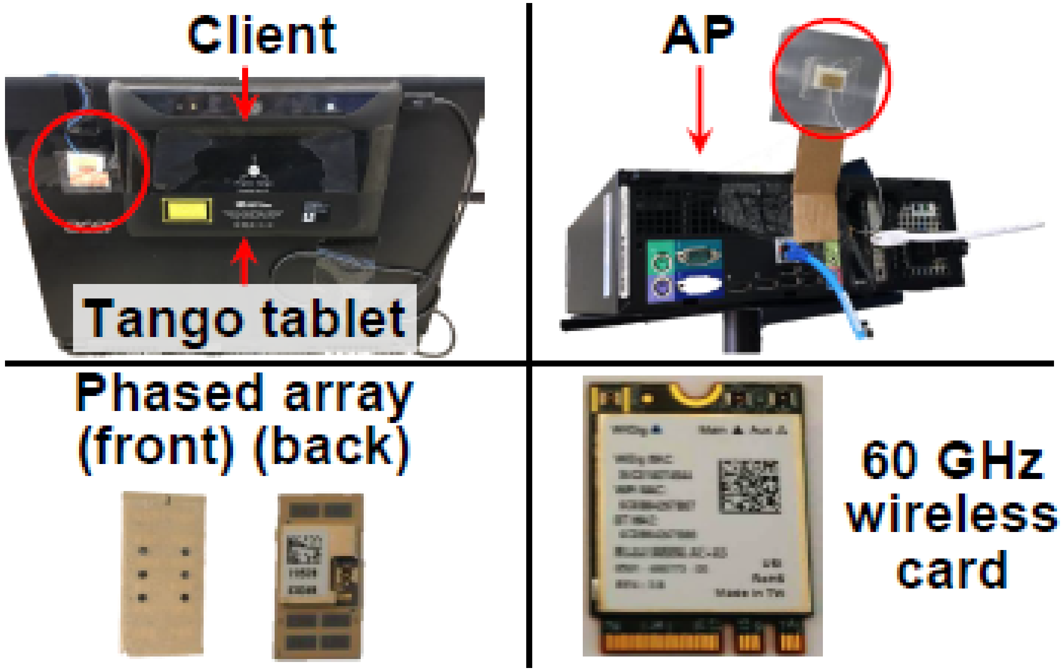

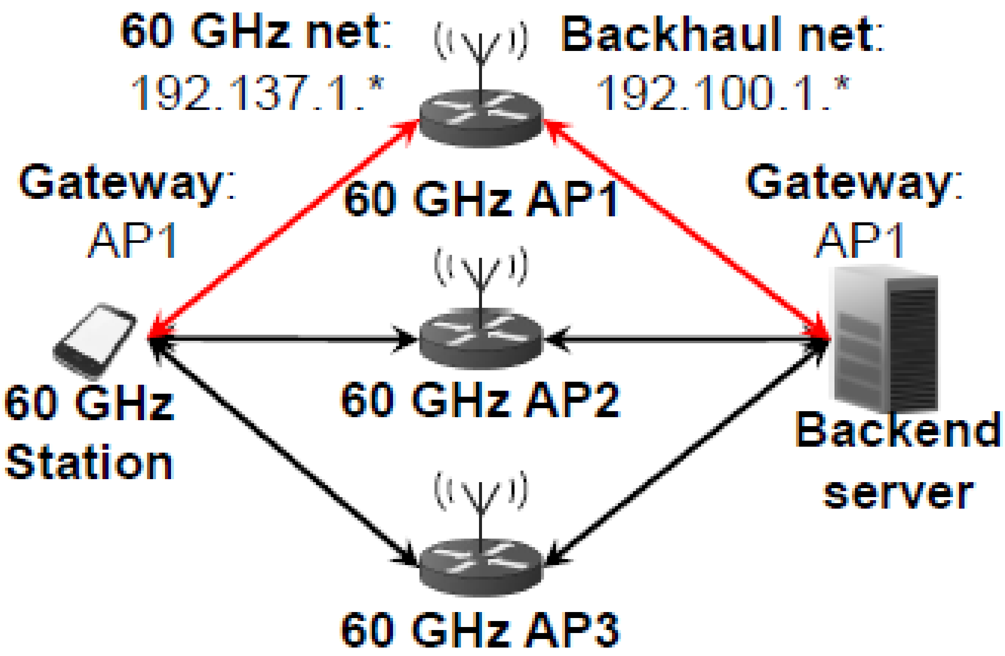

Pia: performance overview

- Link stability

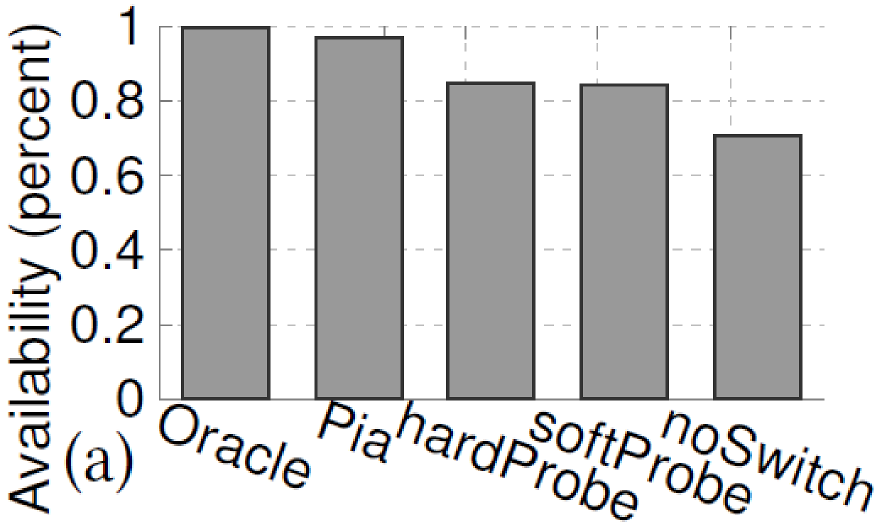

(a) Link availability: percentage of time that throughput exceeds a threshold (1.8 Gbps).

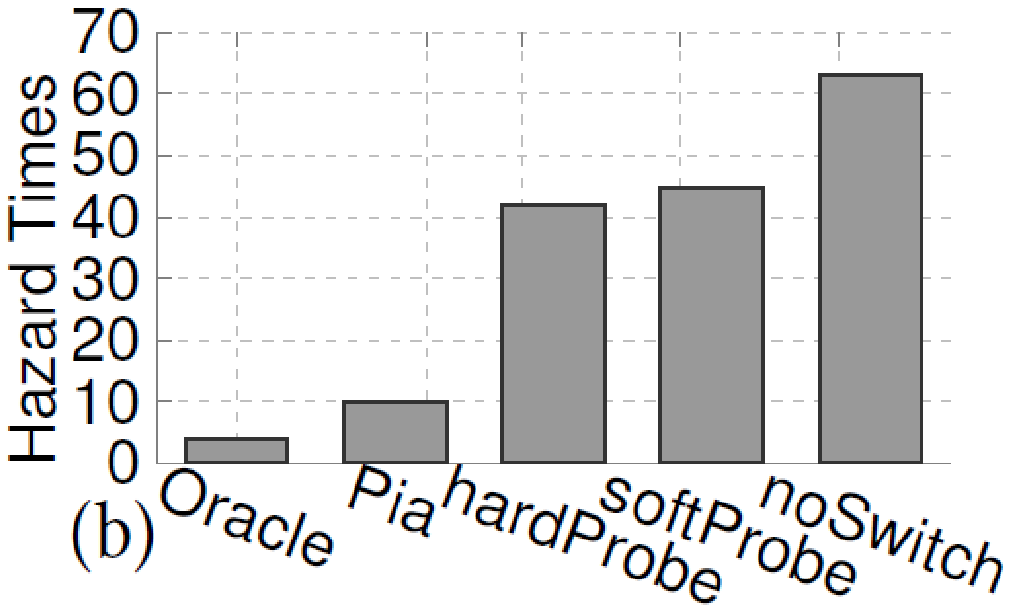

(b) Hazard times: number of occurrences that link throughput drops below the threshold in a 5-minute test.

Pia: performance overview

- Spatial sharing

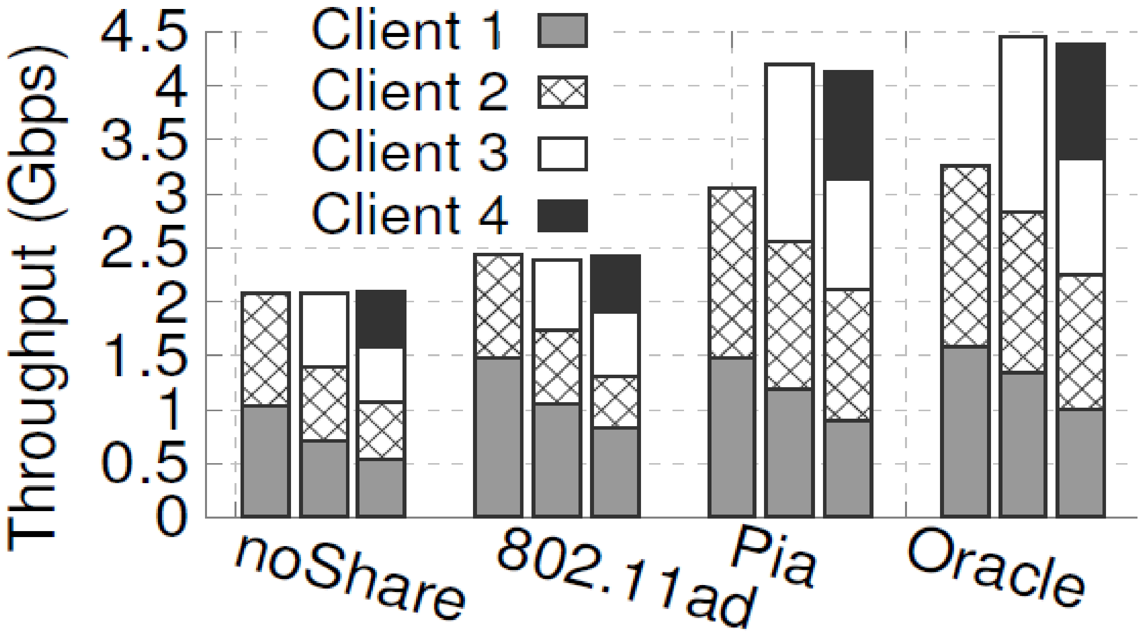

- Why is 802.11ad interference mapping ineffective?
  - Lack optimal mechanism to schedule concurrent transmissions
  - Large overhead esp. in mobile scenarios

Pia: performance overview

- Resilience of AP selection under pose errors

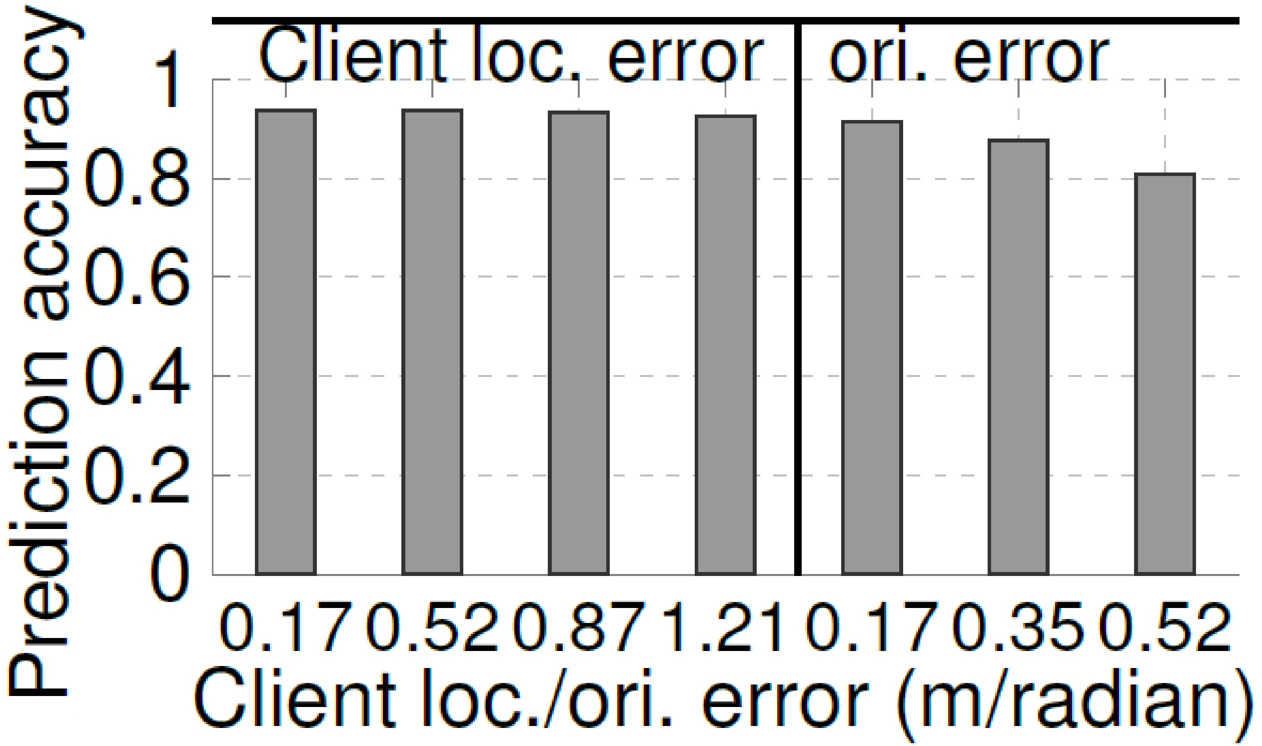

- Only need meter level location precision, and 10+ degrees of orientation precision

MUST: WiFi assisted 60 GHz networks

- Design principle: WiFi as a backup to make 60 GHz network stable
  - Leveraging commodity tri-band 802.11ac/ad radois
  - Predict 60 GHz channel (under mobility) using WiFi CSI
  - Under high risk of low-RSS, proactively switch to WiFi
- Why use WiFi CSI to estimate 60 GHz channel?
  - Much less likely to be blocked
  - MIMO array, instead of phased-array, can estimate channel profile instantaneously (instead of trying all beam directions)

* "*WiFi-Assisted 60 GHz Networks*",
Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, Kyu-Han Kim, ACM MobiCom'17

MUST: alternative design choices

- Why not turn on both 60 GHz and 5 GHz radios?
  - Performance is even worse due to TCP artifacts

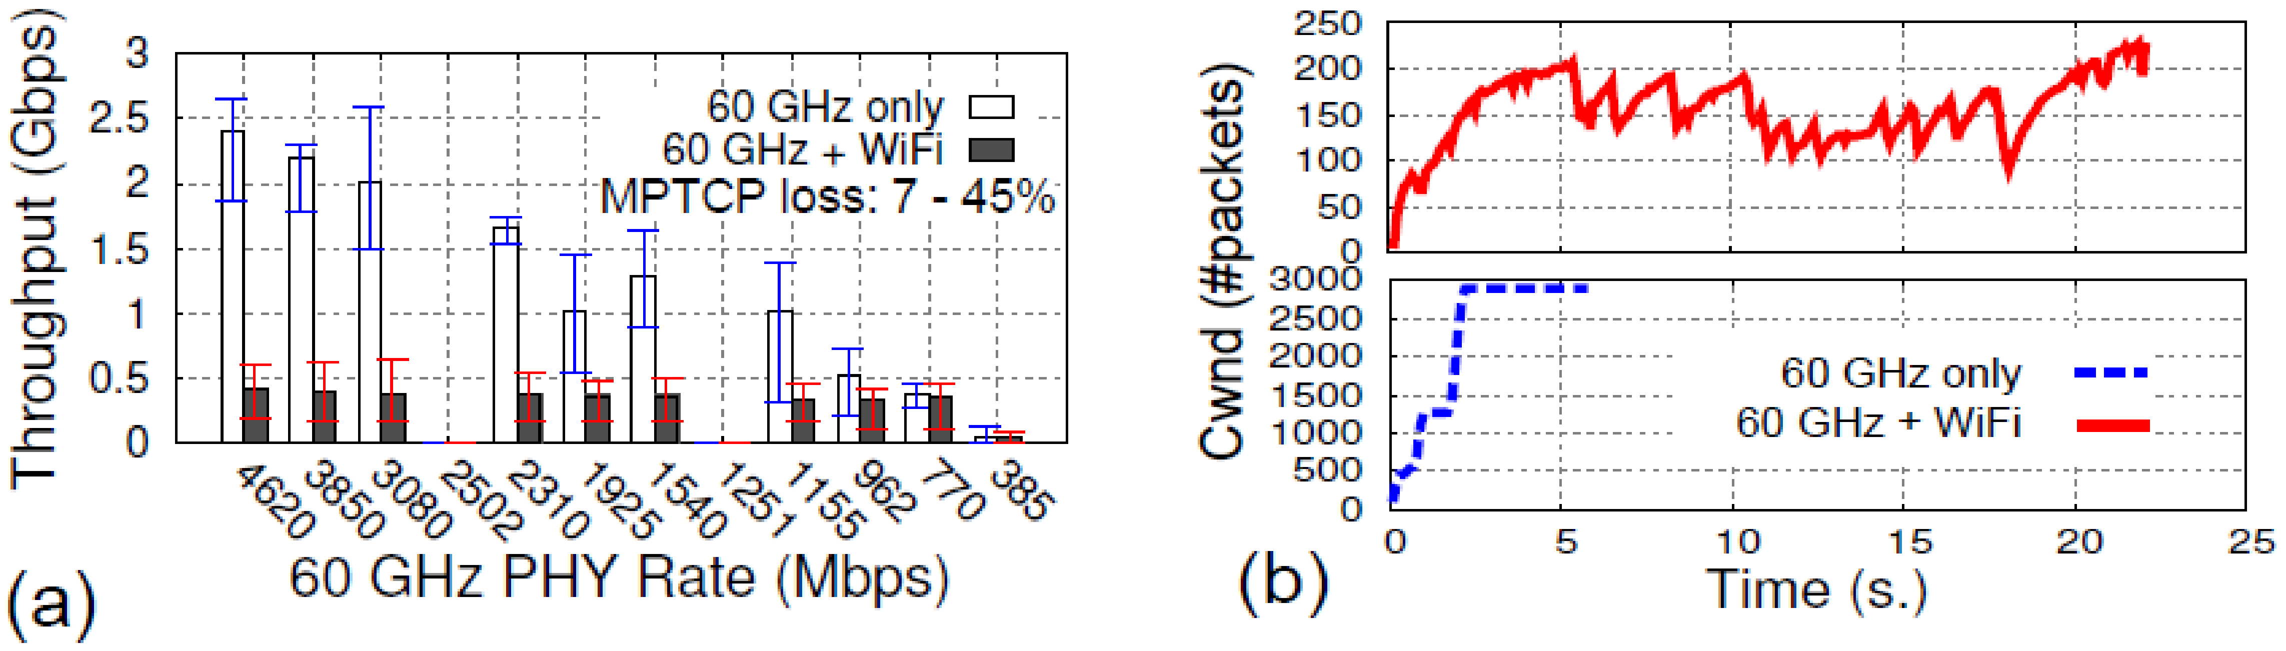

TCP throughput performance.

TCP congestion window size.

MUST: alternative design choices

- Why not react (switch to WiFi) after link outage occurs?
  - Switching latency is long, and amplified at TCP level
  - Non-trivial to determine “when” to switch; non-trivial protocol overhead

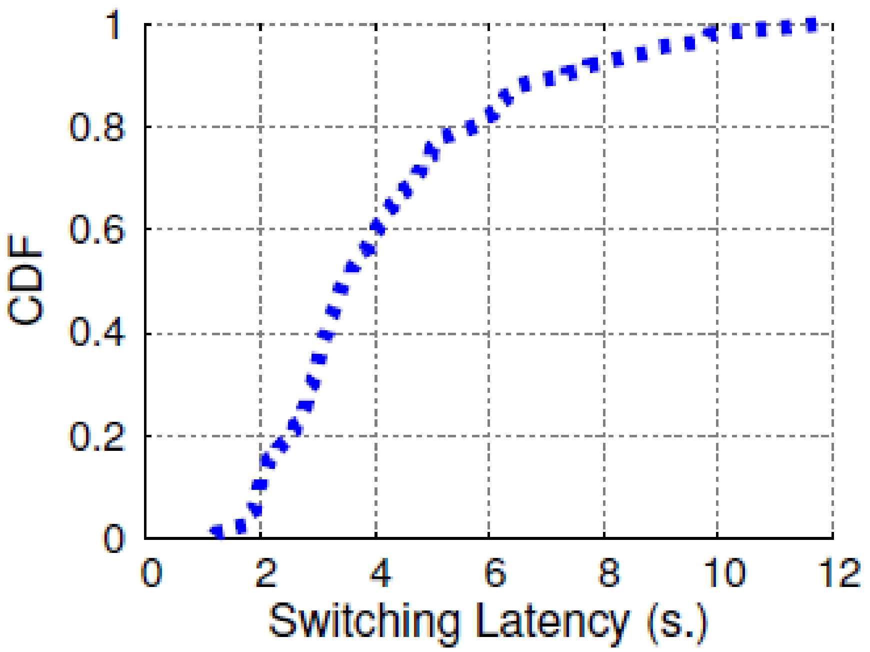

(a) CDF of switching latency on a commodity 802.11ad device.

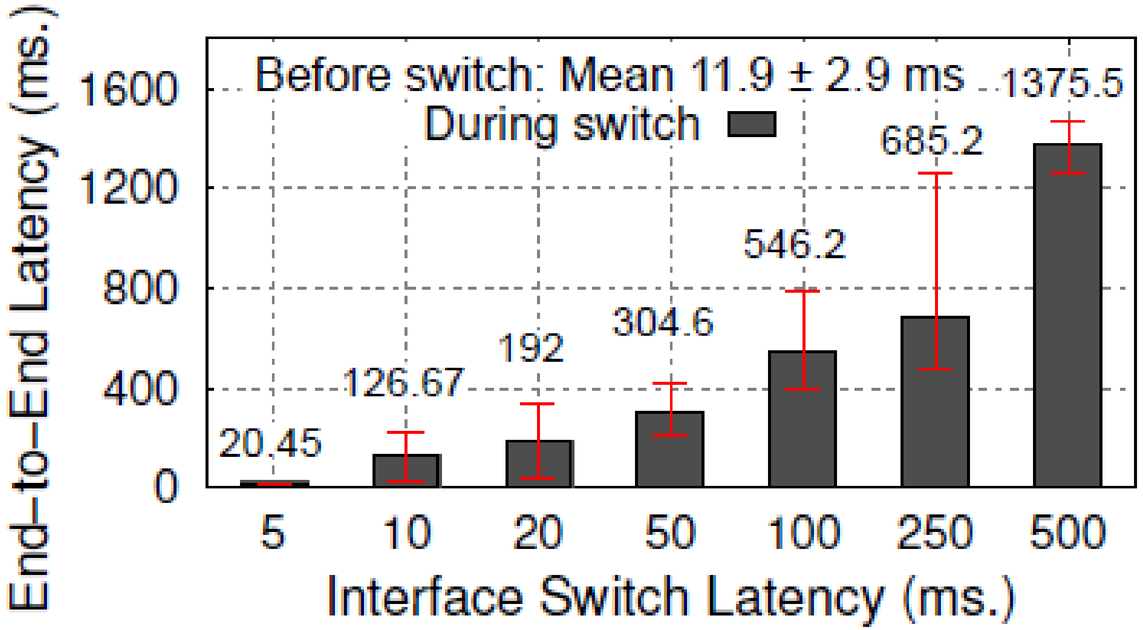

(b) Latency amplified at higher layer.

MUST: predicting 60 GHz channel using WiFi CSI

- MUST work flow

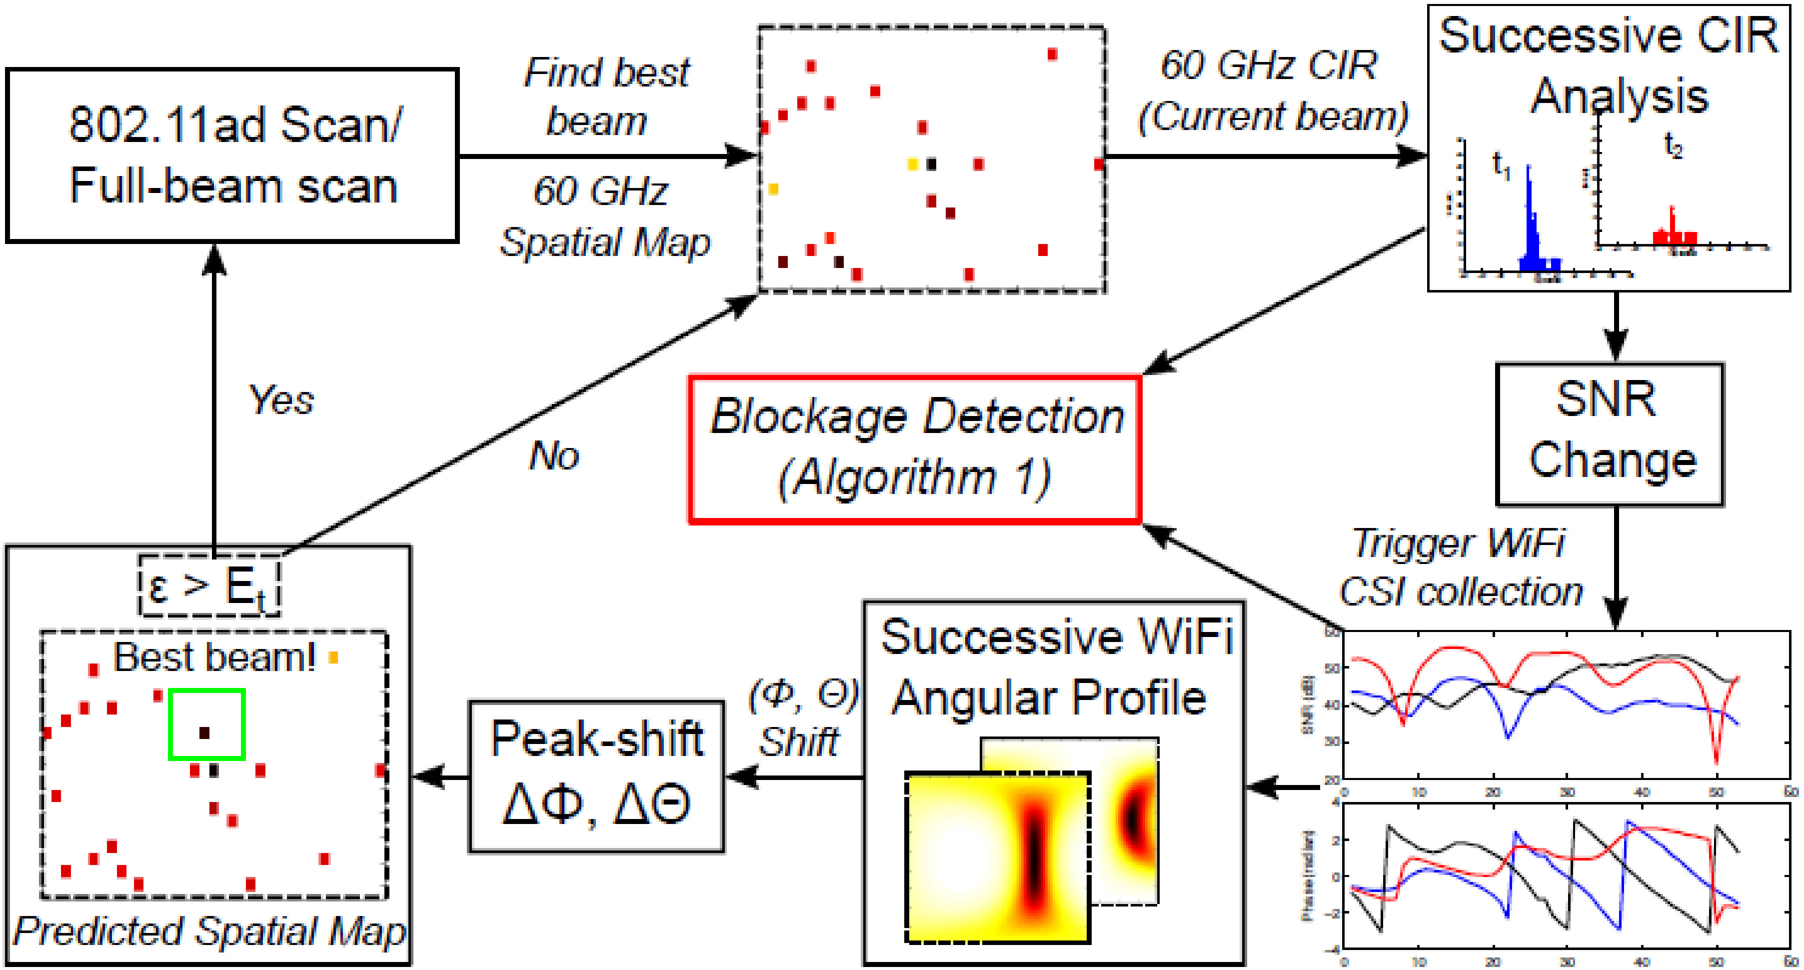

MUST: predicting 60 GHz channel using WiFi CSI

- Identify the angular shift of the 60 GHz dominating path from the successive time-domain spatial snapshots of the WiFi channel

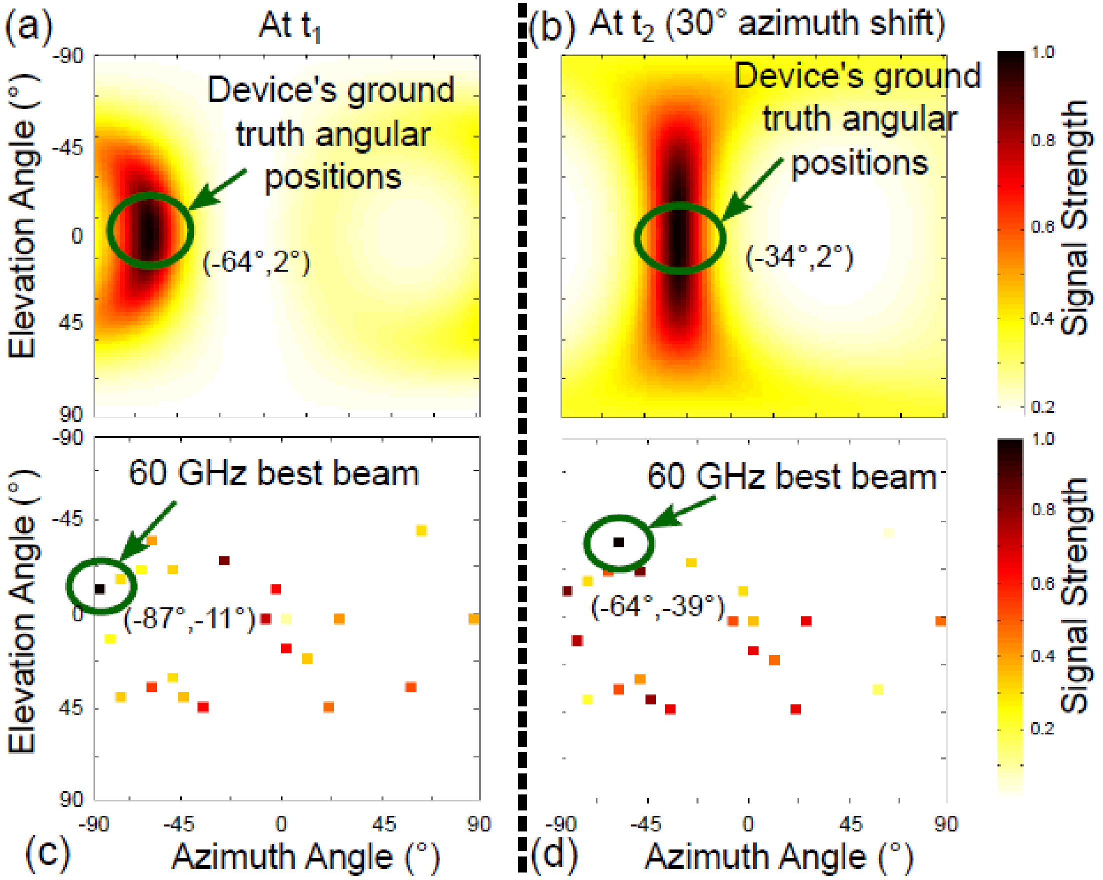

MUST: predicting 60 GHz channel using WiFi CSI

- Denote W1 as WiFi angular profile at t1, and similarly W2. Then the device's angular shift (equivalent to shift of 60 GHz dominating path)

$$\{\Delta\phi, \Delta\theta\} = \underset{\Delta\phi,\Delta\theta}{argmin} \quad |W_1(\phi,\theta) - W_2(\phi+\Delta\phi, \theta+\Delta\theta)|^2$$

- Besides angular change, we need to estimate gain change

$$D'^*(\phi+\Delta\phi, \theta+\Delta\theta) = \underset{D'}{argmin} ||h_k| - |h_k^m||^2$$
$$|h_k^m| = |\textstyle\sum_{\phi,\theta} G_k(\phi,\theta) \cdot D'(\phi+\Delta\phi, \theta+\Delta\theta)|$$

- Straightforward to predict the best beam based on channel prediction (cf. BeamSpy)

MUST: detecting risk of blockage

- Use SNR difference between WiFi and 60 GHz interface as hint to detect potential blockage
  - LOS: constant link budget difference of 27 dB
  - Blocked: large variance of SNR difference

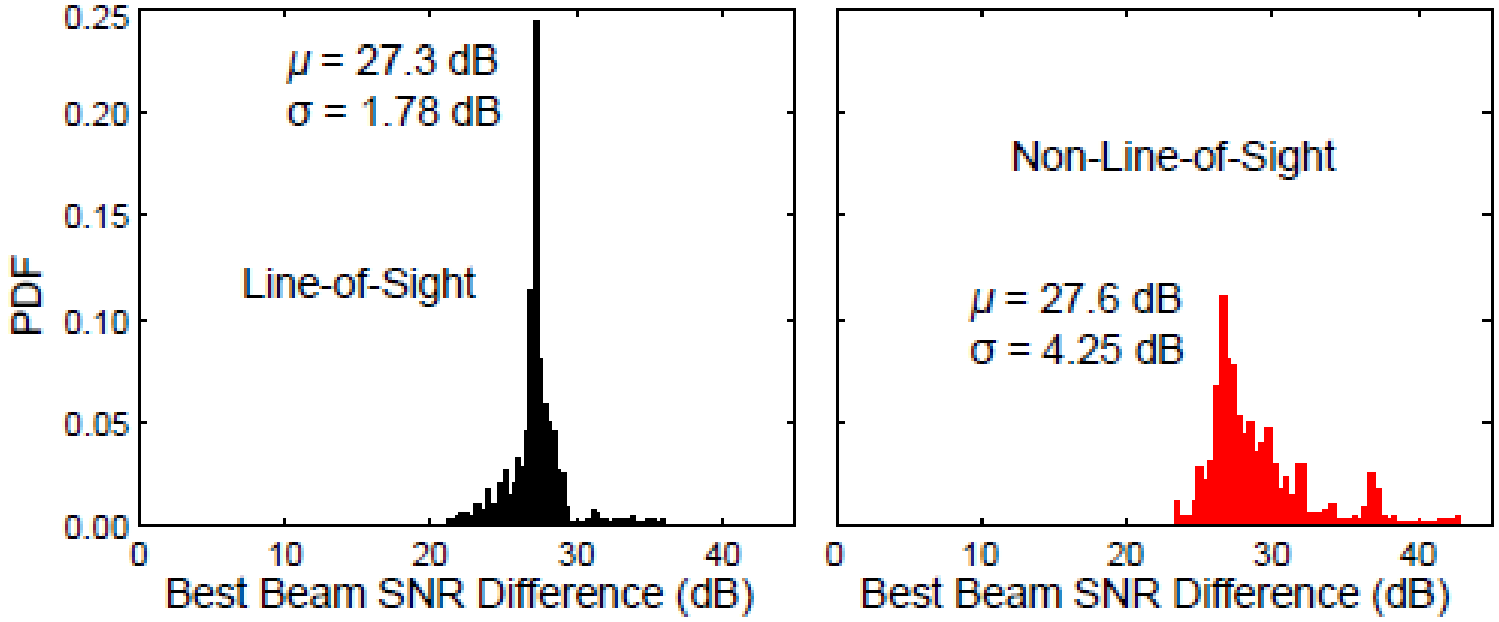

MUST: efficient interface switching

- Implementation and architecture on a tri-band 802.11ad device

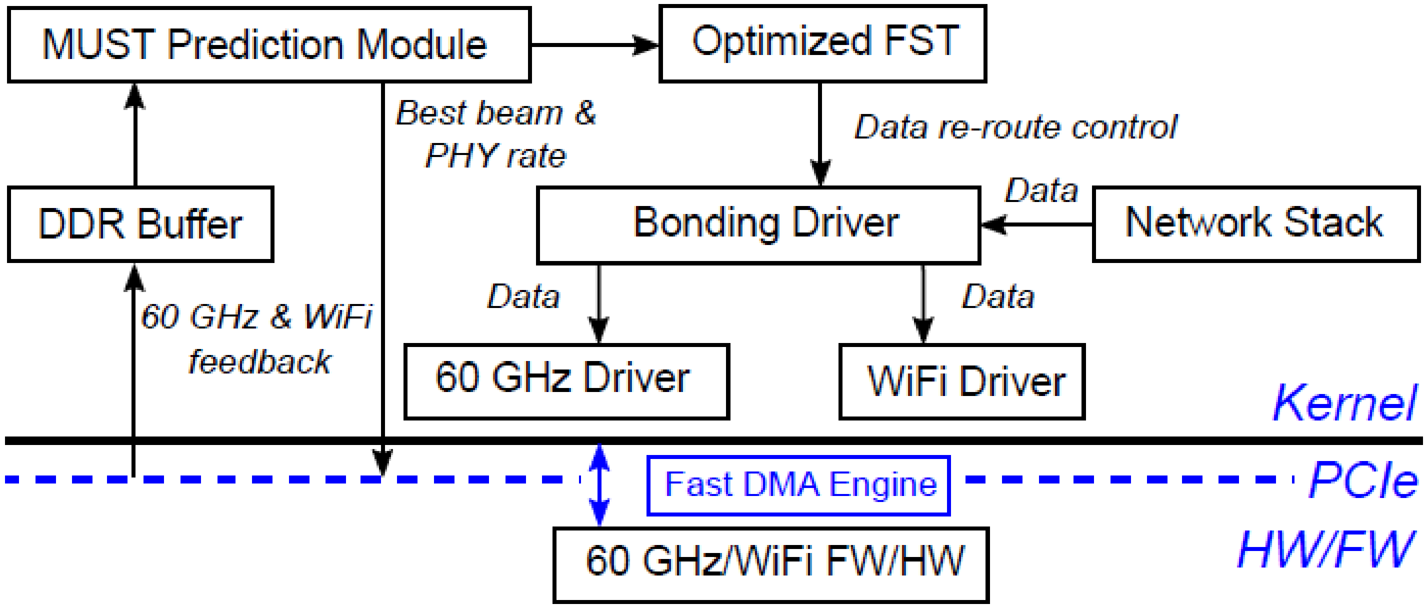

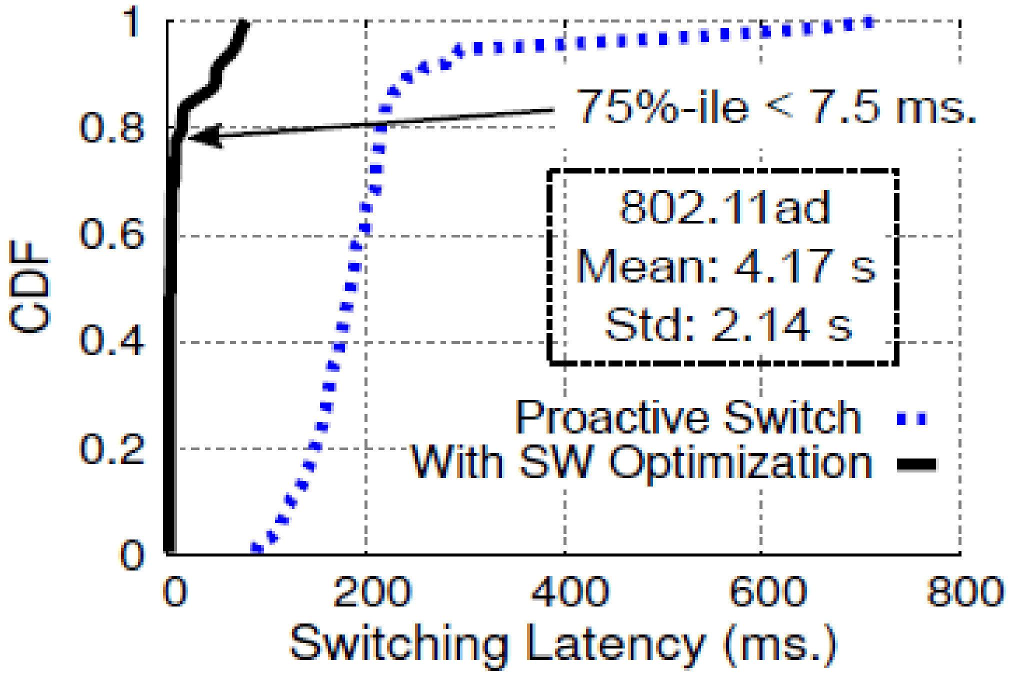

- Optimized software: prioritize FST in kernel; remove unnecessary queuing
- Balanced core affinity: serve 60 GHz and WiFi at different cores, while assigning both IRQ/packet processing of an interface in the same core.

MUST: performance overview

- Link level throughput

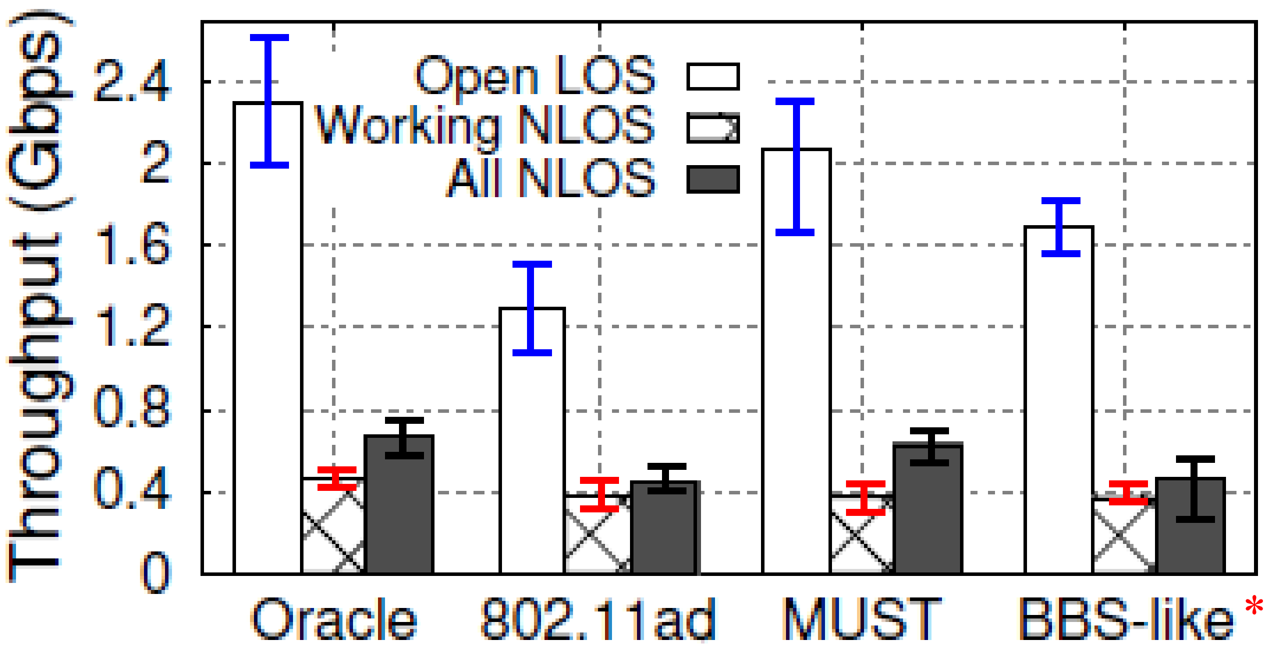

* "*Steering with Eyes Closed: mm-Wave Beam Steering without in-Band Measurement*",
San Thomas Nitsche, Adriana B. Flores, Edward W. Knightly, and Joerg Widmer, IEEE INFOCOM'16

MUST: performance overview

- TCP end to end latency
  - Orders of magnitude reduction

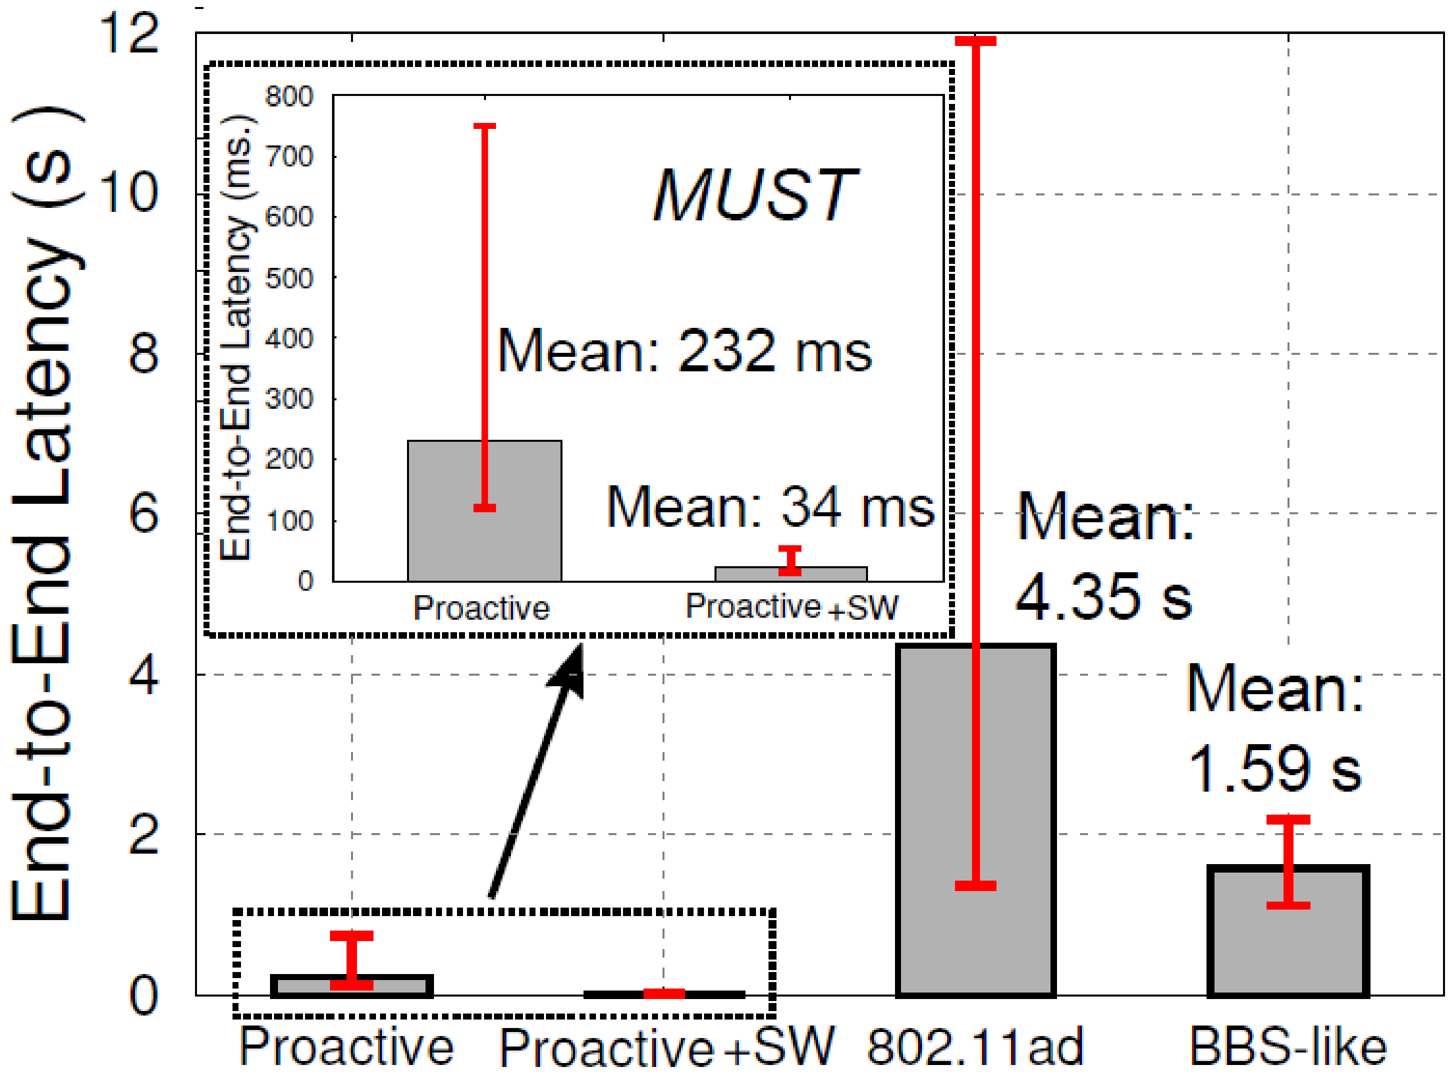

MUST: performance overview

➢ Field trials with mobile users

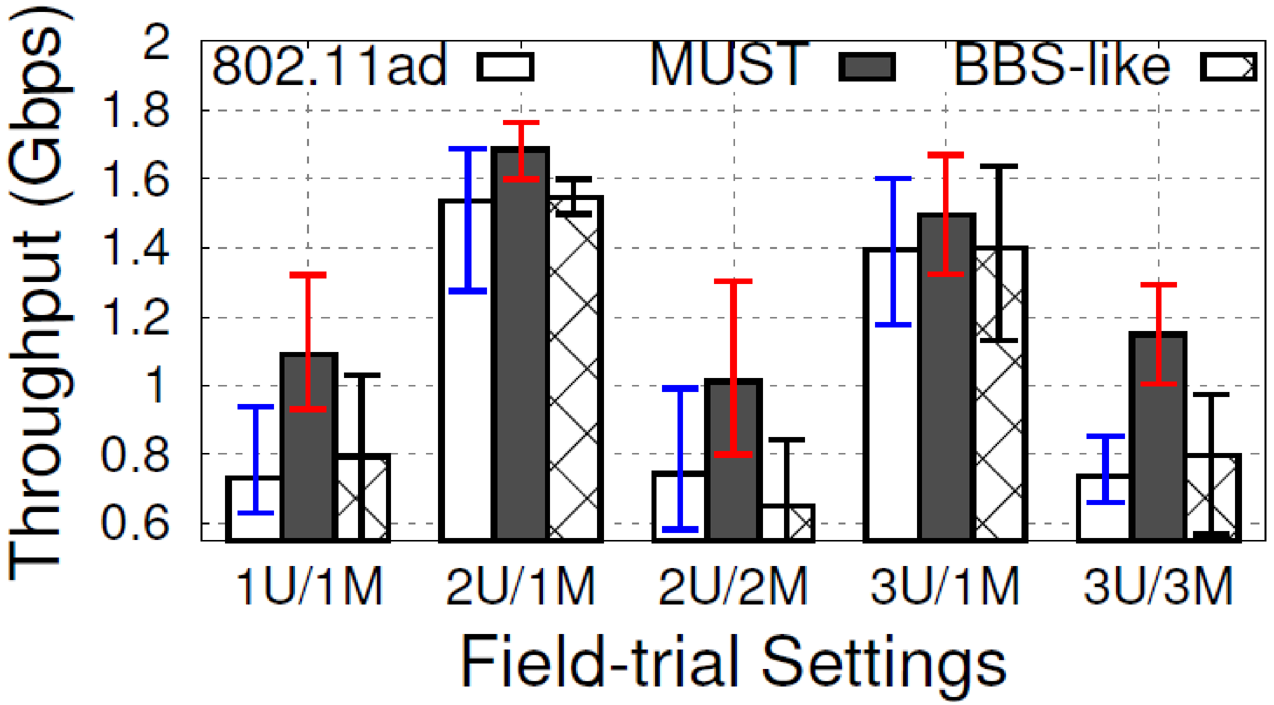

- ~50% gain over 802.11ad and 45% over BBS.
- Higher gain with more mobility.

References

* “*BeamSpy: Enabling Robust 60 GHz Links Under Blockage*”, Sanjib Sur, Xinyu Zhang, Parameswaran Ramanathan, Ranveer Chandra, USENIX NSDI’16

* “*Pose Information Assisted 60 GHz Networks: Towards Seamless Mobility and Coverage*”, Teng Wei, Xinyu Zhang, ACM MobiCom’17

* “*WiFi-Assisted 60 GHz Networks*”, Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, Kyu-Han Kim, ACM MobiCom’17

* “*Steering with Eyes Closed: mm-Wave Beam Steering without in-Band Measurement*”, San Thomas Nitsche, Adriana B. Flores, Edward W. Knightly, and Joerg Widmer, IEEE INFOCOM’16

* “*Blockage and Directivity in 60 GHz Wireless Personal Area Networks*,” S. Singh, F. Ziliotto, U. Madhow, E. M. Belding, and M. Rodwell, IEEE JSAC, vol. 27, no. 8, 2009.

* “*Beam-forecast: Facilitating Mobile 60 GHz Networks via Model-driven Beam Steering*,” Anfu Zhou, Xinyu Zhang, Huadong Ma, IEEE INFOCOM, vol. 27, no. 8, 2017.

* “*Enabling High-Quality Untethered Virtual Reality*,” Omid Abari, Dinesh Bharadia, Austin Duffield, and Dina Katabi, USENIX NSDI, 2017.