

Precise Power Delay Profiling with Commodity WiFi

Yaxiong Xie, Zhenjiang Li, Mo Li,
School of Computer Engineering, Nanyang Technological University, Singapore
{yxie005, lzjiang, limo}@ntu.edu.sg

ABSTRACT

Power delay profiles characterize multipath channel features, which are widely used in motion- or localization-based applications. Recent studies show that the power delay profile may be derived from the CSI traces collected from commodity WiFi devices, but the performance is limited by two dominating factors. The resolution of the derived power delay profile is determined by the channel bandwidth, which is however limited on commodity WiFi. The collected CSI reflects the signal distortions due to both the channel attenuation and the hardware imperfection. A direct derivation of power delay profiles using raw CSI measures, as has been done in the literature, results in significant inaccuracy. In this paper, we present Splicer, a software-based system that derives high-resolution power delay profiles by splicing the CSI measurements from multiple WiFi frequency bands. We propose a set of key techniques to separate the mixed hardware errors from the collected CSI measurements. Splicer adapts its computations within stringent channel coherence time and thus can perform well in presence of mobility. Our experiments with commodity WiFi NICs show that Splicer substantially improves the accuracy in profiling multipath characteristics, reducing the errors of multipath distance estimation to be less than $2m$. Splicer can immediately benefit upper-layer applications. Our case study with recent single-AP localization achieves a median localization error of $0.95m$.

Categories and Subject Descriptors

B.2.1 [Computer-Communication networks]: Wireless Communication

Keywords

Wireless; channel state information (CSI); power delay profile; resolution; bandwidth; channel combination; phase; localization

1. INTRODUCTION

Motivation. The power delay profile gives the power strength of a signal received through a multipath channel as a function of propagation delay, that profiles the multipath arrivals of the signal.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

MobiCom'15, September 7–11, 2015, Paris, France.

© 2015 ACM. ISBN 978-1-4503-3543-0/15/09 ...\$15.00.

DOI: <http://dx.doi.org/10.1145/2789168.2790124>.




Figure 1: Transformation from a channel frequency response to the power delay profile. (a) A channel frequency response, where f_0 , Δf , and B represent starting frequency, frequency sampling resolution, and bandwidth, respectively; (b) Derived power delay profile, where τ_0 and $\Delta\tau$ represent the propagation delay of LoS path and the power delay profile resolution, respectively.

A power delay profile fully characterizes a multipath channel, and has been recently used in various motion- or location-based applications [9, 17, 27, 28, 35, 37, 38, 46] — multipath channel dynamics can be unveiled from consecutive measures of the power delay profile, *e.g.*, tracking the power delay profile changes in a multipath channel can detect an object’s movement [9, 10, 34, 35, 38], like a person’s walking, falling, talking, or gestures, etc. In addition, the exact power level measured from each signal path can also be used to estimate the path length, *i.e.*, ranging between a pair of transmitters [28, 37].

A power delay profile can be measured by directly detecting multipath signals with different arrival times in the time domain, which however requires dedicated hardware of high signal sampling frequency [23, 26]. An alternative way to describe the channel is using the Channel State Information (CSI) which can be obtained from commodity WiFi network interface cards (NICs), *e.g.*, Intel 5300, Atheros 9580, etc. Theoretically, the frequency domain CSI can be transformed lossless to the time domain power delay profile through IFFT (Inverse Fast Fourier Transform). Figure 1 illustrates the process (which will be detailed in §2).

The time resolution of the derived power delay profile from CSI, *e.g.*, $\Delta\tau$ in Figure 1(b), is limited by the bandwidth of the transmitted signal [7, 26], *e.g.*, B in Figure 1(a), and $\Delta\tau = 1/B$. A high resolution power delay profile can differentiate subtle multipath channel changes, and consequently detect tiny activities. For the widely used 20MHz bandwidth in 802.11n [12], the power delay profile resolution is up to $50ns$, which leads to a $15m$ resolution in measuring the multipath lengths. Such a resolution imposes inevitable uncertainty in mobility detection [9, 35, 38, 45, 46], gesture recognition [33], or localization [28, 37]. For a finer grained motion detection, *e.g.*, less than $1.5m$ uncertainty to differentiate slight human body movements, at least 200MHz bandwidth is needed, which is impossible for current commodity WiFi NICs. Some re-

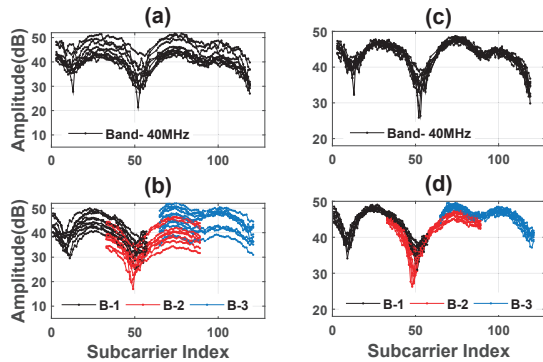


Figure 2: CSI amplitude measurements. Raw CSI amplitudes from (a) one 40MHz 802.11n band and (b) three 20MHz bands; Amplitudes after the offset removal from (c) the 40MHz band and (d) three 20MHz bands.

cent works directly use CSI in replacement of power delay profile to learn the channel dynamics. The CSI description of the channel is, however, essentially limited by the bandwidth. In addition, CSI description is indirect and dependent of hardware uncertainty.

In this paper, we observe that although the width of each individual WiFi band is limited, *e.g.*, 20MHz/40MHz, the total bandwidth allocated to 802.11 WiFi is wide, *e.g.*, more than 200MHz at 5GHz frequency band in 802.11n, which covers 10/5 different 20/40MHz channels. Furthermore, the CSIs measured from these individual WiFi channels can be spliced to derive a finer power delay profile with much higher time resolution. Figure 2 reports the results of our initial measurement study in 802.11n (detailed experiment settings reported in §3.1). The CSI measurements are complex values and hence contain two parts, the amplitude and the phase. Figure 2 (a) presents the CSI amplitude measurements (multiple times) from one 802.11n 40MHz band and Figure 2 (b) presents the CSI amplitude measurements from three 802.11n 20MHz bands that together cover the same 40MHz band in Figure 2 (a). From Figure 2 (a) and (b), we can observe obvious offsets between the measured CSI amplitudes at different times. After removing the amplitude offsets among all measurements, we see that the spliced CSI from 20MHz bands can be very similar in its shape to the one measured from the 40MHz band (Figure 2 (c) and (d)).

Compared with the amplitude splicing, the phase splicing may result much severer errors. Figure 3 depicts the phases of the same CSI traces in Figure 2. The raw CSI phases¹ measured at different times have offsets as well (Figure 3 (a) and (b)). However, even we remove their mutual offsets when splice the traces, the residual phases do not have a common shape, demonstrating diverse phase shifts in different sub-carriers. Consequently, the multiple instances of the 40MHz CSI phase measurements as depicted in Figure 3 (c) do not match each other. The CSI traces from 20MHz bands cannot match the 40MHz measurement neither. To derive an accurate power delay profile, such phase shifts must be precisely compensated because the phase value falls in a small range of $[-\pi, \pi)$, and a slight phase error will result in significant inaccuracy in the power delay profile (as we will demonstrate in §3.1).

Challenges. Removing the CSI measurement error, however, is challenging. The CSI calculated by the commodity WiFi NICs contains the signal distortions due to both channel propagation and

¹A raw CSI phase is in the range of $[-\pi, \pi)$. For a clear representation, we expand the measured CSI phases θ , using $\theta = \theta \pm 2k\pi$, to the range of $[-\infty, +\infty]$ across different channels.

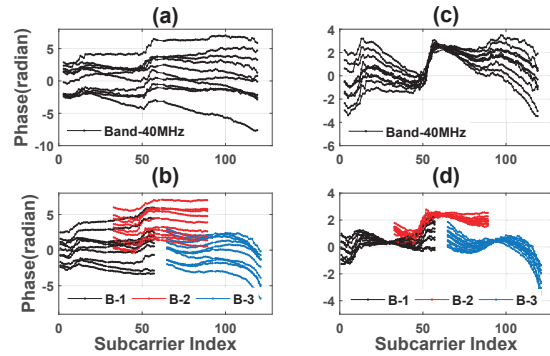


Figure 3: CSI phase measurements. Raw CSI phases from (a) the 40MHz band and (b) three 20MHz bands; Phases after the offset removal from (c) the 40MHz band and (d) three 20MHz bands. The CSI traces in this Figure are the same with Figure 2.

imperfect signal processing on the hardware, *e.g.*, imprecise sampling frequencies at the sender and receiver, shift of the central frequencies, and power control uncertainties. WiFi communication systems do not have to explicitly separate the two sources of signal distortions, because only end-to-end distortion needs to be captured and compensated as a whole in the equalization stage. To derive a precise power delay profile for the channel, however, it requires to precisely separate the channel attenuation part from the mixed signal distortions due to hardware imperfection, which is non-trivial. The sampling clock frequency uncertainty causes frequency-relevant CSI phase measurement errors in each individual channel. The central clock frequency shift and the power control uncertainty further introduce notable phase and amplitude offsets cross different channels, respectively. Based on the raw CSI measures, it is unknown how to compensate those errors for CSI splicing without the knowledge of ground-truth CSI. In addition to above challenges, wireless channels are time-varying, especially in the mobile environment. Few CSI measurements are allowed for scanning the whole WiFi band during a short coherence time. To deal with such a practical limit, we have to devise an effective method to correct and splice CSI measurements with insufficient samples and affordable computation cost.

Contributions. This paper presents a set of key techniques to address above challenges. At the high level, we exercise the observation that the CSIs collected from different frequency bands should lead to the same power delay profile that characterizes the communication channel itself. We propose an efficient method that searches for a CSI manipulation that maximizes the matching between the power delay profiles derived from CSIs obtained at different frequency bands, based on which we can perform a preliminary CSI splicing. However, the power delay profiles used for matching are derived from narrow WiFi bands with limited bandwidth, so the spliced CSI is still of low quality. We devise a wider frequency window and perform a rolling-based calibration on the spliced CSI, based on which we refine the error correction to achieve a precisely spliced CSI. To accommodate the computations in the limited coherence time, we further develop a lightweight scheduler that is able to determine the optimal number of CSIs to measure from each individual WiFi band to strike a trade-off between the error compensation and the total bandwidth that can be afforded for the CSI splicing.

We develop a system, called Splicer, to incorporate above techniques on commodity Athears 9580 NICs. Our benchmark exper-




Figure 4: Power delay profiles derived using two CSI with an amplitude offset of 7dB.

iments show that Splicer can derive high resolution power delay profiles from spliced CSI. We evaluate the derived power delay profile by estimating the distance between the sender and the receiver. According to our experiments, Splicer can reduce the median ranging error from 7.1m to 1.63m compared with using raw CSI traces from NICs. In light of such a high resolution, Splicer can immediately enhance the performance of a plethora of upper-layer applications, *e.g.*, object tracking, gesture recognition, localization, etc., without additional modification to the original application design. We demonstrate this benefit with a case study. We build the recent single-AP localization CUPID [28] on top of Splicer. Our evaluations show that the localization accuracy can be substantially improved. In particular, Splicer improves the CUPID localization accuracy by 71%, with median localization errors about 0.95m.

The rest of the paper is organized as follows. The design principle is stated in §2 and the Splicer design is detailed in §3. We investigate the accuracy of the derived power delay profile using spliced CSIs and evaluate a case study atop Splicer in §4. Related works are reviewed in §5. We conclude in §6.

2. PRINCIPLE OF CSI SPLICING

In this section, we give the theoretical foundation for CSI splicing. According to [7, 26], the channel frequency response $h(f)$ can be expressed by Eq. (1):

$$h(f) = \sum_{l=0}^L \alpha_l \cdot e^{-j \cdot 2\pi \cdot f \cdot \tau_l}, \quad (1)$$

where L is the total number of multipaths, α_l and τ_l stand for the attenuation and the propagation delay of the signal through path l , respectively. Figure 1 (a) depicts a channel frequency response when the channel bandwidth is B , *e.g.*, $f_0 \leq f \leq f_0 + B$. Channel frequency response is reported in the form of CSI in 802.11 WiFi, which is a set of discrete channel frequency response samples. With the sampling rate $F_f = \frac{1}{\Delta f}$, where Δf is the sampling resolution in the frequency domain, a receiver can obtain $M = \frac{B}{\Delta f}$ channel frequency response samples, and each sample contains amplitude and phase information.

To obtain the power delay profile, the CSI can be transformed to the channel impulse response $f(t)$ by IFFT:

$$f(t) = \sum_{l=0}^L \alpha_l \cdot \delta(t - \tau_l), \quad (2)$$

where $\delta(\cdot)$ is the delta function, and L , α_l , and τ_l have the same definitions as they are in Eq. (1). Figure 1 (b) illustrates the channel impulse response transformed from Figure 1 (a). The norm of $f(t)$,




Figure 5: Power delay profiles derived using two CSIs with raw phases and average amplitude of the CSI traces in Figure 4.

$\|f(t)\|_2$, then gives the power delay profile, which describes the power levels of each multipath with different propagation delays.

Feasibility of CSI splicing. According to Eq (1), given one multipath channel, *i.e.*, given each α_n , τ_n , and N in Eq (1), and channel bandwidth B , the channel state information is *deterministic* at each frequency f . We can thus obtain all M CSI samples from either a single measurement covering the entire bandwidth or multiple measurements where each measurement covers a subset of M samples. With the M samples, we can derive a unique power delay profile using Eq. (2) and the norm operation.

Resolution of power delay profile. After the IFFT transformation, we obtain a series of signal samples in the time domain with various delays τ_l in Eq. (2). The norm of each multipath component, $\|\alpha_l \cdot \delta(t - \tau_l)\|_2$, indicates its power level as shown in Figure 1 (b), where the first impulse corresponds to the Line-of-Sight (LoS) path. According to the IFFT theory, the time resolution $\Delta\tau$ of power delay profile is connected to the sampling resolution Δf of the channel impulse response, *i.e.*, $\Delta\tau = 1/(N \cdot \Delta f)$, where N is the IFFT length. As $N \cdot \Delta f = 1/B$, we have $\Delta\tau = 1/B$, where B is the bandwidth. Such a connection indicates that a wider bandwidth CSI leads to a higher resolution of power delay profile.

Given channel bandwidth B , two multipaths of propagation delays τ_1 and τ_2 are not distinguishable if $|\tau_1 - \tau_2| < 1/B$. Hence, all multipaths whose propagation delays differences are less than $1/B$ are viewed as one multipath component in the power delay profile, and the corresponding power level indicates the aggregated power level of those multipaths. As a result, the time resolution $\Delta\tau$ leads to $\frac{c}{B}$ uncertainty in terms of the length difference between non-distinguishable paths, where c is the speed of the signal propagation. In 802.11 WiFi in a 20MHz or 40MHz channel, the path length uncertainty is 15m or 7.5m, respectively, which can merely support coarse mobility tracking and activity recognition.

3. DESIGN

3.1 CSI splicing in practice

We first locate the error sources of CSI splicing in the 802.11 physical layer in §3.1 and then present the design details to address each of them from §3.2 to §3.5.

CSI measurement errors. We perform preliminary CSI measurements to investigate how CSI measurement errors will affect the derived power delay profiles. We use Atheros 9580 NICs that support 802.11n with 20MHz/40MHz channels at the 2.4G/5G frequency band, and modify the driver to extract CSI from the physical

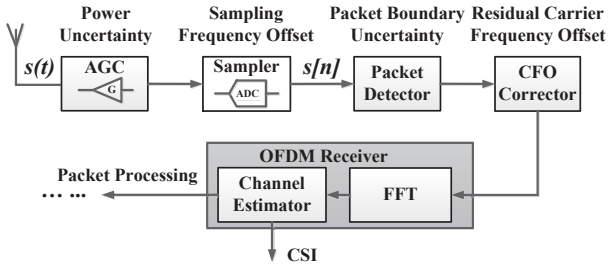


Figure 6: Illustration of signal processing in 802.11.

layer. We configure Atheros nodes to transmit packets with minimum payload to ensure a short transmission delay, *i.e.*, about 0.2ms in our experiment. We collect the CSI traces from one 802.11n 40MHz band as well as three 20MHz bands. The measurement results are reported in Figures 2 and 3, which have been discussed in §1. In this section, we further derive the power delay profiles from the CSI traces and investigate how the CSI amplitude and phase errors may impact the final power delay profiles.

CSI amplitude. In Figure 4, we select two arbitrary CSI traces from the same WiFi band (20MHz-2 in Figure 2) with an amplitude offset of 7dB, and derive two power delay profiles². Figure 4 shows that although two derived power delay profiles have different power levels, *e.g.* the average difference is 7.05 dB, they follow similar shapes. We compute the variance of the power difference for each path to quantify the similarity of two power delay profiles, which is less than 1.0 dB. We observe similar results from other CSI combinations. All of these results indicate that the derived power delay profiles approximately characterize the same multipath channel environment except that their power levels are scaled due to amplitude offsets of the CSI measurements.

CSI phases. In Figure 5, we use the raw phases of two CSI traces in Figure 4 to derive two power delay profiles³. The two profiles in Figure 5 demonstrate opposite results. One power delay profile indicates the existence of the LoS between the transmission pair (*i.e.*, the first multipath component has the strongest power level), however another one indicates that there is no LoS (NLoS) path (*i.e.*, the first arrived signal is much weaker in strength than later arrived signals). In addition, the power levels of each multipath component in these two profiles are very different, *e.g.*, the power difference of the LoS path is more than 10 dB and the variance of the power level differences is up to 4.7 dB. Figure 5 indicates that the CSI phase errors will significantly impact the derived power delay profiles, which completely change both the power loss and the multipath channel features.

As our initial experiment results suggest, the CSIs collected from WiFi NICs are mixed with rich hardware distortions. The raw CSIs cannot derive accurate power delay profiles. We identify the measurement error sources in 802.11 physical layer process and propose solutions to compensate each of them.

Sources of CSI measurement errors. Figure 6 illustrates the wireless signal processing in the 802.11 NICs. An incoming signal from the antenna is down converted to the base band signal $s(t)$ and sampled by Analog-to-Digital (ADC) to derive the digital $s[n]$. The packet boundary detector (PBD) performs correlation between $s[n]$

and a pre-defined 802.11 preamble pattern to confirm an incoming packet. Once the preamble of a packet is detected, the signal central frequency is calibrated by the central frequency offset (CFO) corrector. The OFDM receiver estimates the CSI based on the calibrated $s[n]$ and the CSI is passed to the subsequent equalization module (not shown) to compensate errors prior to the packet decoding. Due to the hardware imperfection, the CSIs measured by NICs introduce the following errors.

Power control uncertainty. Limited by the hardware resolution, Automatic Gain Controller (AGC) cannot perfectly compensate the signal amplitude attenuation to the transmitted power level. The measured CSI amplitude equals to the compensated power level, which is mixed with the power control uncertainty error. According to [14], the CSI amplitude offsets in individual bands can be removed by averaging. However, if the number of CSI measurements on each channel is not sufficient, which may be the usual case due to the stringent delay constraint (§3.5), the averaging cannot perfectly eliminate the power uncertainty. The residual offset between different WiFi bands disallows a direct CSI amplitude splicing.

Sampling frequency offset (SFO). The sampling frequencies of a transmission pair exhibit an offset due to non-synchronized clocks, which can cause $s[n]$ after ADC a time shift τ_o with respect to the transmitted signal. Because clock offsets are relatively stable within a short time, τ_o will introduce near constant errors λ_o to the CSI phases measured from different sub-carriers.

Packet boundary detection (PBD) error. Due to correlator sensitivity of packet detector, the packet detection introduces another time shift τ_b , with respect to the transmitted signal [6, 31]. The timing shift τ_b causes random errors λ_b to the measured CSI phases.

Central frequency offset (CFO). The central frequencies of the transceiver cannot be perfectly synchronized. The central frequency offset is compensated by the CFO corrector, but due to the hardware imperfection, the compensation is incomplete. Signal $s[n]$ still carries residual errors, which can cause the CSI phase offsets β .

The last three error sources cause CSI phase measurement errors. Due to the diverse phase shifts from SFO and PBD, the phases of the overlapped sub-carriers measured from two consecutive bands are inconsistent (different), which impairs the CSI phase splicing. The CSI phases measured from different bands also suffer from notable offsets. In the next section, we introduce our solutions to compensate above CSI errors.

3.2 Phase error correction

We denote S as the number of sub-carriers in one WiFi band. Based on [31, 32], the reported CSI phase value ϕ_k from any sub-carrier k by WiFi NICs can be expressed:

$$\phi_k = \theta_k + k \cdot (\lambda_b + \lambda_o) + \beta, \quad (3)$$

where θ_k is the phase rotation of subcarrier k which is caused by the channel propagation, λ_b and λ_o are phase errors introduced by the packet boundary detection uncertainty and the sampling frequency offset, respectively, β is the phase error caused by the central frequency offset, and $k = 1, 2, \dots, S$. As $\lambda_b + \lambda_o$ is multiplied by the sub-carrier index k in Eq. (3), the phase errors cross different sub-carriers are diverse among different CSI measures as shown in Figure 3. Our target is to obtain the phase value θ_k by eliminating the impact of other parameters, *i.e.*, the λ_b , the λ_o and the β . We focus on the removal of λ_b and λ_o from ϕ_k in the rest of this subsection, and introduce the removal of β when we splice the CSI phases in §3.3.

PBD phase error λ_b removal. Phase error λ_b is caused by the time shift τ_b from the packet boundary detection uncertainty. To investigate the effect of τ_b , we examine the discrete Fourier

²To isolate the impact of the CSI phases, we use the average phases of the two CSI traces such that the derived power delay profiles only differ in the amplitude.

³Similar to the experiment in Figure 4, we use the average amplitude of the two CSI traces to derive power delay profiles to avoid the impact from the amplitude.

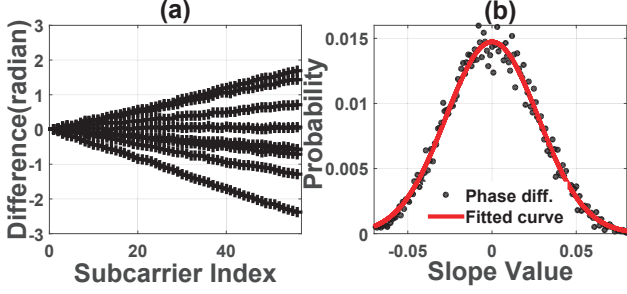


Figure 7: Phase differences cross sub-carriers and the distribution of $\Delta\lambda_b$. (a) The phase differences $k \cdot \Delta\lambda_b$ of 8 randomly CSI pairs; (b) Distribution of slope $\Delta\lambda_b$ calculated from 180 CSIs.

transform of the channel frequency response in Eq. (4):

$$h[k] = \sum_{n=0}^{N-1} f[n] \cdot e^{-j \cdot 2\pi \cdot k \cdot n / N}, \quad (4)$$

where $h[k]$ and $f[n]$ are the discrete versions of $h(f)$ in Eq. (1) and $f(t)$ in Eq. (2), respectively, and N is the IFFT length. With a time shift τ_b in $f[n]$, Eq. (4) can be rephrased as:

$$h[k] \cdot e^{-j \cdot 2\pi \cdot k \cdot \tau_b / N} = \sum_{n=0}^{N-1} f[(n - \tau_b)_N] \cdot e^{-j \cdot 2\pi \cdot k \cdot n / N},$$

where the term $e^{-j \cdot 2\pi \cdot k \cdot \tau_b / N}$ indicates that the time shift τ_b can introduce a phase error, $2\pi \cdot k \cdot \tau_b / N$, in each sub-carrier k . Therefore, $\lambda_b = 2\pi \cdot \tau_b / N$.

To remove λ_b from each ϕ_k , we leverage an observation that the time shift τ_b varies in each packet reception but follows a Gaussian distribution with the zero mean [31]. The error λ_b thus changes accordingly in different CSI measurements and $\lambda_b \sim N(0, \sigma^2)$, where σ is the standard deviation. According to the weak law of large numbers, λ_b can be removed by averaging over the measured CSI phases ϕ_k .

To validate this observation, we perform a trial of experiments in Figure 7. In Eq. (3), λ_b is mixed with λ_o , β , and θ_k in the CSI phase ϕ_k , and we cannot directly investigate its distribution. Therefore, for the collected CSIs from the same WiFi band, we calculate the mutual phase differences of those CSIs and obtain a set of $\Delta\theta_k + k \cdot (\Delta\lambda_b + \Delta\lambda_o) + \Delta\beta = k \cdot \Delta\lambda_b + a$ for each sub-carrier k^4 , where a is a constant. We thus examine the distribution of $\Delta\lambda_b$, because if $\lambda_b \sim N(0, \sigma^2)$, $\Delta\lambda_b$ should be a Gaussian with the zero mean as well. We collect 180 CSIs from a 20MHz channel within a short time interval when the environment is stable. Figure 7 (a) plots the $k \cdot \Delta\lambda_b$ value versus the sub-carrier index k for 8 randomly selected CSI pairs (we omit the constant a for the presentation clarity). Figure 7 (a) shows each line is a straight line and the slopes of those lines are different. The result indicates that λ_b is a constant to each sub-carrier in each individual measure, but varies cross different measures. To further examine its distribution, in Figure 7 (b), we divide the range $[-0.075, 0.075]$ into 100 bins on the x -axis and plot the frequency of $\Delta\lambda_o$ falling into each bin on the y -axis. After the curve fitting, we find that $\Delta\lambda_b$ indeed follows a Gaussian distribution with the zero mean.

According to the λ_b distribution, we can remove it by averaging over the measured CSI phases ϕ_k s. In principle, more measurements lead to a better error removal, but it will prolong the latency

⁴ λ_o is a constant that can be removed by the deduction, which is detailed in the SFO phase error removal.

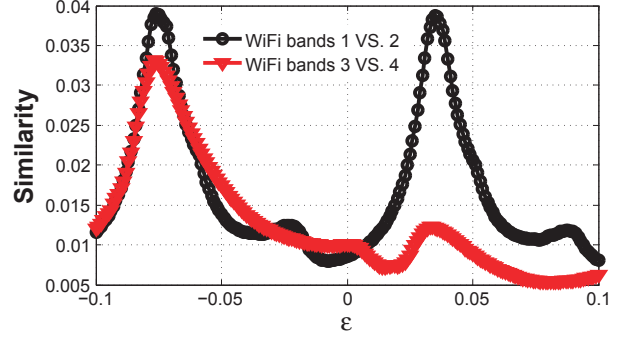


Figure 8: The power delay profile similarities for two pairs of WiFi bands when ϵ varies from -0.1 to 0.1.

to scan each single band. In §3.5, we will determine an optimal number of CSIs collected from each band to balance this trade-off subjected to the stringent channel coherence time. Given the optimal amount \hat{n}_i for any band i , we calculate:

$$\bar{\phi}_k^i = \sum_{j=1}^{\hat{n}_i} \phi_k^i(j) / \hat{n}_i, \quad (5)$$

where $\phi_k^i(j)$ stands for the j -th CSI measure from band i . After λ_b is removed, $\bar{\phi}_k^i = \theta_k^i + k \cdot \lambda_o + \beta$. In the next subsection, we introduce how to remove λ_o from $\bar{\phi}_k^i$.

SFO phase error λ_o removal. Phase error λ_o is caused by the offset of the sampling frequencies of the sender and the receiver, f_s and f_r . We denote $\zeta = \frac{f_s}{f_r} - 1$ as the fractional difference in sampling frequency, and the effect of the sampling frequency offset is to introduce a term, $e^{j \cdot \zeta' \cdot k}$, to the channel frequency response: $h'[k] = h[k] \cdot e^{j \cdot \zeta' \cdot k}$, where ζ' stands for ζ multiplied with a constant and $h[k]$ is the channel frequency response without the sampling frequency offset (Eq. (4)). Hence, $\lambda_o = \zeta'$. As the fractional frequency difference keeps stable in the order of minutes [13], λ_o is a constant during the process of the CSI splicing.

To remove λ_o from $\bar{\phi}_k^i = \theta_k^i + k \cdot \lambda_o + \beta$, obtained from Eq. (5), we leverage an observation that the power delay profiles derived from different WiFi bands should be the same after λ_o is removed (we will show in §3.3 that the phase offset β has no impact on the derived power delay profile), since they characterize the same multipath channel. For any WiFi band i , the CSI phases $\bar{\phi}_k^i$ s from all S sub-carriers form a vector $\bar{\Phi}^i = [\bar{\phi}_1^i, \bar{\phi}_2^i, \dots, \bar{\phi}_S^i]^T$. Therefore, we propose to gradually “rotate” two distinct $\bar{\Phi}^i$ and $\bar{\Phi}^j$ in the frequency domain⁵ and stop when the two derived power delay profiles best match each other. We repeat this process for different pairs of WiFi bands to improve the accuracy. To quantify the likelihood of two power delay profiles, e.g., P_1 and P_2 , we define their similarity as:

$$\rho(P_1, P_2) = \frac{1}{\|P_1 - P_2\|_2}, \quad (6)$$

where the dominator essentially measures the power level differences of each multipath component in the two power delay profiles. A large $\rho(P_1, P_2)$ value indicates that P_1 and P_2 are more similar.

To illustrate this solution, we measure four different 20MHz WiFi bands, and compute the similarities for two channel pairs in Figure 8. We compensate λ_o by $\lambda_o - \epsilon$ and search for the optimal

⁵We multiply $\bar{\Phi}^i$ and $\bar{\Phi}^j$ with $e^{j \cdot \epsilon \cdot k}$, and gradually vary ϵ .

Algorithm 1: SFO Phase Error Compensation

- 1 for each WiFi band pair c do
 - 2 Record the top-two local maximal similarity values and the corresponding ϵ : $\langle \epsilon_1^c, \rho_1^c \rangle$ and $\langle \epsilon_2^c, \rho_2^c \rangle$.
 - 3 Clustering on ϵ and find the cluster with the maximal similarity sum.
 - 4 Return the cluster center as the final ϵ value.
-

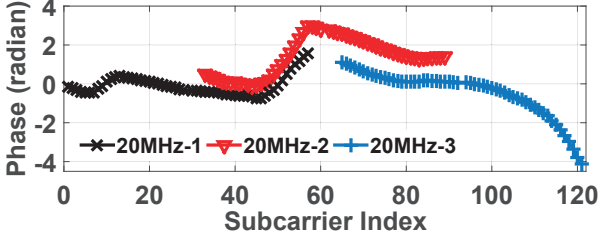


Figure 9: CSI phases after the removal of λ_s and λ_o .

ϵ for both positive and negative directions. For the first pair, we observe four local maximum points, and for the second pair, we observe three local maximum points when ϵ varies in $[-0.1, 0.1]$. According to Algorithm 1, we can determine the final ϵ as the average of ϵ_1 and ϵ_2 in Figure 8. With the optimal ϵ obtained from Algorithm 1, for each $\bar{\Phi}^i$ after the PBD phase error λ_b removal, we can further remove λ_o by $\tilde{\Phi}^i = \bar{\Phi}^i - [\epsilon, 2\epsilon, \dots, S\epsilon]^T$, where S is the number of subcarriers.

Although the ϵ searching introduces extra computational delays, Algorithm 1 does not need to be executed parallel to the CSI sampling in real-time. As a matter of fact, once sufficient CSIs can be obtained subjected to the stringent channel coherence time (§3.5), the latency of Algorithm 1 only impacts the frequency to generate power delay profiles to the upper-layer applications. We evaluate the computational efficiency in §4.

3.3 CSI phase splicing

Figure 9 plots the corrected CSI phases at this stage for the three 20MHz channels in Figure 3 (b). From the result, we observe that after the phase error removals of λ_b and λ_o , the shapes of the overlapped sub-carrier phases from different WiFi bands now become similar and consistent. The only barrier that remains is the offsets. In this subsection, we target to removing offsets to finally enable the phase splicing.

CSI phase offset β removal. Phase offset β is caused by the central frequency offset of the transmission pair. Through our study, we find that for individual WiFi bands, phase offset β has no impact on the derived power delay profile, *i.e.*, given a pair of CSI amplitude (ω) and phase (θ), the power delay profile derived by ω and θ is identical to the one derived by ω and $\theta + \beta$. The reason is that offsets are frequency independent. After IFFT, the error will result in a constant phase rotation term in each α_l of Eq. (2), which will not change the norm operation result. Hence the power level of each multipath component keeps unchanged. According to this observation, we can use the phase measured from any band as a reference and compensate β by calibrating the phases measured from other bands with respect to the reference.

Spliced CSI phase refinement. So far, the CSI phases measured from different channels can be spliced already. However, we find that the CSI phase accuracy can be further improved by leveraging




Figure 10: Spliced CSI phase refinement.

the primarily spliced result. In the SFO phase error λ_o removal, we rely on the similarity of the derived power delay profile to compensate λ_o , but the derived power delay profiles are based on the CSI from single WiFi bands with limited bandwidth. The phase error λ_o thus cannot get fully corrected using low-resolution power delay profiles. To refine the phase information, in the phase splicer, we further manually divide the spliced channel bandwidth into multiple windows illustrated in Figure 10. Each window has a much wider bandwidth, denoted as l , than each single WiFi band. We slide the window and obtain a set of l -wide phase pieces. For each phase pair, we call Algorithm 1 to estimate the SFO phase error compensator ϵ , and use their average value to further compensate the spliced CSI phase.

In Figure 10, the window size l balances a trade-off. A larger l will lead to a higher-resolution power delay profile, which potentially can better compensate the phase error λ_o . However, with a large l , the channel information carried by each divided CSI piece will be highly redundant (more overlapped sub-carriers). Two such power delay profiles will produce a large body of local maximal points in Figure 8, which are dominated by the channel information redundancy, instead of the removal of the phase error λ_o . It thus prevents to find the optimal ϵ to compensate λ_o . Therefore, the length l needs to be carefully selected. In Figure 11, we investigate this trade-off. The experimental setting is detailed in §4, where we use ranging accuracy as the metric for the evaluation. We set the window size l as a fraction of the total bandwidth that the spliced CSI covers. Initially, when we increase the window size, the ranging accuracy improves because λ_o is better compensated by higher-resolution power delay profiles. However, when l is excessive large, *e.g.*, close to 0.95, the accuracy decreases due to the unreliability in the optimal ϵ search. According to Figure 11, we set l to be $\frac{3}{4}$ of the total bandwidth that the spliced CSI covers as default in Splicer.

3.4 CSI amplitude splicing

In Figure 2, we have shown that the amplitudes of raw CSIs also exhibit significant offsets. The reason is that the power control uncertainty [8, 14], which also follow a Gaussian distribution. However, different from the CSI phases, the power uncertainty is frequency band independent, *i.e.*, the amplitudes measured from different bands follow the same distribution. Thus, for CSI amplitudes, there is no need to average them for individual WiFi bands. Instead, we can average the amplitudes after we collect all CSIs for the splicing. The total number of CSIs to collect for splicing is determined by the channel coherence time, which will be discussed and given in the next subsection.

So far, both the CSI amplitudes and phases both can be spliced. Figure 12 depicts the splicing result covering the entire 802.11n WiFi band, 200MHz, obtained from our experiments in §4.

3.5 Battle the coherence time constraint

Wireless channels are time-varying, which enforce a stringent time budget for each round of CSI splicing, since the spliced CSI




Figure 11: Empirical investigation of window size l .

is valid only when the channel condition is relatively stable. In this subsection, we first estimate the minimum number of CSIs to collect from each individual band that can fully compensate phase errors, and then propose an efficient CSI sampling scheduler to balance the trade-off between the error compensation quality in each individual band and the total bandwidth that can be afforded for the CSI splicing within the time budget.

Stringent time budget. The channel coherence time T_c can be expressed as $T_c \approx \frac{1}{2 \cdot f_d}$, where f_d stands for the Doppler shift. Previous works [7] have studied the Doppler shifts caused in typical mobile environments. For instance, when people are walking, the Doppler shifts are usually less than 12Hz, which can be translated into $T_c = 40ms$ with 2.4GHz WiFi. In Splicer, the transmission delay of each packet (without minimum payload) is around 0.2ms, *e.g.*, approximately 200 CSIs can be collected within the time budget. In our current design, we rely on the empirical results from the literature to set T_c . We leave the cooperation with advanced channel coherence time measurement schemes [16] in Splicer as the future work of this study.

CSI measurements for each band. As aforementioned in §3.2, for each WiFi band, we need to collect sufficient CSI measures to fully compensate phase error λ_b , which is caused by the signal boundary detection uncertainty and follows a Gaussian distribution. According to the weak law of large numbers, more CSIs lead to a better compensation. However, we have a stringent time budget T_c to scan the entire WiFi band, which on the other hand limits the number of CSI collected from each WiFi band. To deal with this trade-off, we first investigate the minimum number of CSIs for each band that can achieve a given confidence level.

When we collect n CSIs from one band, for any sub-carrier k , we can calculate the average phase value $\bar{\phi}_k$ by Eq. (5). According to [2], we can define a confidence range $(\bar{\phi}_k - \frac{\sigma}{\sqrt{n}} z_{\alpha/2}, \bar{\phi}_k + \frac{\sigma}{\sqrt{n}} z_{\alpha/2})$, where σ is the standard deviation of λ_b , α is an error rate, and z_j is a normal distribution related parameter that can be obtained from a table [2]. When n increases, the range shrinks, *i.e.*, the confidence increases. The theory in [2] proves that the probability that $E[\phi_k]$ falls into this range is greater than $1 - \alpha$. Therefore, given α and the confidence range length r , the minimal number of CSIs to collect, \hat{n} , can be determined when $\frac{\sigma}{\sqrt{\hat{n}}} z_{\alpha/2} \leq r/2$.

To cover the entire WiFi band, we need to scan multiple (*e.g.*, C) individual WiFi bands, *e.g.*, $C = 5$ for 802.11n at 2.4GHz. However, due to the stringent time budget, it could be infeasible to collect all $C \times \hat{n}$ CSIs within T_c . Therefore, we propose to balance the trade-off between the error compensation quality in each individual




Figure 12: CSI splicing result. (a) The spliced CSI amplitude; (b) The spliced CSI phase.

band and the total bandwidth that can be afforded for the CSI splicing. Through our study, we observe that for any three consecutive WiFi bands, if the CSI phases in the first and the third channels are well compensated, they can serve as two anchors to further calibrate the CSI phase from the second band if its phase error is not fully corrected. In particular, we can rotate the CSI phase from the second band and stop when its phase differences for the overlapped sub-carriers with other two bands are minimized. With this observation, the scheduling of the CSI measurement for each band is as follows. We have three constraints:

- Collecting \hat{n} CSIs from a set of WiFi bands whose indices are odd. They form anchors.
- For each even-indexed band in between, we set a lower bound that the number of CSIs collected from those even-indexed bands is greater than this bound.
- All CSIs are collected within the time budget T_c .

The scheduling objective is to maximize the total number of bands, both odd and even, that can be used for the CSI splicing. After solving this optimization problem, we can determine the optimal CSI collection assignment. When enough CSIs are collected from each band, the time delay handler informs the physical layer to send an ACK such that the sender and the receiver switch to the next band synchronously. As a contingency plan, the sender and receiver will switch back to the first channel if no packets received for a given time-out duration.

Early termination. Due to the wireless channel dynamics, the time budget we adopt from the literature may not always precisely capture the channel coherence time. It is possible that the wireless channel changes dramatically in the middle of CSI collection so that the CSIs from rest WiFi bands will become useless for splicing. To address this issue, we propose an early termination strategy to detect such a case in real time, which leverages the following observation. If the channel is stable, θ_k for any sub-carrier k of a WiFi band in Eq. (3) keeps unchanged. When we collect multiple CSIs from this band and compute their phase differences, we can obtain a set of straight lines as shown in Figure 7. Later, when the channel condition is changed, if we collect another CSI trace from the new channel and compare the phase difference with the CSI from the previous channel, the result demonstrates not a straight line. Therefore, we can apply the linear regression on the phase

difference and use the regression error to indicate the level of the channel condition change.

In light of this, after receiving all CSIs for each round of CSI splicing, the receiver will not directly use them. Instead, for each individual WiFi band involved in the splicing, the receiver will compute the phase difference between the first and the last CSIs collected from this band, and apply linear regression to check the regression error. If the error is greater than a threshold, the receiver will discard all CSIs from this band, as well as all subsequent bands (we will evaluate its effectiveness in §4). CSI splicing only utilizes non-discarded bands.

4. EVALUATION

In this section, we conduct testbed-based experiments to evaluate the performance of Splicer. We introduce our experimental setting in §4.1, evaluate the efficacy of Splicer in §4.2, and report the end-to-end system performance of Splicer-enhanced CUPID in §4.3, which is one of the state-of-the-art indoor localization designs [28].

4.1 Experimental setup

We conduct experiments in a laboratory, which is a typical indoor office environment. We install five APs at five known locations. Each AP is a laptop connected to an Atheros 9580 NIC. The five APs are configured in the monitoring mode as receivers. We select 500 different locations in the laboratory and place another AP, which is set to the RootAP mode as sender, at each of these locations sequentially. At each location, the AP sender conducts multiple rounds of CSI splicing to each of the five AP receivers. To examine the accuracy of the power delay profiles derived by Splicer, we first evaluate the accuracy of the power level measured from the Line-of-Sight (LoS) path. To this end, we measure the distance between the sender and each of the receivers at all 500 different locations as the ground truth. At each location, we compare the derived distance from the power delay profile with the ground truth for evaluation. In addition to the LoS path, we also evaluate the quality of the derived power delay profiles using the power level stability for Non-Line-of-Sight (NLoS) paths.

We first provide a detailed performance analysis of Splicer in §4.2, we disable the sampling scheduler (described in §3.5), and manually control the number of scanned channels. From each channel, we collect 30 CSI traces. After that, in §4.3, we enable the sampling scheduler and evaluate the end-to-end performance with the full-version Splicer in a localization application.

4.2 Results

Phase error correction. To investigate the effectiveness of the CSI error correction designs in Splicer, we first evaluate the ranging performance to estimate the LoS path length d in each single 20MHz WiFi band without using CSI splicing. We evaluate Splicer for three different versions to investigate where the performance gains Splicer achieves come from: the full version with both λ_b and λ_o compensations, as well as two degraded versions — without λ_b compensation and without any phase error compensation. We compare the three versions against the measured ground truth to one randomly selected receiver in Figure 13. In the figure, the x -axis presents different locations and the y -axis illustrates the ranging results of each Splicer version (as well as the ground truth). Each reported accuracy value is the average from 10 measurements. From the result, we observe that without any phase error correction, the ranging performance is highly unreliable. After the phase error λ_b removal, the ranging accuracy Splicer achieves has been dramatically improved, but it is still far away from the ground truth. After the compensations of λ_b and λ_o , we find that the accuracy

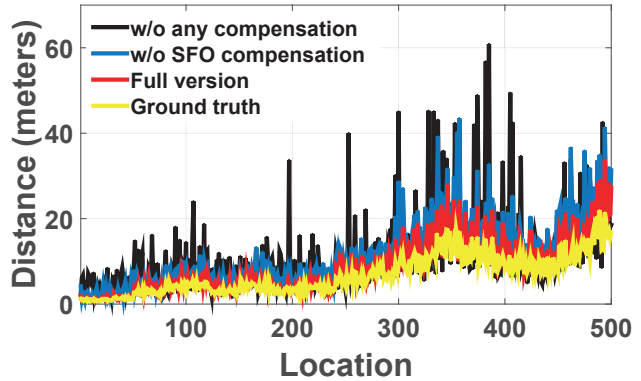


Figure 13: Ranging results at 500 different locations by different versions of Splicer in a single WiFi band.




Figure 14: CDF of ranging errors from a single band.

is very close to the true distance value between the sender and the receiver at different locations.

In Figure 14, we provide detailed statistical results for the performance achieved by different Splicer versions in Figure 13. From the result, we see that when the raw CSI phases are used, the ranging error is 10.7m for 80% of the measurements. The median and the maximum errors are 6.1m and 24m, respectively, compared with the ground truth. After the compensation of phase error λ_b , the ranging error is reduced to 6.3m for 80% measurements. After the removal of both λ_b and λ_o , the ranging error is less than 4.3m for 80% of cases. The improvement is as high as 4.8m on average compared with the traditional ranging performance using the raw phase information.

CSI splicing. In Figure 15, we evaluate the ranging performance of Splicer after we enable CSI splicing from different WiFi bands. As the power delay profile resolution is determined by the channel bandwidth, we select two representative bandwidths after splicing, 200MHz that is the total bandwidth allocated to 802.11n and 120MHz that is in between the entire WiFi band and each single WiFi band. We select all spliced CSIs with these two selected bandwidths to report the the performance. As a benchmark, we also include the performance of the full version Splicer in 20MHz for comparison.

Figure 15 plots CDF of the ranging errors from those three approaches. From the result, we find that the performance of Splicer




Figure 15: CDF of ranging errors with CSI splicing.

using merely a single WiFi band is still limited, even with the CSI error correction. In general, the wider the bandwidth Splicer uses, the smaller error the ranging can achieve. The performance gain stems from more accurate power level measurement of the LoS path. According to the statistics, on average Splicer-120MHz and Splicer-200MHz can outperform Splicer-20MHz by 22.7% and 55.5%, respectively. The performance gain of Splicer stems from both the CSI splicing as well as the CSI phase refinement (§3.2). In Figure 15, we also show the performance gain the phase refinement provides. From the result, we see that the CSI phase refinement can reduce the ranging error by 0.42m approximately.

Non-line-of-sight paths. In Figures 13 to 15, we have evaluated the LoS power accuracy in the derived power delay profiles. In this experiment, we investigate the quality of the derived power delay profiles for all other NLoS paths. In practice, the absolute power level of a NLoS path is not directly useful in applications. Instead, the relative change of the power level of each NLoS path indicates the multipath channel dynamics, which has been used for activity or gesture recognitions [24, 33, 35]. To this end, we randomly select 10 locations in the laboratory. At each location, the transmission pair performs multiple rounds of CSI splicing. We evaluate the stability of the measured power levels for NLoS paths.

Figure 16 depicts the results. In Figure 16 (a), we examine the NLoS path power level stability using a single 20MHz WiFi channel. However, we observe that without the phase error compensation, the power variance is high, *e.g.*, around 16 dB, even the pair of transmitters are static at their locations. Our CSI error compensation designs can reduce the variance to be less than 10 dB on average for single WiFi bands. In Figure 16 (b), due to the CSI splicing, we can derive higher-resolution power delay profiles. Hence, the measured power level for each multipath component is aggregated from fewer non-distinguishable multipaths (§2), which should suffer from even less uncertainty. The result in Figure 16 (b) is consistent to our analysis, which shows that the power level variance for NLoS paths in the derived power delay profiles using 200MHz spliced bandwidth is less than 5.2dB in our experiment.

Computational delay. We evaluate the computational delay of Splicer on commodity WiFi APs to investigate its frequency to generate new power delay profiles to upper-layer applications. We examine the absolute computation delays for Splicer to complete one round of CSI splicing with different aggregated bandwidths in Figure 17. Overall, the computation delay increases when more bands are spliced, since we need to compare more pairs of single bands for the phase error λ_o removal. From Figure 17, we find that the delays increase from 4.1ms to 7.9ms when the total bandwidth varies from 20MHz to 200MHz. We suppose the channel coherence time




Figure 16: Measured power level variance for NLoS path components in the derived power delay profiles. (a) In 20MHz WiFi band; (b) In 200MHz WiFi band.

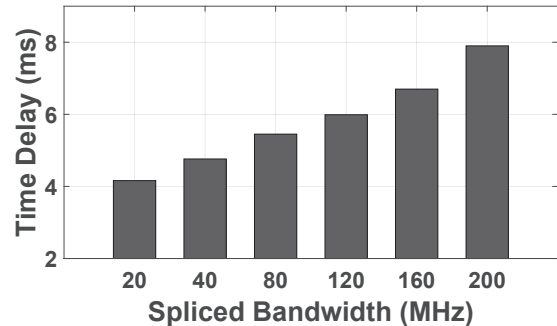


Figure 17: Computation delay of Splicer.

budget is 50ms. With the maximum bandwidth, one round of CSI splicing can be finished within 60ms. Therefore, Splicer can generate 16 new power delay profiles in one second, which fits the needs of most mobile applications [5, 11, 18, 19, 28].

4.3 Case study: indoor localization

In §4.2, we have evaluated the accuracy of the derived power delay profiles by Splicer. With high resolution power delay profiles, the performance of a plethora of upper-layer applications, *e.g.*, localization, object tracking, gesture recognition, etc., can be significantly improved. In this subsection, we take localization as a vehicle to demonstrate this capability of Splicer.

4.3.1 Case study overview

We integrate Splicer into the recent single-AP location approach CUPID [28]. In CUPID, a mobile user can locate its location as follows. The mobile device of the user transmits a packet to an AP with known location. The AP extracts the CSI upon receiving this packet and derives the power delay profile. According to the power level of the LoS path, the AP can estimate this path length. On the other hand, the AP further apply the MUSIC algorithm on the CSI to compute its pseudo spectrum, which can approximate the signal arrival direction from the LoS path. According to the path length

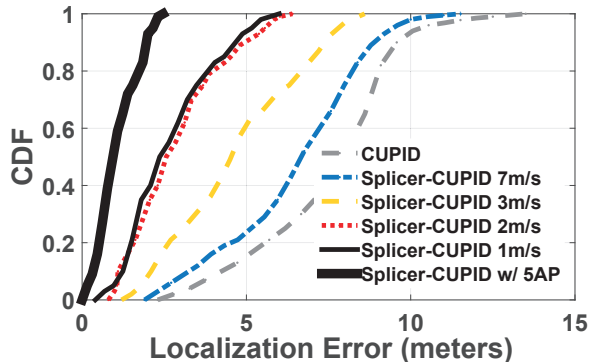


Figure 18: CDF of localization errors using the original CUPID and the enhanced CUPID by Splicer.

and the signal arrival direction, the AP can locate the mobile user and send the localization result back to the user.

To enhance CUPID, we use the spliced CSI as the system input and keep the rest of the CUPID design unchanged. To evaluate the performance, we install one AP in the laboratory with known location and deploy one AP on a robot. The robot is programmed to move along a predefined trajectory with a known speed. According to the time stamp contained in each packet, we can calculate the instant location of the robot when the mobile AP transmits this packet. To evaluate the end-to-end system performance, we enable the sampling scheduler in Splicer.

4.3.2 Results

Localization accuracy. In Figure 18 (a), we compare the performance of the original CUPID and the enhanced CUPID by Splicer, denoted as Splicer-CUPID. We vary the moving speeds of the robot in the experiment. Since the localization in the original CUPID depends on a single CSI measurement, the localization performance of CUPID is not impacted by the moving speed of the robot. For Splicer-CUPID, a higher moving speed leads to a shorter channel coherence time. As a consequence, fewer WiFi bands can be included in each round of CSI splicing and the localization accuracy will decrease. From Figure 18, we find that in general, the localization accuracy of CUPID is not quite accurate, *e.g.*, around 8m for 80% of localizations, which is consistent to the performance reported in [28]. For Splicer-CUPID, with a normal moving speed of a person ($< 2m/s$), the accuracy can be dramatically improved, *e.g.*, the median localization error is 2.3m and 2.5m when the speed is 1m/s and 2m/s, respectively. The localization error is less than 6.4m throughout the experiment. With an even higher moving speed, *e.g.*, 7m/s, Splicer-CUPID still outperforms CUPID.

In [28], the authors propose an AP selection scheme to further improve the localization accuracy by harnessing a dense AP deployment. According to the experiment results in [28], the localization accuracy of Splicer-CUPID using a single AP achieves comparable performance to CUPID with 5 APs, which can significantly improve the usefulness of indoor localizations. In Figure 18, we also leverage such an improvement opportunity. The result shows that the gain from Splicer-CUPID is significant. Localization errors are reduced to 1.75m for 80% cases and the median error is 0.95m when 5 APs are used. Although more APs may improve the localization performance, if the ranging accuracy is not high at the first place, improvement with more APs is limited.

Impact of moving speeds. In Figure 19 (left y-axis), we investigate the total bandwidth that can be spliced using Splicer with

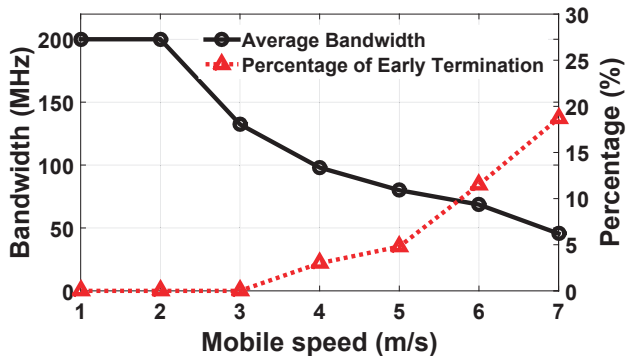


Figure 19: Bandwidth used for localization under different moving speeds. (Left) Spliced bandwidth; (right) Percentages of early terminations.

respect to different moving speeds of transmitters. In general, a higher speed leads to a shorter channel coherence time. As a result, CSI traces are spliced from fewer WiFi channels and the power delay profile resolution is lower. From Figure 19, we see that Splicer can make full use of the 200MHz available WiFi frequency band, when the speed is smaller than 2m/s. The small localization errors observed in this case in Figure 18 is compatible with such an observation. When we accelerate the moving speed from 3m/s to 6m/s, the utilized bandwidth drops from 130MHz to 60MHz, which, however, is still wider than a single WiFi channel, *i.e.* 20MHz or 40MHz. Further more, the bandwidth will drop to 45.5MHz when the speed increases to 7m/s, which is comparable to one single 40MHz channel. Nevertheless, Splicer-CUPID still improve the localization performance since Splicer compensates the CSI measurement errors.

In Figure 18, we further examine the percentages of early terminations (§3.5) occurred in our experiment, in Figure 19 (right y-axis). From the result, we see that the early termination strategy can discover 4.8% to 18.7% rapid channel varying within the channel coherence time budget.

Performance in different environments. In addition to the evaluation in the laboratory environment, we further conduct experiments in other representative environments for localization, including a corridor, a car park, and a lecture hall. Figure 20 plots the median localization error for Splicer-CUPID in comparison with the original CUPID with normal walking speed (1m/s). From the results, we see that the localization error of the original CUPID design can be largely reduced by Splicer-CUPID, *e.g.*, 70.9% in laboratory, 74.0% in corridor, 74.2% in car park, and 76% in the lecture hall. In summary, Splicer achieves general performance gains with different environments.

5. RELATED WORK

Channel sounding. Measuring the wideband channel frequency response requires high-end hardware with high sampling frequencies [4, 20]. The authors in [21, 22] develop systems to measure channel frequency responses from a group of narrow bands to approximate a wideband channel. In [15], the receiver only listens to a few harmonics of a wideband signal each time and then can reconstruct the wideband frequency response. Such a design does not require the modification at the sender. In addition, CSI-SF proposed in [3] can estimate the channel state information for multi-streams using the single stream measurement result. Some ToA-based localization approaches [41, 42] also propose to increase the resolu-




Figure 20: Localization errors in different environments using the original CUPID and the enhanced CUPID by Splicer.

tion using the channel combination, however, only with Software-Defined-Radio. Although Splicer shares a similar principle with those existing works, most of them require tight synchronization between the sender and receiver, *e.g.* devices are connected by the same clock or use GPS and atomic clocks. In this work, however, we meet and address particular challenges due to the hardware imperfection and stringent channel coherence delay constraint, which do not exist in any of existing works. In addition, Splicer can be integrated into commodity NICs without any hardware modification.

CSI phase calibration. Prior works also notice that the CSI traces reported by WiFi NICs contain phase errors introduced by hardware [40, 43, 44]. ArrayPhaser [6] enables the phased array signal processing on commodity WiFi devices. However, ArrayPhaser does not correct any of those phase errors, instead they just treat the phase values measured from one NIC as the reference to calibrate the phase values of other NICs. Hence, they cannot truthfully remove the phase values to derive precise power delay profiles. Prior works [28, 36] try to synchronize the phases from two consecutive received CSIs via a linear transform. After the transformation, if the two measurements are from the same multipath channel, even the collected CSIs are different due to hardware noises, the transformation on these two CSIs leads to the same result, which could be used as fingerprint for localization. Some recent works aim to explicitly correct CSI phase errors, *e.g.*, MegaMIMO [25]. However, MegaMIMO requires both nanosecond-level synchronization and the access to the raw signal at PHY layer, which are not available on commodity NICs. In summary, existing works cannot directly remove measurement errors from CSIs reported by commodity WiFi NICs, and hence cannot address the challenges we met in Splicer.

Power delay profile based applications. At different locations, the received power delay profiles will be different, which can make them a good choice for the fingerprint-based localization design [29, 39]. On the other hand, the Line-of-Sight information can be directly inferred from the power delay profile [36]. The power level of the LoS path can also be used to ranging between a pair of transmitters [28, 37]. Indoor localization based on the ranging results requires no dense AP deployments, no manual fingerprinting site survey, and no sophisticated AP hardware [28, 37]. For activity recognition, although the detailed relation between the multipath channel variance and the different activities is unknown, recent works propose to learn the inner relation. For example, to detect the existence of human beings [46], to count the number of people moving around [38], to detect human falling down in [9], and recognize different types of human activities [1, 30, 35]. Splicer

can benefit all above applications since we can obtain a wider CSI containing more frequency band information to derive a higher-resolution power delay profile, which more precisely describes the multipath channel.

6. CONCLUSION

This paper presents Splicer to derive precise power delay profiles on commodity WiFi devices. The Splicer design leverages the CSIs measured from individual WiFi bands to obtain the CSI of an equivalent wider WiFi band after the CSI splicing, based on which Splicer can derive high-resolution power delay profiles. The major design challenge is that the CSIs collected from commodity NICs do not merely contain the channel information. They are mixed with rich hardware errors. We propose a series of techniques to address the challenge and also battle the stringent channel coherence time constraint. Our experiments on commodity WiFi NICs report high accuracy of the power delay profiles derived by Splicer. The CSI measurement errors stem from several signal-processing modules in the physical layer of standard 802.11 NICs, including the AGC, signal sampler, and packet detector, so Splicer is general for different WiFi chipsets and independent to the hardware. Due to the channel hopping by Splicer, the normal communication of other users to an AP may be interrupted, which is a limitation of the current design. We believe that such a limitation can be mitigated by employing less-active or dedicated APs for Splicer. In the paper, we use that indoor localization application to showcase that Splicer can immediately benefit existing motion- or location-based systems. In the future work, we will study how the high-resolution power delay profile derived from Splicer can be utilized to support more interesting applications like, object tracking, gesture recognition, and so on.

Acknowledgement

We would like to thank the anonymous reviewers and shepherd for their valuable comments and suggestions that improve the quality of this paper. This work is supported by Singapore MOE AcRF Tier 2 Grant MOE2012-T2-1-070 and NTU Nanyang Assistant Professorship (NAP) Grant M4080738.020.

References

- [1] F. Adib, Z. Kabelac, D. Katabi, and R. C. Miller. 3d tracking via body radio reflections. In *Proc. of USENIX NSDI*, 2014.
- [2] A. Azzalini. A class of distributions which includes the normal ones. *Scandinavian journal of statistics*, 1985.
- [3] R. Crepaldei, J. Lee, R. Etkin, S.-J. Lee, and R. Kravets. CSI-SF: Estimating wireless channel state using CSI sampling and fusion. In *Proc. of IEEE INFOCOM*, 2012.
- [4] T. Felhauer, P. Baier, W. König, and W. Mohr. Optimum spread spectrum signals for wideband channel sounding. *Electronics Letters*, 1993.
- [5] Y. Gao, W. Dong, C. Chen, J. Bu, T. Chen, M. Xia, X. Liu, and X. Xu. Domo: passive per-packet delay tomography in wireless ad-hoc networks. In *Proc. of IEEE ICDCS*, 2014.
- [6] J. Gjengset, G. McPhillips, and K. Jamieson. Arrayphaser: Enabling signal processing on WiFi access points. *Proc. of ACM MobiCom*, 2014.
- [7] A. Goldsmith. *Wireless communications*. Cambridge university press, 2005.
- [8] D. Halperin, W. Hu, A. Sheth, and D. Wetherall. Predictable 802.11 packet delivery from wireless channel measurements. In *Proc. of ACM SIGCOMM*, 2010.

- [9] C. Han, K. Wu, Y. Wang, and L. Ni. WiFall: Device-free fall detection by wireless networks. In *Proc. of IEEE INFOCOM*, 2014.
- [10] J. Han, C. Qian, X. Wang, D. Ma, J. Zhao, P. Zhang, W. Xi, and Z. Jiang. Twins: Device-free object tracking using passive tags. In *Proc. of IEEE INFOCOM*, 2014.
- [11] L. He, L. Fu, L. Zheng, Y. Gu, P. Cheng, J. Chen, and J. Pan. Esync: An energy synchronized charging protocol for rechargeable wireless sensor networks. In *Proc. of ACM MobiHoc*, 2014.
- [12] IEEE Standard for Information Technology—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer Specifications, Mar. 2012.
- [13] S. Jana and S. K. Kaser. On fast and accurate detection of unauthorized wireless access points using clock skews. In *Proc. of ACM MobiCom*, 2008.
- [14] V. Jimenez, M. Fernandez-Getino Garcia, F. Serrano, and A. Armada. Design and implementation of synchronization and agc for ofdm-based wlan receivers. *IEEE Transactions on Consumer Electronics*, 2004.
- [15] B. Kempke, P. Pannuto, and P. Dutta. Harmonia: Wideband spreading for accurate indoor rf localization. In *Proc. of ACM HotWireless*, 2014.
- [16] M. Khalid, Y. Wang, I. Butun, H.-j. Kim, I.-h. Ra, and R. Sankar. Coherence time-based cooperative mac protocol 1 for wireless ad hoc networks. In *EURASIP Journal on Wireless Communications and Networking*, 2011.
- [17] H. Liu, Y. Gan, J. Yang, S. Sidhom, Y. Wang, Y. Chen, and F. Ye. Push the limit of wifi based localization for smartphones. In *Proc. of ACM MobiCom*, 2012.
- [18] Q. Ma, K. Liu, X. Xiao, Z. Cao, and Y. Liu. Link scanner: Faulty link detection for wireless sensor networks. In *Proc. of IEEE INFOCOM*, 2013.
- [19] A. T. Mariakakis, S. Sen, J. Lee, and K.-H. Kim. SAIL: Single access point-based indoor localization. In *Proc. of ACM MobiSys*, 2014.
- [20] A. Molina, P. Fannin, and J. Timoney. Generation of optimum excitation waveforms for mobile radio channel sounding. *IEEE Transactions on Vehicular Technology*, 1995.
- [21] D. Molkdar and P. Matthews. Measurements and characterization of the UHF mobile radio channel. part 1: Measurements over the band 853-885 MHz. *Electronic and Radio Engineers, Journal of the Institution of*, 1988.
- [22] K. Pahlavan and A. H. Levesque. *Wireless information networks*. John Wiley & Sons, 2005.
- [23] J. Parsons, D. Demery, and A. Turkmani. Sounding techniques for wideband mobile radio channels: a review. *Communications, Speech and Vision, IEE Proceedings I*, 1991.
- [24] Q. Pu, S. Gupta, S. Gollakota, and S. Patel. Whole-home gesture recognition using wireless signals. In *Proc. of ACM MobiCom*, 2013.
- [25] H. S. Rahul, S. Kumar, and D. Katabi. Jmb: Scaling wireless capacity with user demands. In *Proc. of ACM SIGCOMM*, 2014.
- [26] T. S. Rappaport et al. *Wireless communications: principles and practice*. prentice hall PTR New Jersey, 1996.
- [27] S. Sen, R. R. Choudhury, and S. Nelakuditi. Spinloc: Spin once to know your location. In *Proc. of ACM HotMobile*, 2012.
- [28] S. Sen, J. Lee, K.-H. Kim, and P. Congdon. Avoiding multipath to revive inbuilding WiFi localization. In *Proc. of ACM MobiSys*, 2013.
- [29] S. Sen, B. Radunovic, R. R. Choudhury, and T. Minka. You are facing the mona lisa: Spot localization using PHY layer information. In *Proc. of ACM MobiSys*, 2012.
- [30] L. Shanguan, Z. Yang, A. X. Liu, Z. Zhou, and Y. Liu. Relative localization of rfid tags using spatial-temporal phase profiling. In *Proc. of USENIX NSDI*, 2015.
- [31] M. Speth, S. Fechtel, G. Fock, and H. Meyr. Optimum receiver design for wireless broad-band systems using OFDM. i. *IEEE Transactions on Communications*, 1999.
- [32] J. K. Tan. An adaptive orthogonal frequency division multiplexing baseband modem for wideband wireless channels. Master's thesis, Massachusetts Institute of Technology, 2006.
- [33] G. Wang, Y. Zou, Z. Zhou, K. Wu, and L. M. Ni. We can hear you with Wi-Fi! In *Proc. of ACM MobiCom*, 2014.
- [34] L. Wang, Y. He, Y. Liu, W. Liu, J. Wang, and N. Jing. It is not just a matter of time: oscillation-free emergency navigation with sensor networks. In *Proc. of IEEE RTSS*, 2012.
- [35] Y. Wang, J. Liu, Y. Chen, M. Gruteser, J. Yang, and H. Liu. E-eyes: Device-free location-oriented activity identification using fine-grained WiFi signatures. In *Proc. of ACM MobiCom*, 2014.
- [36] C. Wu, Z. Yang, Z. Zhou, K. Qian, Y. Liu, and M. Liu. Phaseu: Real-time LOS identification with wifi. In *Proc. of IEEE INFOCOM*, 2014.
- [37] K. Wu, J. Xiao, Y. Yi, M. Gao, and L. Ni. FILA: Fine-grained indoor localization. In *Proc. of IEEE INFOCOM*, 2012.
- [38] W. Xi, J. Zhao, X.-Y. Li, K. Zhao, S. Tang, X. Liu, and Z. Jiang. Electronic frog eye: Counting crowd using WiFi. In *Proc. of IEEE INFOCOM*, 2014.
- [39] J. Xiao, K. Wu, Y. Yi, and L. Ni. FIFS: Fine-grained indoor fingerprinting system. In *Proc. of IEEE ICCCN*, 2012.
- [40] J. Xiong and K. Jamieson. Arraytrack: A fine-grained indoor location system. In *Proc. of USENIX NSDI*, 2013.
- [41] J. Xiong, K. Jamieson, and K. Sundaresan. Synchronicity: Pushing the envelope of fine-grained localization with distributed mimo. In *Proc. of ACM HotWireless*, 2014.
- [42] J. Xiong, K. Sundaresan, and K. Jamieson. Tonetrack: Overcoming bandwidth constraints for indoor wireless localization. In *Proc. of ACM MobiCom*, 2015.
- [43] L. Yang, Y. Chen, X.-Y. Li, C. Xiao, M. Li, and Y. Liu. Tagoram: Real-time tracking of mobile rfid tags to high precision using cots devices. In *Proc. of ACM MobiCom*, 2014.
- [44] L. Yang, J. Han, Y. Qi, C. Wang, T. Gu, and Y. Liu. Season: Shelving interference and joint identification in large-scale rfid systems. In *Proc. of IEEE INFOCOM*, 2011.
- [45] L. Zhang, K. Liu, Y. Jiang, X.-Y. Li, Y. Liu, and P. Yang. Montage: Combine frames with movement continuity for realtime multi-user tracking. In *Proc. of IEEE INFOCOM*, 2014.
- [46] Z. Zhou, Z. Yang, C. Wu, L. Shanguan, and Y. Liu. Omnidirectional coverage for device-free passive human detection. *IEEE Transactions on Parallel and Distributed Systems*, 2014.