On the Observability of Urban Wi-Fi Landscapes Through Passive Vehicular Sensing

Giuliano Prestes Fittipaldi*
Sorbonne Université, CNRS, LIP6
Paris, France
giuliano.fittipaldi@lip6.fr

Anne Fladenmuller*
Sorbonne Université, CNRS, LIP6
Paris, France
anne.fladenmuller@lip6.fr

Rodrigo de Souza Couto◇
GTA/Poli/COPPE
Universidade Federal do Rio de Janeiro
Rio de Janeiro, Brasil
rodrigo@gta.ufrj.br

Luís Henrique Maciel Kosmalski Costa◇
GTA/Poli/COPPE
Universidade Federal do Rio de Janeiro
Rio de Janeiro, Brasil
luish@gta.ufrj.br

Marcelo Dias de Amorim*
Sorbonne Université, CNRS, LIP6
Paris, France
marcelo.amorim@lip6.fr

Mostafa Ammar○
Georgia Institute of Technology
Atlanta, USA
ammar@cc.gatech.edu

***Abstract*—The ease of collecting IEEE 802.11 frames with commodity hardware (i.e., a Wi-Fi adapter connected to an embedded platform) enables the deployment of opportunistic sensing nodes in vehicles. While such a system leverages mobility to increase spatial coverage without relying on dense fixed infrastructure, it raises a central question: *What are the observing capabilities of a highly mobile passive Wi-Fi sensing node in characterizing urban wireless activity?* We address this question through a nine-day measurement campaign using a four-antenna vehicular platform, during which we collected over 14 million Wi-Fi frames across approximately 265 km of urban driving. We analyze device discovery, contact patterns, and frame-level characteristics of the observed traffic. Our results show that device discovery follows a stable linear relationship with captured frame volume and is shaped primarily by the spatial structure of the urban environment rather than vehicle speed. We also show that using multiple antennas helps increase device discovery (by up to 66.5% in our experiments). Contact analysis reveals that while most encounters are brief, many persist over more than one kilometer (4,620 contacts), indicating widespread convoy-like co-mobility. At the frame level, over 80% of captured traffic consists of control and management frames, with highly bursty inter-arrival times consistent across routes. These findings provide a quantitative characterization of the observing capabilities and limitations of vehicular passive Wi-Fi sensing, offering a reference point for the design of future urban sensing systems.**

***Index Terms*—vehicular, passive measurement, mobile sensing Wi-Fi, experimentation.**

I. Introduction

Even without actively participating in a network, Wi-Fi devices continuously generate observable IEEE 802.11 frames that can be passively collected in urban environments. This opens the possibility of leveraging *passive* Wi-Fi sensing as a lightweight, non-intrusive tool to study mobility patterns, urban dynamics, and crowd behavior.

In contrast to approaches that require association or infrastructure participation, passive sensing enables observation of Wi-Fi activity without authentication or data-plane access, relying only on protocol headers. A commodity Wi-Fi adapter connected to a small embedded system, such as a Raspberry Pi, is sufficient to collect large-scale wireless traces in the wild. This simplicity allows for the embedding of measurement nodes inside vehicles without requiring specialized infrastructure or dense deployments of fixed sensing devices. Instead of relying on a stationary network of monitors, a small number of mobile vehicular sensors can collect measurements as they move. This is particularly interesting in the context of urban sensing, since prior work has shown that mobility increases spatial coverage compared to stationary deployments [1].

The interaction between vehicular mobility and passive Wi-Fi observability has not been systematically characterized. In particular, it is not well understood how many devices a moving observer can detect, how stable these observations are across space and time, what opportunities the collected data offer, and what structural properties of the wireless environment can be inferred from such measurements. In such a context, several questions remain open: (i) What is the rate of Wi-Fi device discovery from a vehicular passive sensing node? (ii) How consistent are device-level observations across repeated traversals of the same route, and do they exhibit stable spatial structure over time? (iii) What are the spatiotemporal characteristics of device contact events observed along the vehicular route? (iv) How do contact durations vary with vehicle speed, and to what extent do certain devices exhibit repeated co-observation patterns? (v) What are the distributions and temporal characteristics of observed Wi-Fi frames, including frame types, directions, and inter-arrival behavior?

In this paper, we address these questions through a real-world measurement study. We use a vehicle equipped with multiple Wi-Fi antennas and a GPS logging system. Over a period of nine days (five consecutive and four non-consecutive), we collected data over approximately 265 km of urban driving. Our sensing platform passively captured more than 14 million IEEE 802.11 header frames as it traversed diverse environments, including dense residential streets, highways,

979-8-3195-0167-7/26/$31.00 ©2026 IEEE



(a) Capture device. (b) Windshield antenna. (c) Side windows antennas. (d) Back window antenna.

Fig. 1. The hardware within the vehicle.

and a university campus, thereby exposing a range of device densities, mobility conditions, and radio environments.

First, we analyze device discovery patterns from a mobile vantage point, finding that the average number of newly observed devices per 100 m is 60, increasing up to 120 in key transportation zones in the city. We observe that the discovery rate is primarily affected by the measurement region rather than by vehicle speed. We further observe that devices discovered per number of frames exhibit measurable consistency across days on the same route and can be approximated by a linear model, with our two routes yielding a new MAC address after 61 and 70 frames, respectively.

Second, the median contact duration ranges from 22 s to 32 s, while the median contact distance ranges from 229 m to 270 m. When analyzed by speed, contact distance peaks at around 60 km/h, reaching over 600 m, and drops sharply at higher speeds. Contact duration follows a similar trend, peaking between 10 km/h and 30 km/h at around 72 s and remaining above 50 s up to 60 km/h, beyond which it also decreases rapidly. Finally, we identify a subset of devices exhibiting convoy-like behavior, with 4,620 contacts maintained over more than 1 km with our vehicle.

Finally, we analyze frame-level characteristics of the captured traffic. We observe that passive vehicular WiFi sensing primarily observes high-frequency protocol overhead and network maintenance traffic, rather than user-generated payload traffic. Among management, control, and data frames, we find that control frames dominate the observed distribution, with frame-type proportions of around 50%. Regarding transmitter roles, we find that approximately 80% of classified devices behave as stations, while the remaining 20% behave as access points. This suggests that a passive vehicular observer predominantly encounters client devices rather than infrastructure nodes, reflecting both the higher density of user devices in urban environments and the large volume of uplink and discovery traffic generated by modern Wi-Fi clients. Finally, we observe highly bursty traffic patterns, with median inter-frame arrival times below 3 ms across all frame categories.

These results provide a first-order characterization of what can be observed with a highly mobile passive Wi-Fi sensing platform and highlight the opportunities for using vehicular sniffing in large-scale urban sensing applications.

II. Related Work

Several studies have leveraged probe requests and public Wi-Fi infrastructure for large-scale urban monitoring. Alorda-Ladaria et al. used public Wi-Fi telemetry to evaluate social distancing measures during the COVID-19 pandemic and derive city-scale mobility indicators from passive observations [2]. Similarly, Simara et al. used Wi-Fi signals emitted by mobile devices to estimate pedestrian flow and density in public spaces [3]. More recently, Wi-Fi telemetry has also been explored for crowd management applications in smart campuses and public environments [4].

Mobile sensing platforms also yield improved spatial coverage and lower deployment costs. Bertiz et al. integrated a mobile node into an urban wireless sensor network to collect geolocated measurements as it traversed heterogeneous environments, including highways and residential areas [5]. Similarly, Aerts et al. deployed RF exposure sensors aboard postal vehicles to characterize electromagnetic exposure over large urban areas that would be impractical to monitor using fixed installations alone [6].

Distributed opportunistic sensing has also been explored using smartphones and crowdsourced measurements. Villadangos et al. proposed a distributed Wi-Fi mapping system based on opportunistic measurements [7], while Han et al. used crowdsourced fingerprints to construct city-scale radio maps for localization purposes [8]. Although these approaches achieve wide coverage, they require either large numbers of participants or dedicated infrastructure support.

A substantial body of work has studied Wi-Fi measurements for localization, fingerprinting, and wireless sensing applications. Urban fingerprinting systems have demonstrated that Wi-Fi signal characteristics can exhibit considerable spatial structure and temporal stability under repeated traversals [8], [9]. More broadly, recent surveys highlight the growing role of commodity Wi-Fi sensing in smart environments, localization systems, and context-aware applications [10], [11]

Our work combines urban sensing, mobile wireless measurement, and passive Wi-Fi sniffing. Unlike infrastructure-based monitoring systems, we do not require fixed deployments or participation from existing networks. Unlike crowdsourced radio mapping systems, we do not assume large user

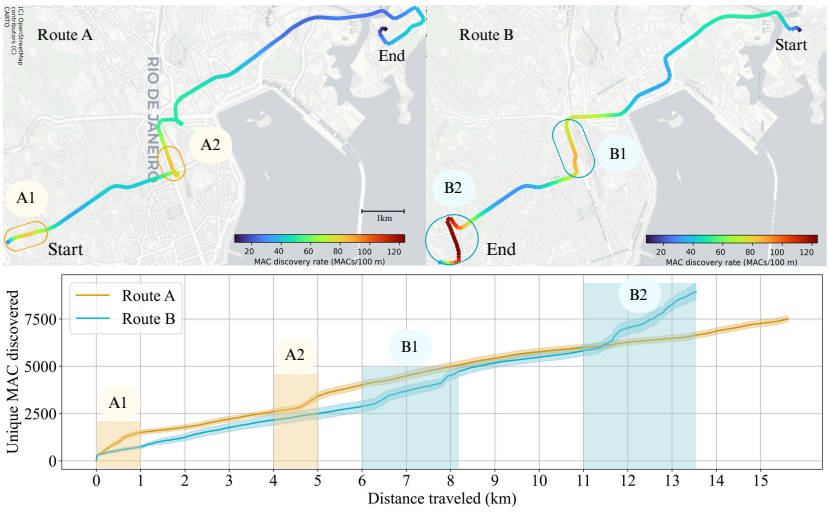


Fig. 2. **Number of unique MAC addresses discovered by distance traveled and rate of discovery along both routes.**

populations or cooperative measurements. Instead, we provide a measurement-driven characterization of what a single vehicular passive sensing node can observe in a real urban environment. Through repeated traversals over multiple days and routes, we study device discovery dynamics, contact behavior, and frame-level properties, providing insights into both the capabilities and limitations of vehicular passive Wi-Fi sensing for future urban sensing systems.

III. Methodology and Experimental Data

Setup. It consists of Raspberry Pi 4 (RPi) Model B boards connected to Alfa AWUS051NH USB adapters with omnidirectional antennas (Figure 1(a)). Traffic is captured in monitor mode using `tcpdump`, while GPS positions are provided by an EVK-5H module (2.5 m CEP). Four antennas are mounted at the vehicle extremities (Figures 1(b)–1(c)) to improve coverage and reduce body-shadowing effects. The collected traces are synchronized, georeferenced, and processed offline.

Vehicle Trajectory. The routes span approximately 16 and 13.5 km, requiring 17-40 minutes to complete, depending on traffic conditions. The routes provide a representative cross-section of urban Wi-Fi infrastructure deployment patterns across different contexts within the city [12]: narrow residential streets (10-30 km/h), faster residential streets (up to 60 km/h), a highway section (above 60 km/h), and the university campus (20-30 km/h). This variety enables analysis of Wi-Fi signal characteristics under different conditions. Residential areas reflect dense access point deployments, highways capture high-mobility signal propagation, and the university campus provides insights into high-density institutional network usage.

Dataset Summary. Our dataset includes both Wi-Fi and GPS measurements collected over multiple experimental runs as summarized in Table I. The dataset contains almost 14 million Wi-Fi frames captured. The total distance covered across all experimental runs is approximately 265 kilometers, with the primary route repeated on 9 separate days over a two-month period to capture temporal variations in network activity. Table II details the specific Wi-Fi frame fields considered.

TABLE I
Dataset characteristics and collection parameters.

Parameter	Value
Total frames captured	14,256,670
Unique MAC addresses	94,227
Unique BSSIDs	20,469
Total distance covered	265 km
Number of experimental runs	18
Collection period	July - August, 2025
Number of monitoring antennas	4
GPS sampling frequency	5 Hz
GPS nominal precision	2.5 m
Antenna gain	5 dBi

TABLE II
Wi-Fi frame contents used in the analyses.

Feature	Description
`frame.time_epoch`	Packet capture timestamp
`wlan.fc.type`	Frame type
`wlan.fc.type_subtype`	Frame subtype
`wlan.ta`	Transmitter address
`wlan.bssid`	Basic Service Set Identifier
`wlan.ssid`	Service Set Identifier
`wlan_radio.signal_dbm`	Received Signal Strength Indicator
`wlan.fc.ds`	Frame direction to or from AP

During collection, raw data is stored locally on each RPi and later transferred via SSH to a central server. For the post-processing step, Wi-Fi captures (pcap files) are processed with `tshark` to extract CSV files, while GPS logs are parsed for coordinates, timestamps, and satellite info. GPS data is projected into a local coordinate system for distance computations. Wi-Fi frames are synchronized with GPS coordinates via nearest-neighbor timestamps. Data cleaning removes MAC addresses observed in fewer than 10 frames and GPS points with fewer than 4 satellite fixes to keep reliability.

IV. Device Discovery

This section investigates the MAC discovery process using the vehicular sensing node. We first characterize how unique

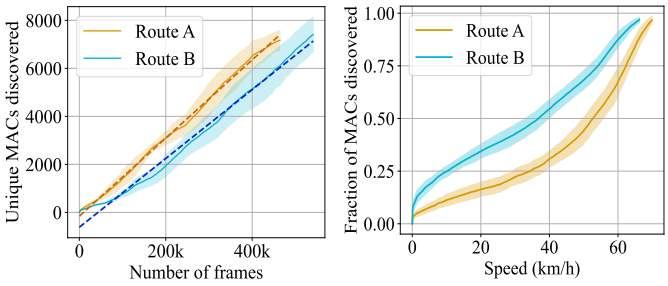
devices are discovered along the route and whether this process exhibits spatial structure. We then analyze the relationship between device discovery and the volume of captured Wi-Fi frames, as well as the influence of vehicle speed. Next, we evaluate the persistence of observed MAC addresses and BSSIDs across days. Finally, we assess the impact of the number of antennas in the capture process.

Figure 2 characterizes the spatial structure of MAC discovery. The bottom plot shows the cumulative number of unique MAC addresses discovered along the routes, averaged across days. The shaded regions represent the 95% confidence intervals. The maps above the curves depict the discovery rate, measured as the number of newly observed MAC addresses every 100 m. For both routes, MAC discovery follows an approximately linear trend, with fewer than 60 new MAC addresses observed per 100 m. However, several route segments exhibit noticeable increases in slope, indicating localized increases in discovery rate. These regions, labeled A1 and A2 for Route A, and B1 and B2 for Route B, correspond to discovery rates of 80-120 MAC addresses per 100 m.

The spatial distribution of hotspots suggests that device discovery is not uniform. Interestingly, A2 and B1 correspond to opposite directions along the same avenue (Av. Presidente Vargas), a major transportation corridor with high population density that connects one of the city's metro lines to several urban and suburban bus lines. Likewise, A1 and B2 are located in residential neighborhoods (Laranjeiras and Flamengo), with higher pedestrian activity and denser residential and commercial infrastructure. These areas likely concentrate a larger number of Wi-Fi-enabled devices, ranging from personal electronics to residential and commercial equipment.

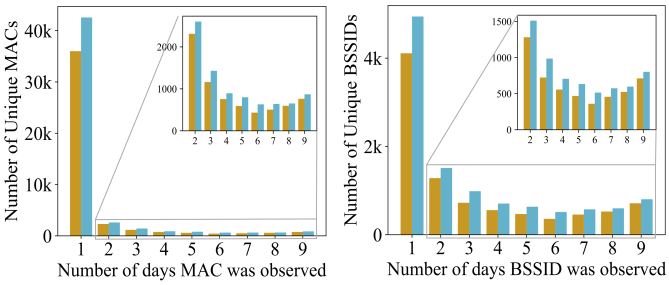
The existence of hotspots raises an important question: are they simply the result of increased frame capture volume, or do they reflect a genuinely larger population of distinct devices? To investigate this relationship, Figure 3(a) presents the cumulative number of unique MAC addresses discovered as a function of the number of captured frames. The curves are averaged across days and accompanied by confidence intervals. Despite the heterogeneous transmission behavior of Wi-Fi devices, a linear model captures the dominant trend of MAC discovery, as shown by the dotted regression lines. For Route A, the fitted slope is 0.016281 MAC/frame $\pm$ 0.000456, corresponding to approximately one new MAC address every 61 captured frames. For Route B, the fitted slope is 0.014330 MAC/frame $\pm$ 0.001003, corresponding to approximately one new MAC address every 70 captured frames. These results indicate a stable relationship between traffic volume and device discovery, with Route A consistently yielding more unique devices per captured frame. This suggests that hotspots produce more frames and more distinct devices in roughly equal proportion, rather than concentrating a disproportionately larger population of devices.

A potential concern is that vehicle speed may influence the device discovery process: at lower speeds, the vehicle spends more time within radio range of nearby devices, which could artificially inflate the discovery rate in slow residential segments relative to faster road sections. To evaluate this effect, Figure 3(b) presents the cumulative fraction of discovered MAC addresses as a function of vehicle speed.



(a) **MACs per captured frames.** (b) **MACs across speed CDF.**

Fig. 3. **(a) MAC discovery per volume of captured frames averaged over 9 days with confidence interval and linear fit. (b) Fraction of MACs discovered across speeds averaged over 9 days with confidence interval.**



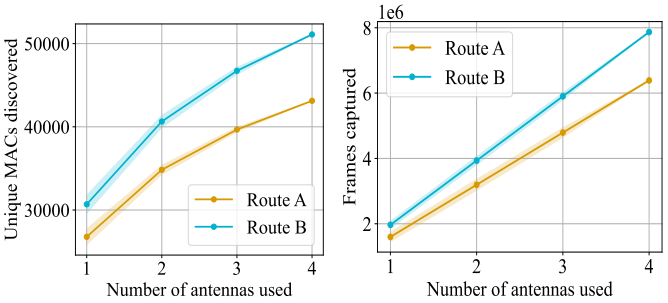
(a) **MAC persistence.** (b) **BSSID persistence.**

Fig. 4. **Discovery persistence.**

For Route B, MAC discoveries are distributed across the entire speed range, with the median discovery occurring at 36 km/h and a maximum average speed of 63 km/h. Route A exhibits a similar behavior, although shifted toward higher speeds, with the median discovery occurring at 51 km/h and a maximum average speed of 70 km/h. This difference reflects the distinct mobility profiles of the two routes: Route A has an average speed of 46 km/h, whereas Route B averages 34 km/h.

Taken together, these results suggest that vehicle speed has limited influence on the probability of discovering new devices. Faster routes reduce the amount of time spent within the coverage area of individual devices and consequently reduce the total number of discoveries per kilometer, as observed in Figure 2. However, within the range of speeds observed in our experiments, speed does not substantially alter the relationship between traffic volume and MAC discovery. Although higher speeds shorten the time spent within the communication range of individual devices, the resulting contacts typically remain long enough to capture the relatively small number of frames required to identify a new MAC address.

Figure 4 evaluates the persistence of observed MAC addresses and BSSIDs across the measurement campaign. Figure 4(a) shows, for each value on the x-axis, the number of distinct MAC addresses observed on exactly that many days out of the nine measurement days, regardless of which specific days they appeared on. Figure 4(b) presents the same analysis



(a) **MACs per antenna.** (b) **Frames per antenna.**

Fig. 5. **(a) MAC discovery and (b) frame capture volume as a function of the number of antennas.** Measuring with four antennas instead of one increases MAC discovery by approximately 61% and 66.5% on Routes A and B, respectively, while frame capture volume scales approximately linearly.

for BSSIDs. The first observation is the substantial difference in cardinality between the two populations: the number of unique MAC addresses is approximately ten times greater than the number of observed BSSIDs. This is expected, as MAC addresses represent individual devices, whereas BSSIDs primarily correspond to network infrastructure.

Most MAC addresses are observed on a single day only. The number of devices decreases sharply as the number of observation days increases, reaching a minimum around six days before increasing again. A similar pattern happens for BSSIDs. This behavior is consistent with the coexistence of two distinct populations: transient devices that appear intermittently along the route and infrastructure-related devices that remain stationary. The latter includes residential routers, enterprise networks, and other fixed Wi-Fi infrastructure.

Finally, we evaluate the impact of the sensing hardware on device discovery. Rather than relying on a single Wi-Fi adapter, our platform employs four independent antennas. Figure 5(a) shows the resulting increase in MAC discovery, while Figure 5(b) presents the corresponding frame capture volume. Using four antennas instead of one increases MAC discovery by approximately 61% on Route A and 66.5% on Route B. As expected, the number of captured frames grows approximately linearly with the number of antennas. These results indicate that additional sensing interfaces translate directly into a larger observable device population, highlighting the importance of hardware configuration when designing vehicular Wi-Fi measurement platforms.

In summary, MAC discovery exhibits a clear spatial structure, with hotspots consistently producing higher discovery rates than elsewhere along the route. The relationship between captured frames and discovered devices is well approximated by a linear model, yielding approximately one new MAC address every 61 frames on Route A and every 70 frames on Route B. While device discovery occurs across the full range of observed speeds with little variation, contact duration and distance are strongly affected by speed, indicating that, for typical urban driving speeds, mobility influences the characteristics of individual contacts more strongly than the probability of discovering new devices. Persistence analysis reveals two distinct populations: a large set of transient devices observed only once and a smaller set of recurring devices and networks consistently detected across days. Finally, increasing the sensing platform from one to four antennas increases MAC discovery by up to 66.5%, demonstrating the direct impact of hardware redundancy on the observable Wi-Fi environment.

V. Device Contacts

Device discovery alone does not fully characterize the sensing capabilities of a vehicular Wi-Fi measurement platform. For many applications, it is equally important to understand how long a sensed device remains observable and how far a sensing node can travel while maintaining contact with it. These properties determine the temporal and spatial window available for exchanging information, tracking mobility patterns, or inferring correlations between devices.

A. Contact Time

We define a contact as the interval between the first and last frame received from a device. Contact duration, therefore, represents the amount of time a device remains observable by the sensing platform. Figure 6(a) presents the CDF of contact durations. Most contacts are relatively short, reflecting the highly dynamic nature of urban mobility. For Route A, the median contact duration is 22.29 s and the mean is 45.23 s. For Route B, the median is 32.01 s and the mean is 55.15 s.

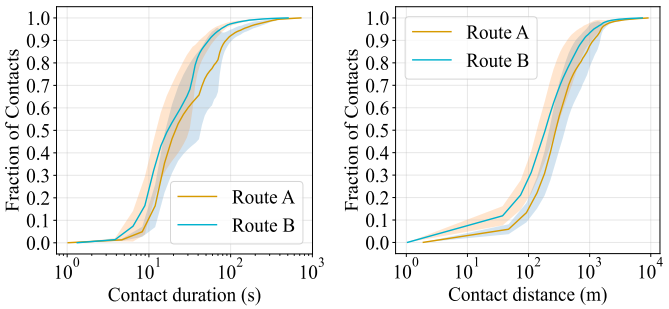
The distributions exhibit a long tail. The longest contacts reach approximately 20 minutes in Route A and 29 minutes in Route B. These extended contacts likely correspond to devices that move along the same transportation corridor as the sensing vehicle and remain within radio range for a substantial portion of the route. This interpretation is plausible given that both routes connect major commuting paths between the city's southern region and the university campus.

The difference between median and mean values highlights the coexistence of two distinct contact regimes. Most contacts correspond to short drive-by encounters, such as pedestrians crossing the street or vehicles traveling in the opposite direction. A smaller fraction of contacts persist for several minutes, significantly increasing the average duration. These longer interactions are relevant to applications that require sustained observability, as they offer substantially more opportunities for communication or sensing than typical drive-by encounters.

B. Contact Distance

While contact duration quantifies how long a device remains observable, contact distance measures how far the sensing vehicle travels before losing contact. This metric provides a complementary perspective on the spatial extent of wireless encounters. Figure 6(b) shows the distribution of contact distances. Similar to contact duration, the distribution is highly skewed. For Route A, the median contact distance is 269.70 m and the mean is 479.33 m. For Route B, the median is 229.70 m, and the mean is 398.13 m.

Most contacts occur over a few hundred meters, indicating that we encounter a large fraction of devices only briefly



(a) **Contact time CDF.** (b) **Contact distance CDF.**

Fig. 6. **the (a) CDF of contact duration between devices and the observation node. (b) CDF of contact distance, with GPS uncertainty of $\pm 14.7\%$.** Both CDFs resemble a normal distribution. The average contact duration is 45.23 for Route A and 55.15 s for Route B. Contacts persist for up to 29 minutes. The average contact distance is 479.33 m for Route A and 398,13 m for Route B, while the longest contact extends nearly 14 km.

as we traverse the city. Nevertheless, a small number of contacts extend for several kilometers, with the longest approaching 14 km. These long-distance contacts are unlikely to result solely from propagation conditions and instead suggest correlated mobility between the observed device and the sensing vehicle. Examples include vehicles traveling in the same traffic flow, passengers following the same commuting corridor, or devices located along extended road segments with uninterrupted visibility. From a sensing perspective, such contacts provide opportunities for repeated observations of the same device across large portions of the route, enabling richer mobility inference than isolated encounters do.

C. Convoy-like Behavior

The contact-distance distribution tail motivates a closer examination of devices that remain observable over extended distances. We define a convoy as a device that we continuously observe for at least 1 kilometer, a conservative threshold chosen to isolate long-lived contacts that are unlikely to result from radio propagation alone. Given the non-ideal propagation conditions in outdoor Wi-Fi environments, maintaining contact over this distance strongly suggests that the device is moving in the same direction as the sensing vehicle and remains geographically close for an extended period.

Figure 7 illustrates these convoy contacts for both routes. Each horizontal line represents a device, while the x-axis indicates the length of the route segment during a contact. Although convoy contacts represent a small subset of the total number of devices discovered in Section IV, they are far from rare. Route A contains 2,627 convoy devices, while Route B contains 1,993. In total, 4,620 contacts persist for more than 1 km, representing 10% of all contacts.

Several convoy trajectories terminate at specific locations along the routes. In Route A, a large number of convoys end near the 5-kilometer mark, while in Route B, similar behavior appears near 8 km. These locations correspond to the opposite ends of the previously identified transportation corridor, regions A2 and B1 in Figure 2. A second concentration of convoy terminations appears near the 11-kilometer mark of Route B, corresponding to the exit of a tunnel where traffic flow disperses toward multiple destinations.

These observations suggest that convoy analysis can reveal structural properties of the transportation network and identify locations where traffic streams merge or diverge. Such information may be useful for applications relying on opportunistic encounters, where prolonged co-mobility directly affects communication and sensing opportunities.

D. Impact of Speed

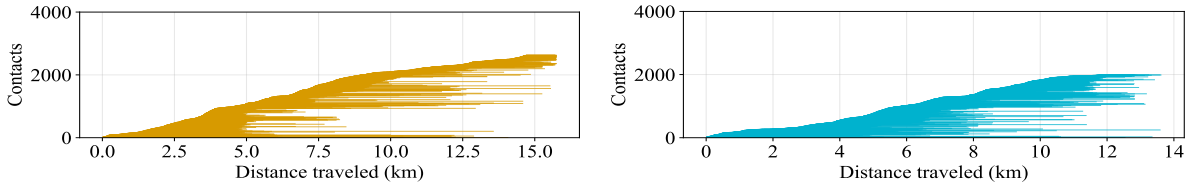
The previous analyses characterized contact behavior independently of vehicle speed. However, speed is an important operational parameter for vehicular sensing systems, as it determines both the rate at which the environment is sampled and the time available to maintain contact with nearby devices. Figure 8(a) shows the fraction of contacts observed across different speeds. Contacts are distributed across almost the entire speed range up to $\approx$70 km/h. This behavior is consistent with the driving profile of the measurement campaign, where most of the route was traversed below this threshold. Contacts observed at higher speeds likely correspond to devices traveling alongside the sensing vehicle, such as smartphones and infotainment systems in nearby cars or buses.

Figure 8(b) shows the average contact distance as a function of speed. Contact distance increases up to approximately 60 km/h, indicating that the sensing platform can maintain contact reliably within this range. Beyond 60 km/h, contact distance decreases substantially, and the variability of the experiment increases. This suggests that, at higher speeds, contacts increasingly resemble short drive-by encounters and that the practical observation range of Wi-Fi devices in our dataset saturates near 600 m. Figure 8(c) provides a complementary view through contact duration. Up to 60 km/h, average contact durations remain above 40 s, reaching approximately 73 s near 20 km/h. These longer contacts are consistent with the traffic conditions commonly observed in residential and congested urban areas. Above 60 km/h, contact duration decreases sharply, reaching values as low as 12 s.

These results indicate that the sensing platform is effective over a broad range of urban speeds but that high-speed segments fundamentally alter the nature of observed contacts. While device discovery remains possible, contacts become increasingly transient and provide a much smaller temporal window for observation. Consequently, applications requiring sustained interaction with nearby devices should account for this limitation when operating on high-speed road segments.

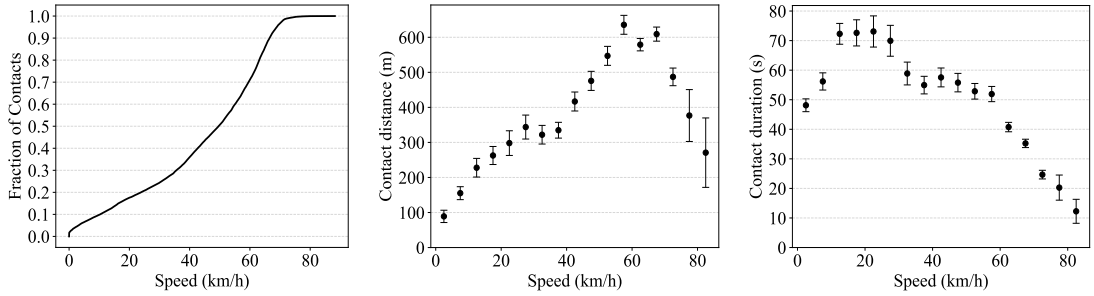
VI. Frame Type Analysis

In this section, we investigate the nature of the WiFi traffic observed by a passive vehicular sensing platform. In particular, we seek to understand which types of devices are discovered, how the vehicle interacts with them, and which types of traffic dominate the collected traces. These characteristics directly determine the types of analyses that can be supported by



(a) Route A. (b) Route B.

Fig. 7. **Contacts lasting more than 1 kilometer, characterizing convoy-like behavior.** A total of 4,620 contacts remain observable for more than 1 kilometer, indicating sustained co-mobility between the sensing node and the observed device.



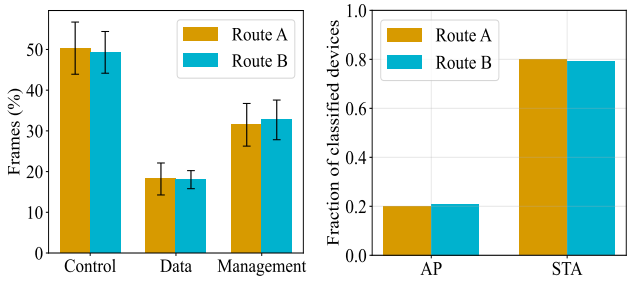
(a) Fraction of Contacts vs. Speed CDF. (b) Contact Distance vs. Speed. (c) Contact Time vs. Speed.

Fig. 8. **Impact of vehicle speed on contact characteristics across the 265 km of driving.**

mobile Wi-Fi sensing, such as infrastructure discovery, traffic inference, or user mobility analysis.

Figure 9(a), where the y-axis represents the average percentage of captured frames per route, and Table III show that the composition of captured traffic is remarkably consistent across routes, suggesting that the observed traffic mix is independent of the urban environments traversed by the vehicle. Approximately half of all captured frames are control frames, which mainly consist of RTS, CTS, ACK, and Block ACK exchanges used for medium access coordination and reliable delivery, but which contain little semantic information regarding users. Management traffic accounts for roughly one third of the captured frames and is dominated by beacons and probe responses, making passive vehicular sensing particularly effective for infrastructure discovery and characterization. In contrast, only about 18% of the frames carry user payload, indicating that the traffic observed by the sensing platform is dominated by protocol maintenance and network discovery activity rather than user data transmissions.

Figure 9(b) shows the fraction of transmitter MAC addresses classified according to their inferred device role. Devices transmitting beacon frames are labeled as access points, while devices transmitting probe requests are labeled as stations. For devices that do not transmit either subtype, we infer their role using the DS flags observed in their data frames, where "To DS" indicates a station transmitter and "From DS" indicates an access point transmitter. Despite this classification relying on the presence of beacon, probe request, or data transmissions, the method successfully classifies 93% and 92% of observed transmitters on routes A and B, respectively. Across both routes, the observed wireless environment is dominated by station devices, which account for 79.9% and 79.1% of classified transmitters, while access points represent 20.1% and 20.9%. This result suggests that a passive vehicular observer primarily encounters client devices rather than infrastructure nodes, reflecting both the higher density of user devices in

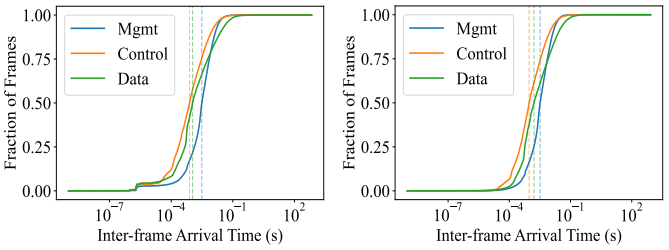


(a) Frame types. (b) Device roles.

Fig. 9. **(a) Percentage of frame types in the data. (b) Device role indicated by data type and frame direction.**

TABLE III
FRAME SUBTYPE DISTRIBUTION.

Subtype	**Route A (%)**	**Route B** (%)
Beacon	22.31	25.04
CTS	13.31	13.05
RTS	12.96	11.69
ACK	11.96	13.05
Probe Response	9.00	8.88
Block ACK	8.19	6.99
Data	5.26	5.46
Null Data	4.63	4.14
QoS Data	3.97	3.87
QoS Null	2.90	2.07
Probe Request	2.34	2.11
Block ACK Request	1.60	2.32
Action	0.87	0.63
CF-End	0.57	0.45
Authentication	0.08	0.12



(a) Route A. (b) Route B.

Fig. 10. **Interframe arrival time.** Passive Wi-Fi frames observe, on average, 10 ms between frames.

urban environments and the large volume of probe and uplink traffic generated by modern Wi-Fi clients.

Figure 10 depicts the inter-frame arrival time distributions, per MAC, distinguishing management, control, and data traffic. Control frames exhibit the shortest inter-arrival times, with median values below 1 ms in both routes (0.78 ms and 0.93 ms), reflecting the rapid exchange of acknowledgments and channel access coordination messages required by the IEEE 802.11 MAC protocol. Management frames arrive less frequently, with median intervals close to 3 ms and effective rates of approximately 300 frames/s, consistent with periodic beacon transmissions and discovery procedures. Data traffic presents intermediate median inter-arrival times (1.10 ms and 1.63 ms) but significantly heavier tails, with 99th percentile delays exceeding 120 ms, revealing the bursty and application-dependent nature of user traffic. Despite the different environments traversed by the vehicle, both routes exhibit similar inter-arrival distributions, suggesting that protocol dynamics play a major role in shaping the temporal structure of the observed Wi-Fi traffic. Together with the frame type distributions, these results indicate that passive vehicular WiFi sensing primarily observes high-frequency protocol overhead and network maintenance traffic, while data frames account for a smaller fraction of the observed frame count and exhibit more intermittent transmission patterns.

VII. Conclusion

In this paper, we presented a measurement-driven characterization of passive vehicular Wi-Fi sensing based on more than 14 million IEEE 802.11 frames collected over 265 km of urban driving in Rio de Janeiro. Through analyses of device discovery, contact dynamics, and frame-level traffic characteristics, we identified some key observability properties of a highly mobile Wi-Fi sensing platform.

Our results show that device discovery is closely tied to captured traffic volume and is driven primarily by the spatial structure of the urban environment rather than by vehicle speed. While most observed devices are encountered only briefly, thousands of contacts persist over distances exceeding one kilometer, revealing substantial opportunities for mobility inference and the identification of co-moving devices. We also show that increasing the number of sensing antennas directly expands the observable device population, highlighting the importance of hardware redundancy in the design of mobile sensing platforms.

At the frame level, we find that passive vehicular sensing predominantly captures management and control traffic rather than user-generated data. This observation suggests that such platforms are well-suited to characterizing wireless infrastructure, device presence, and mobility patterns, but less appropriate for inferring application-level communication behavior. The consistency of these observations across routes and days further indicates that many properties of the urban Wi-Fi landscape exhibit stable large-scale structure despite the highly dynamic nature of individual wireless encounters.

Taken together, our findings establish both the capabilities and the limits of passive vehicular Wi-Fi sensing. Beyond providing a quantitative reference for future measurement studies, they offer practical guidance for the design of urban sensing systems. For example, the observed stability of device discovery and recurring wireless landmarks can support infrastructure mapping and roadside environment characterization, while long-duration convoy contacts provide opportunities for identifying co-moving vehicles in applications such as opportunistic networking and vehicular federated learning.

Acknowledgments

This work was financed in part by the CAPES-COFECUB “Convergencia” project (88887.987121/2024-00, 88881.987122/2024-01, and Ma1074/25).

References

[1] B. Liu, P. Brass, O. Dousse, P. Nain, and D. Towsley, “Mobility improves coverage of sensor networks,” in *ACM Mobihoc*, May 2005.

[2] B. Alorda-Ladaria, M. Ruiz-Pérez, and V. Ramos, “A near real-time monitoring system using public wi-fi data to evaluate covid-19 social distance measures,” *Electronics*, 2022.

[3] E. Šimara, P. Kilnarová, J. Palacký, and R. Vašut, “Using wi-fi signals from mobile devices to determine characteristics of pedestrian behavior in public spaces,” *Transportation Research Record*, 2020.

[4] A. Variskhanov and T. Arreeras, “Early peak badges from wi-fi telemetry: A field feasibility study of lunchtime crowd management on a smart campus,” *Urban Science*, 2026.

[5] C. A. S. Bertiz, J. J. F. Lozano, J. A. Gómez-Ruiz, and A. García-Cerezo, “Integration of a mobile node into a hybrid wireless sensor network for urban environments,” *Sensors (Basel, Switzerland)*, 2019.

[6] S. Aerts, G. Vermeeren, M. V. D. Bossche, R. Aminzadeh, L. Verloock, A. Thielens, P. Leroux, J. Bergs, B. Braem, A. Philippron, L. Martens, and W. Joseph, “Lessons learned from a distributed rf-emf sensor network,” *Sensors (Basel, Switzerland)*, 2022.

[7] J. Villadangos, F. Falcone, A. Lopez-Martin, J. J. Astrain, P. Sanchis, and I. R. Matías, “Distributed opportunistic wireless mapplicationing system towards smart city service provision,” *2021 IEEE Sensors*, 2021.

[8] D. Han, A. Sahar, J. Berkner, and G. Lee, “City radio map construction for wi-fi-based citywide indoor positioning,” *IEEE Access*, 2019.

[9] X. Tong, H. Wang, X. Liu, and W. Qu, “Mapfi: Autonomous mapping of wi-fi infrastructure for indoor localization,” *IEEE Transactions on Mobile Computing*, 2021.

[10] Y. Yang, M. Chen, Y. Blankenship, J. Lee, Z. Ghassemlooy, J. Cheng, and S. Mao, “Positioning using wireless networks: Applications, recent progress, and future challenges,” *IEEE Journal on Selected Areas in Communications*, 2024.

[11] F. Wang, T. Zhang, W. Xi, H. Ding, G. Wang, D. Zhang, Y. Cui, F. Liu, J. Han, J. Xu, and T. X. Han, “A survey on wi-fi sensing generalizability: Taxonomy, techniques, datasets, and future research prospects,” *IEEE Communications Surveys & Tutorials*, 2025.

[12] B. Motte-Baumvol and C. D. Nassi, “Immobility in Rio de Janeiro, beyond poverty,” *Journal of Transport Geography*, 2012.