Documents in the Air

Dragoș Ștefan Niculescu, Laura Ruse, Costin Carabaș
University Politehnica of Bucharest, Romania
Department of Computer Science and Engineering
ORCID: 0000-0001-7590-2015, 0000-0002-3756-6142, 0000-0003-4078-202X

2022 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops) | 978-1-6654-1647-4/22/$31.00 ©2022 IEEE | DOI: 10.1109/PerComWorkshops53856.2022.9767318

***Abstract*—$AirDocs$ is a middleware system that allows placing and retrieving objects or documents at different indoor locations without requiring a positioning system. We show how existing WiFi/BLE infrastructure can be used to create unique place signatures that can serve for indexing a collection of documents. The middleware enables many location centric applications, and relies only on an app installed on the smartphone, and a server placed in the intranet or cloud.**

***Index Terms*—WiFi, fingerprint, indoor location, context, middleware, $AirDocs$.**

I. Introduction

As environments become very cluttered in schools, homes, and institutions, one of the challenges will be in how we manage these collections of Internet connected objects. Operating and maintaining physical object databases comes with significant challenges in management of devices, topologies, interoperation, security, privacy, portability, and context awareness. This concept of **context** is heavily overloaded, but is in fact quite central for pervasive computing and IoT, so that ongoing research still requires extensive surveying effort and building of taxonomies [1]–[4].

Position is certainly a context, but obtaining it, especially indoors requires nontrivial effort. Indoor location will no doubt play an important role in the quest for context awareness, but research in locating and tracking devices, which has ramped up significantly in the last decade [5], [6], has shown that obtaining indoor location is costly in several ways: necessary infrastructure, specialized hardware(UWB [7], LiDAR [8]), instrumentation of the environment, low accuracy (WiFi, BLE), effort (training and maintaining location databases), battery consumption (GPS, WiFi, 4G methods), maintenance of privacy (Google tracking). Even if GPS were fully available indoors, its continuous daily use would place significant battery usage. In contrast, this project proposes a method to **associate data to indoor locations without the need to obtain and maintain positions** in a Cartesian space. Instead, the system relies on wireless measurements (WiFi, 3G-5G, Bluetooth) are monitored permanently by all mobile devices as part of their normal functioning, requiring no additional gathering effort. Collecting this information in a **signature of the context**, which includes the WiFi fingerprint and other context specific information, enables retrieval of information based on signatures.

Work supported in part by the Romanian National Authority for Scientific Research and Innovation, UEFISCDI project PN-III-P2-2.1-PED-2019-5413

$AirDocs$ is a middleware system that uses this context specific signature to allow placing and retrieving objects or documents at different indoor locations without requiring actual position. It relies on WiFi/Bluetooth infrastructure existing in most homes and institutions, requires only a single additional intranet/cloud server, and a mobile app that can be installed on any Android mobile device. $AirDocs$ enables many applications that involve natural placing and retrieving of documents at locations.

The system is akin to augmented reality with the users having the illusion of the documents being spread in the physical environment, visible only at certain locations. Leaving a document "in the air" allows for a natural way to use it as a wireless post-it for museums explanations, maps directions in airports and malls, advertising, lab door announcements, restaurant menus, office pin-boards, refrigerator post-its, general reminders, and notices around the house and office. Many of these applications would usually require location, but if the $AirDocs$ service is available, the functionality of placing and retrieving documents can be used right away requiring only the installation of the Android application.

Contact tracing [9], [10] has recently seen a surge of interest, and has similar requirements with our system: no additional infrastructure, and simple operation with existing smartphones. $AirDocs$ explores the same idea of proximity based on dissimilarity, and can be used as support for a contact tracing app, since the 1m-4m proximity detection is within range of current health advisories.

II. System Architecture

The $AirDocs$ architecture is represented in Figure 1. The middleware provides an API for scanning Wi-Fi APs, cellular networks, Bluetooth Low Energy (BLE) devices, GPS information, and sound, in order to build signatures. Also, it includes methods for sending documents to the server along with the associated signature, and for retrieving documents from the server for a recently collected signature. This middleware can then be used by actual applications in order to store and retrieve documents depending on their specific.

The server is responsible with storing documents and their associated signatures, and also with identifying the appropriate document for a certain signature. It does this by comparing the collected signature with other signatures stored in the database, by using a (dis)similarity function. The most similar

Authorized licensed use limited to: Polytechnic University of Bucharest. Downloaded on July 04,2023 at 10:58:06 UTC from IEEE Xplore. Restrictions apply.

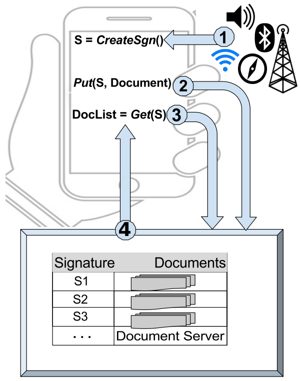

Fig. 1: Put/Get documents on the server based on radio signatures that are location specific, but without necessarily using Cartesian locations.

signature is identified and the associated document is retrieved and displayed in the application.

The unique rich signatures can then be used to manage a document collection without mapping them to geographical locations, but in fact obtaining an association between a document and its location in the building, not necessarily a Cartesian position. The data-structure obtained provides many functionalities of a location indexed database. The middleware (Figure 1) offers three main primitives to the applications:

- $S = CreateSgn()$ collects a location specific signature from the phone sensors(WiFi, BLE, 4G, sound, etc) and creates a multidimensional signature that is unique from any other location signature in the building;
- $Put(S, document)$ stores a document on the server associated with the signature S; The signature is created by a phone, but the indexing of the signatures and the document storage happen on a cloud/intranet server.
- $Get(S)$ - a phone harvests its current signature, and asks the server for a list of documents that have similar signatures, that have been stored at nearby locations. The server **searches in the signature space, and real geographical coordinates are never needed**.

On the server, documents are indexed by their signatures, based on the similarities between them. For a given signature query, the server may 1. Retrieve the document with the lowest dissimilarity with the query, or 2. Retrieve all documents with dissimilarities below a threshold. Since real physical positions are not known, the database of signatures on the server needs to be organized using clustering and labeling methods.

The institution owning the WiFi infrastructure may deploy $AirDocs$ services on a known port so that smartphone apps can easily discover instances of the servers with documents. In fact, as long as the WiFi/BLE infrastructure is stable, the system could be used with a server placed in the cloud, and the app could access a community or application specific server that stores signatures obtained from a given physical space, but without being administratively tied to that space.

A. Dissimilarity measure

Generally, positioning using fingerprints uses some function of distance in signal space, with Euclidean used in the RADAR paper [11], and many others tested in the literature. [12] tests Minkowski, cosine, Pearson correlation, and Shepard, finding that Euclidean and Pearson correlation provide the best results. In other studies, Mahalanobis is found to have the best performance, but for our setup it cannot be applied, since Android only gives one RSSI scan every 3 seconds, therefore a covariation matrix between RSSI distributions of different APs cannot be obtained without a long wait. Torres-Sospedra et al. [13] explore many others distances ad dissimilarities used in the literature, and found Sørensen (BrayCurtis coefficient) to perform best. Of these, we decided to test Minkowski ($p = 2$, Euclidean), cosine, Pearson correlation and Bray-Curtis. We chose to use Bray-Curtis on the basis of providing better monotonicity in cross validation against our dataset. In addition, we adopted some other improvements proposed in [13]: zero-to-one normalized representation (equation 1) of a RSSI value x_i in dBm:

$$X_i = normalized(x_i) = \alpha(1 - \frac{x_i}{min})^e \quad (1)$$

We chose the scale value α so that the range of x_i observed values -95 dBm .. -30 dBm get mapped to X_i in the interval [0,1]. The purpose of this normalization is double: it maps negative power reading in dBm to positive values that are needed by some similarity measures, but also discounts more differences between low power readings. The latter means that differences in stronger signals are penalized, for example a -90 dBm to -85 dBm difference is less important than a -40 dBm to -35 dBm difference, as RSSI readings are known to be much noisier at low power values. The Bray-Curtis dissimilarity (Sørensen distance), deemed the most performant in [13], relies APs common between the two fingerprints:

$$BCurtis(X, Y) = \frac{\sum_{i=1}^{c} |X_i - Y_i|}{\sum_{i=1}^{c} (X_i + Y_i)} \quad (2)$$

For APs missing between the two signatures, we consider them visible at -100dBm (-99dBm is the minimum observed value in datasets in this article), so that they contribute to the dissimilarity.

B. Measurements

We collected a dataset of 85 points in our own office building with collection points uniformly spread along a square shaped corridor. For each measurement point, 4 directions were collected, rotating 90^o after each scan. Two Android devices (Google Pixel 4A and Redmi Note 8) were held at face level away from the body at one step (0.6m) resolution. The building has an infrastructure WiFi, and a measurement point receives a median of 32 APs (minimum 20APs, 95% = 49APs). For the following numbers we use data collected with one phone for querying against the database collected with the other device.

Authorized licensed use limited to: Polytechnic University of Bucharest. Downloaded on July 04,2023 at 10:58:06 UTC from IEEE Xplore. Restrictions apply.

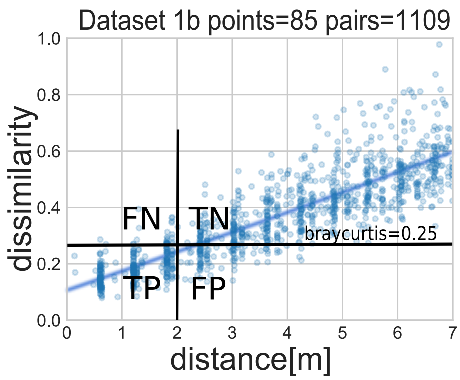

Fig. 2: Classification using a circle: simple thresholding in signal space produces false negatives and false positives. TP=true positives, FP=false positives, FN=false negatives.

As shown in Figure 2, the relation between dissimilarity and real distance is plotted with one blue dot for each pair of points which are at a distance up to 7 m. As expected, points with large dissimilarities are at larger distances to each other, but the linear fit still leaves some residuals.

An user may harvest its radio signature for his current location, and request either the the closest document in the database, or all documents in an area around him. If requesting the closest document, 70% of the queries return the exact same location, and distances to closest points are bounded by 0.6m and 1.67m (50% and 95% respectively).

If requesting documents in a circular range, the simplest scheme is to threshold dissimilarity values to limit the area in which documents are returned. Even if $AirDocs$ doesn't use real locations, users would have the feeling of discovering documents in a sphere around them that is sized relative to human dimensions. Using a threshold as indicated in Figure 2, the given query may produce some false negatives (FN) and false positives(FP). If we limit FN so to obtain a recall($\frac{TP}{TP+FN}$) of 0.99, the threshold used and the precision($\frac{TP}{TP+FP}$) obtained are presented in Table I.

TABLE I: Threshold classification of dissimilarity with 99% recall.

Radius	dissimilarity	precision	95% of results below
2.0m	0.32	0.57	3.8m
2.5m	0.35	0.63	4.2m
3.0m	0.37	0.66	4.4m
4.0m	0.45	0.75	5.4m
5.0m	0.54	0.78	6.2m

The third column indicates that false positives are returned up to a potentially larger radius, which defines the resolution of the system. While these preliminary results are promising, there are a number of considerations:

AP density affects the performance of the system in more than one way: on one hand more APs means more ways to discriminate between close locations, on the other hand unstable AP picked up by Android scanning introduces additional noise in the dissimilarity. More study is needed to understand which APs contribute positively to the signatures.

Whitelists: in many setups there will be temporary APs, or APs that change location. If these are a small fraction of the total, their effect will not be visible, but for low densities of stable APs the system should only use APs in a whitelist with MAC addresses that belong to the infrastructure. Also, most modern APs create virtual SSIDs, so the same physical card would broadcast under MAC addresses differing by one byte (Cisco), therefore a whitelist would be beneficial in unifying these readings.

Increased resolution: One of the first method to obtain increased resolution for the dissimilarity of the signatures is to use of additional sensors besides WiFi. BLE infrastructures are not as prevalent as WiFi, but all the issues explored in this paper for WiFi apply directly when beacons are available. For WiFi, a substantial increase in accuracy could come from collecting WiFi beacons at a higher rate than 0.3Hz offered by Android phones. Most laptops in monitor mode can collect the 10 beacons per second emitted by regular SSIDs, a value which would drastically improve both dissimilarity accuracy, and collection time. Unfortunately, the use of laptops would decrease the accessibility of the project.

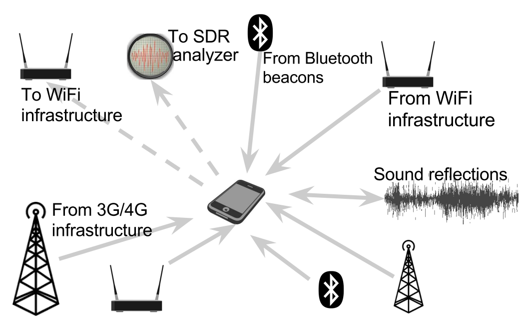

Fig. 3: a location specific signature is built using other sensors available the smartphone.

Figure 3 shows several possible sources of data that could make signatures more rich, and thus more discriminate with respect to location. 4G/5G has a rather low positioning accuracy, but is available in all smartphones, and could be used to speed up the searching structures in the server. Sound reflections (as used in project EchoTag [14]) are another source of enriching the signature that does not require deploying of additional infrastructure. In summary, any context information that is *stable, available, and easily collectable* by the phone can become part of the signature.

III. Related Work

Using location as context for enhancing human computer interaction has been proposed more than two decades ago [15], [16], and in $AirDocs$ we are revisiting some of those visions in the context of newly available functionalities (Wi-Fi, BLE, magnetic). Wireless propagation indoor has a hard to predict behaviour because of the heterogeneity of the environment, furniture and people, therefore many positioning systems require extensive training and updating to maintain a positioning service [17]. In contrast, $AirDocs$ proposes

Authorized licensed use limited to: Polytechnic University of Bucharest. Downloaded on July 04,2023 at 10:58:06 UTC from IEEE Xplore. Restrictions apply.

management of documents in a context aware fashion, but not linked to geographic locations which are natural contexts. Association of documents with locations has been explored — web documents are being geo-tagged and geo-referenced [18], and in the database community there are efforts to formalize searches for objects distributed in space [19].

POIs (points of interest) are associated to GPS maps to assign data to physical locations. They are used in the realm of outdoor activities to associate documents to locations. A project called Digital Graffiti [20] from University of Linz, Austria aims at associating data to outdoor locations, but it requires access to GPS, and to the Internet to access the data.

The EchoTag project [14] uses the microphone and speaker of the mobile phone to create a sound signature specific to the location. We plan to explore this direction with the purpose of creating an even richer signature for $AirDocs$.

Augmented reality is an emerging technology that "supplements the real world with virtual (computer-generated) objects that appear to coexist in the same space as the real world" [21]. [22] mentions projects spawned from MIT Media Lab's project sixth sense, that achieves a form of augmented reality by requiring the user to carry a projector and a camera to recognize hand gestures. $AirDocs$ is an enabler of augmented reality in the sense that documents are embedded in physical space, but without requiring positioning, head mounted displays, or instrumentation of the environment.

Dousse et. al [23] develop a purely fingerprint-based $placelearning$ method. Its core is a density-based clustering algorithm that works directly on the raw WiFi fingerprints. They also study the behavior of fingerprints with respect to space and time, but their focus on learning about stationary $places$ by using 60s sampling, manually labeled sets, and an unspecified spatial resolution of these places. Also, locations are visited for more than 5 minutes, in contrast with $AirDocs$, which aims for a more fluid user experience.

IV. Conclusion

We propose $AirDocs$, a system that makes use of signatures composed of stable information about the location that is easily collectable by smartphones, so that documents are managed spatially, but without the use of a location system, which usually requires extra infrastructure, training, or crowd-sourcing of measurements. $AirDocs$ works without any setup of the mobile or of the environment, and relies on a single server (intranet/cloud) to manage the placing and retrieving of signatures, and possibly the storage of the documents as well. We explore the use of WiFi fingerprints as the main component of a location dependent signature, and define a measure of dissimilarity that is mostly monotonic with real distance. We show that typical WiFi deployments enable efficient retrieval of documents for two popular queries: closest document, and all documents in a radius.

References

[1] Rathin Chandra Shit, Suraj Sharma, Deepak Puthal, and Albert Y Zomaya. Location of things (lot): A review and taxonomy of sensors localization in iot infrastructure. *IEEE Communications Surveys & Tutorials*, 20(3):2028–2061, 2018.

[2] Everton de Matos, Ramão Tiago Tiburski, Carlos Roberto Moratelli, Sergio Johann Filho, Leonardo Albernaz Amaral, Gowri Ramachandran, Bhaskar Krishnamachari, and Fabiano Hessel. Context information sharing for the internet of things: A survey. *Computer Networks*, 166:106988, 2020.

[3] Ritika Lohiya and Ankit Thakkar. Application domains, evaluation data sets, and research challenges of iot: A systematic review. *IEEE Internet of Things Journal*, 8(11):8774–8798, 2021.

[4] Omer Berat Sezer, Erdogan Dogdu, and Ahmet Murat Ozbayoglu. Context-aware computing, learning, and big data in internet of things: A survey. *IEEE Internet of Things Journal*, 5(1):1–27, 2018.

[5] Microsoft indoor localization competition, 2020. https://github.com/location-competition/indoor-location-competition-20.

[6] Dimitrios Lymberopoulos and Jie Liu. The microsoft indoor localization competition: Experiences and lessons learned. *IEEE Signal Processing Magazine*, 34(5):125–140, 2017.

[7] Antonio Ramón Jiménez Ruiz and Fernando Seco Granja. Comparing Ubisense, Bespoon, and Decawave UWB location systems: Indoor performance analysis. *IEEE Transactions on instrumentation and Measurement*, 66(8):2106–2117, 2017.

[8] Qin Zou, Qin Sun, Long Chen, Bu Nie, and Qingquan Li. A comparative analysis of lidar slam-based indoor navigation for autonomous vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 2021.

[9] C.T. Nguyen and et al. A Comprehensive Survey of Enabling and Emerging Technologies for Social Distancing—Part I: Fundamentals and Enabling Technologies. *IEEE Access*, 8:153479–153507, 2020.

[10] C.T. Nguyen and et al. A Comprehensive Survey of Enabling and Emerging Technologies for Social Distancing—Part II: Emerging Technologies and Open Issues. *IEEE Access*, 8:154209–154236, 2020.

[11] Paramvir Bahl and Venkata N Padmanabhan. Radar: An in-building rf-based user location and tracking system. In *Proceedings IEEE INFOCOM 2000*, volume 2, pages 775–784. Ieee, 2000.

[12] Giuseppe Caso, Luca De Nardis, and Maria-Gabriella Di Benedetto. A mixed approach to similarity metric selection in affinity propagation-based wifi fingerprinting indoor positioning. *Sensors*, 15(11):27692–27720, 2015.

[13] Joaquín Torres-Sospedra, Raúl Montoliu, Sergio Trilles, Óscar Belmonte, and Joaquín Huerta. Comprehensive analysis of distance and similarity measures for wi-fi fingerprinting indoor positioning systems. *Expert Systems with Applications*, 42(23):9263–9278, 2015.

[14] Yu-Chih Tung and Kang G Shin. Echotag: Accurate infrastructure-free indoor location tagging with smartphones. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, pages 525–536, 2015.

[15] Jason Pascoe. The stick-e note architecture: extending the interface beyond the user. In *Proceedings of the 2nd international conference on Intelligent user interfaces*, pages 261–264, 1997.

[16] Fredrik Espinoza, Per Persson, Anna Sandin, Hanna Nyström, Elenor Cacciatore, and Markus Bylund. Geonotes: Social and navigational aspects of location-based information systems. In *International Conference on Ubiquitous Computing*, pages 2–17. Springer, 2001.

[17] Xiaoqiang Zhu, Wenyu Qu, Tie Qiu, Laiping Zhao, Mohammed Atiquzzaman, and Dapeng Oliver Wu. Indoor intelligent fingerprint-based localization: Principles, approaches and challenges. *IEEE Communications Surveys Tutorials*, 22(4):2634–2657, 2020.

[18] Ingmar Poese, Steve Uhlig, Mohamed Ali Kaafar, Benoit Donnet, and Bamba Gueye. Ip geolocation databases: Unreliable? *ACM SIGCOMM Computer Communication Review*, 41(2):53–56, 2011.

[19] Dingming Wu, Gao Cong, and Christian S Jensen. A framework for efficient spatial web object retrieval. *The VLDB Journal*, 21(6).

[20] Heinrich Schmitzberger and Wolfgang Narzt. Leveraging wlan infrastructure for large-scale indoor tracking. In *2010 6th International Conference on Wireless and Mobile Communications*.

[21] Dimitris Chatzopoulos, Carlos Bermejo, Zhanpeng Huang, and Pan Hui. Mobile augmented reality survey: From where we are to where we go. *Ieee Access*, 5:6917–6950, 2017.

[22] Maria V Sanchez-Vives and Mel Slater. From presence to consciousness through virtual reality. *Nature Reviews Neuroscience*, 6(4):332–339, 2005.

[23] Olivier Dousse, Julien Eberle, and Matthias Mertens. Place learning via direct wifi fingerprint clustering. In *2012 IEEE 13th International Conference on Mobile Data Management*, pages 282–287. IEEE, 2012.

Authorized licensed use limited to: Polytechnic University of Bucharest. Downloaded on July 04,2023 at 10:58:06 UTC from IEEE Xplore. Restrictions apply.