Sensing Social Interactions through BLE Beacons and Commercial Mobile Devices

Michele Girolami[1 a], Fabio Mavilia[a], Franca Delmastro[b]

[a]*Institute of Information Science and Technologies, National Research Council (ISTI-CNR), Pisa, Italy*
{name.surname}@isti.cnr.it
[b]*Institute of Informatics and Telematics, National Research Council (IIT-CNR), Pisa, Italy*
franca.delmastro@iit.cnr.it

Abstract

Wearable sensing devices can provide high-resolution data useful to characterise and identify complex human behaviours. Sensing human social interactions through wearable devices represents one of the emerging field in mobile social sensing, considering their impact on different user categories and on different social contexts. However, it is important to limit the collection and use of sensitive information characterising individual users and their social interactions in order to maintain the user compliance. For this reason, we decided to focus mainly on physical proximity and, specifically, on the analysis of BLE wireless signals commonly used by commercial mobile devices. In this work, we present the SocializeME framework designed to collect proximity information and to detect social interactions through heterogeneous personal mobile devices. We also present the results of an experimental data collection campaign conducted with real users, highlighting technical limitations and performances in terms of quality of RSS, packet loss, and channel symmetry, and how they are influenced by different configurations of the user's body and the position of the personal device. Specifically, we obtained a dataset with more than 820.000 Bluetooth signals (BLE beacons) collected, with a total monitoring of over 11 hours. The dataset collected reproduces 4 different configurations by mixing

[1]Michele Girolami is corresponding author, National Council of Research ISTI-CNR Italy, michele.girolami@isti.cnr.it

two user posture's layouts (standing and sitting) and different positions of the receiver device (in hand, in the front pocket and in the back pocket). The large number of experiments in those different configurations, well cover the common way of holding a mobile device, and the layout of a dyad involved in a social interaction. We also present the results obtained by SME-D algorithm, designed to automatically detect social interactions based on the collected wireless signals, which obtained an overall accuracy of 81.56% and F-score 84.7%. The collected and labelled dataset is also released to the mobile social sensing community in order to evaluate and compare new algorithms.

Keywords: Social Interactions; Human Proximity; Bluetooth Low Energy; Wearable sensors.

---

1. Introduction

Human social interactions represent complex behaviour to be described, and they cannot be easily connected with sensing data, even though we are currently able to collect a huge amount of data from our personal and mobile devices. Most of the works in the literature rely on video signals and processing, from which identifying specific conditions of social interactions. However, users are even more skeptical of being continuously monitored, especially in terms of video recording, generating thus not spontaneous behaviour and resulting in biased information. More recent works focus on the use of sensing data derived from wearable sensors, but also in this case they often need customised hardware to be worn in specific conditions to optimise the wireless signal performances. In addition, there are very few datasets available in the community to test and evaluate detection algorithms in this field.

Since we consider proximity and face-to-face meetings as fundamental information to detect social interactions, we propose a mobile framework, called SocializeME, based on the use of commercial devices (i.e., smartphones), and designed to collect and to analyse data from wireless BLE signals, commonly emitted by wearable devices. We conducted an experimental campaign by involving several students from a high school to test our framework and to execute specific interaction experiments in order to collect the wireless data and the appropriate labels related to the different interaction phases: non-interaction, approaching, interaction. Tests have been designed to reproduce heterogeneous situations typical of human interactions by using commercial devices. We analysed the collected data with particular attention to the differences in the emitted signals, in terms of frequency and RSS, generated by the different configurations and uses of the mobile device (e.g., standing face-to-face with smartphone in one hand, standing with the phone in the front/back pocket, sitting, etc.). The main goal is the emulation of real-life scenarios and the collection of realistic data. Specifically, we study the wireless channel both in terms of quality and symmetry. As a first step, we focus on the RSS variations during the different stages of a social interaction. In addition, we show how RSS is attenuated depending on the position of the receiving device, which is witnessed also by the beacon loss rate resulting from the experiments. Finally, we present the performance of the SocializeME Detector (SME-D) algorithm, originally presented in [1]. SME-D is designed to automatically classify interactions or non-interactions according to some features of beacon's RSS. In particular, we measure the performance of SME-D in terms of Accuracy and F-Score metrics with different tests.

Our experience provides some important lessons, both in terms of the realistic usage of commercial mobile devices for social interaction monitoring and the characteristics of the BLE wireless signals in realistic scenarios. The encountered technical limitations, together with a detailed description of the experiments and the collected dataset, represents an important contribution for the entire community aiming at defining and evaluating new algorithms for human social interactions from mobile sensing. Furthermore, we argue that the technologies we adopted for collecting our dataset, enables the possibility of modelling and predicting social gatherings among people with high temporal resolution, which is a crucial asset for managing the epidemic diffusion of diseases in the modern society such as the recent COVID-19 pandemic.

The paper is organised as follows. Section 2 describes the related works on the automatic detection of social interaction through sensing units. We also review the existing datasets specifically collected for this purpose. Section 3 describes the SocializeME mobile framework in terms of selected wearable devices and the mobile application running on commercial smartphones. Section 4 describes the experimental campaign and the data analysis, in terms of quality and symmetry of the BLE channel. Section 5 presents the main features of SME-D algorithm, and its performances in different configurations.

2. Related Works

The study of human social interactions is traditionally approached with questionnaires and diaries [2, 3] periodically compiled by monitored subjects. Such tools represented for a long time an essential source of information, however they require the explicit user intervention, which can impact on the accuracy of the experiment and the period of involvement of voluntary people. In order to overcome this limitation, sometimes an external observer has been involved to track the subjects and take note of the time, people involved in the meeting, and the type of social interaction (e.g. occasional, recurrent, or intimate interaction). This approach introduces an additional possible bias to the collected data, both related to the involvement of additional persons and on the natural way people interact, perceiving the presence of external subjects.

With the introduction of mobile sensing technologies, the previous approach can be further improved. Researchers are now investigating how objective measures derived from the available technologies are able to provide information on the user social interactions in the physical world. As a first step, we can distinguish between the available technologies, the type of collected data and the processing required to extract meaningful information.

Video recording and wireless signal analysis are the most used technologies for this purpose, they can be associated with the recognition of the physical activity or the emotions. The target is to define the appropriate processing to detect *when* and *how strong* people are engaged in a face-to-face interaction.

In addition, embedding the sensing technologies in mobile and wearable devices (i.e., smartphones and smart watches) allows the collection of huge amount of data, reducing the obtrusiveness and improving the quality of the collected dataset. However, it is important to focus on the user compliance in participating in the monitoring, especially related to the collection of this personal and sensitive information.

The video-based techniques [4] are widely adopted in this area. They exploit the analysis of frames derived from fixed or mobile cameras deployed in the environment. However, the interactions can only be tracked in the monitored physical area and people can feel observed, thus not moving and interacting in a natural way. Authors of [5] propose a system based on camera recordings from a first-person video. Specifically, an actor records videos with a camera capturing scenes from his point of view. Authors analyze the collected video clips during a one-day social event, with the goal of classifying the interactions in 3 categories: dialogues, discussion, and monologue. The analysis is based on location and orientation of the involved people with respect to the recording person, in order to detect if they were interacting or not. The work presented in [6] is also based on video recordings from a surveillance system. Authors adopt a technique based on the motion features of the recorded subjects with the goal of detecting if two subjects are close enough to interact, introducing the concept of proxemic [7]. Similarly, a more recent work focuses on the recognition of the social interactions based on the pictures captured by a worn camera [8]. The authors adopt cameras based on a low frame rate in order to capture long-lasting recording sessions.

Datasets collected in the previous works are not publicly available, also due to privacy concerns. In fact, it is not possible to anonymize this type of data without loosing their significance, and generally users are really skeptical of being visually recorded during their daily activities for monitoring purposes, generating data that are often biased by a not natural behaviour. However, several works in the literature rely on this sensitive information, like SALSA [9], which collected an audio-visual dataset combined with sensing information obtained by Sociometric badge[2] [10, 11, 12]. The badge is a wearable device equipped with a microphone, an accelerometer and an infrared and Bluetooth sensor. The microphone is used to assess the talkativeness of subjects in noisy condition, with the goal of inferring if subjects are speaking or not. The accelerometer provides an indication about the subjects' movements, while infrared and Bluetooth are used to detect proximity among the subjects. The experiment conducted by SALSA collects data during an indoor social event in which 16 subjects are recorded for 60 minutes. Moreover, in order to share the dataset, the authors provide annotations to the captured scenes. The annotations describe, for each subject, the position, and the orientation of head and body with a time resolution of 0.3 Hz. In addition, the monitored subjects were asked to fill in a questionnaire before joining the experiment. The questionnaire allows the authors to define a profile of the subjects in terms of 5 personal traits: Extraversion, Agreeableness, Conscientiousness, Emotional Stability, and Creativity. As a result, the SALSA dataset provides sensing information collected during the experiment and a description of the involved subjects.

From the collected data the authors were able to infer users' proximity based on the IR signals, talkativeness through the audio signal of the microphone, and body motion through the accelerometer data. This type of data are useful to reveal complex dynamics of human interactions with a fine-grained time scale, however the users need to wear the proposed device and accept to be visually and audio recorded.

The authors of [13] exploit audio and video signals in order to detect groups of conversing people. More specifically, 12 wireless microphones were deployed among the participants and 3 cameras recorded the scene. Participants worn a sensor pack around the neck composed by: a triaxial accelerometer, a proximity sensor, and an indoor positioning device. An interesting aspect of this work consists in the detailed questions that the authors aim to respond: who is speaking and who are the people involved. Authors aim also to understand the

[2] https://hd.media.mit.edu/badges/information.html

quality of the social interaction, by detecting if people are enjoying themselves in crowded scenarios. The authors collected data from 32 volunteer students from different universities. Authors annotated the dataset by analysing the video recordings and detecting the F-Formations [14].

In order to define a less invasive system, which does not require video and audio signals, a different approach has been proposed in the literature, consisting in inferring the existence of an interaction between a pair of subjects by only detecting their co-location. The term co-location represents the co-presence of two or more people in the same place at the same time. As discussed in [15] co-location can be used as a proxy to reveal a face-to-face meeting. As a key-observation, the more people lie in proximity, the more likely they are involved in meaningful social interactions. Works based on the co-location analysis have flourished in the recent years, thanks to the diffusion of wireless interfaces such as WiFi Direct, Bluetooth Low Energy and legacy Bluetooth, as well as the upcoming LTE Direct technology. All these interfaces are pervasive in our mobile devices (smartphones, smartwatches, wristband and small pocket devices), making much easier and cheaper to setup a long-lasting data collection campaign. The technique often used to infer an interaction from the co-location consists in broadcasting wireless signals to a relative short distance. Therefore, only devices in the nearby can hear them. In turn, such signals are collected and analyzed in order to infer when a pair of devices is close enough to assume that their owners are involved in a social interaction. Of course, such technique might lead to false positive and negative answers but, on the long-term, the analysis provides valid and reliable results.

Authors of [16] represents one of the pioneering work in this field, exploiting the Bluetooth scans of Nokia 6600 devices in order to discover devices in the nearby with a time resolution of 5 minutes. The dataset comprised 75 students, and the authors analyzed proximity, the visited locations, and a pattern of encounters among volunteers on a daily basis. The dataset collected offers a first outlook to study dynamics of the social interactions, however we argue that the approach followed cannot be applied for large-scale deployments, since all the volunteers use the same device model.

A more recent work is represented by SocioPattern[3] research platform, which provides several datasets gathered during various social events. SocioPattern has been adopted in several remarkable works addressing the goal of dynamics' analysis of the social interactions in a natural context[17, 18, 19]. The platform consists on customised wearable badges based on RFID sensors emitting a signal in a range of 1 to 1.5 meters. Badges are worn by people, while receivers are installed in the environment (generally on the ceiling) in order to collect and store the acquired data. The SocioPattern platform represents a state-of-the-art solution for the purpose of understanding dynamics of the social interactions in crowded areas.

Finally, the Copenhagen Networks Study [20, 21] is another interesting project aimed at studying the social interactions of people by means of data collected with a mobile app for smartphone. The application captures multiple signals, including WiFi scans, locations and Bluetooth scans. The data collection campaigns (SensibleDTU 2012 and 2013) involve a high number of volunteers. The organizers distributed the smartphones to the users therefore, differently from our experience, they did not face with the extreme heterogeneity of device manufacturing. Social ties are identified with Bluetooth, WiFI and other channels. Concerning Bluetooth, the organizers of the Copenhagen Networks Study adopted Bluetooth scan in order to discover devices in proximity. We consider that such technology offers an interesting approach, however the power of emission of the Bluetooth scan can not be configured in order to limit the devices seen in proximity. Furthermore, the frequency of emission of Bluetooth scans is also limited. Bluetooth beacons, differently, allow to tune both of the parameters (power and frequency) giving rise to a even more accuracy of devices detected.

Table 1 presents a summary of the technologies and systems presented in the literature and described above, highlighting the main advantages and draw-

[3] http://www.sociopatterns.org/

Table 1: Comparison of Technologies used for detecting social interactions.

Technology	References	Pros & Cons
Audio-Video Recordings	[4, 5, 6, 8]	+ camera can detect multiple subjects + camera can record gestures, proximity, talkativeness - tuning of the camera - subjects might feel uncomfortable with the system and collected data can be biased by non natural behaviour
Wireless signal analysis	[12, 15, 17] [18, 19, 20, 22]	+ interactions are inferred by detecting proximity between subjects - interactions tracked only in the monitored area - subjects generally have to wear a customised device
Mixed Solutions	[9, 13, 20, 21] [10, 11, 16]	- subjects generally wear a customised device - multiple sensing units + interactions are inferred by detecting multiple sociological markers

backs.

All the presented works focus on physical social interactions, while in the last few years researchers also investigated this social phenomenon associated with the cyber world, and especially to the virtual interactions though Online Social Networks (OSN). In this case, data collected from OSNs (i.e. Facebook, Twitter) are used to build up the ego-network of one or more users in order to identify his/her contact graph, and to evaluate the correlation among virtual and physical social interactions [23] . OSN data are also used to estimate the trust of the users we are in contact with in order to define trustable communications [24] or to select users with which exchange contents or transactions. However, even in this case it is fundamental to identify physical social interactions and to define unobtrusive systems able to automatically detect them. For this reason, in this paper we present a system designed to monitor and detect social interactions through commercial mobile devices, and in particular smartphones and low cost BLE tags, and a wide dataset collected during several experimental campaigns involving young students. Specifically the system, called SocializeME, consists of an Android mobile app that receives beacon messages sent by wearable BLE tags in proximity. All the collected data are sent to a remote server where the wireless signals are analysed in order to identify the features that characterise social interactions among two or more users, depending also on their posture and the position of the receiving device. The main advantage of SocializeME with respect to the previous solutions is represented by the use of commercial mobile and wearable devices, widely adopted and not designed specifically for the presented goal, and the different configuration settings that can be used. However, even if the use of commercial devices increases the acceptability of a sensing campaign, they can also introduce some technical limitations, making even more challenging the overall goal. To this aim, we will also present the technical limitations derived from our real world experience, the characteristics of the wireless signal and how the posture of the subjects and position of the devices can dramatically influence the quality of the data and the symmetry of the radio channel. Finally, we present a description of the available dataset, which can be used to reproduce our results and for further investigation by the research community.

3. BLE Beacons and Commercial Mobile Devices: the SocializeMe framework

As previously introduced, we designed a mobile solution aimed at collecting data related to proximity of personal mobile devices in order to infer face-to-face interactions among their users. The target users of the research project are young students of a high school specialised in computer science, which have been directly involved in the design and testing of the proposed solution before participating in the experiment. This activity has been carried out in the framework of SocializeME project, which provides the same name to the proposed software framework.

The goal of the project is to design software and analytic tools able to detect face-to-face interactions without adopting customised hardware, and to provide the users with a non-invasive technological solution, both in terms of hardware and software interactions. In fact, we mainly rely on commercial smartphones, continuously used especially by the younger categories. In this section we present the technical aspects of SocializeME solution, including the technical limitations we encountered during its development, which arose important real-world challenges for its deployment for a large-scale experimentation.

As a first step, we identified Bluetooth technology and its Low Energy evolution as the ideal communication paradigm on which implementing a mobile app aimed at collecting information about other users and devices in proximity. In fact, one of the main Bluetooth Low Energy (BLE) functionalities is the *advertising procedure*. It allows mobile devices to broadcast information to announce themselves and define their intentions. Specifically, BLE defines a single packet format for both advertising and data transmissions composed of four elements: preamble (1 octet), access address (4 octets), Protocol Data Unit – PDU (2-257 octets), and Cyclic Redundancy Check – CRC (3 octets). PDU contains a 16-bit header and a variable size payload. The header contains the packet type (i.e., advertising or data packet) and the length of the payload. A mobile app can specify the content of the advertising payload to announce a service running on the device and the device itself through a unique identifier. Advertising packets are periodically broadcasted on each advertising channel. The time interval between sent packets is composed of a fixed part and a random delay to avoid collisions. Devices can receive advertising packets only when they are in a scan phase.

Considering these characteristics, we decided to develop the first version of SocializeME App as a mobile app aimed at advertising the presence of the local device and its user to the other devices in proximity (i.e.,in the range of few meters) and to store the information of the received advertising packets from third parties. The app should have worked both in foreground and in background, with a minimum interaction with the final user. However, only Android devices supports advertising and scan modes in background, while iOS automatically interrupt advertising and scanning operations while the app moves to background. Therefore, we decided to focus only on Android, but also in this case we identified some minimum technical requirements: the minimum version of Android OS (5.0), and the presence of a BLE physical interface based on Bluetooth standard 4.0, at least.

(a) The SocializeME app for collecting beacons.

(b) The BLE tags used for emitting beacons.

Figure 1: Hardware and software setting of the SocializeME project.

Since the app collects personal data of the user, it is designed in accordance with the GDPR and the user register only after the acceptance of the privacy statement. From the technical point of view, the app broadcasts the advertising packet with a certain frequency and transmission power. These two parameters are not configurable by the application with a numerical value, but only through predefined variables (i.e., high, medium, low and ultra low transmission power; high, balanced and low frequency). Each device type and model has its own settings for these threshold values, and the application has no control on these parameters, which highly affect the app performances.

We ran a set of preliminary experiments with some volunteer students to evaluate the performances of the system, and we encountered several issues: (i) most of the commercial devices do not support advertising mode even satisfying the basic technological requirements (this feature is referred to as peripheral mode); (ii) advertising and scan operations have a big impact on the device power consumption. As far as device heterogeneity is concerned, we tested more than 15 different devices and 42% of them did not support the minimum application requirements. A detailed list of tested device is presented in [1].

To avoid these limitations, we decided to modify the reference architecture from a dedicated mobile app, to a wider system including a set of commercial BLE Tags produced by Global Tag[4] aimed only at transmitting the advertising packets to announce the presence of a user, while the mobile app is dedicated to maintain the scan phase active and to collect the received advertisements. In Figure 1a we present a screenshot of the app, which presents a user interface defined as a monitoring tool of the app itself (i.e., Receiver ID, number of received beacons, start and end time of the encounter), and a picture of the tags, which are not invasive and can be easily worn as bracelets. In addition, they are low cost devices, easy-to-configure and fully compliant with Bluetooth v4 protocol stack. They support both the Eddystone and iBeacon beaconing protocol, and the advertisement rate and the transmission power can be tuned, ranging from 1 Hz to 10 Hz and from -23 dBm to 4 dBm, respectively. The tag configuration is managed through a specific software used to pair a tag with the user device.

It is important to highlight that the mobile app has been installed directly on the users' mobile devices after they signed an informed consent, according to the current GDPR in terms of privacy and management of personal data. In addition, the app requires to configure the tags before the first use, to make the system able to distinguish between the tag worn by the local users, and those available in proximity. Finally, all the collected data are stored on a remote server for offline analysis. For all these reasons the app cannot be released publicly, but all the presented results can be verified and are completely reproducible by analysing the associated dataset [25].

[4]http://www.global-tag.com

3.1. Technical limitations

During the development and testing of SocializeME app, we encountered some technical issues mainly related to the continuous monitoring requirement and how the user interacts with the device. Specifically, to maintain continuously active BLE scan mode is not completely automatic. Originally, Android apps had the ability to use the "wakelock" procedure to prevent the phone from going into a power-saving deep sleep mode. However, since several apps exploited this mechanism to be continuously active, highly impacting on the battery lifetime, Android 6 introduced a new mode called Doze, which reduces the use of CPU and the network access to those apps using a wakelock mechanism in case the phone is in idle state, stationary and not connected to the power. In case of Android 7.0 version, this condition has been further restricted, not considering the stationary condition, therefore the phone can enter the doze mode even if the phone is idle while moving together with the user. This condition highly impacts on SocializeME performances in terms of received beacons and detection of social interactions. To reduce the impact of this limitation, SocializeME app has been updated as a foreground service, maintaining an icon in the status bar, so that the user is aware of the running service.However, this update didn't solve the problem completely. Even though the app was not affected by doze mode, BLE scans stopped after a period of 20 min of the background running. This time has been estimated with several stress tests executed in laboratory before the app release. In these cases the experiment has been set up to maintain the app active for 12 hours, receiving beacons from an increasing number of tags. In the heaviest case, we had a smartphone on a table and 25 tags configured to send beacons with a frequency of 5Hz and -6dBm power. In five tests with the same configuration the app stopped receiving beacons after 20 minutes. Therefore, we assumed that the operating system implemented a control on the usage of BLE interface to avoid an interface overload. In fact, by implementing a periodic software reset of the BLE interface, indicatively every 30 min, SocializeMe App was able to maintain scan operations active and receive beacon messages indefinitely. In addition, we also made additional stress test configuring the beacons with 2Hz frequency and -6dBm power and there was no interruption but we experienced some inactive periods in the experiments, and consequently the loss of advertising packets. Therefore, considering the high school environment as the reference for our solution, we expect to have a not negligible packet loss while maintaining the ability to detect face-to-face meeting. In order to support this assumption, we present the results obtained by a specific set of experiments, enriched with the ground truth provided by the involved users. In this way, we will be able to analyse the obtained dataset and the accuracy of possibile detection algorithms.

4. Experimental analysis

In this section we present the experimental sessions we conducted in order to analyse the impact of the users' configuration, with respect to their mobile devices, on the signal quality and, consequently, on the correct detection of face-to-face interactions in a realistic context.

- number of involved subjects;
- posture of the subjects: Standing (ST) or Sitting (SI);
- position of the receiving device with respect to the user body: Hand (H), Front Pocket (FP) and Back Pocket (BP);
- layout of the involved subjects: dyad, trio, foursome, and group of five.

During a test, the subjects mimic a face-to-face interaction characterized by three stages differentiated by the physical distance between the participants [7]: Non interaction (3 to 3.5 meters), Approaching (3 to 2.5 meters) and Interaction (2.5 to 1 meter).

Table 2 provides an overview of the experiments. The table reports the identifier of each session, the number of volunteers, the number of different smartphones models, the number of collected beacons and the overall duration of the session's tests. The number of collected beacons is obtained as the sum

Table 2: Experiments overview.

Session	#Volunteers	Smartphone Models	#Beacons	Duration [min]
1	8	4	53375	111
2	9	6	87467	114
3	10	4	205152	111
4	8	5	198603	130
5	8	3	247776	130
6	6	4	27886	73
Total	**-**	**-**	**820259**	**669**

of all the beacons received by all the smartphones during each of the session's tests. Differently, the duration column provides the average time spent by the volunteers in order to complete a specific session. For example, the 8 volunteers of Session 1 gathered 53.375 beacons with 4 different types of smartphones. The 4 pairs of volunteers completed the tests with an average time of 111 minutes.

We can note that the sessions differ remarkably in two aspects: (i) the number of collected beacons, and the duration of the session. Such differences are mainly caused by the number of volunteers participating in the experiments. As expected, the lower the volunteers, the lower the number of the collected beacons and the session's duration. As a general result, we can say that we obtained a dataset with more than 820.000 beacons collected over 11 hours of monitoring.

The experiments have been designed to reproduce some meaningful combinations of posture and position of the receiver, as reported in Table 3. In summary, we tested 4 different configurations (C_1 - C_4), by mixing two posture's layouts, standing (ST) or sitting (SI), different positions of the receiver device (FP, BP and H) as detailed previously. Such variations of the physical configurations of the experiments well cover the common way of holding a mobile device, and the layout of the dyad involved in a social interaction.

Table 3: Experimental configurations and the reference sessions.

Configuration	Posture	Position of Receiver	Sessions/Tests
C_1	ST	FP	$S_1/1$ $S_2/1$ $S_3/1$ $S_4/1$ $S_5/1,2,3,4$ $S_6/1.3$
C_2	ST	BP	$S_1/3$ $S_2/3$ $S_3/3$ $S_4/2$ $S_5/1,2,3$ $S_6/1,2$
C_3	ST	H	$S_4/3$ $S_5/1,3,4$ $S_6/2,3$
C_4	SI	FP	$S_1/2$ $S_2/2$ $S_3/2$ $S_4/2$

4.1. Wireless communication performances

In order to analyse the system performances, we firstly analyse the characteristics of BLE wireless communication in terms of the RSS distribution during the sessions and the beacon loss rate. We identify these two parameters as an evaluation of the wireless communication performances aimed at identifying social interactions.

In fact, it is well-known in the literature that most of the wireless communication channels are influenced by several factors: the distance between the emitter and the receiver, the presence of obstacles in between, as well as the material of the obstacles, the surrounding environment (i.e., indoor or outdoor) but also the presence of humans in the line of sight between the emitter and the receiver [26, 27]. BLE beacons are characterised by similar conditions and, being related to smaller distances, they can be highly influenced by the posture and the position of the receiver with respect to the user's body. As a meaningful example, we show the RSS values recorded by 3 different receivers in three configurations (see Table 3): C_1 = (ST, FP), $C2$ = (ST, BP) and C_3 = (ST, H), as shown in Figure 2. For each of the three configurations, Figure 2 shows

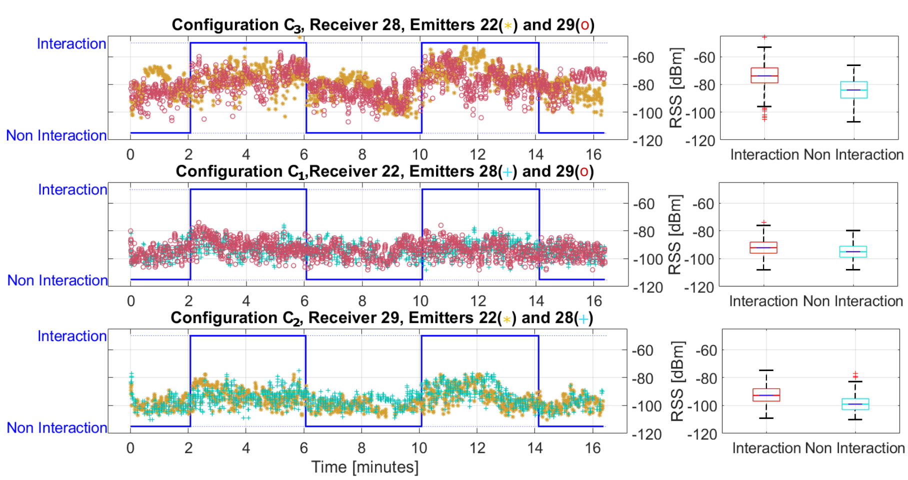

Figure 2: RSS variation with 3 configurations: C_3, C_1 and C_2.

on the left the time series of the RSS values, and on the right the distribution of the RSS values during the Interaction and Non interaction stages as box plots. The blue line shows the ground-truth, namely the time intervals during which the subjects are actually engaged in a social interaction.

Configuration C_3 shows an evident pattern of the RSS values during the Interaction and Non Interaction stages. In this case, the receiver is hold in hand (H) on the same line of sight of the emitter, therefore the RSS values well reflect time periods of Interaction with respect to time periods of Non Interaction. In particular, during the Interaction stage the RSS median value is -74 dbm, the $75th$ percentile is -68 dbm and the $25th$ percentile is -79 dbm. While, during the Non Interaction stage, the median value is -84 dbm, the $75th$ percentile is -78 dbm and the $25th$ percentile is -90 dbm.

The situation is slightly different in case of configuration C_1, in which the receiver is positioned in the front pocket (FP). The RSS values give rise to a soft, but still present pattern: during the Interaction stage the RSS values tend to increase, while during the Non Interaction stage such values decrease. In particular, during the Interaction stage the RSS median value is -92 dbm, the $75th$ percentile is -88 dbm and a $25th$ percentile is -96 dbm. During the Non Interaction stage, we recorded lower values of such statistics, specifically the median value is -95 dbm, the $75th$ percentile is -91 dbm and the $25th$ percentile is -99 dbm. Therefore, a first qualitatively observation is that positioning a receiver in the FP causes an attenuation of the RSS values both during the Interaction and Non Interaction stages. This is also clear from the two box plots of C_1 in Figure 2 which show very similar values of the median during Interaction and Non Interaction stages (-91 dbm vs -95 dbm with a $\delta = -3$ dbm).

Similar considerations also apply for configuration C_2 in which the receiver is put in the back pocket (BP). In this case, the body attenuation produces even lower RSS values. The interaction pattern is still present (increase of the RSS values during Interaction and decrease of the RSS values during Non Interaction stages) but with lower statistical values with respect to C_1. In particular, during the Interaction stage the RSS median value is -93 dbm, the $75th$ percentile is -88 dbm and the $25th$ percentile is -97 dbm. During the Non Interaction stage, we recorded lower values of such statistics, specifically the median value is -99 dbm, the $75th$ percentile is -95 dbm and the $25th$ percentile is -103 dbm. Configuration C_2 differs from C_1 as the box plots show a median of the Interaction and Non Interaction of respectively -93 dbm vs -99 dbm with a difference of $\delta = -6$ dbm.

The consideration given so far are based on a select dyad of one of the experiments, in order to highlight the distinguishing features of the RSS. We now extend our analysis to all the beacons collected and in particular, we analyze the RSS distributions for each of the 4 configurations described in Table 3: C_1 to C_4. The goal is to show quantitatively how the RSS distributions are affected by the combination of the posture and the position of the receiver. Figure 3 shows the probability density function of the RSS values in the 4 configurations divided by the 3 stages: Interaction, Approaching and Non Interaction. We further present in Figure 4 the RSS distribution as box plots with median (mid line), 25_{th} and 75_{th} percentile, max and min RSS.

For each of the 12 distributions in Figure 3, we report the median (μ) and the standard deviation (σ) of the received beacons. As a general trend, we observe that the mean value μ of the signal strength reduces moving from Interaction, Approaching to the Non Interaction stages. In fact, the distance between emitters and receivers increases that, in turn, reflects on the signal strength of the beacons. We also provide a trend of the distribution by computing the KDE (Kernel Density Estimator) reported as a blue line on each distribution.

We observe that C_3 configuration provides the highest μ values. This represents the best case since all the tests are executed with the devices hold in hand and no barriers between the emitter and the receiver, resulting with higher RSS values. Concerning the Interaction stage, we observe that C_3 (ST, H) and C_1 (ST, FP) provide better results with respect to the other two configurations. Specifically, in C_3 $\mu = -69$ and $\sigma = 10.05$, while in C_2 $\mu = -78$ and $\sigma = 11.92$. Differently, the signal strength of the received beacons remarkably reduces with configurations C_4 and C_2. Concerning C_4 (SI, FP), we can observe that the sitting position slightly degrades the signal strength. With such configuration,

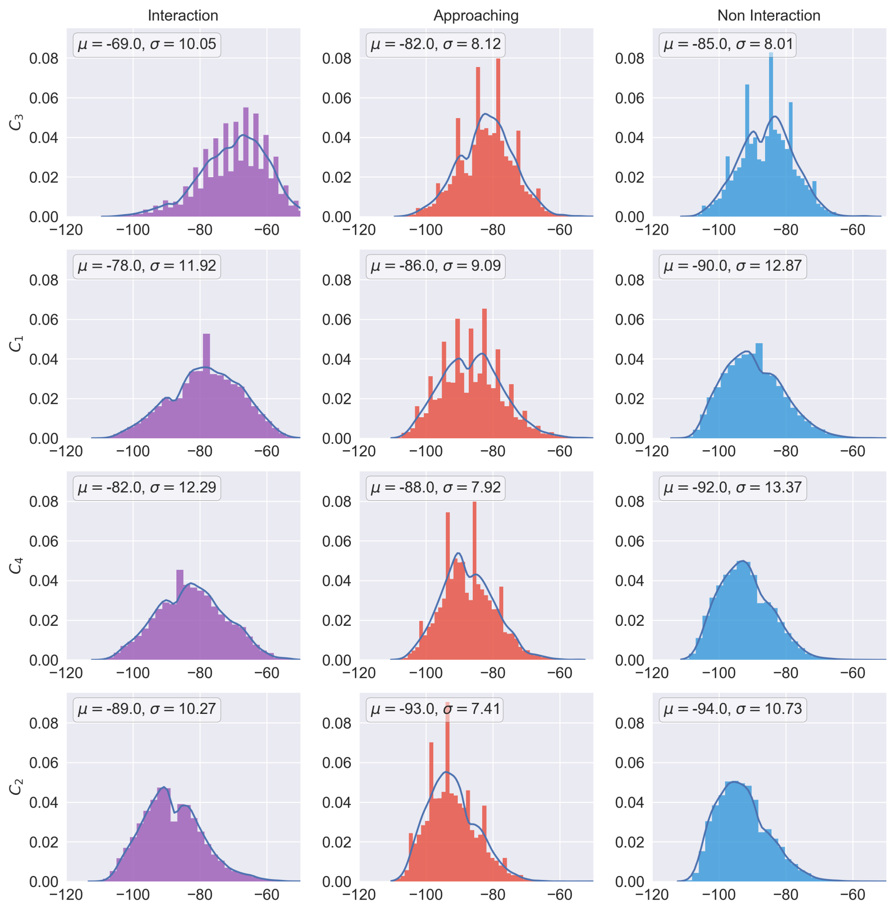

Figure 3: Probability density function of the RSS values.

we measure $\mu = -82$ and $\sigma = 12.29$, 13 dbm lower than C_3 values. The worst case is C_2 (ST, BP). In this case, the tests are executed with devices put in the back pocket demonstrating that the body attenuation highly degrades the signal strength. In fact, in this configuration we measure $\mu = -89$ and $\sigma = 10.27$, 20 dbm lower than that of C_3.

The Approaching stage provides generally mid values between Interaction and Non Interaction stages as shown by μ values. Also in this stage, C_3 tests represent the best case, while C_1 tests are the worst one. Moreover, we observe that the shape of the distributions across all the configurations is generally

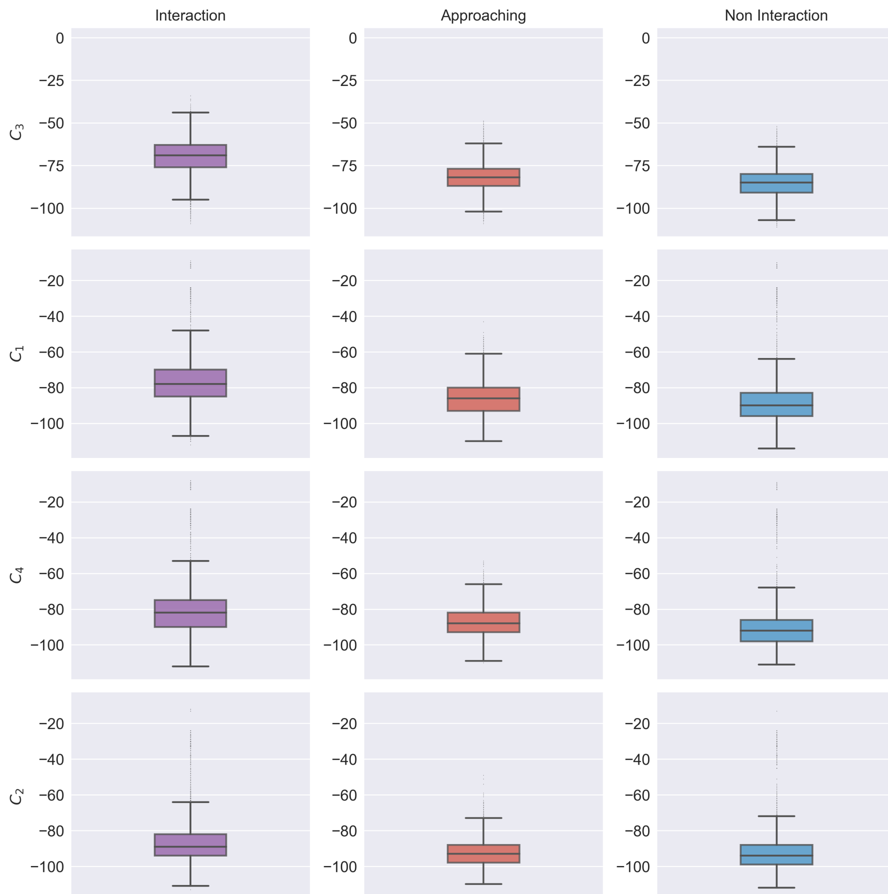

Figure 4: Box Plot of the RSS values with 4 configurations and 3 stages.

more irregular than the other stages. The RSS distributions show spikes in correspondence to several RSS values, which are intrinsic in the Approaching movement of the users and the devices.

The Non Interaction stage shows a more evident effect of the body attenuation. As a general trend, we note that μ values are lower than those measured during the Approaching and Interaction stages. Also in this stage, C_3 represents the best case, while C_1 confirms the worst configuration. As a further confirmation, we observe that during the Interaction stage of C_2, the mean $\mu = -89$ dbm, which is still lower than that obtained in C_3 during the Non Interaction stage ($\mu = -85$ dbm). Therefore, in C_2 even if two devices are 1 meter distant (Interaction stage), the mean value of the RSS is lower than that of C_3 in which the devices are hold in hand at about 3 meters distance.

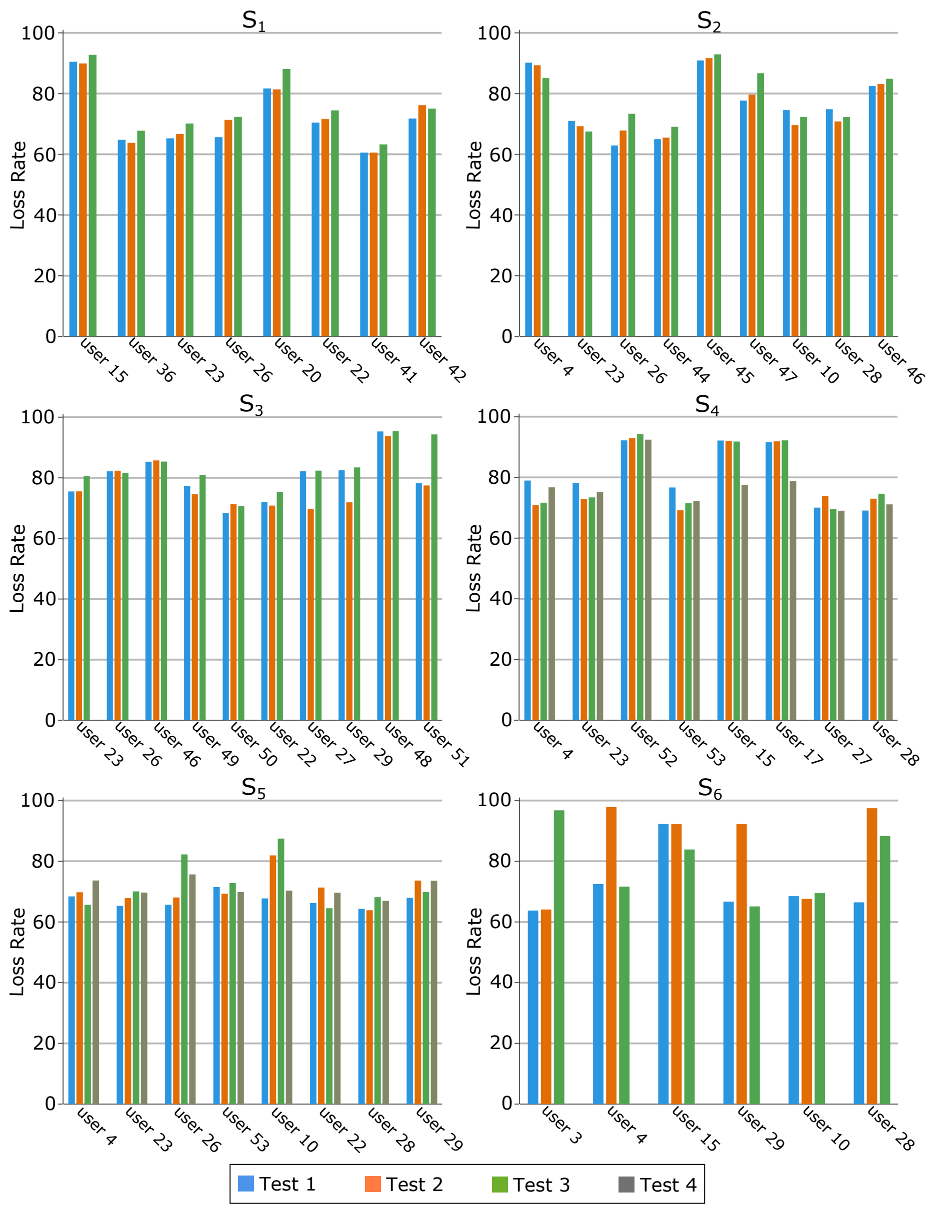

Figure 5: Beacon loss rate for all the sessions.

Another important parameter of our experimental analysis is represented by the beacon loss rate experienced in the different configurations. We focus on the Interaction stage of all the configurations since it represents the best opportunity for the receivers to collect beacons from their partners. Differently, during the Approaching and Non Interaction stages, we expect to receive a minor number of beacons due to the increased distance and the minor RSS values. The beacon loss rate is computed for each user and for each completed test. For example, test 1 of S_1 is organized as pairs, e.g. user 1 vs user 2. Each device emits beacons with a frequency of 5 Hz, and the Interaction stage lasts for 5 minutes. Therefore user 1 will receive beacons only from user 2 for 5 minutes. The beacon loss rate for user 1 in test 1 is given by $x/1500$ where x is the number of received beacons from user 1, and 1500 is the number of expected beacons for 5 minutes at 5 Hz.

Figure 5 shows the beacon loss rate as a set of distinct bar plots. The figure shows, for each of the 6 experimental sessions (S_1 to S_6), the users involved and the beacon loss rate measured by their devices. We report the beacon loss rate for each of the tests performed by the users with the goal of showing significant differences among the tests. Sessions S_1, S_2, S_3 and S_6 are composed by 3 tests each, while Sessions S_4 and S_5 by 4 tests

The sessions involve different device models and different number of users. Details of the beacon loss rate divided by device models and number of users are reported in Appendix Appendix A.

It is worth noting that the beacon loss analysis does not aim to rank the performance of the device models we experienced, rather it aims to highlight the heterogeneity of the devices in a real-world experiment, and how it impacts on the final results. Unfortunately this represents a limitation for the deployment of a large scale experimentation with commercial mobile devices. In fact, the beacon loss rate can be due to multiple reasons such as:

- the different versions of the OS;
- energy saving apps that can shutdown the Bluetooth interface or kill high-

computational apps used for our tests;

- the posture of the user and the receiver position;
- any environmental interference.

The combination of such factors can decrease, even remarkably, the overall beacon loss rate. In our case, we experienced values ranging from a minimum of 60.53% to a maximum of 95.45%. We found some device models performing better during some tests with respect to the others. As an example, the Nexus 6 based on Android 7.0, generally, provides beacon loss rate in the range between 60% and 80%, while Huawei P9 Lite or Samsung Galaxy S experience about 90%.

4.2. Symmetry of the Channel

We finally analyze the symmetry of the channel. This analysis focuses on measuring the differences between the RSS mutually estimated by a pair of devices. The goal is to measure how much the RSS values estimated by a devices' pair are similar during the Interaction stage. Similarly to the previous analysis, we focus on the Interaction stage, since the Approaching and Non Interaction stages are supposed to provide non comparable results. As a general example, the symmetry between user 1 and user 2 is given by the absolute value of the difference of the RSS values estimated by user1's device with respect to user2's device. The lower the difference, the higher their symmetry during the Interaction stage. The process we followed for measuring the symmetry is reported in Figure 6. Given a pair of devices, the first step is to slice the time

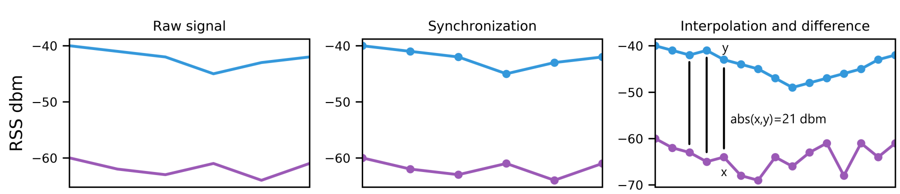

Figure 6: Re-sampling process of raw signals.

series of the RSS recorded by the pair at regular intervals. The second step, is to synchronize the time series, so that the RSS values are synchronized with respect to the time. Synchronization is obtained by aligning the RSS samples temporally, e.g. given the RSS sample (time, dbm): $(10.3s, -64dBm)$, it is synchronized to: $(10, -64)$. After the synchronization step, the time series of the device pair match with respect to the time. The third step is to interpolate any missing RSS sample and finally to build their differences. The interpolation estimates a RSS value by considering the previous existing samples; in our case we applied a linear interpolation by exploiting the previous RSS sample in order to estimate the value of the missing ones. The interpolation process is required since every device might record a different number of RSS values (as discussed with the beacon loss rate analysis). Finally, once the time series are both synchronized and interpolated, we measure their symmetry by computing the absolute value of the differences of the RSS values.

We measure the symmetry of the channel by considering all the possible combinations of configurations that we tested, in particular we distinguish between homogeneous symmetry: $C_1 - C_1$, $C_2 - C_2$, $C_3 - C_3$, $C_4 - C_4$ with respect to heterogeneous symmetry: $C_1 - C_2$, $C_1 - C_3$, $C_2 - C_3$. In the first case, we measure the symmetry only between subjects with the same posture and position of the receiver. For example, hand to hand, front pocket to front pocket or back pocket to back pocket. In the second case, we measure the symmetry between subjects with different postures and holding the device not symmetrically. The results concerning the symmetry are given as a distribution. In particular, we measure the symmetry for all the possible pairs of devices with homogeneous and heterogeneous combinations, as shown in Figure 7. The figure reports, for each combination of a configuration, a violin plot which combines the box plot with the KDE density estimation in order to better visualize the distribution's trend. On the left side of Figure 7, we report the homogeneous combinations (subjects with the same configurations), while on the right side we report the heterogeneous ones (subjects with different configurations). Each violin shows the median (the white dot), the $25th$ and $75th$ percentile (lower and upper bound of the inner box) of the distribution. The symmetry is generally higher with the homogeneous combinations with respect to the heterogeneous combinations, as shown by the median values and by the fat shape of the violins. In these cases, the median value of the symmetry ranges in the interval $[5-10]$ dbm, meaning that the RSS values recorded mutually by a pair of device differ at most by 10 dbm. However, we observe a high number of outliers in the homogeneous combinations, as shown by the sharp whiskers whose values are above —80— dbm. We further investigated such aspect and we found that one of the devices that we tested estimates the RSS of the beacons with higher values with respect to any other device but we are not able to clearly identify the cause except for the fact that it is the only one running version 5 of Android OS. Differently, with the heterogeneous combinations, the violin plots are more stretched, meaning that the symmetry is lower with respect to the homogeneous combinations. The symmetry ranges in a wider interval with respect to the homogeneous combinations, in particular $[5-30]$ dbm. Therefore, the RSS values of a devices' pair can differ by 30 dbm. As expected, subjects with different configuration can result with divergent RSS values collected by their devices. A meaningful example is given by the combination $C_2 - C_3$ in which one device is hold in BP and the other device is hold in H. In this case, the symmetry is the lowest we experienced during all our tests and the violin plot is stretched in a wide interval with a median value of approximately 20 dBm.

5. Detecting Social Interactions with SME-D Algorithm

As a first work towards the automatic detection of social interactions from commercial mobile devices and, in particular, through SocializeME framework, we designed an algorithm based on the analysis of the beacon messages collected in our experimental campaign. Specifically, we used the beacon loss rate and the RSS value experienced by each pair of users involved in an interaction as the main parameters of our algorithm called SocializeME Detector (SME-D). The algorithm has been originally proposed in [1] and we report here only its

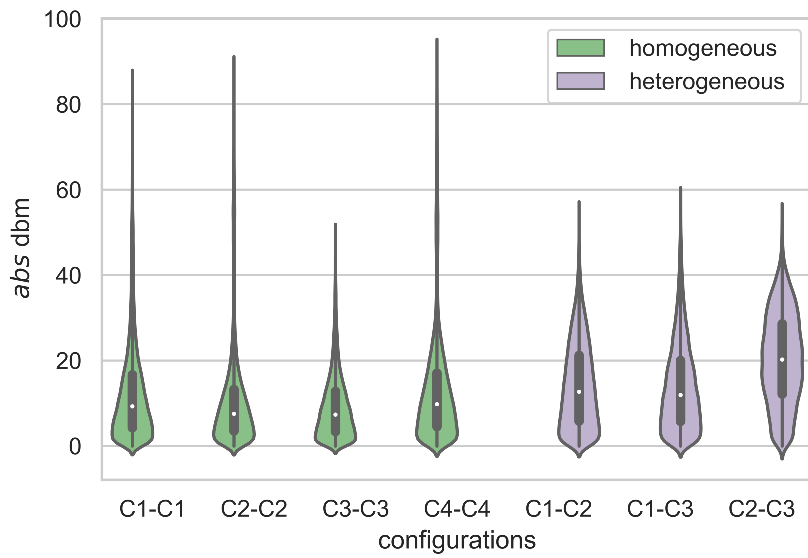

Figure 7: Symmetry of the channel with different configurations.

main characteristics and the preliminary results. We are currently working on possible enhancements. SME-D analyses the time series of beacon messages received by each dyad by using a sliding time-window of predefined duration (Δ_{up}), and in that window it evaluates the following conditions to identify the starting time of a social interaction (i.e., opening conditions):

- to receive at least $p\%$ of the expected beacons;
- the RSS of the received beacons is greater or equal to a threshold value T_{rss}

We consider that a social interaction is starting in that time window if the two conditions are verified at least in one of the two directions of the dyad. Then, we assume that the interaction is active until the closing condition is detected: the time interval between the last received beacon with $RSS \geq T_{rss}$ is greater than or equal to Δ_{down}.

We evaluated SME-D performances in terms of accuracy and F-score when applied to the experimental sessions obtained in C_1, C_2 and C_3 configurations and for different values of p and T_{rss}. This has been possible due to the ground-truth provided by the students while performing all the tests.

For each configuration we analysed the SME-D performances for p values equals to $3\%, 7\%, 11\%, 19\%$ and T_{rss} in the range $[-90, -70]$dbm. Values outside these ranges negatively affect the algorithm performances. In C_1 and C_2 configurations the accuracy increases as T_{rss} increases to reach the maximum value for -82 dbm for C_1 and -84 dbm for C_2, for all p values indicated above. Then the accuracy decreases. The situation is different in C_3 (ST-BP) since the body attenuation highly affect the overall performances. In this case there is no local maximum, but both accuracy and F-score decreases as T_{rss} increases.

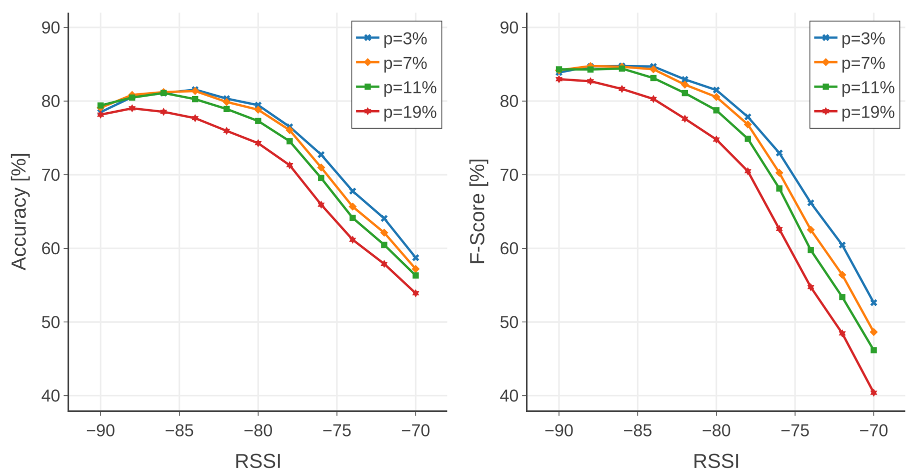

Figure 8: Performance of SME-D algorithm in terms of Accuracy and F-Score.

Since we cannot assume a priori the configuration of the users and their devices, to provide a more realistic analysis of the proposed algorithm we combined the results of all the configurations and we obtained a general trend similar to that obtained in C_1 and C_2 configurations. Specifically, SME-D obtains the best performances with $p = 3\%$ and $T_{rss} = -84$dbm. In Figure 8 we present the accuracy and F-score values obtained with this setting in the three configurations and the general one. Moreover, table 4 provides a summary of the evaluation

Table 4: Summary of the performance of SME-D in Session S_3.

	Test 1	Test 2	Test 3	All Tests
Accuracy [%]	85.32	87.19	72.21	81.56
F-Score [%]	88.54	89.28	75.67	84.7

metrics of SME-D. In particular, SME-D presents an overall accuracy of 81.56% and F-score 84.7%, which are promising results. To further evaluate SME-D in real scenarios, we should be able to collect large-scale traces in real conditions, which is not completely feasible at the moment considering the technical issues we encountered in the experimental campaign. However, the collected dataset, presented and detailed in [25], can be useful to evaluate possible enhancements of the proposed solution and to compare them with other algorithms proposed for the same purpose.

6. Conclusions

Social interactions represent a complex phenomena to be captured and studied with traditional methods, such as diaries and self-reported questionnaires. We exploit mobile sensing technologies, and specifically BLE beacon messages, to collect proximity information and detect face-to-face interactions in real-life scenarios. To this purpose, we conducted an experimental campaign involving students of a high school to use SocializeME framework on their personal mobile devices while emulating interactions in specific configurations. The students provided also the ground-truth of each experiment, including technical limitations and errors. This allowed us to identify limitations and challenges in using BLE messages as proximity and interaction information, and to define a dataset [25] to be available for the community. The experiments reproduce interactions by varying the posture of the subjects as well as the way they hold the device. We analyzed the quality of the wireless signals in terms of RSS distribution and beacon loss rate. We also studied the symmetry of the channel as a measure of similarity among the RSS mutually estimated by the involved devices.

The experience we gained in this study lead us to derive some considerations. Firstly, we adopted commercial mobile devices to collect BLE beacons in order to demonstrate their effectiveness for a future massive collection campaign. Such choice represents a challenging task since we experienced very different behaviors of the heterogeneous devices, which cannot be controlled but only mitigated, especially in terms of beacon loss rate. In addition, even though BLE is commonly adopted in all commercial devices, several of them do not support the advertising mode, not allowing the smartphone to send its own beacons, but only to be used as a receiver of BLE emitting tags.

Even though BLE tags are low-cost devices and highly configurable, we expect in the future that the advertisement mode will be homogeneously adopted in the commercial devices, and that the reduced signal quality, due to the emitter position, could be supported by additional information derived from the embedded sensing devices.

Finally, the user's posture and the position of the mobile device are two crucial factors in the system performances. They have been often ignored by the current literature by proposing customised hardware to be worn in the optimal position for the signal acquisition. However, they represent important characteristics of the collected data, especially to define the detection algorithm, as shown in the evaluation of the SME-D algorithm.

Appendix A. Details of the Experimental Sessions

We report on this appendix all the sessions' details. In particular, we report for each of the 6 sessions the IDs of the users, the device models, the version of the Android OS and the loss rate of each of the session's tests. We finally report in this appendix the box plots of the RSS values with 4 configurations: C_1 to C_4 and 3 stages: Interaction, Approaching and Non Interaction.

Table A.5: Session 1

Group ID	User ID	Device Model	Android Version	Loss Rate [%]		
				Test 1	Test 2	Test 3
1	15	Samsung Galaxy S7	7.0	90.47	89.93	92.77
	36	Nexus 6	7.0	64.77	63.78	67.76
2	23	Nexus 6	7.0	65.25	66.72	70.13
	26	LG G4	6.0	65.66	71.34	72.33
3	20	Huawei P9 Lite	6.0	81.69	81.39	88.11
	22	Nexus 6	7.0	70.43	71.63	74.47
4	41	Nexus 6	7.0	60.53	60.53	63.26
	42	Nexus 6	7.0	71.78	76.18	75.04

Table A.6: Session 2

Group ID	User ID	Device Model	Android Version	Loss Rate [%]		
				Test 1	Test 2	Test 3
1	4	Honor 8	7.0	90.17	89.34	85.12
	23	Nexus 6	7.0	70.99	69.27	67.48
	26	LG G4	6.0	62.89	67.82	73.35
2	44	Nexus 6	7.0	65.03	65.49	69.05
	45	Samsung Galaxy S5	6.0.1	90.92	91.71	92.95
	47	Samsung Galaxy J3	5.1.1	77.68	79.65	86.73
3	10	Nexus 6	7.0	74.59	69.64	72.32
	28	Nexus 6	7.0	74.86	70.82	72.29
	46	Huawei P9 Lite	7.0	82.51	83.18	84.87

Table A.7: Session 3

Group ID	User ID	Device Model	Android Version	Loss Rate [%]		
				Test 1	Test 2	Test 3
1	23	Nexus 6	7.0	75.50	75.52	80.54
	26	LG G4	6.0	82.15	82.31	81.60
	46	Huawei P9 Lite	7.0	85.30	85.75	85.35
	49	Nexus 6	7.0	77.38	74.60	80.94
	50	Nexus 6	7.0	68.37	71.33	70.71
2	22	Nexus 6	7.0	72.07	70.84	75.34
	27	Nexus 6	6.0	82.17	69.74	82.38
	29	Nexus 6	7.0	82.52	71.95	83.43
	48	Samsung Galaxy S5	6.0.1	95.30	93.78	95.45
	51	Huawei P9 Lite	7.0	78.25	77.48	94.33

Table A.8: Session 4

Group ID	User ID	Device Model	Android Version	Loss Rate [%]			
				Test 1	Test 2	Test 3	Test 4
1	4	Honor 8	7.1	78.96	70.88	71.64	76.75
	23	Nexus 6	7.0	78.17	72.88	73.41	75.21
	52	Samsung Galaxy S8	7.0	92.22	92.97	94.25	92.43
	53	Nexus 6	7.0	76.69	69.18	71.51	72.24
2	15	Samsung Galaxy S7	7.0	92.16	92.06	91.82	77.51
	17	Samsung Galaxy S7 Edge	7.0	91.66	91.89	92.23	78.78
	27	Nexus 6	7.0	70.05	73.84	69.62	69.00
	28	Nexus 6	7.0	69.09	72.99	74.62	71.14

Table A.9: Session 5

Group ID	User ID	Device Model	Android Version	Loss Rate [%]			
				Test 1	Test 2	Test 3	Test 4
1	4	Nexus 6	7.0	68.42	69.80	65.65	73.68
	23	Nexus 6	7.0	65.30	67.90	70.08	69.68
	26	LG G4	6.0	65.72	68.07	82.27	75.68
	53	Nexus 6	7.0	71.48	69.32	72.80	69.86
2	10	OnePlus 5	8.0	67.77	81.93	87.46	70.33
	22	Nexus 6	7.0	66.21	71.34	64.54	69.67
	28	Nexus 6	7.0	64.31	63.85	68.16	66.98
	29	Nexus 6	7.0	67.97	73.64	69.86	73.61

Table A.10: Session 6

Group ID	User ID	Device Model	Android Version	Loss Rate [%]		
				Test 1	Test 2	Test 3
1	3	Nexus 6	7.0	63.77	64.11	96.80
	4	Honor 8	7.1	72.50	97.86	71.63
2	15	Samsung Galaxy S7	7.0	92.27	92.25	83.88
	29	Nexus 6	7.0	66.69	92.25	65.13
3	10	OnePlus 5	8.0	68.52	67.61	69.55
	28	Nexus 6	7.0	66.47	97.50	88.33

Acknowledgement

This work is partially supported by the SocializeME project co-funded by "Fondazione Cassa di Risparmio di Lucca", a collaboration between CNR-ISTI/IIT and Polo Scientifico Tecnico Professionale E. Fermi – G. Giorgi.

References

[1] M. Girolami, F. Mavilia, F. Delmastro, E. Distefano, Detecting social interactions through commercial mobile devices, in: 2018 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops), IEEE, 2018, pp. 125–130 (2018).

[2] J. De Jong Gierveld, T. Van Tilburg, The de jong gierveld short scales for emotional and social loneliness: tested on data from 7 countries in the un generations and gender surveys, European journal of ageing 7 (2) (2010) 121–130, 20730083[pmid] (Jun 2010). `doi:10.1007/s10433-010-0144-6`. URL `https://www.ncbi.nlm.nih.gov/pubmed/20730083`

[3] Q. Chang, F. Sha, C. H. Chan, P. S. F. Yip, Validation of an abbreviated version of the lubben social network scale ("lsns-6") and its associations with suicidality among older adults in china, PLOS ONE 13 (8) (2018) 1–11 (08 2018). `doi:10.1371/journal.pone.0201612`. URL `https://doi.org/10.1371/journal.pone.0201612`

[4] B. Schiele, N. Oliver, T. Jebara, A. Pentland, An interactive computer vision system dypers: Dynamic personal enhanced reality system, in: Computer Vision Systems, Springer Berlin Heidelberg, Berlin, Heidelberg, 1999, pp. 51–65 (1999).

[5] A. Fathi, J. K. Hodgins, J. M. Rehg, Social interactions: A first-person perspective, in: 2012 IEEE Conference on Computer Vision and Pattern Recognition, 2012, pp. 1226–1233 (June 2012). `doi:10.1109/CVPR.2012.6247805`.

[6] P. Rota, N. Conci, N. Sebe, Real time detection of social interactions in surveillance video, in: A. Fusiello, V. Murino, R. Cucchiara (Eds.), Computer Vision – ECCV 2012. Workshops and Demonstrations, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 111–120 (2012).

[7] E. Hall, C. P. C. L. of Congress), The Hidden Dimension, Anchor books, Doubleday, 1966 (1966).
URL `https://books.google.it/books?id=zGYPwLj2dCoC`

[8] M. Aghaei, M. Dimiccoli, P. Radeva, With whom do i interact? detecting social interactions in egocentric photo-streams, in: 2016 23rd International Conference on Pattern Recognition (ICPR), 2016, pp. 2959–2964 (Dec 2016). `doi:10.1109/ICPR.2016.7900087`.

[9] X. Alameda-Pineda, J. Staiano, R. Subramanian, L. Batrinca, E. Ricci, B. Lepri, O. Lanz, N. Sebe, Salsa: A novel dataset for multimodal group behavior analysis, IEEE Transactions on Pattern Analysis and Machine Intelligence 38 (8) (2016) 1707–1720 (Aug 2016). `doi:10.1109/TPAMI.2015.2496269`.

[10] T. Choudhury, A. Pentland, Sensing and modeling human networks using the sociometer, in: Proceedings of the 7th IEEE International Symposium on Wearable Computers, ISWC '03, IEEE Computer Society, Washington, DC, USA, 2003, pp. 216– (2003).
URL `http://dl.acm.org/citation.cfm?id=946249.946901`

[11] J.-P. Onnela, B. N. Waber, A. Pentland, S. Schnorf, D. Lazer, Using sociometers to quantify social interaction patterns, Scientific Reports 4 (2014) 5604 EP –, article (Jul 2014).
URL `https://doi.org/10.1038/srep05604`

[12] J.-I. Watanabe, H. Atsumori, M. Kiguchi, Informal face-to-face interaction improves mood state reflected in prefrontal cortex activity, Frontiers in human neuroscience 10 (2016) 194–194, 27199715[pmid] (May 2016). `doi:`

10.3389/fnhum.2016.00194.
URL https://www.ncbi.nlm.nih.gov/pubmed/27199715

[13] H. Hung, G. Englebienne, L. Cabrera Quiros, Detecting conversing groups with a single worn accelerometer, in: Proceedings of the 16th International Conference on Multimodal Interaction, ICMI '14, ACM, New York, NY, USA, 2014, pp. 84–91 (2014). doi:10.1145/2663204.2663228.
URL http://doi.acm.org/10.1145/2663204.2663228

[14] T. M. Ciolek, A. Kendon, Environment and the spatial arrangement of conversational encounters, Sociological Inquiry 50 (3-4) (1980) 237–271 (1980). arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1475-682X.1980.tb00022.x, doi:10.1111/j.1475-682X.1980.tb00022.x.
URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1475-682X.1980.tb00022.x

[15] M. Génois, A. Barrat, Can co-location be used as a proxy for face-to-face contacts?, EPJ Data Science 7 (1) (2018) 11 (2018). doi:10.1140/epjds/s13688-018-0140-1.
URL https://doi.org/10.1140/epjds/s13688-018-0140-1

[16] N. Eagle, A. (Sandy) Pentland, Reality mining: Sensing complex social systems, Personal Ubiquitous Comput. 10 (4) (2006) 255–268 (Mar. 2006). doi:10.1007/s00779-005-0046-3.
URL http://dx.doi.org/10.1007/s00779-005-0046-3

[17] J. Fournet, A. Barrat, Contact patterns among high school students, PLOS ONE 9 (9) (2014) 1–17 (09 2014). doi:10.1371/journal.pone.0107878.
URL https://doi.org/10.1371/journal.pone.0107878

[18] C. Cattuto, W. Van den Broeck, A. Barrat, V. Colizza, J.-F. Pinton, A. Vespignani, Dynamics of person-to-person interactions from distributed rfid sensor networks, PLOS ONE 5 (7) (2010) 1–9 (07 2010).

doi:10.1371/journal.pone.0011596.
URL https://doi.org/10.1371/journal.pone.0011596

[19] R. Mastrandrea, J. Fournet, A. Barrat, Contact patterns in a high school: A comparison between data collected using wearable sensors, contact diaries and friendship surveys, PLOS ONE 10 (9) (2015) 1–26 (09 2015). doi:10.1371/journal.pone.0136497.
URL https://doi.org/10.1371/journal.pone.0136497

[20] V. Sekara, S. Lehmann, The strength of friendship ties in proximity sensor data, PLOS ONE 9 (7) (2014) 1–8 (07 2014). doi:10.1371/journal.pone.0100915.
URL https://doi.org/10.1371/journal.pone.0100915

[21] A. Stopczynski, V. Sekara, P. Sapiezynski, A. Cuttone, M. M. Madsen, J. E. Larsen, S. Lehmann, Measuring large-scale social networks with high resolution, PLOS ONE 9 (4) (2014) 1–24 (04 2014). doi:10.1371/journal.pone.0095978.
URL https://doi.org/10.1371/journal.pone.0095978

[22] K. Katevas, K. Hänsel, R. G. Clegg, I. Leontiadis, H. Haddadi, L. N. Tokarchuk, Finding dory in the crowd: Detecting social interactions using multi-modal mobile sensing, CoRR abs/1809.00947 (2018). arXiv:1809.00947.

[23] V. Arnaboldi, M. Conti, A. Passarella, F. Pezzoni, Analysis of ego network structure in online social networks, 2012 (09 2012). doi:10.1109/SocialCom-PASSAT.2012.41.

[24] C. A. Kerrache, N. Lagraa, R. Hussain, S. H. Ahmed, A. Benslimane, C. T. Calafate, J. Cano, A. M. Vegni, Tacashi: Trust-aware communication architecture for social internet of vehicles, IEEE Internet of Things Journal 6 (4) (2019) 5870–5877 (2019).

[25] M. Girolami, F. Mavilia, F. Delmastro, Bluetooth low energy dataset for the analysis of social interactions with commercial devices, Data in Brief under review.

[26] P. Barsocchi, Position recognition to support bedsores prevention, IEEE Journal of Biomedical and Health Informatics 17 (1) (2013) 53–59 (Jan 2013). doi:10.1109/TITB.2012.2220374.

[27] P. Barsocchi, G. Oligeri, F. Potorti, Measurement-based frame error model for simulating outdoor wi-fi networks, IEEE Transactions on Wireless Communications 8 (3) (2009) 1154–1158 (March 2009). doi:10.1109/TWC.2009.060475.