



HAL
open science

Souk: Spatial Observation of hUman Kinetics

Marc-Olivier Killijian, Roberto Pasqua, Matthieu Roy, Gilles Trédan,
Christophe Zanon

► **To cite this version:**

Marc-Olivier Killijian, Roberto Pasqua, Matthieu Roy, Gilles Trédan, Christophe Zanon.
Souk: Spatial Observation of hUman Kinetics. *Computer Networks*, 2016, 111, pp.109-119.
10.1016/j.comnet.2016.08.008 . hal-01372329

HAL Id: hal-01372329

<https://hal.science/hal-01372329v1>

Submitted on 27 Sep 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

SouK: Spatial Observation of hUman Kinetics

Marc-Olivier Killijian, Roberto Pasqua, Matthieu Roy*, Gilles Tredan*, Christophe Zanon

LAAS-CNRS, Université de Toulouse, CNRS, Toulouse, France

Abstract

Simulating human-centered pervasive systems requires accurate assumptions on the behavior of human groups. Recent models consider this behavior as a combination of both social and spatial factors. Yet, establishing accurate traces of human groups is difficult: current techniques capture either positions, or contacts, with a limited accuracy.

In this paper we introduce a new technique to capture such behaviors. The interest of this approach lies in the unprecedented accuracy at which both positions and orientations of humans, even gathered in a crowd, are captured. The open-source software pipeline we developed to exploit captured data allows extraction of several metrics on movement and social contacts, and permits study of their respective interrelationship. From the mobility to the topological connectivity, this framework offers a layered approach that can be tailored, allowing to compare and reason about models and traces.

We demonstrate the accuracy and validity of our approach on social events and calibration runs in which we captured the motions of humans. In particular, we introduce an open-access trace of 50 individuals and compare it against random waypoint models that have the same global characteristics. Our fine-grain analyses, that take into account social interactions between users, show that the random way point model does not provide accurate predictions for socially-induced motion; to model human kinetics, new group- and interaction-based models should be developed. From the computer science point of view, these models are required to fully exploit the power of human-centered mobile computing, crucial for ubiquitous computing, and referred to as Short Range Communication Systems, Mobile Opportunistic Networking, or Mobile Networking in Proximity.

Keywords: Human-centered computing, Reliability and robustness, Human mobility modeling, Ubiquitous computing, Mobile computing.

1. Introduction

During the past two decades, the problem of understanding human mobility has received a growing attention from the research community. Thanks to the widespread use of mobile handheld devices, large scale datasets have been produced and successfully exploited to provide a wealth of results characterizing human mobility patterns. Applications are multiple: from

*Corresponding authors: roy@laas.fr, tredan@laas.fr

street planning to epidemics modeling, every hint about how humans move is a powerful ally for designing tomorrow’s information society.

Indeed, one of the most crucial parameters of a ubiquitous system that relies on the network formed by users’s devices —often referred to as Mobile Opportunistic Networking, or Mobile Networking in Proximity (MNP)— is its users contact model. Available mobility traces are usually coarse grained and do not allow to precisely emulate short range communication topologies. Contact traces exist, but these are usually established using wireless technologies themselves. The problem of such approaches is genericity: how to simulate a Bluetooth communication topology using an RFID contact trace, and vice-versa? Due to the wide variety of wireless technologies and their rapid evolution, it is of prime importance to establish datasets that are technology independent.

The idea that users’ mobility and social contacts are connected has recently given rise to the development of mobility models taking these two dynamics in consideration. But in the absence of traces capturing both interactions and movements, such models remain only partially validated.

One of the fundamental question that is left unanswered is “*what is a good analytical model for crowd connectivity, required to implement the MNP paradigm?*”, and, as a corollary, “*how to validate models?*”.

To that end, we present Souk— Social Observation of hUman Kinetics. This platform allows the precise capture, in real time, of both the *position* and the *orientation* of individuals in a dense region. To achieve this, each individual is equipped with two lightweight wireless tags that are localized with a 15 cm accuracy using a network of sensors. More precisely, we present Souk as a means to test the realism of existing and future human mobility models.

Because of Souk’s ability to capture users’ localization, orientation and interactions in dense crowds, social events (*e.g.* meetings, cocktails, concerts) constitute perfect use cases for Souk. More precisely, the level of accuracy attained when capturing localization and orientation data allows Souk to infer interactions between users in a crowd. The dynamic network of social interactions arising during the social event can thus be computed, exposed in real-time and logged for off-line analysis. To the best of our knowledge, this is the first time that both social interactions and movements are assessed at such a fine granularity and at such a scale.

Contributions. In this paper, we introduce Souk, a platform to capture the behavior of a crowd at an unprecedented scale and resolution. Traces captured using this platform allow access to both the position and the social contacts of individuals, allowing precise simulation of any Short Range Communication topology.

We provide a set of tools to compare models and reality using a wide variety of metrics, in a layered approach. We believe this tool chain could be of prime importance to develop realistic models and compare them to reality.

We showcase our approach by comparing the random waypoint model against an experimental deployment of the platform on 50 individuals during a social event. Notice that the drawbacks of random waypoint models have already been pinpointed before, and the comparison is only presented to demonstrate how the proposed approach allows to easily compare a given model against reality.

The paper is organized as follows: the following section discusses related works. Then, we present in Section 3 the Souk experimental platform. Section 4 describes different experiments we conducted and associated results. The last section concludes the paper and exposes some trails for future work. A preliminary version of this article appeared in [17], and described initial

experiments conducted in 2012 only. The current article aims at providing a comprehensive and complete view of our framework, on the software and hardware points of view, including calibration and assessment of accuracy.

2. Typical Mobility and Interaction Analyses

Several mobility data collection campaigns that have been conducted and are published in the CRAWDAD project¹. Table 1 presents a selection of datasets that include some flavor of mobility in captured data. These campaigns use Off-The-Shelf hardware, such as smartphones, to capture information, thus their localization source is either a GPS system or based on wireless interfaces.

Compared to the datasets we capture, the major difference lies in the scale: Souk’s dataset has a smaller scale (i.e., building-wise vs. town-wise, and short term vs. long-term) but provides a higher accuracy (i.e., in the order of 10cm vs. 10 – 100m) and includes users orientation, thus enabling a precise capture of social interactions between users.

Since Souk aims at capturing micro-mobility whereas all other datasets in Table 1 are interested in macro-mobility, these two types of datasets are complementary: on the one hand, understanding micro-mobility and fine-grained interactions between users requires the accuracy offered by a localization system similar to the one we used, and allows a better understanding of connectivity patterns within a pedestrian-carried mobile system. On the other hand, long-term evolution of systems and recurrent behaviors can only be captured on platforms that are large scale both in terms of time and space [2].

Dataset	Users	Duration	Resolution
Yellow cabs	100	1 month	GPS: 1/min
Reality Mining	100	9 months	BlueTooth: 12/h
UIUC-UIM	28	3 weeks	BlueTooth: 1/min, WiFi: 1/h
Nokia	200	1 year	GPS: user-defined frequency
Yonsei/Lifemap	8	2 months	GPS, WiFi 12 – 30/h
Souk	50	1.5h	Ultra-Wide Band: 10cm at 1Hz

Table 1: Typical available mobility/interaction datasets

Indoor positioning has been a very active area of research in ubiquitous and pervasive computing. Much effort has been spent on developing indoor localization technologies: from the original Cricket system [25] that used both radio and ultrasonic signals, to more recent systems using power lines [23], Ultra-Wide Band signals [1], digital cameras and SLAMs [18], CDMA mobile phone fingerprinting [28], resonant magnetic coupling [24], etc.

The study, and modeling, of the relationship between human mobility and social aspects of human behavior has recently gained a lot of attention. In particular, much effort is spent in developing socially inspired mobility or propagation models [14, 20, 6, 5, 22, 26]. In these works, positioning is not necessarily of primary interest but, rather, access to data concerning contacts or proximity between the individuals is necessary. Several technologies and methods have been used to collect or infer social contacts. Bluetooth and WiFi networks were used

¹The CRAWDAD project: crawdad.cs.dartmouth.edu

in various environments : in offices [19, 10], conferences [15], during a roller trip [3], on a campus [11] or in shopping malls [12]. The main limitation of these experiments lies in the fact that contacts are inferred when two devices are co-located or in communication range. Accuracy of this inference can be questioned. In the same manner, dead reckoning can be used to estimate relative positions between individuals [16] but has the same drawbacks. RFIDs may be used to record contacts when individuals are engaged in face-to-face interaction [9, 26], without knowing users positions — notice that tags have to really face each other for a long enough period of time and thus, some interactions can be missed.

Recent works have focused on analyzing interactions in crowds [8, 7]. Yet, to the best of our knowledge, this is the first time that such an accurate and precise framework for capturing *both positions* and *contacts* is produced in the context of dense populations.

3. Souk: the Experimental Platform

Souk consists of three parts: an experimental platform to capture the position and orientation of mobile individuals through wireless *tags*, a framework to develop mobility models, and a software system that exploits the output of either the capture process or model-generated traces. In a nutshell, both mobility models and the experimental platform can feed a database that is then accessed by the software pipeline. The use of a database between the production and the exploitation of positions ensures the repeatability of the experiments, and a certain degree of genericity: any model or positioning system can be used, for real-time exploitation of data or for later off-line analysis.

3.1. Position capture platform

The experimental platform relies on two types of elements: a sensor network, and a set of tags. Currently, Souk relies on a Ultra Wide Band (*UWB*) based localization system developed by Ubisense [27]. The whole system can be deployed within a day, and the *UWB* technology allows Souk to work seamlessly in crowded environments. Since any system providing a similar localization accuracy would fit, the precise whereabouts concerning the hardware system are out of the scope of this paper. However we briefly present the system and its impact on the experiment.

3.1.1. Tags

A tag is localized by the system in a three dimensional space. Each participant is equipped with two tags, so that along with the position of the participant, we are able to compute the orientation of his body.

Each tag dimension is approximately $4 \times 4 \times 1.5$ cm and weights 25 g, hence users tend to forget about them and the impact of wearing the tags on users behaviors is rather limited.

Tags communicate with sensors using *UWB* wave trains, and although these wave trains can traverse the human body to a certain extent, we chose to locate the tags on participants' shoulders to limit obstacles, as shown on Figure 1.

Moreover, since tag localization accuracy is about 15 cm (symbolized by the red halos on Figure 1), it is important to put sensors as far apart as possible from each other, in order to attain a sufficiently reliable orientation.

3.1.2. Sensors

Each tag periodically beacons a UWB pulse train. Sensors, which have fixed positions in the environment, locate the tag using both the angle of arrival and the time difference of arrival principles. Therefore, sensors need to be tightly synchronized together, and precisely positioned and oriented in space. The practical range of a sensor is a cone of 90° and 25 m depth. At least two sensors are required to provide an accurate position, but, in practice, the more sensors that compute the localization, the more precise and robust the position is (in our current setting, we use a set of 6 sensors).

Sensors use a Time Division Multiple Access (*TDMA*) channel access method. Indeed, both UWB and traditional WiFi signals are used to provide synchronization in the system. WiFi is used to negotiate the TDMA slots allocated to tags: every tag communicates using 802.11 signals with a *master* sensor that provides it with a unique TDMA slot. As soon as a tag gets allocated a TDMA slot, it starts to beacon pulse trains in its slot in the UWB range.

The number of TDMA slots per second is a fixed parameter given a set of sensors: let c be the number of communication slots per second, n the number of attendees in the monitored event (requiring $2 \times n$ tags), and f the position refresh frequency, in Hertz. All these parameters are bound by the relation: $f \leq \frac{c}{2n}$. In our setup, the highest rate available on the system is $c = 128$. As a consequence, to obtain a position refresh rate of 1 Hz, we limit to at most 64 participants. Notice that breaking this 64 participants barrier is just a matter of resources: duplication of sensors doubles the systems capacity; more generally, the *cost of the system is linear* with respect to the number of equipped participants. As of 2016, the price of the position capturing system for n users is roughly $\lceil n/64 \rceil \times 20k\$$.

3.2. Mobility model framework

When devising a mobility model for humans, it is fundamental to question its accuracy. To that aim, we provide researchers with a large generic toolbox of metrics that can characterize both models and experiments. Souk allows using this toolbox regardless of the model-generated or experimental nature of the data, allowing a fair comparison between models and experiments.

In the framework, it is simple to develop a mobility model to use it as an input of the analysis process. It is thus possible to check generated traces against real, captured ones. More interestingly, it is possible to use model-generated traces as an input of the analysis pipeline, and make comparisons on higher level metrics, as we show in the next subsection, rather on low-level traces that cannot be easily compared.

Practically, developing a new mobility model is as simple as extending the provided random way-point model.

3.3. Analysis pipeline

This part constitutes the heart of Souk. It consists of a set of software bricks that dynamically retrieve data from the database, filter results, infer contacts, and maintain various statistics. Figure 2 illustrates the information processing pipeline used in Souk. The layered approach used allows partitioning the analysis in different bricks that are detailed below.




Figure 1: Each participant carries two lightweight tags, clamped on clothes on her/his shoulders.




Figure 2: Overview of the Souk processing pipeline

3.3.1. Snapshot creation

The localization system produces a flow of tag positions (more precisely, $(\text{tagId}, x, y, z, \text{timestamp})$ tuples) that need to be transformed into participants position/orientation couples. This is the role of the snapshot creation module, that periodically polls the database for new data, associates sensor pairs to each participant, and computes participants orientation.

Snapshot is an important abstraction in Souk. Intuitively, it provides a series of “pictures” of the system as if it were periodically frozen. Such a discretization of system states dramatically simplifies data processing.

An important feature of this mechanism is to handle possible missed position reports. Indeed, due to the high contention on the wireless medium, or to possible obstructions, some position reports might be missed. As a result, one of the two tags locating a participant may have no position known at a given round. The snapshot creation module implements a synchronization and filtering policy to handle such cases: periodically, when a snapshot is created, the last known position is used for every tag, mitigated by a classical aging policy — when a tag has no position known for Δ time units, it is considered to be out of the system. From our first experiments, we found that setting $\Delta = 3\text{s}$ provides good filtering properties while still guaranteeing sufficient temporal and spatial accuracy.

3.3.2. Contact extraction

Every snapshot is then transferred to the contact extraction module. Its role is to decide whether two given participants (or the devices they carry) a and b are potentially engaged in a contact. This decision is based on the analysis of participants respective positions and orientations, and depends on the nature of the contact that one seeks to capture, and therefore on the finality of the study.

In this paper, we present two different contact extractions, corresponding to two different use cases for Souk. One extracts the communication topology of *devices* carried by users and communicating wirelessly at short range. The other extracts social interaction between *users*. Although defining what is a social contact is far from being trivial —several definitions can coexist— Souk modular architecture allows to implement and test heuristics corresponding to those different definitions. In the next section, we showcase different techniques: a device-based unit-disc graph wireless contact extraction, and a two user-based social contact extraction techniques.

3.3.3. Interaction extraction

Once contacts have been extracted, they are exploited by the interaction module to decide whether a and b are engaged in an interaction. The interaction module keeps an history of the previous contact graphs, and uses information filtering techniques to transform contact graphs into an interaction graph.

This step is very important for two reasons. First, recall that information about tags positioning is noisy, and therefore contact graphs might be biased. Therefore this information filtering step allows to remove spurious interactions caused, for instance, by inaccurate participant orientation. Second, a contact is not necessarily the beginning of an interaction: two people crossing each other because they travel opposite directions in the same corridor will appear for a short period of time as being in contact, but this scenario cannot be considered as an interaction.

Finally, each pipeline step can be decorated with observers that perform additional functions. We used this feature to export topologies to a database and to the ns-2 simulator, to stream the graph to a graph rendering software in real-time for visualization purposes, and to compute statistics.

4. Experimentations and Results

4.1. Calibration of the Souk Platform

Extracting social interactions from a set of $(position, orientation)$ couples is not a trivial task. This is particularly the case when one considers the noisy nature of the dataset, and the social nature of the observed process: as demonstrated by Hall [13], the interaction distance is highly impacted by multiple social parameters, such as the interacting subjects gender, their originating country, or the size of the room.

To parameterize the different contact detection methods, we therefore defined the following experimental process. A group of 20 volunteers was equipped with tags and colored badges. These colored badges, clearly visible from anyone, defined 4 different classes of participants. Participants were asked to alternatively walk around for a given period, and then chat only to members of their own color class, and then walk again, and so on. For each of the 15 experiments, a subset of the participants was asked to evolve in a specific area of the experimental space. Notice that participants were not asked to discuss with all the other members of their color class, they could for instance form smaller subgroups.




Figure 3: Illustration of the contact detection parameters

A social interaction between a and b is modeled as a sustained visual contact between participants. The “visual contact” is parametrized by an angle α , which is the maximal tolerated deviation between a ’s (resp. b ’s) shoulders orientation and the line-of-sight from a to b (resp. b to a). Due to the high density of users in a crowd, we also require a and b to be close enough, namely closer than a given threshold D . We require this contact to be observed during at least C snapshots over the last W snapshots. Figure 3 summarizes this detection process.

Given the a priori knowledge of which interactions should take place (assuming participants correctly respected the interaction color classes rules), we were able to run our contact detection model on the experimental dataset to assess the impact of the model parameters. For each tuple of α , D , W and C values, we extracted the detected social interactions and classified them using the following metrics:

- **me**: movement error: a contact was detected between two participants, while they were moving.
- **ie**: interaction error: a contact was detected between participants that belong to two different color groups.
- **ic**: interaction – correct: a contact was detected between participants that belong to the same color class during a chat phase.

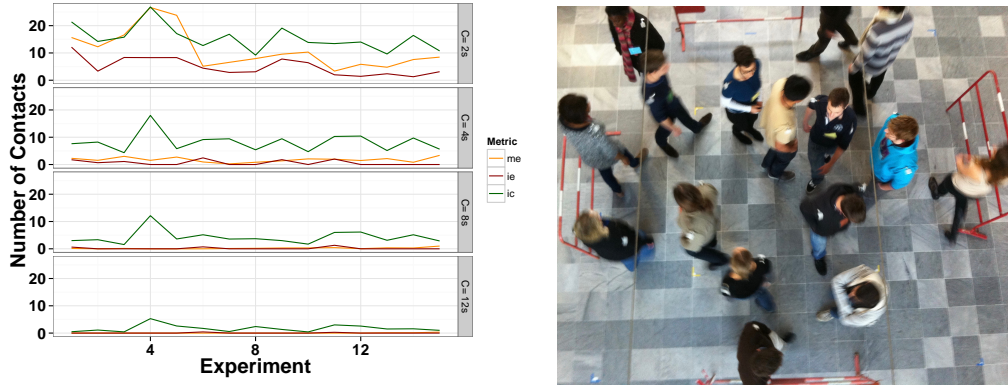


Figure 4: *Left*: Impact of C on detection quality, for each experiment. It shows the number of detected contacts per snapshot, averaged on the whole experiment. *Right*: Picture of experimental setup

Figure 4 illustrates the results obtained during this process. It represents the evolution of ic , ie and me metrics for $\alpha = 60^\circ$, $W = 12$ s and $D = 1$ m.

Each experiment is composed of 10 alternations of two phases: 30 seconds of movement and 15 seconds of interaction in variable sizes enclosures. During the movement phase, participants are all required to move randomly in the enclosure. During the interaction phase, participants are required to engage in a conversation, but only with a member of their group, materialized as color tags (see Figure 4). In experiments 4,5,9,10,14,15, a total of 16 participants are randomly assigned to 4 groups, in the other experiments, 12 participants are randomly assigned to 3 groups. Experiments 1 – 5, 6 – 10 and 11 – 15 were conducted in respectively $4m^2$, $16m^2$ and $100m^2$ square enclosures.

Experiment 4 constantly yields to many more detected contacts as it was the experiment with highest number of participants in the smallest space, therefore leading to much more contacts. Interestingly, the number of correctly detected contacts is constantly higher than both ie and me . While me is often higher than ie , especially with smaller contact threshold C , it can also be considered as a more benign error, as detecting immobility can be done efficiently. Although the details of such an immobility detection are beyond the scope of this study, we managed to detect immobility with 97% accuracy on this dataset without relying on other sources of data.

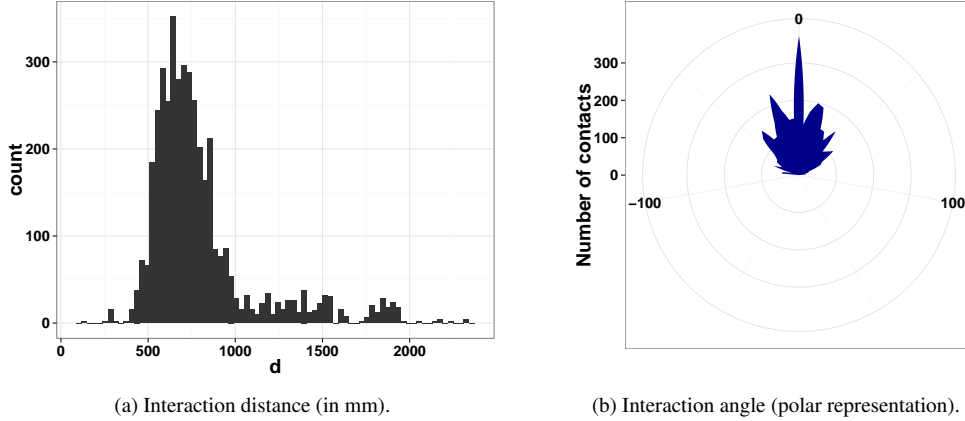


Figure 5: Interaction parameters

Given the good confidence we have in the correctly detected links, we can now explore the spatial characteristics of the social interactions (again, these results may apply solely in France, for educated young volunteers — these results are likely to vary in other contexts).

Figure 5a represents the interaction distance distribution, for $W = 12$ s, $C = 4$ s, and $\alpha = 60^\circ$. For instance, it shows that 185 contacts happened between 50 cm and 58 cm. The vast majority of the interactions happened between 50 cm and 1 m.

Figure 5b represents the distribution of the interaction angle for correctly detected contacts, using $W = 12$ s, $C = 4$ s, and $D = 1$ m. Interestingly enough, one can notice 3 distinct spikes in the distribution, which correspond to interacting dyads (face-to-face, null angle), triads ($-\pi/6$ or $\pi/6$ angle) and tetrads ($-\pi/4$, 0 or $\pi/4$ angle): angles in parenthesis correspond to the angle between a node looking at the barycenter, or center of mass, and the direction of another node in regular polyhedra of 2, 3 and 4 nodes respectively.

Finally, in order to attain both a good detection probability and a relatively low false-positive rate, we set the following parameters: $W = 12$ s, $C = 4$ s, $\alpha = 60^\circ$ and $D = 1$ m.

4.2. Experimental setup

We present here two measurement studies, $M1$ and $M2$, that were conducted on July, 4th and July, 5th 2012 during two receptions following the inauguration of a new building. The attendance of both receptions was a mix of scientists, journalists, and representatives of local institutions. More than 0.6 million position reports were collected during these events².

In the first experiment, $M1$, we deployed 116 tags, equipping 58 out of around 100 participants (the number of participants was continuously evolving during the experiment, 100 participants is an estimation of the maximum number of attendees at the beginning of the experimentation). In the second experiment, $M2$, we deployed 126 sensors, equipping 63 out of approximately 70 participants.

²Both datasets can be downloaded freely from <http://projects.laas.fr/souk>. They are distributed with a creative commons licence for further reuse.

In both $M1$ and $M2$ we collected approximately 90 minutes of data. We provided any volunteer with a pair of tags and briefly explained them the scientific aim of this experiment. In $M2$, only one participant refused to take part to the experiment for privacy concerns, and the difference between $M1$ and $M2$ tag counts can be explained by the number of available functional tags. For privacy reasons, we did not keep track of who picked which tags.

In the following, we focus on experiment $M2$, because it provides the best coverage in terms of equipped attendees.

The room used for experiments is approximately a $10\text{m} \times 10\text{m}$ square zone that is represented in Figure 6a. The room has 4 exits, and the South-East corner is occupied by a staircase that was unreachable during the reception. Attendees were able to get food and drinks from two large buffets (marked “Buffet” on the figure). Since these experiments also had the objective of demonstrating some of the research activity led in the lab, a live representation of the collected data was continuously exposed on a large screen (marked “Visualization”). The coordinate system is indicated with the arrows on the figure (z axis omitted).




Figure 6: Capture environment: physical map and noise map

4.3. System accuracy and results sanitizing

Before analyzing captured data, we first present some results related to the accuracy of collected data, and explain simple sanitizing operations performed on the database.

First of all, we monitored the position of several immobile tags scattered in the room. The distribution of the collected positions for each tag provides an interesting way to evaluate Souk’s accuracy. The measurement error follows a normal distribution. Figure 6b represents the standard deviation of the tags position depending on their position in the room. One can observe that *i*) positioning is more precise in the center of the room, and *ii*) global accuracy is within specification (the median standard error is 16.3 cm, and 3rd quartile is at 26.1 cm).

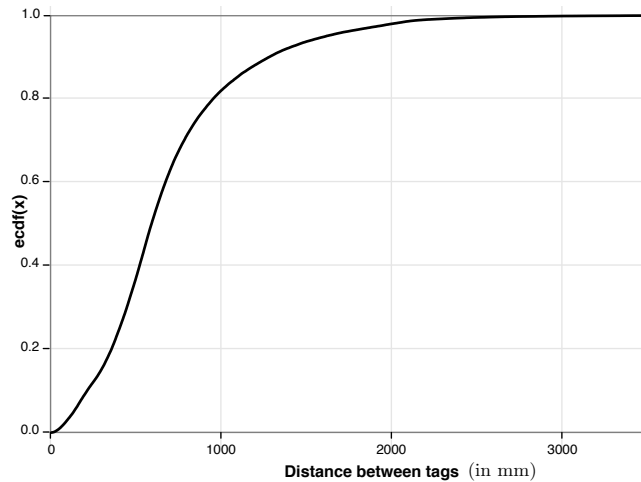


Figure 7: Expected cumulative distribution function of inter-tag distance

To get a better idea of tags’ location accuracy, we can observe how left and right tags are relatively located. Figure 7 represents the Expected Cumulative Distribution Function (ECDF) of tags distance. As a back-of-the-envelope approximation, Leonard’s *Vitruvian Man* sets the ideal shoulder width to a quarter of body height. In our country, males average height is ≈ 177 cm, and therefore we should observe an average distance of 44 cm. The observed average distance, however, is around 70cm and the median is 59 cm. This deviation can be explained by the measurement process: since both tags report individually, and since position reports might get lost, there can be a temporal shift of up to 3 s between left and right tags position reports. Thus, when subjects move, their shoulder width is measured bigger than it actually is.

Assuming all attendees have shoulders at the same height, and since tags report individually, tags height can be used to assess the precision of the system. The vast majority of position reports exhibit a very small tag pair height inconsistency- 40 mm for 92% of the pairs. The remaining 8% pairs are discarded as outliers. After these filtering steps, we retain 49 traces that meet our criteria of length and accuracy. The remainder of this paper only considers two-dimensional data.

4.4. Movement patterns: exploiting snapshots

In this section, we showcase possible exploitation of Souk’s snapshots. As stated before, a snapshot is an instant picture of attendees positions. Transforming the continuous flow of positions delivered by the hardware platform into a series of time-discrete snapshots is an operation that simplifies a lot the following processing operations, both computationally and conceptually, at the price of some accuracy loss. However, since Souk captures the raw output of the hardware platform, the impact of different snapshot creation policies can be tested and compared offline. In the following, a snapshot is produced every 3s.

4.4.1. Movement patterns

The first question one might ask is “do we all move identically?”. Figure 8a partly answers this question by representing the attendees sorted by traveled distance. The median walked distance is about 150m, with 1st and 3rd quartiles being respectively 115m and 200m.

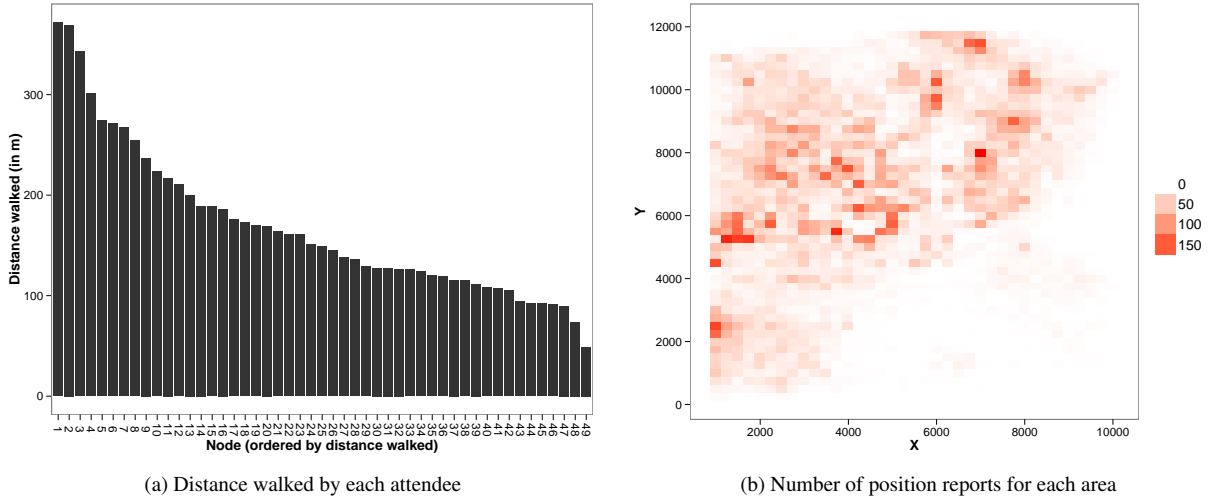


Figure 8: Movement patterns

Another interesting question concerns how people spread out in the experimental room. Figure 8b answers this question by representing the spatial density of position reports. The uneven nature of the distribution is the first striking difference with the outcome of a random waypoint mobility model, that would exhibit a more uniform positions density —of course, this is only an illustration, since the presence of points of interest would require a more refined discussion on density of users, see e.g. [4].




Figure 9: Number of pauses by attendee

The temporal aspect of the movement behavior can be explored using snapshots and some

basic signal processing. For instance, one can compare the position of attendees using a sliding window, considering that an attendee is immobile if she/he spends 7 out of the last 10 snapshots within a 40cm radius range. We call this a “pause”. It turns out that people’s pausing behavior shows a high degree of variability, as illustrated by Figure 9. It represents the number of such detected pauses by attendee. The median number of pauses is 17. If one had to parameter a random way point model to represent such behavior, the average behavior would be “walk for 9.3 snapshots” and “pause for 45.6 snapshots”. Yet, all the extracted indicators related to the pausing behavior demonstrate that this average is not representative of any attendee.

4.4.2. Destinations

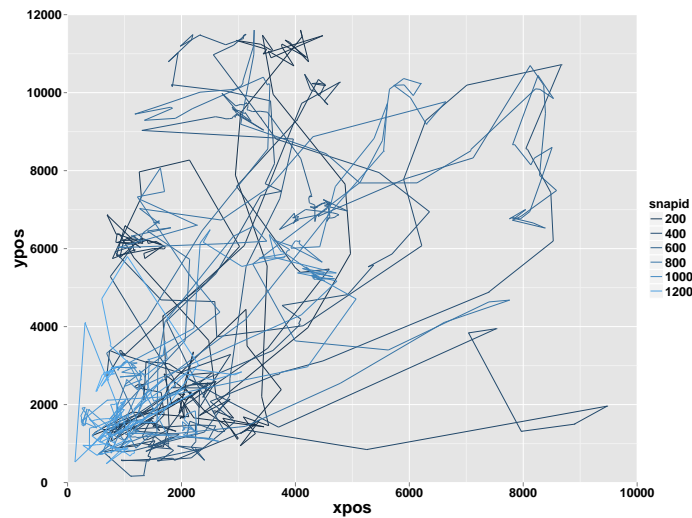


Figure 10: A sample trajectory. Color depicts time, black being the time origin and light blue the end of the experiment.

A simple exploitation of the snapshot flow is to obtain individual trajectories, like the one displayed on Figure 10. This figure represents the trajectory of one of the authors during *M2*, with snapshots taken every 3 seconds. These trajectories can then be exploited in many application-dependent ways. As an example, it is possible to measure the time spent by each attendee in precisely defined zones, in order to quantify attendees’ habits and points of interests.

Figure 11 represents the time spent by each attendee on each buffet (“Buffet” zones on Figure 6a, using the same color scheme). Line 17 can be read as: “on the total buffet time spent by this participant, 63% was in Buffet 1, and 37% in Buffet 2”. Surprisingly we can clearly observe non uniform preferences between participants: whereas the affluence of both buffets is rather even, individual results suggest that each attendee has a preferred buffet, even though both buffets were providing the exact same set of food and beverages. Among the extreme profiles that spent most of their time at Buffet 1 or Buffet 2, two are waiters that were participating in the experiment.

Overall, all these observations underline non-uniform processes. None of the observed distributions are likely to be observed using a random way-point model, arguably at the exception of the walked distance. These results suggest that (i) crowds spread not uniformly, but around some points of interests (POIs), (ii) the attractivity of POIs varies from one individual to the other and (iii) behavior in between POIs has a huge variability. Although devising a mobility model is not




Figure 11: Time spent by each attendee on Buffet 1 and Buffet 2

the aim of this paper, it raises the question whether patterns of mobility of different users are independent.

4.5. Contacts

By analyzing snapshots produced by Souk, we can use a device-based model to explore the topology of a system consisting of devices carried by users, and we can use a user-based model to dig some information about the interpersonal relations that took place during the experiments. As we show hereafter, the latter is important to model the former.

We developed three simple models of what a contact is. Defining social contacts is far beyond the scope of this paper, but we provide these examples in order to illustrate the potential uses of the platform. Incorporating a new contact model to Souk is a matter of minutes, the only limitation being that such a contact model has to be solely based on position and orientation of individuals.

4.5.1. Device-based

The simplest model considers that a and b are attendees' devices and detects whether these are within wireless contact range r . In this case, the simplest approach is to decide upon the distance between them, using a *unit disc* wireless communication model (*i.e.* we count one link iff $d(a, b) < r$).

Figure 12 represents the number of wireless links obtained using this simple model. Apart from a connectivity drop around snapshot 540 noticeable at all ranges, the link count is relatively stable over time.

4.5.2. User-based, cone

Alternatively, one can consider that a and b are the attendees themselves and that their awareness is limited by a cone in which social interactions can happen. Therefore, each attendee i has




Figure 12: Number of wireless contacts, for various wireless ranges

a “social cone” of $2 \times \alpha$ in front of him, with a range of $2m$. Everybody in i ’s cone is potentially interacting with i . If j is in i ’s cone, and i is in j ’s cone, *i.e.*, they face each other and their distance is lower than $2m$, we consider them as interacting with each other.

4.5.3. User-based, Voronoï

The third model we exploit relies on Voronoï diagrams. Conceptually, it is based on the hypothesis that interacting people form circles, or polygons, and that all people composing this circle and facing it are interacting. In this context, the Voronoï cell of each attendee represents the area that is closer to this attendee than to any other. We consider that people facing each other while sharing contiguous cells are interacting. Practically, we first compute the Voronoï cell of each user i (*i.e.*, the polygon containing all the points that are closer to i than to any other attendee). Then we consider that i and j are interacting if (1) they have neighboring cells and (2) they face each other, again with an angle of maximum $2 \times \alpha$. This approach, despite being algorithmically more complex, has a lower computational complexity than the cone model (n against n^2 , where n is the number of attendees). Moreover, it does not require a “social” distance parameter. Such a parameter is hard to calibrate, as it is affected by both cultural factors and environmental factors such as local people density, as studied by Hall [13].

Figure 13a illustrates the impact of the contact detection model by representing the number of detected social links over time for both user-based detection techniques. Two main parameters impact the number of detected links: the detection strategy (cone, or Voronoï) and the maximal deviation angle α . Voronoï and cone detection techniques roughly detect the same amount of links for a fixed α , although Voronoï always detects less links than the cone method. This is probably an impact of the “line-of-sight” effect of Voronoï: consider 3 attendees i, j, k aligned, i and k can be in contact using the cone method (provided $d(i, k) < 2m$) whereas the Voronoï method will never detect an interaction between them. Interestingly, the Voronoï method leads to a more stable link count over time. The grayed zones surrounding each curve represent the standard deviation of the smoothing applied. Analyzing these suggest that increasing α decreases the stability of the results, and that Voronoï method always provides more stable results.




Figure 13: Evolution of contacts and CC (RWP: Random Way Point, UB: User-Based)

4.6. Interactions

Figures 13b, 14a and 14b exploit the extracted interactions from a graph perspective by exploring the “knitting” of the structure. Figures 13b and 14a both compare traces obtained from experiment *M2* (User-based/UB, dark colors) and from a random waypoint model (RWP, light colors) tailored to copy the observed behavior (identical attendee speed, pause duration and pause probability).

Figure 13b represents the evolution of the largest connected component size in the contact graph derived from traces using various radio ranges. Computing the size of the largest connected component provides an upper bound of the number of devices able to exchange messages using a multihop communication scheme at a given moment in time. One can notice a striking difference between results obtained using synthetic and real traces: the connected component size of synthetic traces is constant over time, and always overestimate the number of connected attendees. This figure also illustrates the dramatic impact of range: above 3m all attendees are connected. A 1m-range never allows to connect more than 12 devices, while a 2m-range allows to reach nearly everyone.

Figure 14a presents another striking difference between *M2* and the output of a random waypoint model. Again, we compare both UB and RWP traces. From these traces, we counted the number of wireless contacts made by each pair of devices assuming a range of 2 m. In other words, we compute the weights of a wireless contact graph for both traces. Figure 14a presents the distribution of these weights. It is interpreted as follows: in the UB trace, around 28 devices pairs were in contact between 150 and 160 times. The striking difference is that RWP trace provides a (not surprisingly) normal distribution centered around 50, whereas the real trace exhibits a heavy tailed distribution: some devices pair connect very often while some others nearly never connect.

Figure 14b partly explains this striking difference: it represent layouts of the obtained final *social* interaction graph —as extracted by Souk in *M2*—, when considering only most frequent links (more than 50 snapshots, approximately 2.5 minutes). Each link is weighted proportionally to the amount of time its endpoints spent together. Colors represent communities found using a classical community detection algorithm [21]. It is interesting to observe the variety of contact

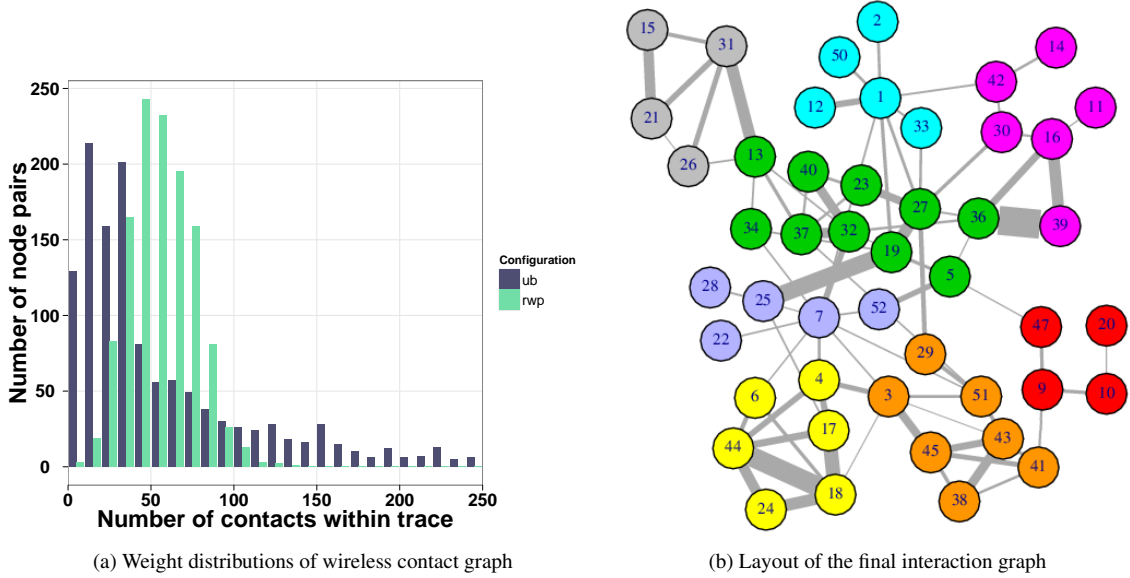


Figure 14: Digging into social interactions

patterns: whereas some attendees only have few but very strong connections, such as node 39 for example, others like node 7 have many links of lesser importance. The modularity (see [21] for a reference) of this graph is 0.51, which is relatively high.

The conclusion we draw from these two observations is that random waypoint models have no chance of correctly emulating human micro-mobility because they ignore the primordial social dimension of our behavior. As we move to meet our friends (and avoid our foes) we drastically bias the connection pattern of the devices we carry. Even if the limits of random waypoint models are already known, we obtain these results by analyzing the social structure of the underlying interaction network, to showcase the possible use of the Souk platform.

5. Conclusion

This paper presents a framework to capture and analyze mobility data, with the long-term goal of refining mobility models or deriving new ones. Instead of using raw mobility data or abstract mobility models to test the impact of mobility on human-carried devices, we seek to study and characterize crowd mobility using the presented framework. We argue this strategy will enable assessment of the level of realism and generality of models and traces, allowing us to better understand and simulate human-centered Short Range Communication-based systems.

As a first step towards this goal, we present the results obtained during the first experimental deployments of the platform. To the best of our knowledge the dynamics of such a dense crowd had never been assessed that precisely before. Analysis reveals that crowd behavior is all but random, demonstrating the need of a better model toolbox to design and test mobility-resilient software systems.

On the practical side, the Souk platform, which is fully open source, has been designed as a scalable solution towards analysis of large crowds: although the results presented here are illustrated on an experiment involving 50 persons, both the hardware cost and the complexity of software analysis (using Voronoi-based approach) are linear with respect to the number of tracked individuals.

This preliminary study opens many research directions. One of the first is to design mobility models that reflect more closely crowd connectivity dynamics. Another direction is to study this evolving connectivity following a graph-theoretic perspective: like most observed *static* interaction networks that exhibit scale-free properties, we conjecture that the dynamics of interactions in human-carried mobile networks should have common properties tied to the geographical and social distribution of users. Finally, whereas this paper focused on micro-mobility, mostly as a means to observe social interactions, our platform is technology agnostic and allows studying mobility at any scale. We believe that a major challenge lies in a careful study of this scaling behavior, to find the relationships that exist between micro-mobility and macro-mobility.

References

- [1] Ryan Aipperspach, Tye Rattenbury, Allison Woodruff, and John Canny. A quantitative method for revealing and comparing places in the home. In *Procs. of UbiComp'06*, pages 1–18. Springer, 2006.
- [2] Adam J. Aviv, Micah Sherr, Matt Blaze, and Jonathan M. Smith. Evading cellular data monitoring with human movement networks. In *Proc. 5th USENIX Workshop Hot Topics on Security*, 2010.
- [3] Farid Benbadis, Jeremie Leguay, Vincent Borrel, Marcelo Amorim, and Timur Friedman. Millipede: a rollerblade positioning system. In *Procs. of ACM WiTECH'06*, pages 117–118. ACM, 2006.
- [4] Christian Bettstetter, Giovanni Resta, and Paolo Santi. The node distribution of the random waypoint mobility model for wireless ad hoc networks. *IEEE Transactions on Mobile Computing*, 2(3), 2003.
- [5] Chiara Boldrini, Marco Conti, and Andrea Passarella. The sociable traveller: human travelling patterns in social-based mobility. In *Procs. of ACM MobiWAC'09*, pages 34–41. ACM, 2009.
- [6] Vincent Borrel, Franck Legendre, Marcelo Dias De Amorim, and Serge Fdida. Simps: using sociology for personal mobility. *IEEE/ACM Trans. Netw.*, 17(3):831–842, June 2009.
- [7] Chlo Brown, Christos Efstratiou, Ilias Leontiadis, Daniele Quercia, and Cecilia Mascolo. Tracking serendipitous interactions: How individual cultures shape the office. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW 2014)*.
- [8] Chlo Brown, Christos Efstratiou, Ilias Leontiadis, Daniele Quercia, Cecilia Mascolo, James Scott, and Peter Key. The architecture of innovation: Tracking face-to-face interactions with ubicomp technologies. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp 2014)*.
- [9] Ciro Cattuto, Wouter Van den Broeck, Alain Barrat, Vittoria Colizza, Jean-Francois Pinton, and Alessandro Vespignani. Dynamics of person-to-person interactions from distributed rfid sensor networks. *PLoS ONE*, 5(7), 07 2010.
- [10] Augustin Chaintreau, Pan Hui, Jon Crowcroft, Christophe Diot, Richard Gass, and James Scott. Impact of human mobility on opportunistic forwarding algorithms. *IEEE Transactions on Mobile Computing*, 6(6):606–620, June 2007.
- [11] Nathan Eagle and Alex (Sandy) Pentland. Reality mining: sensing complex social systems. *Personal Ubiquitous Comput.*, 10(4):255–268, March 2006.
- [12] Adriano Galati and Chris Greenhalgh. Human mobility in shopping mall environments. In *Proceedings of the Second International Workshop on Mobile Opportunistic Networking, MobiOpp '10*, pages 1–7. ACM, 2010.
- [13] E.T. Hall. *The hidden dimension*. Doubleday Anchor Books. Doubleday, 1966.
- [14] Klaus Herrmann. Modeling the sociological aspects of mobility in ad hoc networks. In *Procs. of ACM MSWIM'03*, pages 128–129. ACM, 2003.
- [15] Pan Hui, Augustin Chaintreau, James Scott, Richard Gass, Jon Crowcroft, and Christophe Diot. Pocket switched networks and human mobility in conference environments. In *Procs. of WDTN'05*, pages 244–251. ACM, 2005.
- [16] Daisuke Kamisaka, Takafumi Watanabe, Shigeki Muramatsu, Arei Kobayashi, and Hiroyuki Yokoyama. Estimating position relation between two pedestrians using mobile phones. In *Procs. of Pervasive Computing'12*, pages 307–324. Springer Berlin Heidelberg, 2012.
- [17] Marc-Olivier Killijian, Matthieu Roy, Gilles Trédan, and Christophe Zanon. Souk: social observation of human kinetics. In *UbiComp*, pages 193–196, 2013.

- [18] Moritz Köhler, Shwetak N. Patel, Jay W. Summet, Erich P. Stuntebeck, and Gregory D. Abowd. Tracksense: infrastructure free precise indoor positioning using projected patterns. In *Procs. of Pervasive computing'07*, pages 334–350. Springer-Verlag, 2007.
- [19] Vincent Lenders, Jörg Wagner, and Martin May. Analyzing the impact of mobility in ad hoc networks. In *Procs. of ACM RealMan'06*, pages 39–46. ACM, 2006.
- [20] Mirco Musolesi and Cecilia Mascolo. Designing mobility models based on social network theory. *SIGMOBILE Mob. Comput. Commun. Rev.*, 11(3):59–70, July 2007.
- [21] M.E.J. Newman and M. Girvan. Finding and evaluating community structure in networks. *Physical review E*, 2004.
- [22] Andrea Passarella, Marco Conti, Chiara Boldrini, and Robin I.M. Dunbar. Modelling inter-contact times in social pervasive networks. In *Proc. ACM MSWiM '11*, pages 333–340. ACM, 2011.
- [23] ShwetakN. Patel, KhaiN. Truong, and GregoryD. Abowd. Powerline positioning: A practical sub-room-level indoor location system for domestic use. In *Procs. of UbiComp'06*, pages 441–458. Springer Berlin Heidelberg, 2006.
- [24] Gerald Pirkel and Paul Lukowicz. Robust, low cost indoor positioning using magnetic resonant coupling. In *Procs. of ACM UbiComp'12*, pages 431–440. ACM, 2012.
- [25] Nissanka B. Priyantha, Anit Chakraborty, and Hari Balakrishnan. The cricket location-support system. In *Procs. of ACM Mobicom'00*, pages 32–43. ACM, 2000.
- [26] Juliette Stehlé, Nicolas Voirin, Alain Barrat, Ciro Cattuto, Lorenzo Isella, Jean-Francois Pinton, Marco Quaggiotto, Wouter Van den Broeck, Corinne Régis, Bruno Lina, and Philippe Vanhems. High-resolution measurements of face-to-face contact patterns in a primary school. *PLoS ONE*, 6(8), 08 2011.
- [27] Ubisense. Series 7000 ip sensors. ubisense.net/en/media/pdfs/factsheets_pdf/83188_series_7000_ip_sensors.pdf.
- [28] Waqas ur Rehman, Eyal de Lara, and Stefan Saroiu. Cilos: a cdma indoor localization system. In *Procs. of ACM UbiComp'08*, pages 104–113. ACM, 2008.