An Experimental Platform for Quantified Crowd

Mateusz Grabowski*, Michal Marschall*, Wojciech Sirko*, Maciej Debski*, Marcin Ziombski*,
Przemyslaw Horban*, Szymon Acedanski*, Marcin Peczarski*, Dominik Batorski†, and Konrad Iwanicki*

*Faculty of Mathematics, Informatics and Mechanics, University of Warsaw, Poland
†Interdisciplinary Center for Mathematical and Computational Modeling, University of Warsaw, Poland
E-mails: {mg291540,mm291693,ws292444,md319428,mz292739,ph262940}@students.mimuw.edu.pl
{accek,marpe}@mimuw.edu.pl   db@uw.edu.pl   iwanicki@mimuw.edu.pl

Abstract—Quantified crowd, a vision in which on-body sensors of nearby people collaborate to detect various phenomena within a crowd and produce feedback, is an emerging research area. One of the issues that impedes progress in this new area is a lack of a broadly applicable experimental platform, such as the platforms that enabled research on wireless sensor networks.

In this paper, we aim to address this issue by presenting such an experimental hardware-software platform. Not only does the platform introduce custom badge-form low-power devices that feature a number of sensors relevant to quantified crowd, but it also ensures that these devices can interoperate with commercial off-the-shelf smartphones, wristbands, and other devices for self-quantification. In effect, it constitutes a powerful experimental instrument, as we show in a preliminary real-world evaluation.

I. INTRODUCTION

The progress in miniaturization, energy efficiency, and wireless communication has made small personal devices, such as smartphones or activity monitors, a popular technology. Equipped with diverse sensors and specialized software, such devices allow for continuous, fine-grained monitoring of various aspects of peoples' lives and for giving personalized feedback aimed to improve these aspects. This feedback loop is the essence of the vision referred to as *quantified self* [1].

There is mounting evidence that self-quantification devices will become even smaller and commonplace, embedded in clothes, accessories, and everyday objects. Combining this with the fact that they are increasingly capable of short-range wireless communication, one may envision scenarios in which the devices of various people exchange information directly, depending on their hosts' co-presence in physical spaces. In other words, the devices collaborate to perform sensing and coordinate to provide feedback not only for an individual, but also for collocated groups or even crowds of people. This vision, which we refer to as *quantified crowd*, has numerous potential applications. Let us thus give just a few examples.

1) Detecting crowd dynamics: To start with, equipped with proximity, direction, acceleration and other sensors, quantified crowd devices worn by a large number of people could communicate to collaboratively detect patterns of coordinated movement of nearby persons and to assign a higher-level semantic meaning to these patterns. For instance, they could recognize formation of queues, emergence of flows in pedestrian lanes, intersections of such flows, and other problematic situations, such as clogging [2]. This, in turn, could help improve designing public spaces, planning train schedules, provisioning staff, and avoiding disasters during mass events.

2) Maximizing event experience: Quantified crowd devices could also be beneficial in trade fairs, conferences, and multi-stage concerts [3], [4], [5], [6]. They could detect interpersonal interactions and provide the participants with each others' contact information from online social networks. They could guide a participant to interesting peers or exhibitions. They could also track what music performances a group of friends enjoyed (e.g., by dancing or jumping together), so that the songs could be obtained after the concert. In short, the technology could help maximizing the overall experience of an event.

3) Mitigating exclusion: Another problem the quantified crowd technology could address is the exclusion of disabled people. In particular, while blind people are often proficient in navigating throughout a city, when two or more of them want to arrange a meeting in a large, open, crowded public space, they can have major difficulties finding each other.[1] Quantified crowd devices could thus again be used to facilitate the navigation through the crowds toward each other.

A. Problem Statement and Related Work

Recently, a lot of research progress facilitating the aforementioned and similar applications has been made. To begin with, it was shown that physical proximity, measured based on wireless beacons broadcast by wearable devices, can be used as a crude proxy for detecting social interactions between the people wearing the devices [3], [6], [7], [8]. Using infrared instead, one can also detect whether nearby people are facing each other, which increases the likeliness of them having an actual social interaction [4], [5], [9], [10]. Acceleration and orientation are other modalities that can potentially allow for detecting the physical activities of individuals [3], [11], [12] or even for clustering individuals into groups behaving similarly or moving together [13], [14]. Likewise, audio can be used to detect a human conversation, including, for instance, recognizing the speakers' mood [15], [16], [17]. Being able to sense these modalities has allowed researchers to study real-world social interactions in workplaces or during conferences [4], [8], [9], to investigate techniques for identifying communities and groups of influence [18], [19], [20], or to analyze basic patterns in the behavior of crowds in physical spaces [2], [14],

---

[1]Personal communication.

[21], even giving birth to a new field of science: Computational Social Science [22]. Finally, new decentralized algorithms for highly mobile networks, in particular, for opportunistic forwarding [23], in-network aggregation [24], and density estimation [25], have started being investigated to close the feedback loop in the futuristic quantified crowd applications.

However, despite this progress, there are still numerous research problems that must be solved before the quantified crowd vision becomes a reality. One such basic issue is *a lack of a broadly applicable experimental platform*. To illustrate, instead of quantified crowd, consider the now mature wireless embedded sensing. The immense development in wireless sensor networks was possible in the last years partially due to the versatile hardware-software platforms, such as the Mica2 [26] and TelosB [27] motes together with the TinyOS environment [28]. They allowed people with varying expertise to innovate at different levels and to validate novel ideas in the real world. In contrast, for quantified crowd, no such platforms exist yet.

A popular solution is to use smartphones [3], [6], [7], [14], [17]. They are easy to program, contain various sensors, and are relatively widespread, thereby being a perfect platform for validating simple ideas quickly and at scale. Yet, as commercial devices with use cases different than the futuristic quantified crowd applications, they are not well suited for more advanced studies of novel techniques. In particular, they are missing sensors and interfaces that, despite no commercial value today, could be useful for quantified crowd in the future. In addition, their form factor does not match many experimental scenarios, they consume relatively much power, are heterogeneous, and difficult to extend with custom hardware.

To cope with these problems, an opposite approach is often taken: designing devices dedicated for studying a given phenomenon or technique [2], [4], [5], [8], [9], [10], [11], [13]. An advantage of this approach is that the devices can be fully customized. A major disadvantage is in turn that they are usually aimed at few experimental scenarios. As such, they may not support any wireless communication, focusing on sensor data acquisition for offline analysis [11], [13]; alternatively, they may have wireless interfaces, but those are specialized, precluding interoperability with, for instance, smartphones [2], [4], [5], [8], [9], [10]. Furthermore, the devices often disregard energy consumption, assuming short-term operation. Finally, in many cases, they are not open, which makes reproducing results of other groups laborious.

B. Contributions and Paper Organization

In this paper, we introduce a new experimental platform for quantified crowd that addresses the aforementioned issues. The platform consists of hardware and software that together bridge the two previous experimental approaches. The platform is also open, thereby allowing other groups to adapt it to their own experimental needs without repeating our work.[2]

Drawing from the aforementioned research experiences, the hardware has a form of name badge; the software, in turn,

---

[2]For the relevant hardware specs and software, please, refer to http://www.mimuw.edu.pl/~iwanicki/projects/SocSenSys/.

comprises a programming environment for the badges and utility applications for mobile personal devices. Together, they offer to their users numerous sensors that proved useful to date, as well as sensors that seem appealing for measuring some aspects of quantified crowd, but have not been evaluated yet. In other words, even on its own, the platform enables experiments with virtually all modalities recognized as relevant to quantified crowd and in many different scenarios, as we demonstrate in preliminary evaluation. Equally important, however, it supports three different wireless communication interfaces that enable interoperation with both smartphones and wireless sensor nodes. In particular, it can exploit smartphones' computational resources or their Internet connectivity. Finally, despite these features, it minimizes power consumption, allowing for weeks of operation without recharging.

The rest of this paper is organized as follows. Section II discusses the platform in more detail. Section III demonstrates its utility in preliminary experiments. Section IV concludes.

II. ARCHITECTURE OF THE PLATFORM

The goal of our platform is facilitating experimentation in as many quantified crowd scenarios as possible. To this end, it combines both new dedicated low-power hardware—the badges—and software developed not only for the hardware but also for commodity devices, notably smartphones. As one of the consequences, it allows dedicated and commodity devices to interoperate, opening numerous experimental possibilities.

A. Hardware

In line with this goal, we wanted the hardware to ensure the following two properties. First, on its own, it should enable experimentation in as many scenarios as possible. In particular, its form factor, sensing capabilities, and power consumption should be sufficient for a wide range of quantified crowd experiments. Second, it should allow for extending the platform with external devices. This implies that at the very minimum, it should interoperate with commodity devices; preferably, it should also bridge specialized wireless sensing technology with commercial off-the-shelf personal devices. Let us discuss how its design ensures these properties.

1) Form factor: Building upon the aforementioned research experiences, our hardware has a popular form of name badge (cf. Fig. 1). As such, it has a well defined position on the body of its wearer. This simplifies detecting if the wearer is facing another person with a badge or reasoning about the wearer's current physical activity. Likewise, the form factor facilitates small-scale prototyping compared to tinier personal devices.

2) Local sensing: Drawing from quantified self, quantified crowd relies heavily on sensing various modalities of individual people. Accordingly, our badge hosts a number of sensors. A tri-axis accelerometer is meant for determining the current physical activity of the badge wearer, for example, whether the wearer is standing, walking, running, jumping, and so on. A magnetometer enables detecting the badge wearer's orientation. Combined with a gyroscope, the accelerometer and magnetometer allow for dead reckoning, that is, navigating
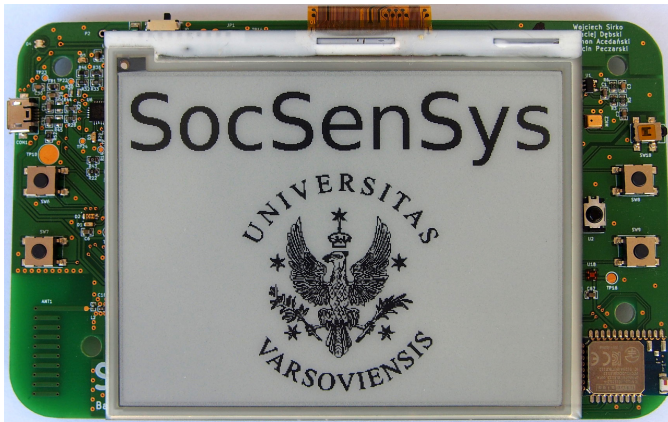
Fig. 1. Our badge without a casing.

the wearer without additional positioning. A microphone that amplifies frequencies corresponding to human speech is meant for detecting conversation as well as finer-grained sensing, such as inferring the sex and mood of the speakers. A highly sensitive barometer can give context information, such as the elevation at which the badge resides. Light and temperature sensors provide further context information, for example, if the badge is outdoors or indoors. Overall, the suite of sensors available on a badge covers virtually all modalities employed in the aforementioned experiments conducted to date.

3) Communication and crowd texture sensing: Furthermore, what differentiates quantified crowd from quantified self is communication between the devices of various people. In quantified crowd, nearby devices must coordinate their sensing and collaborate when analyzing the data and providing feedback. In addition, by communicating, they must determine spatial interdependencies between them, notably proximity and mutual orientation, which is crucial for inferring the spatial structure—the texture—of a crowd [2]. For these reasons, our badge hosts three different communication interfaces.

First, a low-power radio (RF) is meant as the basic communication interface as well as proximity detector. We chose a popular radio chip for the sub-gigahertz ISM bands with dynamically configurable output power, and coupled it with a custom on-board antenna. The chip is compatible with several wireless sensor network platforms and some personal devices, like eZ430-Chronos smart watches [29], which enables extending the platform with other devices. The configurable output power allows in turn for changing the transmission range, which, combined with the selected band and antenna, enables fine-grained sensing of the distance between badges.

Second, being more directional than the radio, an infrared (IR) transceiver is meant to detect if badge wearers face each other. Again, we designed the interface so that it can be adjusted: the output power can be selected at runtime to enable controlling the transmission range; the radiation cone can in turn be regulated by the casing of the badge. Finally, infrared can also be used for communicating information: the physical bit encoding and frame format are determined by software, so again the badges can interoperate with other infrared devices.

Third, our badge is equipped with a novel Bluetooth Low Energy (BLE) transceiver. In effect, it can communicate with modern smartphones and commercial devices for self-quantification, such as the Fitbit wristband [30]. Considering that BLE is gaining popularity, the number of external smart devices that can extend our platform is growing as well. Although BLE can also be used for sensing proximity, it is not as versatile in terms of configuration and communication protocols as the low-power radio, so we decided to have both.

4) Input-output: Apart from the communication interfaces, our badge also has several input-output elements. To show the name and affiliation of its bearer but also arbitrary dynamic information, it has a 4.41-inch, 400×300-pixel e-paper display. We have chosen e-paper over other technologies as it requires energy only when changing the displayed image and the image does not degrade in sunlight, which was problematic for other technologies used in previous experiments, both by us and others. Other I/O elements on our badge include buttons, LEDs, a buzzer for audio feedback, external flash memory and an additional SD card slot for logging data, and a USB socket for charging batteries and post-mortem log acquisition.

5) Energy consumption: Finally, both individual components of a badge and the badge as a whole minimize energy consumption, which is crucial for experimental scenarios that require badges to work unattended for weeks. In particular, as the core of our badge, we have chosen an ARM Cortex-M MCU, which is relatively powerful and yet ultra-low power. Other component choices, such as BLE instead of the classic Bluetooth or e-paper instead of another display technology, were also motivated by the power requirements; so was the lack of power hungry components, such as a GPS receiver or a WiFi transceiver, which can be used indirectly if necessary, by having a badge communicate with a smartphone. Last but not least, the way the selected components are integrated as well as the way their software drivers work (discussed in the next section) strongly emphasize low energy consumption.

6) Summary: To sum up, our hardware builds upon previous research activities but addresses the drawbacks of the instruments used in those activities, often because we experienced these drawbacks ourselves (e.g., display). The form factor of a badge, its low energy consumption, and the rich suite of sensors enable using our badges alone in a wide range of experimental scenarios. The three low-power communication interfaces in turn allow for increasing the capabilities of the badges with external devices, both custom wireless sensors and commercial personal devices. To the best of our knowledge, these properties of the badges make them truly unique.

B. Software

The software of the platform consists in turn of two main parts: the programming environment for the badges and utilities for mobile devices. Like the hardware, the software is tailored to be applicable in as many experimental scenarios as possible. Unlike the hardware, which is now stable, the software is continuously extended, which hopefully further increases its utility. Let us give an overview of its two parts.

1) Badge software: The programming environment for the badges is based on the popular operating system for wireless sensor networks, TinyOS [28]. We selected TinyOS for a few reasons. First, its layering promotes rapid innovation, even without expert knowledge about its internals. Second, its driver architecture facilitates minimizing energy consumption. Third, its code base contains many libraries that can be useful also for experimenting with quantified crowds. Finally, it supports many wireless sensor node platforms that, due to the radio chip employed by us, can interoperate with our badges.

To provide the support for our badge in TinyOS, we have created a completely new TinyOS platform, called *badge*. In the process, we had to face a number of challenges.

To begin with, the ARM MCU used by our badges is 32-bit and has a few unique features compared to the MCUs supported by TinyOS, which are mostly 8- or 16-bit. As a result, when implementing the support for our MCU, we had to handle these new features within the boundaries of the TinyOS architecture, which was sometimes problematic. For instance, we discovered and fixed a bug in the low-level operating system code [31], which existed because apparently the code had never before been run on the 32-bit ARMs.

Furthermore, many sensor types available on our badges as well as elements such as the IR transceiver had not been supported in TinyOS. Consequently, we had to develop their drivers at all the system layers. In addition, some of the components, such as the BLE interface or the e-paper, had not even had designs of a software architecture. We had to create these designs from scratch because the operation of these components differs significantly from the components that are supported by TinyOS. Finally, we also adapted some of the TinyOS services for the badges, like the timers library or over-the-air reprogramming; likewise, we implemented services customized for the badges, for instance, for acquisition of logged data. In effect, we have obtained a complete, fully-functional TinyOS platform with additional services useful for experimentation in scenarios relevant to quantified crowds.

2) Mobile device software: In addition to the software for the badges, we have developed software for commercial personal devices. The software is written in Java as we targeted devices running the Android OS. It is tailored mostly to individual experimental scenarios we considered. For example, since in some quantified crowd scenarios, such as conferences or trade fairs, it is beneficial to know the profiles and/or contacts of individuals, we created utilities for extracting such information from online social networks. Other examples include utilities for scanning nearby Bluetooth devices.

One piece of the Android software deserves more attention: software for communicating with the badges. More specifically, for testing badge prototypes, we implemented utility code that enables exchanging information with a badge over BLE. In the future, similar code can allow a badge to harness the computational power of a smartphone (e.g., to perform sensor data analysis that is beyond the capabilities of the badge's MCU) or to utilize the Internet connectivity of the smartphone (e.g., to download user profiles from online social networks or upload changes to such profiles). Developing this code was not straightforward, as some BLE modes were (and likely still are) poorly supported in Android.

3) Summary: To sum up, the layered TinyOS programming environment allows developers to write experimental code accessing the badge hardware at various levels and using various services out of the box, depending on the developers' expertise. At the same time, the Android software can enable using the badges in combination with mobile off-the-shelf devices, such as smartphones, wristbands, pedometers, and so on. All in all, our software may facilitate building even complex distributed systems for quantified crowd experiments.

III. Preliminary Evaluation

We have started experimenting with the platform to study various phenomena relevant to quantified crowds. Although we do provide sample results from such experiments, covering all of them is out of the scope of this paper. Instead, the goal of the paper is illustrating the potential of our platform.

A. Energy consumption

To start with, Table I presents the average standby- and active-state current of each major hardware component of our badge, as given by the manufacturer and validated empirically by us. The leakage current is not included and is less than $1\ \mu$A per component. In short, the table shows that the energy consumption of our badge can be ultra low, allowing for weeks of operation on a default 3.7-V, 2400-mAh battery. This facilitates using the badges in a wide range of experimental scenarios that require long operation without charging.

TABLE I
AVERAGE CURRENT CONSUMPTION OF EACH HARDWARE COMPONENT

Component	Standby [μA]	Active [mA]
MCU	1.7	6.1
RF transceiver	0.2	16.9 (RX), 34.2 (TX @ 12 dBm)
Infrared transceiver	0	0.9 (RX), 72 (TX @ power level 3)
BLE module	0.4	25 (RX), 36 (TX)
E-compass	2	0.44
Gyroscope	1	5
Barometer	0.5	0.025
Light sensor	4	0.13
Microphone	0	1.8
External flash	100	20
SD card	0	depends on a card (e.g., 100)
E-paper display	0	8

B. Sensing crowd texture with RF

As we mention previously, an important modality in such experiments is proximity of people. One of the proximity sensors on our badge is the radio with the specifically selected transceiver chip and antenna. More precisely, by dynamically adjusting the radio output power, one can control the transmission range, based on which proximity can be inferred.

As a demonstration, we conducted experiments in which two people wearing badges were positioned at a given distance from each other with line-of-sight visibility: one of the badges being a transmitter, the other one, a receiver. We varied the distance between the people, and for each distance, we studied a set of radio output power values. For each (distance, power)

configuration, we measured the packet reception rate (PRR), received signal strength indicator (RSSI), and link quality indicator (LQI) over a thousand transmitted packets. Table II presents the results for PRR for selected output power levels.

Dist. [m]	TX power [dBm]							
	-30.2	-27.7	-20.5	-15.7	-11.0	-5.0	-0.3	3.7
1	74.8	85.3	99.0	99.7	99.4	99.3	99.7	99.6
2	14.0	65.7	99.7	99.2	99.0	99.6	99.3	98.4
3	2.0	0.0	98.7	98.9	99.3	99.6	99.2	99.3
5	14.9	0.0	99.4	99.1	99.6	99.4	97.9	99.3
10	0.0	0.0	36.0	96.0	99.7	99.9	99.0	99.3
20	0.0	0.0	0.0	14.9	98.3	99.9	99.3	99.5
30	0.0	0.0	0.0	0.0	23.6	35.0	99.0	99.1
40	0.0	0.0	0.0	0.0	20.0	22.0	68.5	99.5
50	0.0	0.0	0.0	0.0	0.0	1.0	11.0	98.8
60	0.0	0.0	0.0	0.0	0.0	0.0	23.4	98.1
70	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

The results indicate that, despite the unpredictability of low-power wireless communication, PRR does allow for reasoning about proximity of badges. Using output power levels from the table, as well as intermediate ones not depicted in the table, one can distinguish shorter distances with an accuracy of down to a few meters; higher distances can in turn be detected with an accuracy of about ten meters (the maximal distance is around 100 m). Moreover, it is also possible to exploit the raw physical signal quality metrics, RSSI and LQI (not shown due to space constraints), to further improve proximity inference.

C. Sensing crowd texture with IR

Like RF, IR can also be used for sensing proximity. To this end, we conducted the same experiments as for RF but employing IR instead. Their results are shown in Table III.

Dist. [m]	TX power [level]			
	0	1	2	3
1	100.0	99.0	93.0	88.0
2	74.0	67.0	99.0	88.0
3	0.0	55.0	33.0	51.0
4	0.0	0.0	50.0	53.0
5	0.0	0.0	0.0	30.0
6	0.0	0.0	0.0	0.0

In contrast to the radio, which is a single chip, the IR transceiver is a custom circuit. We designed the circuit so that it supports four power levels (0–3), derived from earlier experiments. The table shows that these power levels allow for inferring proximity with an even better accuracy than RF. What is more, the accuracy can be further tuned by changing the length of transmitted packets, which however is out of the scope of this paper given the available page limit.

In addition, IR is more directional than RF. In particular, Table IV shows how the orientation of the receiver with respect to the transmitter affects PRR. PRR is the highest when the badges face each other and becomes lower when they rotate. It is the lowest when they have their backs to each other, though it is not zero, because the experiment was conducted indoors,

where IR reflections can occur. Note also that the badges were without casings. With an appropriate casing, one can regulate the angle of the transmission cone, which may allow for detecting the mutual orientation of badges more precisely.

Angle [°]	0	30	60	90	120	150	180	210	240	270	300	330
PRR [%]	98	99	89	46	18	26	4	5	23	7	94	99

D. Sample illustrative experiment

The previous experiments suggest that IR can be more accurate than RF for sensing proximity. Yet, the maximal IR bandwidth on our platform is two orders of magnitude lower than that of RF, and IR transmissions require more energy than RF ones. Therefore, when it comes to inferring proximity, these two modalities should be combined. Likewise, mutual orientation can be inferred from a combination of IR and magnetometer. These are just two examples highlighting the need for multi-modal sensing in quantified crowd applications.

Therefore finally, to illustrate a tiny sample of the multi-modal sensing possibilities offered by our platform for such applications, we present briefly one of our very first experiments aimed at classifying interpersonal situations. More specifically, we wanted to detect and distinguish between four types of pairwise situations involving people wearing badges: (1) no social interaction at all (denoted N), (2) two people passing each other (P), (3) two people standing and having a chat (C), and (4) two people walking together (T). To this end, we first conducted 100 micro-benchmark experiments for each of the situations, in each micro-benchmark recording data from RF transmissions at three low output power levels, IR at the minimal power level, the accelerometer and microphone. We used the resulting data as a training set for a machine learning toolkit, Weka [32], to obtain classification heuristics. The heuristics were then deployed on the badges and used in an experiment involving three participants moving freely in a foyer of a building with other people present. Figure 2 depicts the resulting pairwise classifications (a) obtained at the badge of participant 1 and the actual situations recorded by an external observer and verified based on a video stream (b).

Although a more detailed discussion of the experimental results is out of the scope of this paper, the figure already suggests that the four situations can be automatically recognized to a large extent. Note also that this was one of the very first experiments with the platform, involving few modalities and simple, automatically synthesized heuristics. Our recent research indicates in turn that pairwise situation classification can be improved with more advanced decision algorithms and other modalities sensed by the badges. What is more, a variety of additional external devices, which our platform interoperates with, may improve the classification even further.

IV. CONCLUSION

All in all, our hardware-software platform does have the potential to be a broadly applicable experimental instrument for quantified crowd. Even on its own, it allows for a wide

(a) the interactions of participant 1 with participants 2 and 3, as detected by the badge of participant 1

(b) the actual interactions of participant 1 with participants 2 and 3, as recorded by an external observer and later verified
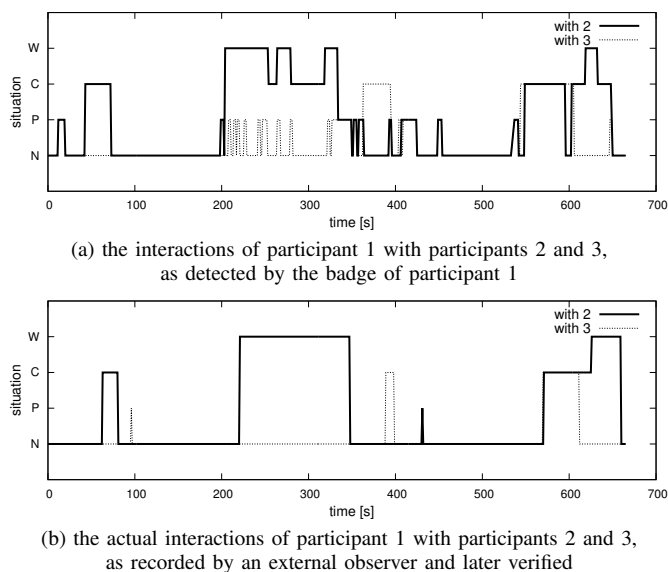
Fig. 2. Sample experiment for pairwise interaction classification based on multi-modal sensing with our badges.

range of real-world experiments including studying human activity, crowd textures, social interactions, collective behavior, and the like, which we are confirming in our present research. What is more, however, it can interoperate with both existing wireless sensor node platforms and commercial devices for self-quantification, thereby bridging these two worlds. As a result, it can be extended with new hardware and software to handle scenarios that we may not even envision yet. Since to date such a platform has been lacking, we hope that our platform will become for quantified crowd what the aforementioned Mica2, TelosB, and TinyOS solutions were for wireless sensor networks: an instrument that eliminates the need for redesigning everything, instead allowing different groups to focus solely on research falling within their expertise, which promotes innovation.

ACKNOWLEDGMENT

REFERENCES

[1] K. Kelly, "What is the quantified self?" web: http://quantifiedself.com/2007/page/3/, October 2007.

[2] C. Martella, M. van Steen, A. Halteren, C. Conrado, and J. Li, "Crowd textures as proximity graphs," *IEEE Commun. Mag.*, vol. 52, no. 1, pp. 114–121, 2014.

[3] E. Miluzzo, N. D. Lane, K. Fodor, R. Peterson, H. Lu, M. Musolesi, S. B. Eisenman, X. Zheng, and A. T. Campbell, "Sensing meets mobile social networks: The design, implementation and evaluation of the CenceMe applications," in *ACM SenSys '08*, 2008, pp. 337–350.

[4] K. Ara, N. Kanehira, D. O. Olguín, B. N. Waber, T. Kim, A. Mohan, P. Gloor, R. Laubacher, D. Oster, A. S. Pentland, and K. Yano, "Sensible organizations: Changing our businesses and work styles through sensor data," *J. of Inf. Processing*, vol. 16, pp. 1–12, 2008.

[5] J. A. Paradiso, J. Gips, M. Laibowitz, S. Sadi, D. Merrill, R. Aylward, P. Maes, and A. Pentland, "Identifying and facilitating social interaction with a wearable wireless sensor network," *Pers. Ubiquit. Comput.*, vol. 14, no. 2, pp. 137–152, 2010.

[6] T. Do and D. Gatica-Perez, "Human interaction discovery in smartphone proximity networks," *Pers. Ubiquit. Comput.*, vol. 17, no. 3, pp. 413–431, 2013.

[7] N. Eagle and A. S. Pentland, "Reality mining: Sensing complex social systems," *Pers. Ubiquit. Comput.*, vol. 10, no. 4, pp. 255–268, 2006.

[8] C. M. Angelopoulos, C. Mouskos, and S. Nikoletseas, "Social signal processing: Detecting human interactions using wireless sensor networks," in *Proc. ACM MobiWac '11*. ACM, 2011, pp. 171–174.

[9] T. Choudhury and A. Pentland, "Sensing and modeling human networks using the sociometer," in *IEEE ISWC '03*, 2003, pp. 216–222.

[10] M. A. White, "SmartBadge: An electronic conference badge using RF and IR communications," Master's thesis, University of Canterbury, Christchurch, New Zealand, 2006.

[11] P. Siirtola, P. Laurinen, E. Haapalainen, J. Roning, and H. Kinnunen, "Clustering-based activity classification with a wrist-worn accelerometer using basic features," in *IEEE CIDM '09*, 2009, pp. 95–100.

[12] G. Uslu, O. Altun, and S. Baydere, "A Bayesian approach for indoor human activity monitoring," in *Proc. HIS '11*, 2011, pp. 324–327.

[13] D. Roggen, M. Wirz, G. Tröster, and D. Helbing, "Recognition of crowd behavior from mobile sensors with pattern analysis and graph clustering methods," *Net. Heterog. Media*, vol. 6, no. 3, pp. 521–544, 2011.

[14] M. B. Kjaergaard, M. Wirz, D. Roggen, and G. Troster, "Detecting pedestrian flocks by fusion of multi-modal sensors in mobile phones," in *Proc. ACM UbiComp '12*, 2012, pp. 240–249.

[15] P. Ruvolo, I. Fasel, and J. Movellan, "Auditory mood detection for social and educational robots," in *IEEE ICRA '08*, 2008, pp. 3551–3556.

[16] F. Burkhardt, T. Polzehl, J. Stegmann, F. Metze, and R. Huber, "Detecting real life anger," in *IEEE ICASSP '09*, 2009, pp. 4761–4764.

[17] H. Lu, D. Frauendorfer, M. Rabbi, M. S. Mast, G. T. Chittaranjan, A. T. Campbell, D. Gatica-Perez, and T. Choudhury, "StressSense: Detecting stress in unconstrained acoustic environments using smartphones," in *ACM UbiComp '12*, 2012, pp. 351–360.

[18] P. Hui, E. Yoneki, S.-Y. Chan, and J. Crowcroft, "Distributed community detection in delay tolerant networks," in *Proc. ACM MobiArch '07*, 2007.

[19] H. Kim, J. Tang, R. Anderson, and C. Mascolo, "Centrality prediction in dynamic human contact networks," *Comput. Netw.*, vol. 56, no. 3, pp. 983–996, 2012.

[20] D. Fay, J. Kunegis, and E. Yoneki, "Centrality and mode detection in dynamic contact graphs; a joint diagonalisation approach," in *Proc. ACM/IEEE ASONAM '13*, 2013, pp. 41–48.

[21] M. Wirz, D. Roggen, and G. Tröster, "A methodology towards the detection of collective behavior patterns by means of body-worn sensors," in *Proc. the Ubicomp in the Large Workshop at Pervasive '10*, 2010.

[22] D. Lazer, A. S. Pentland, L. Adamic, S. Aral, A.-L. Barabási, D. Brewer, N. Christakis, N. Contractor, J. Fowler, M. Gutmann, T. Jebara, G. King, M. Macy, D. Roy, and M. V. Alstyne, "Computational social science," *Science*, vol. 323, no. 5915, pp. 721–723, 2009.

[23] N. Wang and E. Yoneki, "Impact of social structure on forwarding algorithms in opportunistic networks," in *Mobile and Wireless Networking (iCOST), 2011 International Conference on Selected Topics in*, 2011.

[24] M. Gregorczyk, T. Pazurkiewicz, and K. Iwanicki, "On decentralized in-network aggregation in real-world scenarios with crowd mobility," in *IEEE DCOSS '14*, 2014, pp. 83–91.

[25] M. Cattani, M. Zuniga, A. Loukas, and K. Langendoen, "Lightweight neighborhood cardinality estimation in dynamic wireless networks," in *ACM/IEEE IPSN '14*, 2014, pp. 179–189.

[26] J. Hill and D. Culler, "Mica: A wireless platform for deeply embedded networks," *IEEE Micro*, vol. 22, no. 6, pp. 12–24, 2002.

[27] J. Polastre, R. Szewczyk, and D. Culler, "Telos: Enabling ultra-low power wireless research," in *IPSN '05*, 2005, pp. 364–369.

[28] P. Levis, S. Madden, J. Polastre, R. Szewczyk, K. Whitehouse, A. Woo, D. Gay, J. Hill, M. Welsh, E. Brewer, and D. Culler, "TinyOS: An operating system for sensor networks," in *Ambient Intelligence*. Springer Berlin Heidelberg, 2005, pp. 115–148.

[29] Texas Instruments, "eZ430-Chronos development tool: User's guide," web: http://www.ti.com/lit/ug/slau292f/slau292f.pdf, 2009.

[30] Fitbit Inc., "Fitbit user's manual," web: http://fitbit.com, 2009.

[31] W. Sirko, "[Bug] BusyWaitCounterC – missing explicit cast," web: https://github.com/tinyos/tinyos-main/issues/247, 2013.

[32] I. H. Witten, E. Frank, and M. A. Hall, *Data Mining: Practical Machine Learning Tools and Techniques*, 3rd ed. Morgan Kaufmann, 2011.