Capturing Interactions in Face-To-Face Social Networks

Francesco Ficarola and Andrea Vitaletti
Department of Computer, Control, and Management Engineering,
Sapienza University of Rome, Via Ariosto 25, 00185 Rome, Italy

Keywords: Mobile, Interactions, Social Networks, RFID, Proximity, Face-To-Face.

Abstract: Online social networks, formed by cyber interactions between users, are nowadays explored in a number of papers. In this work, we present our experimental activity on Face-To-Face (F2F) social networks tracing physical interactions of humans in real-world scenarios. We briefly present the technologies to observe F2F social networks focusing on the SocioPatterns platform that we have employed in our real-world experiments. Motivated by the requirements of heterogeneous applications, we discuss how to tune the SocioPatterns collection protocol parameters in order to better capture fast interactions between users; carried out experiments confirm the effectiveness of such tuning.

1 INTRODUCTION

The participation in online social networking is continuously growing and more and more people use websites such as Twitter, Facebook or Google+. This new online phenomenon has led to a greater attention of the research community to social networks, also motivated by relevant application scenarios, such as reputation management (Yu et al., 2010), recommendation systems (Adomavicius and Tuzhilin, 2005; Walter et al., 2008) or information sharing platforms such as Quora. However, a new trend involving our "real life" relationships is becoming more and more contemporary due to the continuous development of new devices capable to track physical interactions. In this paper we are interested in Face-To-Face (F2F) social networks, namely networks made of evolving graphs in which nodes are humans and edges between nodes dynamically appear whenever a F2F interaction between humans takes place. Over the last decade, some papers such as (Hui et al., 2005), have focused on tracking physical proximities, however the size of those experiments is relatively small and some of them were deployed employing unsuitable technologies; so far, the main reason behind the lack of analysis of real-world interactions is mainly due to hardware limitations (e.g., bluetooth is too inaccurate to obtain an efficient face-to-face measure between people). One of the first attempt to effectively track face-to-face interactions has been made by the SocioPatterns researchers (soc, nd; Barrat et al., 2008) using the Radio Frequency IDentification (RFID) technology. We believe SocioPatterns tags are among the most promising devices to track F2F interactions in real-world scenarios. For this reason, our experiments have been conducted using such devices with the main purpose of tuning the multiple access control (MAC) (Wu and Pan, 2008) protocol programmed into the tags to better suit heterogeneous application scenarios. Notice that, standard MAC protocols for Wireless Sensor Networks (WSNs) are not suitable in real-world social scenarios where the dynamics of interactions are extremely fast and unpredictable when compared to the WSNs ones.

Contribution. The ability of capturing F2F interactions opens the possibility of quantitatively exploring the way in which humans interact in the real world. In this paper, we compare [1] the **`SOCIOMAC`** protocol proposed by the SocioPatterns initiative and already used in monitoring activities, with the new **`PROXMAC`** protocol. Our experimental results show that a proper tuning of the protocol parameters can better support the rich heterogeneity of F2F real-world scenarios and guarantees an effective trade-off between interaction resolution and network life-time.

Roadmap. In the next section we introduce more formally F2F social networks, presenting some of their applications and the technologies employed to track the interactions. In section 3 we discuss the issues in tracking F2F interactions in heterogeneous applica-

[1]This work has been possible thanks to the cooperation with the SocioPatterns collaboration (soc, nd).

Ficarola F. and Vitaletti A..
Capturing Interactions in Face-To-Face Social Networks.
DOI: 10.5220/0005432206130620
In *Proceedings of the 11th International Conference on Web Information Systems and Technologies* (WEBIST-2015), pages 613-620
ISBN: 978-989-758-106-9

tion scenarios and the corresponding MAC protocols. In section 4 we show the results of the experiments conducted on F2F social networks and on a simulation environment to test the performance of the proposed MAC protocols. Conclusions are discussed in section 5.

2 F2F SOCIAL NETWORKS

A Face-To-Face (F2F) social network is a dynamic network made by linking nodes (i.e., humans) that interact at short range (e.g., 1-1.5 meter of distance) for a sufficient amount of time to potentially exchange meaningful information. More formally, a F2F social networks is a dynamic evolving network **G** made of a finite sequence $G_0, G_1, G_2, \ldots, G_t$ of t static networks over the same vertex set V and a variable edge set E_i (Avin et al., 2008); a link connecting nodes $\langle u, v \rangle \in E_i$ iff at least a F2F interaction between u and v occurred in the interval in time Δ between G_i and G_{i+i}. We call Δ the *resolution* of the network.

Technologies. At present, there are a limited number of solutions able to track interactions between individuals in a distributed way that include sensors and wireless technologies, such as Bluetooth, WiFi and RFID. One of the first experiments to collect information from a real group of people was made in (Hui et al., 2005), where 54 individuals participating in a conference were given Intel Imote devices, equipped with a Bluetooth radio to discover nodes in proximity. However, Bluetooth does not allow a fine-grained recording of social interactions because of two reasons: 1) the discovery process to identify potential F2F neighbors is slow and 2) since the radio range is $\approx$ 5-10 meters, it is possible to record as "social interaction" the simple fact of being in the same room, even if the users are not interacting in any way.

Jayagopi et al. (Jayagopi et al., 2010) deployed some experiments involving 24 groups of 4 members each. Every participant was asked to wear a sociometric badges capable of recognizing speech activity and line-of-sight presence. However, all the experiments performed in (Jayagopi et al., 2010) are based on relatively small groups in which interactions are limited inside the group, while we are interested in a technology to study the dynamics of larger groups in which members are free to move.

SocioPatterns is an interdisciplinary research collaboration that supports the development of the *SocioPatterns Sensing Platform*, an infrastructure including new experimental RFID sensors that can be worn by humans in order to track their mobility and F2F interactions in real-world scenarios. The SocioPatterns platform is made of two main entities: active RFID tags and RFID readers. When two persons are in proximity within 1-1.5 meter (see Figure 1), their tags exchange proximity packets containing their IDs (process 1) so that each tag is able to know who is talking to. Eventually, tags send the received proximity packets to close-by readers (process 2), which in turn will forward those messages to a central server running a UDP logger (process 3). Proximity ranges can be controlled via firmware by setting the transceiver's transmission at 4 different levels of powers: 0, -6, -12, -18 dBm. Low-power transmissions entail lower ranges of proximity. The SocioPatterns platform has proved to work well in a number of F2F application scenarios that are discussed in the next section. For this reason the experiments presented in this paper have been implemented using such a platform.

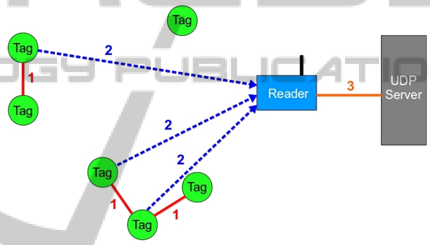

Figure 1: The communication process.

Applications. SocioPatterns performed several installations in different social contexts analyzing the obtained dynamics, such as the study of relationships between attendees in conferences (Barrat et al., 2010; der Broeck et al., 2010) or the spread of disease in hospitals and schools (Isella et al., 2011; Vanhems et al., 2013; Stehlé et al., 2011). One of the most important work is certainly Live Social Semantics (LSS) (Barrat et al., 2010), which is an applications that relates virtual and real interactions among individuals during conferences. Furthermore, since sensing human behaviors in real-world scenarios opens new frontiers in ubiquitous areas, SocioPatterns have started studying and analyzing characteristics in this kind of social networks (Barrat et al., 2013). Analogous experiments to the SocioPatterns' ones were deployed by Chin et al. (Chin et al., 2012) giving each person an active RFID badge during the course of a conference. They were interested in realizing a system able to find and connect people to each other. Analyzing such an experiment, they discovered that more proximity interactions result in an increased probability for a person to add another as a social connection. Finally, Becchetti et al. deployed two large-scale social experiments (Becchetti et al., 2012), one at the Department of Computer Engineering in Rome and another one at the MACRO museum during the opening of the NEON exhibition. They collected data coming from active RFID tags worn by around 120 volunteers normally moving and interacting in the monitored areas.

3 CAPTURE F2F INTERACTIONS

A number of protocols nowadays exists for wireless sensor networks applications (Demirkol et al., 2006), opportunistic networks contexts (Conti et al., 2010) or mobile ad hoc networks (Taneja and Kush, 2010). Nevertheless, most of those protocols are not well suited to track fast-changing F2F networks where interactions are very dynamic and potentially might change the structure of the network at every time-step. We stress that the main purpose of our experiments is to accurately track F2F interactions rather than to allow users to exchange messages. In this perspective, the proposed MAC protocols are primarily designed to quickly exchange small packets to build a neighborhood map rather than to exchange relatively big packets.

In this section we introduce **SOCIOMAC** and **PROXMAC**, two MAC protocols specifically designed to track F2F interactions, and we discuss how the proper tuning of the protocol parameters can allow us to effectively meet the needs of heterogeneous F2F tracking scenarios. In particular, depending on the expected dynamics of the underlying social graph (e.g., the interaction time), the protocols have to strike a balance between accuracy and network life-time.

SocioPatterns tags were initially provided with **SOCIOMAC** (soc, 2014) that has been designed for experiments characterized by long-lasting F2F interactions. **SOCIOMAC** allows to identify nodes in the proximity and to deliver the position and the list of encountered nodes to the readers. Communication is encrypted to support user privacy and packets are broadcast to the readers in a best-effort fashion with no acknowledgment mechanism to minimize the overhead of the communication. The three types of packets used by the protocol are: *contact* (or *proximity*), *beacon* (or *sighting*), *report*. Contact packets are used to identify close-by tags. This information, as well as the position of the node, is then delivered to the reader using the report and beacon packets, respectively.

Figure 2 illustrates an entire phase of the default **SOCIOMAC** protocol. Each phase, composed of 8 cycles, represents the main policy of the protocol. Every cycle starts with a random sleeping period `S`, between

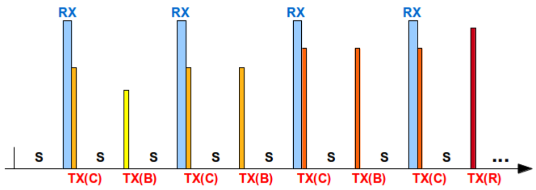

Figure 2: The **SOCIOMAC** protocol.

50 and 91 ms, to avoid periodic collisions, then it is concluded alternatively by the transmission of contact packets `TX(C)` or by the transmission of beacon packets `TX(B)`. Before performing a `TX(C)`, tags go in receiving mode `RX` for 5 ms so as to receive contact packets from possible neighbors. A phase is concluded with the 8th cycle, when a report packet `TX(R)` is sent to the reader. It can be computed that the average duty cycle of **SOCIOMAC** is less than 10% making this protocol suitable for long-lasting deployments; however the probability P_{succ} of exchanging a contact packet between two neighbors in a phase "only" ranges between 10% and 17%.

According to the analysis and evaluation given in (Cattuto et al., 2010), a resolution of $\Delta = 20$ seconds is considered sufficient to accurately track F2F interactions and to consequently assess proximity. More specifically, a contact between two individuals is established with a probability in excess of 99%, namely there is a link in G_i, as long as at least one of them receives a contact packet from the other every 20 seconds. However, such a time scale could be not sufficient in some cases. In particular, we consider a scenario to investigate the so called "wisdom of crowds" (Surowiecki, 2004), where a group of people is asked to answer a set of questions in two separate phases. In the first phase participants have to guess the answer without having any kind of relationship with others, while in the second phase they are allowed to interact and exchange their opinions. Since users are already aware of the questions asked in the first phase, in the second round fast interactions, such as "What did you answer the first question?", are likely to happen. In this case, a higher resolution (i.e., a lower time scale, such as 5 seconds) may be required in order to capture all those kinds of interactions, consequently it is crucial a fine-tuning of the **SOCIOMAC** parameters to increase P_{succ}.

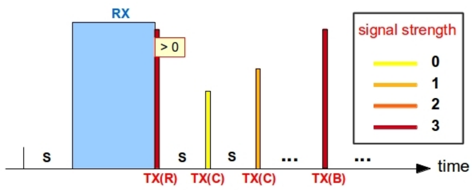

Figure 3: The **PROXMAC** protocol.

Since in short-lasting deployments energy constraints are less demanding and interactions tend to be very fast and dense, at the cost of a higher duty cycle, the first tuning that had to be made was the increase of the receiving interval so as to collect with a higher probability incoming packets from fast interactions. We recall that the **SOCIOMAC** protocol sets the `RX` interval to 5 ms; because the transmission of a packet approximately lasts the same order of magnitude, the protocol is not always able to capture packets in their first exchanges. For that reason, we decided of redesigning the main policy of the protocol, named **PROXMAC** and shown in Figure 3, so that a fine-tuning of the parameters would have been achievable. In the next section we will show how P_{succ} increases, making this protocol more suitable for detecting fast dynamics. Values of the sleeping period and the receiving interval can be configured according to the scenario which is going to be deployed.

A sketch of the phase is shown in Figure 3. Notice that, in this new policy, the report packet can be sent after every receiving operation only if the tag has previously received some contact packets, otherwise this transmission is skipped. This feature guarantees communications broadcasting only useful packets and reducing contingent collisions towards the readers. Furthermore, to better discriminate F2F links, `TX(C)` are transmitted with signal strengths 0 and 1, namely -18 and -12 dBm, while the default **SOCIOMAC** protocol uses values of 1 and 2, making it more prone to false positives. Finally, a tag running **PROXMAC** sends a *beacon* packet only every t phases, so as to avoid to incessantly flood the channel. The t parameter is usually configured to receive at least one beacon packet at each time-step. That ensures a continuous monitoring of the tag inside the network. Both the beacon and report packets are sent using a signal strength of 3 in order to ensure with a good probability the collection of the packet by close readers. Some other minor optimizations of the protocol have been implemented but will not be discussed in this paper for the sake of space.

Of course, depending on the chosen parameters, we can experience with different duty cycles and probabilities P_{succ} of successfully receiving a packet. As sample of setting, also used in the experiments discussed in the next section, we pick out a sleeping period ranging in $[20,30)$ ms and a receiving interval in a range between 30 and 39 ms. Such values, as well as other settings, were tested in our labs and, after a long time of measurements, we found that such a configuration was a good compromise in terms of performance for our purposes and social experiments. Indeed, we were interested in deploying social experiments where a high number of iterations of the main phase in a time-step was required to capture fast interactions as much as possible in dense real F2F social networks. Other values of `S` and `RX` are clearly available and allowed; they only depend on the purpose of the deployment. However, you should take into account the fact that the minimum value of the receiving interval should always be equal or greater than the maximum value of the sleeping period, so that the **PROXMAC** policy is still guaranteed. Indeed, if that condition is not fulfilled, then it may happen that a whole receiving interval completely falls in the middle of the sleeping period, thus compromising P_{succ}.

As we will better see in the next section, using those parameters, **PROXMAC** can almost always ensure the reception of at least one proximity packet every 2 or 3 time-steps if two tags are facing each other at 1-1.5 meter. Therefore, people having very fast interaction can be now tracked and logged. Unlike **SOCIOMAC** where the recommended resolution is $\Delta = 20$ seconds and the maximum signal strength for collecting contacts is 2, **PROXMAC** can exploit a more fine-grained Δ and better manage false positives thanks to a reduced signal strength. An improvement of the performance even depends on the new length of the phase being very short, so that the protocol can iterate the main cycle several times within a time-step, which is the most fine-grained resolution.

4 EXPERIMENTS AND EVALUATION

In this section we evaluate the **SOCIOMAC** and **PROXMAC**'s performance under different conditions in real-world testbeds.

4.1 Preliminary Experiments

The first experiment measures the number of contact packets exchanged in 10 minutes between two F2F tags 1 meter apart and then collected by a reader. This kind of setting is used to simulate an interaction as happens in real-world scenarios. Of course, in real-world testbeds, communications are more challenging with respect to this preliminary experiment because of overlapping transmissions and dense networks. However, this first experiment allows us to figure out how protocols perform in an ideal setting, so that we can understand when to use determined settings in real-world social networks in which performance measures are too difficult to explore and analyze in detail. The plots depicted in Figure 4 show the results of the experiments. The x-axis represents the time of the testbed, while the y-axis the total number of distinct proximity packets collected by both the protocols. Notice that, the slope of both the plots is less than 45 degrees. This means that for both the protocols, a resolution of 1 second (i.e., the minimum time-step in real-world scenarios) cannot be achieved. However, as expected, **PROXMAC** provides better performance in terms of proximity collection. In this simple setting it is able to support an average resolution of $\Delta = 1.8$ seconds, while **SOCIOMAC** supports a resolution of $\Delta = 2.9$ seconds.

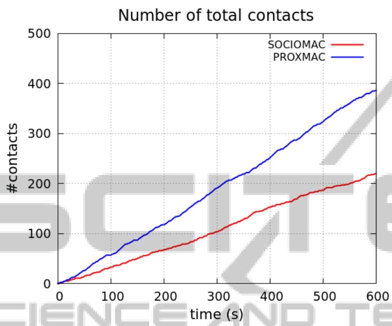

Figure 4: All the packets exchanged over time.

The second experiment is shown in Figure 5 and is used to evaluate the number of false positives reported by both the protocols, namely all those contact packets exchanged by two tags *not* facing each other. In this experiment two couples of nodes are placed at a distance of 2.5 meters. As in the previous experiment, the two tags in a couple are 1 meter apart. In principle, a tag should detect only the other one in the couple (i.e., a F2F interaction). However, due to the proximity of the two couples, false positives may occur. Of course, similar behaviors occur in real-world scenarios whenever there are high network densities. For instance, it can happen that two independent groups of people at close range wrongly exchange contact packets.

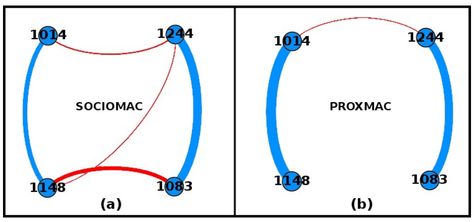

Figure 5: Correct contacts and false positives.

As shown in Figure 5(b), most of the **PROXMAC** contact packets are correctly exchanged (blue arcs) between the tags in the couple. On the contrary, that is not the same for **SOCIOMAC** (see Figure 5(a)), where a number of false positive occurs (red arcs). The thickness of the edges is proportional to the number of contact packets exchanged between the tags. A ticker edge indicates a longer interaction, while a thin arc shows a short-lasting interaction. This feature helps to visualize what kinds of contacts were common. In the **SOCIOMAC** case, the formed network is the result of 5 edges, including two regular interactions and three false positives. The two more thin false positives could be excluded using a filter that analyzes the structure of the network enforcing a certain threshold, but the third arc has a thickness pretty similar to one of the two regular. Therefore, there is a significant problem in selecting which arc to keep and which arc to discard. The main cause of this issue in **SOCIOMAC** depends on the signal strength used by the protocol and the corresponding policy. As already described in section 3, **SOCIOMAC** allows tags to exchange contact packets using three level of power, including the signal strength 2. This brings to a wider range of communication catching farther tags which cannot be considered face-to-face. This issue is rather reduced using **PROXMAC**. Indeed, the only false positive registered during the experiment was the thin edge between the nodes 1014 and 1244. However, when the thickness of false positives is quite thin with respect to other regular edges, then that arc can be easily rejected using some techniques of filtering in post-processing.

4.2 Real-World Social Experiments

WSDM 2013 Conference. On February 2013 we deployed a real-world social experiment in Rome during the WSDM Conference, where 69 attendees agreed to wear our tags running the **SOCIOMAC** protocol. The experiment lasted more than 1 hour, where 50 minutes were allocated for the social interaction, thus collecting data from a large area, including several rooms, the corridor and common spaces. The main purpose of the experiment was to collect data to study how F2F interactions can possibly influence the wisdom of a group of people, which is usually called the "wisdom of crowds" phenomenon (Lorenz et al., 2011). People were not constrained to talk with a restricted group of other participants, but they were allowed to interact with anyone they wanted.

At the end of the experiment we collected more than 23000 single proximity packets. Figure 6 shows the social network built aggregating all graphs collected over time. Each of them was built with a resolution of $\Delta = 3$ seconds and an edge threshold of $\theta = 5$ proximity contacts. This means that an edge was built only if at least 5 proximity packets were collected, each of which had to be captured within a 3-second time interval. Such a resolution was chosen according to the results obtained in section 4.1, while that θ empirically gives a good chance to record a real interaction. Indeed, assuming to choose a lower θ, an edge may be established by the simple fact that two users can bump into each other just for a moment. Darker and bigger nodes of the graph represent people that in the whole experiment accumulated the bigger number of F2F interactions, while the thickness of an edge is proportional to the number of interactions observed between two nodes over the whole time.

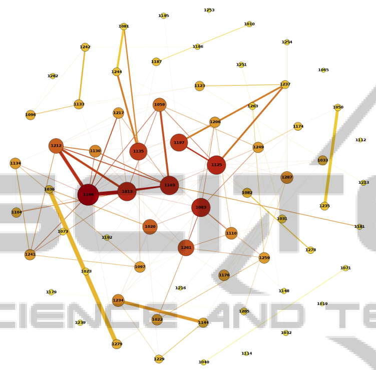

Figure 6: The social graph of the network formed during the experiment in Priverno.

The graph is made of 69 nodes and 133 undirected edges. The average degree is 3.855, while the network diameter and the average path length are 7 and 3.222, respectively. Only 6 nodes are isolated, while the maximum degree is 13. Finally, the average density of the graphs in the evolving network is 0.003.

Priverno's Country Fair. Similarly to the previous experiment and with the same modalities, on May 2014 we deployed another real-world social experiment in Priverno (LT), Italy, during a country fair, recruiting 60 volunteers. As usual, each participant wore a tag, but this time running the **PROXMAC** protocol. They were free to move within a large monitored area (around 10 x 15 meters) in a green park. The total experiment lasted around 30 minutes to study the same social phenomenon of WSDM. The interaction part of the experiment lasted around 10 minutes, time in which we collected more than 11000 single proximity packets. The graph in Figure 7 depicts the aggregated result of the social interaction. Since the preliminary experiments showed that **PROXMAC** was able to collect almost twice the number of proximity packets than **SOCIOMAC**, in this experiment we leave a resolution of $\Delta = 3$ seconds, but we set an edge threshold of $\theta = 10$. The same resolution allowed us to compare this experiment to the previous one, while a double

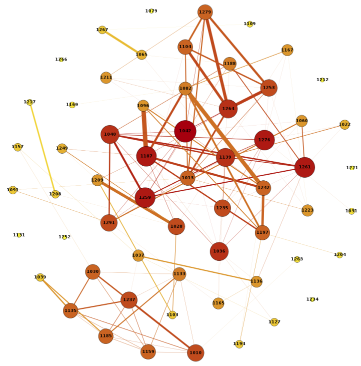

Figure 7: The social graph of the network formed during the WSDM experiment.

edge threshold tested the performance of **PROXMAC** measured in the ideal experiments.

The graph formed of 60 nodes and 128 undirected edges has an average degree of 4.267, a network diameter and the average path length of 9 and 3.713, respectively. 7 users result isolated, while the maximum degree is 10. Finally, the average density of the graphs is 0.012. Despite this network is quite similar in terms of size with respect to the WSDM's one, we can notice how the Priverno's density is an order of magnitude greater than the WSDM's one. This fact, besides a likely higher participation, is also due to **PROXMAC** which could collect many more proximity packets. Clearly, these small real-world social networks are only a starting point to test such collection protocols. For this reason, in the next section, we simulated the behavior and the performance of both the protocols in larger graphs and higher densities.

4.3 Simulation on Larger and Denser Graphs

A simulation on graphs having many more nodes and higher densities allows us to analyze what kind of situation we may encounter in future and bigger real-world scenarios. First of all, we remind that **SOCIOMAC** sends report packets to the readers only after 8 cycles, while **PROXMAC** can report after every receiving operation if new contact packets have been received. Due to this policy, in dense networks **SOCIOMAC** can more easily experience a buffer overflow (tags have room for only 4 packets) and consequently discard new incoming contacts. We set up a simulation for different graph densities and sizes for a total of 1200 time-steps. For each step, we generate a random graph based on the Erdos-Renyi model (Erds and Rnyi, 1960), with $n = \{60, 100, 200, 500\}$ and $p = \{0.01, 0.02, 0.05, 0.1\}$, where n is the num-

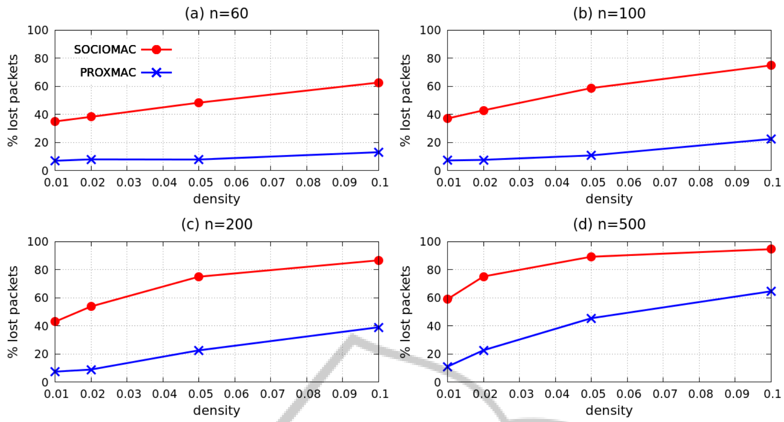

Figure 8: The percentage of lost packets for different graph densities and number of nodes.

ber of nodes and p is the probability of including an edge in the graph. Notice that for a sufficient large amount of time-steps, the probability p approximates the average density of the whole evolving network. Figure 8 shows the results of our simulation.

The equivalent case to the real-world experiment deployed in Priverno is depicted in Figure 8(a), where the number of nodes is 60 and the density is 0.01. **PROXMAC** turned out to be very efficient in this case, where the percentage of lost packets is around 7%. Vice versa, **SOCIOMAC** was not able to collect almost the 40% of the total possible packets. A similar behavior can be observed also in the other cases considered in the simulation where **PROXMAC** is able to better support denser graphs. Indeed, we recall that **PROXMAC** has a possible forwarding action after every receiving operation, while **SOCIOMAC** forwards its contact packets only every 8 cycles. Also, the whole phase lasts less than the **SOCIOMAC**'s one. This brings to a faster actions cycle in **PROXMAC**, which is able to execute each operation several times in every time-step. Still considering the plot with $n = 60$, **PROXMAC** has a similar percentage of lost packets for higher densities as well, while the **SOCIOMAC**'s trend slightly increases. Their behavior changes little for the graph having 100 nodes, but increases in graphs with $n = 200$ and $n = 500$. However, for a density of 0.01, **PROXMAC** was able to maintain a very low level of loss. On the contrary, **SOCIOMAC** starts with a level of loss equals to 40% and 60%, respectively. Then, their trends increase and reach more than 80% for **SOCIOMAC** with a density 0.1, while **PROXMAC** amounts to 40% in $n = 200$ and 60% in $n = 500$. We recall, however, that very high densities are unlikely in real-world scenarios, as already observed in our social experiments.

5 CONCLUSION AND FUTURE WORK

The ability of capturing F2F interactions opens the possibility of quantitatively exploring the way in which humans interact in real-world scenarios. The heterogeneity of such scenarios, requires a proper tuning of the collection protocols in order to achieve the most effective trade-off between resolution and network life-time.

In this paper we compared the **SOCIOMAC** protocol proposed by the SocioPatterns initiative and already used in monitoring activities, with the new **PROXMAC** protocol, specifically designed to support the monitoring of fast-changing interactions. The experimental results, performed both in real-world social experiments and in denser and bigger simulated networks, confirm that **PROXMAC** can collect twice the number of proximity packets than **SOCIOMAC**, thus providing a better resolution for capturing fast interactions. Thanks to an optimized protocol policy, **PROXMAC** delivers report packets more efficiently than **SOCIOMAC**, which suffers in denser networks.

In our vision, the collection of reliable F2F interactions will support a better understanding of the so-called “wisdom of crowds” phenomenon (Surowiecki, 2004). Indeed, all the proposed models in this research area (such as (DeGroot, 1974; Friedkin and Johnsen, 1990; Bindel et al., 2011)) rely on the availability of interaction graphs.

In the future, we plan to extend the real-world testbeds considering a higher number of individuals in order to study bigger and denser networks.

REFERENCES

(2014). The SOCIOMAC patent. https://www.google.com/patents/US8660490.

(n.d.). SocioPatterns. http://www.sociopatterns.org/.

Adomavicius, G. and Tuzhilin, A. (2005). Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *Knowledge and Data Engineering, IEEE Transactions on*, 17(6):734–749.

Avin, C., Koucký, M., and Lotker, Z. (2008). How to explore a fast-changing world (cover time of a simple random walk on evolving graphs). In *Automata, Languages and Programming*, pages 121–132. Springer.

Barrat, A., Cattuto, C., Colizza, V., Gesualdo, F., Isella, L., Pandolfi, E., Pinton, J.-F., Rav, L., Rizzo, C., Romano, M., Stehl, J., Tozzi, A., and Broeck, W. (2013). Empirical temporal networks of face-to-face human interactions. *The European Physical Journal Special Topics*, 222(6):1295–1309.

Barrat, A., Cattuto, C., Colizza, V., Pinton, J.-F., den Broeck, W. V., and Vespignani, A. (2008). High resolution dynamical mapping of social interactions with active rfid. *CoRR*, abs/0811.4170.

Barrat, A., Cattuto, C., Szomszor, M., den Broeck, W. V., and Alani, H. (2010). Social dynamics in conferences: Analysis of data from the live social semantics application. In *Proceedings of the 9th International Semantic Web Conference (ISWC 2010)*.

Becchetti, L., Bergamini, L., Ficarola, F., and Vitaletti, A. (2012). Population protocols on real social networks. In *Proceedings of the 9th ACM Symposium on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks*, PE-WASUN '12, pages 17–24, New York, NY, USA. ACM.

Bindel, D., Kleinberg, J., and Oren, S. (2011). How bad is forming your own opinion? In *Proceedings of the 2011 IEEE 52Nd Annual Symposium on Foundations of Computer Science*, FOCS '11, pages 57–66, Washington, DC, USA. IEEE Computer Society.

Cattuto, C., Van den Broeck, W., Barrat, A., Colizza, V., Pinton, J.-F., and Vespignani, A. (2010). Dynamics of person-to-person interactions from distributed rfid sensor networks. *PLoS ONE*, 5(7):e11596.

Chin, A., Xu, B., Wang, H., and Wang, X. (2012). Linking people through physical proximity in a conference. In *Proceedings of the 3rd international workshop on Modeling social media*, MSM '12, pages 13–20. ACM.

Conti, M., Giordano, S., May, M., and Passarella, A. (2010). From opportunistic networks to opportunistic computing. *Communications Magazine, IEEE*, 48(9):126–139.

DeGroot, M. H. (1974). Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121.

Demirkol, I., Ersoy, C., and Alagoz, F. (2006). Mac protocols for wireless sensor networks: a survey. *Communications Magazine, IEEE*, 44(4):115–121.

der Broeck, W. V., Cattuto, C., Barrat, A., Szomszor, M., Correndo, G., and Alani, H. (2010). The live social semantics application: a platform for integrating face-to-face presence with on-line social networking. In *First International Workshop on Communication, Collaboration and Social Networking in Pervasive Computing Environments (PerCol 2010)*.

Erds, P. and Rnyi, A. (1960). On the evolution of random graphs. In *Publication of the Mathematical Institute of the Hungarian Academy of Sciences*, pages 17–61.

Friedkin, N. and Johnsen, E. (1990). Social-influence and opinions. *Journal of Mathematical Sociology*, 15(3-4):193–205.

Hui, P., Chaintreau, A., Scott, J., Gass, R., Crowcroft, J., and Diot, C. (2005). Pocket switched networks and the consequences of human mobility in conference environments. In *Proceedings of ACM SIGCOMM first workshop on delay tolerant networking and related topics*.

Isella, L., Romano, M., Barrat, A., Cattuto, C., Colizza, V., Van den Broeck, W., Gesualdo, F., Pandolfi, E., Ravá, L., Rizzo, C., and Tozzi, A. E. (2011). Close encounters in a pediatric ward: Measuring face-to-face proximity and mixing patterns with wearable sensors. *PLoS ONE*, 6(2):e17144.

Jayagopi, D. B., Kim, T., Pentland, A. S., and Gatica-Perez, D. (2010). Recognizing conversational context in group interaction using privacy-sensitive mobile sensors. In *Proceedings of the 9th International Conference on Mobile and Ubiquitous Multimedia*, MUM '10, pages 8:1–8:4, New York, NY, USA. ACM.

Lorenz, J., Rauhut, H., Schweitzer, F., and Helbing, D. (2011). How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences*, 108(22):9020–9025.

Stehlé, J., Voirin, N., Barrat, A., Cattuto, C., Isella, L., Pinton, J.-F., Quaggiotto, M., Van den Broeck, W., Regis, C., Lina, B., and Vanhems, P. (2011). High-resolution measurements of face-to-face contact patterns in a primary school. *PLoS ONE*, 6(8):e23176.

Surowiecki, J. (2004). *The Wisdom of Crowds*. Anchor Books.

Taneja, S. and Kush, A. (2010). A survey of routing protocols in mobile ad hoc networks. *International Journal of Innovation, Management and Technology*, 1(3):2010–0248.

Vanhems, P., Barrat, A., Cattuto, C., Pinton, J.-F., Khanafer, N., Rgis, C., Kim, B.-a., Comte, B., and Voirin, N. (2013). Estimating potential infection transmission routes in hospital wards using wearable proximity sensors. *PLoS ONE*, 8(9):e73970.

Walter, F., Battiston, S., and Schweitzer, F. (2008). A model of a trust-based recommendation system on a social network. *Autonomous Agents and Multi-Agent Systems*, 16(1):57–74.

Wu, H. and Pan, Y. (2008). *Medium Access Control in Wireless Networks*. Nova Science Publishers.

Yu, H., Shen, Z., Miao, C., Leung, C., and Niyato, D. (2010). A survey of trust and reputation management systems in wireless communications. *Proceedings of the IEEE*, 98(10):1755–1772.