Have You Also Seen That? Collaborative Alert Assessment in Ad Hoc Participatory Sensing

Fátima Castro-Jul(✉), Rebeca P. Díaz-Redondo, and Ana Fernández-Vilas

Universidade de Vigo, Campus Universitario s/n, 36310 Vigo, Spain
{fatima,rebeca,avilas}@det.uvigo.es

Abstract. Due to their flexibility and immediacy, participatory sensing systems based on ad hoc networks have become a valuable source of proximity-based information about events and incidents in the city. However, easy contribution may result in the transmission of larges amount of data, some of them being duplicate or not relevant. We have proved that collaboratively filtering the information to be disseminated improves network efficiency while maintaining the system effectiveness.

Keywords: Ad hoc networks · Smart cities · Collaboration · Alert dissemination

1 Introduction

The proliferation of sensor-rich handheld devices has allowed users to constantly monitor their surroundings. The fact that they can easily share their findings with their peers has given rise to participatory sensing schemes [1], where their contributions form collective knowledge that can be used to retrieve any kind of local information. Thus, participatory sensing provides an opportunity to collaboratively detect and stay aware of what is going on anytime anywhere. This is specially interesting when it comes to alert about incidents in urban areas, where device density allows reliance on participatory strategies.

Once devices detect an event there are mainly two ways in which they can disseminate the information: social and ad hoc networks. Contribution to both is straightforward and available for any user ready to take part on them. On the one hand, this enriches and makes the detection process flexible. On the other hand, easy contribution can result in big amounts of irrelevant or duplicate data. As a result, it is necessary to establish mechanisms able to deal with user-generated information in order to filter received contributions.

Alert detection in social networks is performed using data mining techniques [2], such as *keyword burst* [3] or supervised classification [4]. Thus, it consists on a centralized system that processes large amounts of data and detects new events by finding trends or similarities between different messages. As a consequence, these systems benefit from receiving abundant data and do not worry about duplicate or irrelevant information, which is dealt with through data processing.

C.R. García et al. (Eds.): UCAmI 2016, Part II, LNCS 10070, pp. 125–130, 2016.
DOI: 10.1007/978-3-319-48799-1_15

However, in ad hoc alert systems the situation is quite different. There exist no central entity in charge of data processing and, as a result, there is no filter for user-generated alerts. This can result in a large amount of unnecessary transmissions that may provoke network congestion or even collapse. Research on network efficiency has been extensive but the attention has been aimed mainly at routing and dissemination techniques [5–7], instead of at the alert content.

There exist, however, proposals that assess message content for message dissemination. One possibility is to establish hierarchical structures, whose head is in charge of alert aggregation and transmission [8,9]. Nevertheless, this involves a setup cost that decreases network flexibility. An alternative consists on filtering on the receiver side. Nodes receive first metadata about available new messages in order to decide whether they want to receive them or not taking into account their interests and energy constraints [10,11]. Although this reduces processing overhead and energy expenditure, including a negotiation phase on every hop of the data dissemination process is much more resource-consuming than applying a starting filter, where the overhead falls only on the initial neighbor nodes.

In this paper, we introduce a collaboration strategy that involves neighbor nodes in alert assessment. Thus, we include a filtering step in order to reduce irrelevant or duplicate alert transmissions in ad hoc networks. And we do so in a distributed manner, without adding any centralized or hierarchical structure. Our evaluation of the strategy, using network simulations on urban scenarios, demonstrates the outcome of our approach.

The paper is organized as follows. In Sect. 2 we start by explaining the urban scenario we have considered and our collaborative assessment proposal. Then, we detail the simulation process (Sect. 3). Finally, we outline our conclusions and current line of research (Sect. 4).

2 Scenario and Collaborative Assement Proposal

We consider an urban scenario with a high density of pedestrians, equipped with sensor-rich handheld devices that allow them to monitor their surroundings and inform about their findings. Devices include short-range radio technologies, such as WiFi or Bluetooth, that allow them to communicate with each other directly. As a consequence, no dedicated infrastructure is required. Connections build a flat network, where no hierarchies or clusters are formed since the high pedestrian mobility does not make setup costs worth the effort. As a result, connections are flexible and proximity-based, which makes it simple for any device to join and contribute to the detection and dissemination process. Our system is targeted at general interest alerts, such as traffic incidents, public transport issues, crowds or unsafe situations, like robberies or fights. Thus, we ensure users' concern and ability to discern them.

We propose a collaboratively-triggered dissemination mechanism composed of three steps: user alert, collaborative assessment and alert dissemination.

The alert process may be triggered manually by users or automatically by devices. In the latter case, devices rely on their sensors data in order to start or assess an alert. For instance, a car accident or a big crowd can be detected using sound sensors. When the incident is detected, a broadcast notification is sent to neighbor nodes in order to be assessed and, if approved, disseminated.

Our contribution focuses on how to assess alert relevance and prevent alert duplication by using ad-hoc device cooperation. Neighbor nodes receiving the assessment request must decide whether to confirm the alert or not by sending an assessment reply to the requester, either positive or negative. This kind of assessment can be employed in scenarios where nodes' support requires users' intervention or where it is performed autonomously. Since we employ proximity-based communications, collaboration is restricted to physically close devices. Thus, only nodes with actual knowledge of the area, and therefore able to confirm or refute the alert, take part in the assessment process. This represents an advantage regarding social network approaches, where it is necessary to verify alert senders' location. Devices remember alerts they have been required to assess so they will not start a new dissemination if they detect the same event. However, if an alert has been received through flooding dissemination, instead of a direct request for assessment, the node will start its own assessment process as the alert has been originated outside its local area.

If the alert is backed up by neighbors, that is to say, most of the received replies in a certain time are confirmations; then it is disseminated through the whole network using simple flooding. Since we aim our attention to the assessment stage, developing an efficient strategy in the dissemination phase is out of the scope of this paper. Alert verification takes part exclusively in the first sender's neighborhood, no negotiation or agreement is required in its further dissemination. Negotiation on every hop is suitable for interest-based dissemination but not for alerts, which are considered of general interest and for whose verification we rely only on devices in the incident area.

Figure 1 details the alert process described in this paper. A pedestrian alerts about an event in its surroundings. The alert is first shared with its neighbors in order to be assessed (green arrows). Then, if they agree with the alert content, the alert is disseminated through the whole area using flooding (red arrows). In order to implement collaborative assessment, it is necessary to establish when several alerts can be considered the same. Alert description consists only on the incident code, which results in easier message processing. Categories must be narrow enough to be able to give significant information about certain incidents but broad enough not to make the detection process too complex. The list of possible categories can be adapted to different systems' purposes.

3 Simulation

In order to perform a simulation based on a real-life urban scenario, we have chosen the city center of Vigo (Spain). The chosen region, shown in Fig. 1, comprises an area of 1.46 km^2 and includes commercial streets and Vigo's old town. As a result, most of the streets are either pedestrian-only or include wide sidewalks, which makes the area suitable for pedestrian simulation.

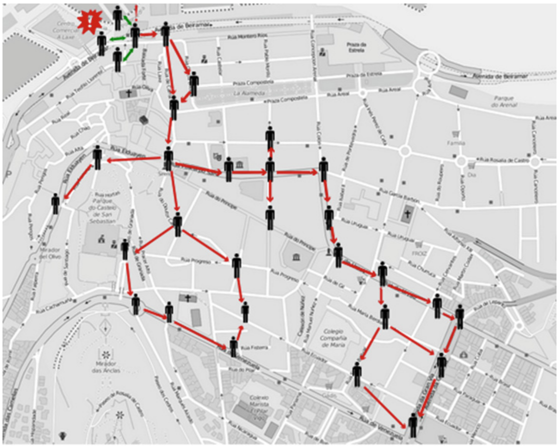

Fig. 1. Urban area used in simulation. (Color figure online)

We have exported the selected area from OpenStreetMap [12] and populated it with random pedestrian routes using SUMO urban mobility simulator [13]. Then, we have exported pedestrian traces to ns-3 network simulator [14], where we have implemented the assessment strategy and carried out the simulation. We have compared it with a simple flooding mechanism in situations where several neighbor devices detect the same incident and alert about it.

We have implemented the assessment strategy proposed in Sect. 2. We have set a time to wait for assessment replies of 100 ms. We have assumed that all the alerts are relevant and therefore, that they will be all backed up by neighbors. Besides, nodes reply automatically once they receive the assessment request, without users' intervention.

Transmission technology chosen for the implementation was ad-hoc WiFi. Reasons for this election include its availability in current smartphones and its transmission range ($\simeq$100 m), which makes it appropriate for urban communications.

We have carried out simulations with 100, 200, 500 and 1000 pedestrians in the area in order to assess the collaborative assessment behavior with different device density. In each one of those, we have simulated situations where 1, 2, 5, 10 or 20 nodes alert about the incident. Physically close nodes have been chosen as alert senders and time interval between their alerts has been set to one second. The following network parameters were considered as performance metrics:

- Message efficiency. We have measured the total number of messages that are necessary for an alert dissemination, including assessment requests, replies and disseminated alerts.
- Delivery ratio. We have measured the percentage of nodes in the area that receive the alert. Due to the randomness of the positioning and the characteristics of the urban area, some of the pedestrians are placed on the edges and cannot be reached. For this reason, delivery ratio never achieves 100 %.

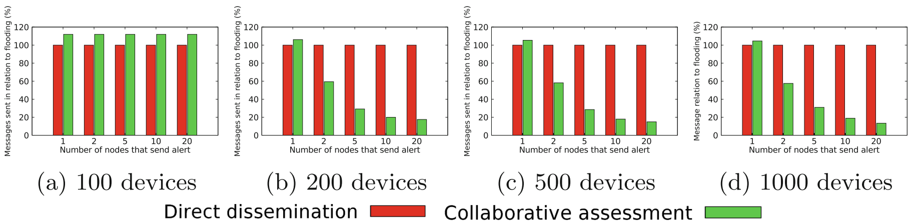

Fig. 2. Percentage of messages sent in relation to flooding.

3.1 Evaluation

We have concluded that our assessment mechanism outperforms direct flooding dissemination as long as more than one device alerts about the incident and there is enough device density. 100 devices have been proved to be too few to cover the considered area. As shown in Fig. 3(a), delivery ratio is always below 50 %. Distance between nodes prevents further dissemination while maintaining the quantity of messages sent constant. Moreover, assessment stage messages make the collaborative approach more inefficient than flooding, as shown in Fig. 2(a). From 200 devices on we can observe a clear trend on the system behavior. For one alert sender message efficiency of flooding outperforms the collaborative (Fig. 2) as it does not require any extra messages apart from the dissemination ones. For two alert messages, collaborative assessment requires about half of the messages flooding does. As the number of the nodes that send alert increases, so does the difference between flooding and collaborative assessment. With 20 alert senders, it requires less than 20 % of the messages flooding employs. And that happens with little difference in the delivery ratio, which is about 90 % with 200 and 500 devices and around 98 % for 1000 (Fig. 3). It could be argued that flooding is an inefficient dissemination technique that can easily be outperformed, but it has to be taken into account that flooding is also the dissemination mechanism used in the collaborative approach. As a result, comparison is made on equal terms.

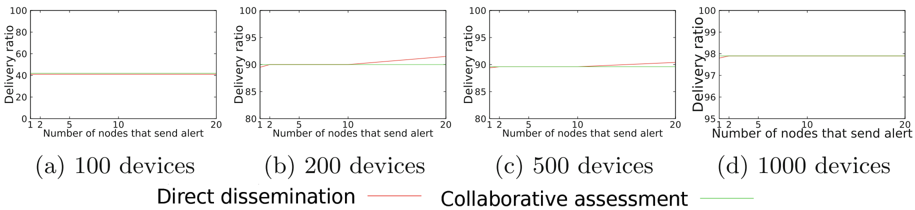

Fig. 3. Percentage of nodes in the area that receive the alert.

4 Conclusion and Future Work

To the best of our knowledge, no other work has been presented targeted at collaboratively filtering what is worthy to be disseminated in an ad-hoc alert system. Our proposal has been proved able to reduce network load while maintaining a good delivery ratio. Moreover, it does not require setup costs or negotiations on every step. The quantity of messages employed in the assessment phase has been shown to be small and therefore, the strategy cost is minor. However, we believe it could still be reduced. Replying to every alert proposal they receive may imply a significant overhead for mobile devices in crowded scenarios, where many alerts may be generated. Therefore, our current work focuses on developing a new collaborative assessment strategy that reduces this burden.

Acknowledgments. This work is funded by: the European Regional Development Fund (ERDF) and the Galician Regional Government under agreement for funding the Atlantic Research Center for Information and Communication Technologies (AtlantTIC) and the Spanish Ministry of Economy and Competitiveness under the National Science Program (TEC2014-54335-C4-3-R).

References

1. Campbell, A.T., Eisenman, S.B., Lane, N.D., Miluzzo, E., Peterson, R.A.: People-centric urban sensing. In: WICON 2006, pp. 18–31 (2006)
2. Imran, M., Castillo, C., Diaz, F., Vieweg, S.: Processing social media messages in mass emergency: a survey. ACM Comput. Surv. **47**(4), 67:1–67:38 (2015)
3. Mathioudakis, M., Koudas, N.: Twittermonitor: trend detection over the twitter stream. In: SIGMOD 2010, pp. 1155–1158 (2010)
4. Ritter, A., Mausam., Etzioni, O., Clark, S.: Open domain event extraction from twitter. In: SIGKDD 2012, p. 1104 (2012)
5. Goyal, D., Tripathy, M.R.: Routing protocols in wireless sensor networks: a survey. In: ACCT 20122, pp. 474–480 (2012)
6. Kokuti, A., Simon, V.: Location based data dissemination in mobile ad hoc networks. In: TSP 2012, pp. 57–61 (2012)
7. Tonguz, O., Wisitpongphan, N., Parikh, J., Bai, F., Mudalige, P., Sadekar, V.: On the broadcast storm problem in ad hoc wireless networks. In: BROADNETS 2006 (2006)
8. Saleet, H., Basir, O.: Location-based message aggregation in vehicular ad hoc networks. In: GLOBECOM, pp. 1–7 (2007)
9. Sun, B., Wu, K., Pooch, U.W.: Alert aggregation in mobile ad hoc networks. In: WiSE 2003, pp. 69–78 (2003)
10. Heinzelman, W., Kulik, J., Balakrishnan, H.: Adaptive protocols for information dissemination in wireless sensor networks. In: MobiCom 1999, pp. 174–185 (1999)
11. Rehena, Z., Roy, S., Mukherjee, N.: A modified SPIN for wireless sensor networks. In: COMSNETS, pp. 1–4 (2011)
12. OpenStreetMap (2015). http://wiki.openstreetmap.org/
13. Krajzewicz, D., Erdmann, J., Behrisch, M., Bieker, L.: Recent development and applications of SUMO simulation of urban mobility. Int. J. Adv. Syst. Meas. **5**(3&4), 128–138 (2012)
14. NS - 3 network simulator (2015). https://www.nsnam.org/