Leveraging Human Mobility in Smartphone Based Ad-Hoc Information Distribution in Crowd Management Scenarios

Tobias Franke*, Sascha Negele†, George Kampis* and Paul Lukowicz*

*DFKI; Kaiserslautern, Germany; Email: {*firstname.lastname*}@dfki.de
†University of Kaiserslautern; Kaiserslautern, Germany; Email: s_negele@cs.uni-kl.de

***Abstract*—We propose a novel approach for Ad-Hoc WiFi based distribution of information within large crowds of mobile users. The work is motivated by civil protection scenarios where infrastructure based communication often breaks down in cases of emergency. We follow a basic opportunistic networking approach by making use of the smartphones' built-in WiFi hotspot functionality which in combination with the devices switching between access point and client modes facilitates the propagation of messages on a multi-hop basis. We make three contributions with respect to previous work on this topic. First, we empirically determine core boundary conditions given by the performance of modern smartphones. To maximize system performance under such circumstances we propose novel heuristics for a mode switching strategy based on client mobility instead of random strategies that have mainly been utilized so far. Finally, we compare its performance to a random role switching strategy in a large-scale simulation based on a real dataset consisting of movement traces from 28'000 people during a three day festival in Zurich. Within the simulation we investigate the influence of various parameters on the system's behavior.**

I. Motivation

Managing crowds at large scale events and during everyday life in urban areas is a complex and highly dynamic problem. Up to date knowledge about the current state of a crowd (e.g. with respect to its density) is absolutely vital for operational personnel. This includes the earliest possible detection of potential hot spots and the means to communicate with the crowd in an efficient manner to solve problems before they turn into critical situations.

Over the past years, we therefore developed a smartphone app-based crowd management system, which analyzes smartphone sensor data voluntarily contributed by visitors of large scale events and creates a real time overview about crowd conditions such as crowd density, crowd pressure and crowd turbulences [1], [2]. More importantly from a communication perspective, the system also provides the means for event organizers and emergency forces to communicate with the crowd in a targeted fashion utilizing location based messages [3].

Until today, this system has been deployed during numerous international events and was consequently used by over 100.000 users who contributed well over 50 million data points in total. The deployment during the Zurich festival [4] resulted in the creation of the – according to our knowledge – largest GPS based mobility dataset collected at a public event at that time.

A key concern that has emerged during virtually all deployments was the ability to deal with network outages. Such outages must be expected during emergency situations and are also not uncommon at large enough public gatherings under normal circumstances. Clearly, a network outage during an emergency would mean that the system functionality (sending messages containing safety advice) would become unavailable at the very moment when it is most needed. As a consequence we have extended the app part of the system with the ability to use the users' smartphones to create an Ad-Hoc network within the crowd. We follow a basic opportunistic networking approach by making use of the smartphones' built-in WiFi hotspot functionality which in combination with the devices switching between access point and client modes facilitates the propagation of messages on a multi-hop basis. The use of the hotspot mode instead of the various peer-2-peer networking capabilities built into modern smart phone operating systems is motivated by the need to be compatible with as many devices as possible. While the P2P capabilities differ between devices, the ability to connect to a WiFi hotspot is virtually universal. While a number of publications (see related work) have dealt with this topic from a networking point of view before, there is currently little insight into the specific needs of a highly mobile smartphone-driven scenario.

In our work we present the boundary conditions imposed by the capabilities of smartphones (which differ greatly from standard networking devices) on the design of an opportunistic communication protocol. As our use case involves large crowds of people, we also consider the influence of the human body on the system's performance. We furthermore introduce a novel mobility based mode switching strategy for our opportunistic networking approach which is more fit for use in crowd management scenarios than traditional random strategies. Finally, we present a large-scale simulation based on a dataset consisting of movement traces from 28'000 people. Within this simulation, the influence of various parameters on the system is investigated and a random mode switching strategy is compared to our approach.

II. Related Work

The task of establishing communication networks in the absence of a fixed infrastructure (i.e. building ad-hoc networks) has been researched for decades and was initially motivated

978-1-4799-9923-1/15/$31.00 ©2015 IEEE

by military and disaster recovery motives. Especially the field of mobile ad-hoc networks (MANETs) has gained increased interest in the 1990s [5].

Traditional Ad-Hoc approaches work on the TCP/IP stack and require its modification on almost every layer. Most modern smartphones (e.g. iPhone and Android devices) on the other hand, do not grant access for developers to those parts of the operating system to make the necessary modifications. However, given that our app-based scenario requires a system that can be used on standard off-the-shelf hardware, it is no option to pursue an approach which requires a rooted device. To solve this problem, the WiFi Alliance introduced the WiFi-Direct standard which facilitates direct communication between WiFi-enabled devices. Unfortunately, the vast majority of smartphones currently available on the market does not support this standard [6], so it is not feasible for our work.

Trifunovic et al. introduced WiFi-Opp [8] which lays the basis of our work and is designed as a delay tolerant opportunistic network. WiFi-Opp relies on the fact that modern smartphones have a built-in WiFi-hotspot functionality. The underlying principle is that if a certain amount of devices take over the role as access points and the others are trying to connect to those, information can be exchanged between the connected devices. If those modes are switched periodically, the information is propagated to all participating devices. To assign modes to devices, WiFi-Opp employs a random strategy which is further refined in [9]. [8] also shows that WiFi-Opp is more energy efficient than the WiFi-Direct standard.

Fet et al. presented a model for WiFi signal attenuation of the human body [10] in 2013. Since our scenario usually involves large (and mostly dense) crowds of people, we consider their findings within our system performance parameters.

It is worth mentioning that none of the related work has been performed on the iOS platform. Our work is no exception as we focus on Android, the mobile operating system which has currently the largest market share. The reason for focusing on Android is to be found in the iOS API which restricts access to vital system level functionalities. Having said that, our approach will of course also work on iOS devices if Apple should decide to provide the necessary APIs one day.

III. Contribution

Previous works have presented smartphones as a simple "tool" to enable opportunistic networking without actually determining what conditions have to be met by the overall system in order for the smartphones to be utilized to their full capacity from a network performance point of view. We therefore provide the results of an empirical analysis of the capabilities and limitations of smartphones, that are currently available on the market, with respect to opportunistic networking. These findings serve as a basis for the development of a smartphone based opportunistic networking approach.

This approach employs a novel mode switching strategy based on node mobility. We show that using this strategy leads to information being spread amongst a crowd with a lower amount of devices acting as access points at a comparable or even greater speed compared to traditional random mode switching approaches. Since our switching strategy is heavily based on device mobility it could be assumed that it is not working in static scenarios. Therefore, we also present the results of a real world experiment under worst case conditions – i.e. in a scenario where all devices are immobile. It will be shown that information can be spread at decent speeds even under those circumstances.

Finally, we present the findings of a large scale simulation of an event where a temporary network outage affects a limited area of the event region. The simulation is used to fine-tune the parameters of our switching strategy. The virtual agents in the simulation follow a movement model that was extracted from a real dataset containing data from 28.000 people and can therefore be assumed to be realistic.

IV. Boundary Conditions Imposed by Smartphones

In order to design a communication protocol and a mode switching strategy optimal for our scenario, an understanding about the capabilities and limitations of smartphones from an opportunistic networking perspective is needed. Given that to our knowledge no such research has been carried out up to now, we performed a series of experiments with a host of smartphones. The set of devices was assembled based on the most popular smartphones manufactured between 2010 and 2014 thereby providing a good overview about the market. In the following subsections we present the outcome of our experiments.

A. Device Type

To learn if older smartphones perform worse with respect to opportunistic networking than the latest flagship models, we performed a series of tests where a fixed number of messages (each 100 characters long) was transmitted between an access point device and a client device. We chose the Google Nexus S (released 12/2010), the Samsung Galaxy S3 (released 05/2012) and the Google Nexus 5 (released 10/2013) as reference devices. Our findings clearly show that the oldest device performs broadly on the same level as the newest device, with the "middle aged" device surprisingly outperforming both. We therefore concluded that the device type should not be regarded as a factor for either the mode switching strategy or the communication protocol. For details, please see the top left plot in Fig.1.

B. Payload Size

In order to find out what impact they payload size has on the transmission speed, we performed a series of tests where messages of varying size were exchanged between two devices. The payload size ranged between 100 and 10'000 characters. Our findings show that while there is a linear dependency between message size and transmission time, it is very small and therefore does not have any meaningful impact on the communication protocol. The top right plot in Fig. 1 demonstrates that the average transmission time only grows by 32% when the message size is increased by a factor of five (280ms vs. 370ms).

C. Packaging Strategy

One key design decision for our communication protocol was whether to exchange all messages within one large packet

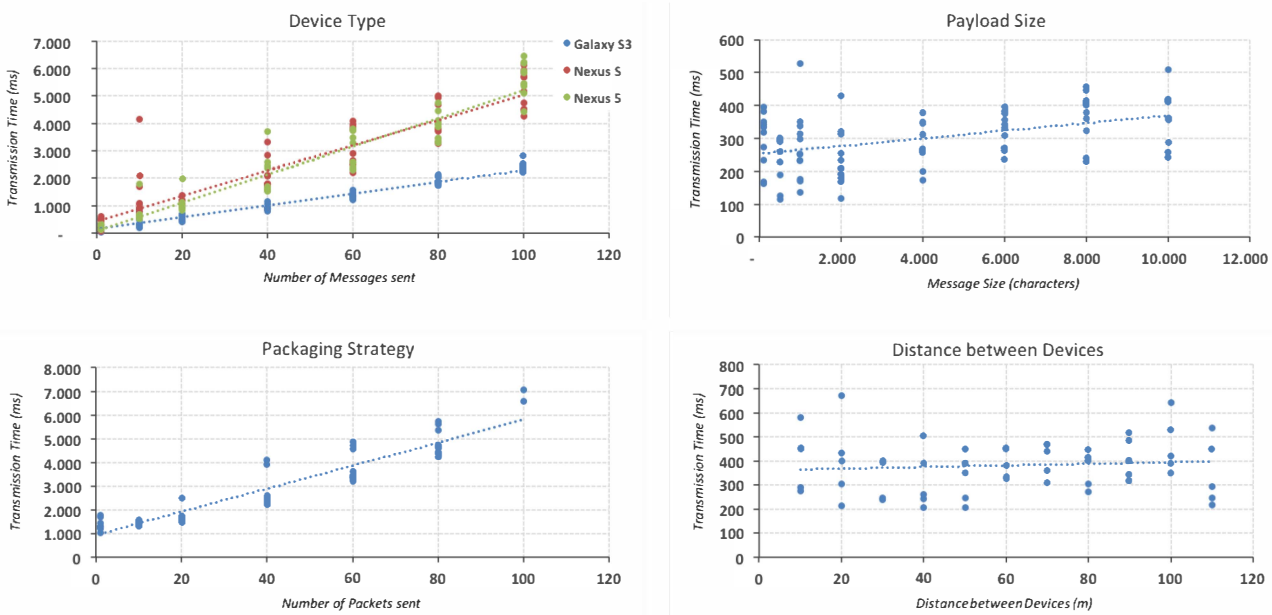


Fig. 1: Results of Smartphone Evaluation

or if each message should be transmitted as a single small packet. We therefore performed a series of experiments where we altered the packaging strategy for each iteration. The first iteration consisted of a single 100'000 character packet being sent while the last iteration consisted of 100 packages each 1'000 characters long. Since our findings showed that the former performed about 7 times better than the latter (see bottom left plot in Fig. 1), we made the decision to exchange messages between devices in bulks.

D. Distance between Devices

A further experiment was performed to find out if the distance between two devices had a large impact on transmission times – especially in scenarios with a lot of mobility this would be an important factor. We therefore performed outdoor tests where messages were exchanged between two devices at varying distances. The distance was increased from 10m to 100m in 10m increments. Our findings clearly show that under line-of-sight (LOS) conditions, the distance between the devices can be disregarded (see bottom right plot in Fig. 1).

However, the situation looks completely different when there is no LOS between the two devices. Especially, in indoor scenarios, a smartphone's ability to act as an access point over longer distances can be severely weakened due to dense multi-path effects in the environment and reflection, diffraction and scattering of the WiFi signals [15]. While our scenario is focused purely on outdoor situations, one still shouldn't make the mistake to assume that smartphones will be able to communicate via opportunistic networking over distances of 100+ m in the average case as it's very likely that there are no LOS conditions (please also see the next subsection for details on this topic).

E. The Effect of dense Crowds

WiFi signals are mainly transmitted in the 2.4 GHz range – this happens to be more or less exactly the same as the resonance frequency of water [15]. Incidentally, the human body is made up of up to 72% water [16]. Therefore, people must have an effect on the transmission of WiFi signals – a fact that is even more relevant in our scenario consisting of (potentially very dense) crowds. [17] establishes that attenuation due to the human body on the WiFi signal strength is stronger when a person is closer to an access point than when it is further away. [10] quantifies these findings: when a person is standing 1m away from an access point facing the device, the RSSI of the WiFi signal drops over 40 dBm when receiving the signal behind that person compared to receiving it in front of it. In a distance of 10m however, there is no noticeable difference anymore. In practice, this means that smartphones which are carried in a pocket by their users are "weak" access points and we cannot assume that they will cover the same distance that we observed during our LOS experiments where the devices were hand-held. Tests in crowded open air environments have shown that the smartphones carried inside of trouser pockets were still able to communicate over a distance of approximately 50m but would sporadically lose their connection when the test subjects moved away even further.

F. Exchange Strategy

The communication within a micro-network consisting of one device in access point mode and several devices in client mode can be described best as a synchronization between all participating devices which is in our case controlled by the access point device. The question is consequently, if synchronization should happen via flooding (i.e. all clients send all their messages to the access point which in turn sends the combined set back to all clients) or if it makes sense for the synchronization to take place in a filtered way. The latter would introduce an additional stage taking place before the actual synchronization where each client sends the access point a message containing the IDs of all the messages the client has stored. The server would return a list containing only those message IDs that are new to him so that the amount of transferred data can be optimized. Under ideal conditions (i.e.

no duplicate messages exist in the micro-network), flooding is about 15% faster than the filtering approach. Since such an ideal scenario is not realistic under real-world conditions, we decided to implement the filtering mechanism in our communication protocol in order to be able to cope with situations where there might be many large duplicate messages existing on the connected devices.

G. Access Point Performance

During our experiments we realized that smartphone vendors actually limit their devices when it comes to the number of clients that can connect to an access point device simultaneously. This limit does not seem to depend on hardware factors alone as we discovered that identical phones running different operating systems (stock Android vs. the Cyanogen Mod for example) imposed different limitations on the access point feature. After having analyzed 14 different phones in total, we came to the conclusion that it is safe to assume that each device can handle 5 clients when in access point mode – an important number when designing the mode switching strategy. However, it needs to be said that some devices we tested were able to handle as much as 13 connected clients, so in practice there will be micro-networks with more than just six participating devices.

H. Mode Switching Duration

Critical performance parameters are the time it takes an access point device to provide clients with a WiFi hotspot and the time it takes a client to connect to an access point. Since our scenario is very mobile, devices potentially don't have much time to connect to each other and synchronize before they are out of range again. Our experiments showed that it can take a device anything between 5 and 10s to create a usable WiFi hotspot. On the other hand, client devices took a maximum of 5s to connect to an access point. Based on the dataset recorded at the Zurich festival we found out that the average speed of a pedestrian at an event is rarely above 1 m/s under normal conditions. Assuming an average access point range of 50m and a total micro-network setup time of 15s, this leaves enough time for information synchronization to occur even when the two devices are moving in opposite directions.

V. Mobility-Based Mode Switching Strategy

Our scenario assumes that a section of the communication network collapses under heavy load as it can happen during large scale events. Therefore, some parts of the visitors are in a "blackout area" unable to connect to the internet and thus unable to receive messages via our app based system. The idea is now that the app realizes that there's no connectivity which leads to it enabling the AdHoc communication feature trying to exchange messages in a peer-to-peer manner. Messages sent by the event organizers or emergency forces are received by people outside the "blackout area" and are carried into the area as those people move around. Once inside the affected area, the apps of those people can inject the messages into the AdHoc network where they are propagated.

As mentioned earlier, our approach for delivering messages to those parts of a crowd currently affected by a network outage takes the works of [8] as a starting point. The underlying idea is that some smartphones enable their built-in WiFi hotspot mode, while others connect to those access points as clients. Information can then be exchanged within this micro network. If the modes of access point and client are frequently changed, it leads to an eventual propagation of the information to all participating devices – data is consequently being transmitted in a multi-hop fashion.

The crucial element defining the efficiency of an opportunistic networking implementation is the mode switching strategy. A good strategy should ensure that the devices communicate in a way that maximizes the amount of new messages being exchanged within each micro-network. Hence, the case where the same devices are part of the same micro-network in consecutive mode switching iterations should be avoided if possible. Taking these goals as a basis, we present a novel mobility based mode switching strategy which is based on the following two assumptions:

1) Devices with a high mobility carry potentially more new messages than those with a low mobility. The reasoning is intuitive: if a device traveled a long distance, it potentially carries messages from a source that's physically further away from it's current location. Since opportunistic networking is mainly about bridging physical distances as quickly as possible, device mobility should clearly be regarded as an important factor. We therefore introduce the quality factor *device mobility* as DM and the *device mobility threshold* as DM_{thresh}.
2) Ideally, each micro network consists of devices that haven't been connected for a long time. Again, the reasoning is rather intuitive: the likelihood of new and therefore relevant messages being exchanged between devices in an opportunistic network is higher when the devices haven't been in touch with one another for a long time, since they will have had the chance to collect new messages in the meantime. We therefore introduce the quality factor *last seen* as LS and the *last seen threshold* as LS_{thresh}.

Our mode switching strategy follows a strictly asynchronous approach: as soon as a device loses connectivity for a certain amount of time, it enters into opportunistic networking mode and initiates a so-called *decision phase* where it decides whether it becomes an access point or a client without complying with a global schedule like it is done in [13]. Within the decision phase, the following main steps are taken:

1) The device collects information about all available access points that are within reach. Access point devices broadcast their DM value as part of their SSID.
2) All access points with $DM \geq DM_{thresh}$ are inserted into a list L_{DM} ranked according to their DM values.
3) All access points with $LS \geq LS_{thresh}$ are inserted into a list L_{LS} ranked according to their LS values. Each device keeps a local record of when it was last connected to another device so that this ranking can be created individually for each device.
4) An overall ranking is computed based on each access point's ranking within L_{DM} and L_{LS}. Both rankings can be weighted (see section "Large Scale

Simulation" for details on the optimal weight factors DM_{weight} and LS_{weight}).

5) The device goes into client mode and connects to the access point with the highest overall ranking.
6) If no suitable access point could be found (because either none was available or none exceeded DM_{thresh} or LS_{thresh}), the device collects information about all available devices that are within reach.
7) If there are at least n devices within reach, the device goes into access point mode and waits for clients to connect. Please note that due to the limitations of some devices discovered during our experiments (see subsection "Access Point Performance" for details), the initial value of n should not be greater than 5.

The devices also need to be able to make a decision about their modes, when the parameters DM_{thresh}, LS_{thresh} and n are not met by their surrounding peers. In those cases the decision phase is repeated, this time with lower values for DM_{thresh}, LS_{thresh} and n, effectively relaxing the requirements regarding the quality factors a little bit. Furthermore, with each iteration of the decision phase, a random factor is added to the decision process so that after a large enough number of iterations, the decision basically defaults to a random mode switching strategy. In those extreme cases, 20% of the devices become access points while the rest goes into client mode. The reasoning for this 20/80 distribution is again found in the limitations of some smartphones with respect to their access point performance. We chose this upper bound instead of a distribution with a higher access point ratio for battery drainage reason (devices in access point mode cause a higher battery drain than those in client mode [14]). Section "Large Scale Simulation" will show that from a performance point of view, more access points ensure a faster spreading of information – however, for our fallback random strategy we decided to rank battery drain higher than overall performance.

Furthermore, an approach is needed for defining the actual communication within the micro-networks. While the main parameters of our communication protocol have been defined in the section "Boundary Conditions Imposed by Smartphones", an integral element of the actual mode switching strategy is to define when a device currently acting as a hotspot should quit its current mode and go back into the decision phase. Our approach makes hotspots responsible for breaking up connections – clients stay connected to an access point until that access point decides to go back into a decision phase. This happens when a certain amount of time has passed without any data being exchanged within the micro network – we made the decision against a pre-defined global duration for the existence of access points as it can always be the case that new clients connect to the access point and deliver potentially new data to the micro network. Only when no data exchange has taken place for a set duration U is the micro network dissolved by the access point going back into a decision phase which triggers the clients to do the same. For further elaboration on the value of U please see section "Large Scale Simulation". Clients are of course automatically disconnected from the micro-network when they move out of the access point's range.

Since our mode switching strategy is heavily based on device mobility we wanted to assess how it would fare under worst case conditions – i.e. in a scenario where all devices are immobile. We therefore laid out 14 devices over an area of approximately 500 m^2 and let them each create random messages in a way that resulted in a total number of 200 messages each 100 characters long. The parameter U was set to a length of 3 minutes for this test (a rather high value which would be well suited for a scenario with high device mobility). Table 1 summarizes our findings.

	avg	med	std
Reception Duration	309.43 s	219.00 s	301.09 s
Full Msg. Coverage	774.26 s	750.45 s	187.03 s
Hop Count	3.2	3.0	1.2
Max Hop Count	4.5	4.0	0.9

TABLE I: Results of live test without device mobility.

It can be seen that it took an average of roughly 3 minutes for each message to make it to another device and almost 13 minutes for all 200 messages to be available on all 14 devices. While those values might seem high at first glance, they can be easily explained by the fact that each decision phase needed to be iterated numerous times as all devices were stationary – i.e. there were no usable DM parameters and LS couldn't be used as intended, either. Furthermore the fact that U was set to 3 minutes – a value that makes sense in a moving crowd when there is a chance of new clients connecting to an access point – slowed the entire test down considerably, as it meant that 3 minutes needed to pass before a new micro-network could be established. In summary, we come to the conclusion that our approach even works under conditions that are actually destructive to the mode switching strategy's underlying concepts.

VI. Large-Scale Simulation

To test the effect of various parameters on our mode switching algorithm, and to study operational domains lying outside those of current experimental devices, we have developed and analyzed a large-scale agent-based simulation model (ABM) using the framework NetLogo. The simulation example is based on a real world dataset obtained during the 2013 Zurich festival and the map of Zurich where pedestrians are randomly located with a desired density and motion parameters. This way we define abstract pedestrians that are "as good as real" and allow us to study realistic motion patters as well as their effect on communication.

The simulation allows for broader parameter sweeps we have obtained using NetLogo's own BehaviorSpace module. We developed 2 varieties of the model and both have been parametrized using real data. One version uses city street maps (shape files) for pedestrian motion and the other a free area to reduce complexity and to obtain a general model. The shape file based version is domain independent and has been tested on various city layouts (e.g. Torino and Budapest) and is here studied at the example of Zurich. A rectangular area can be defined to define a signal outage, i.e. an area within which our AdHoc communication approach is simulated. Pedestrians can be put on streets randomly or by mouse commands to define "hot spots" which are common during large scale events. Using our simulation model, various scenarios can be tested.

The simulation unites the motion model with the communication model. For the former, see the next subsection; the latter directly implements the algorithm given in Section V.

A. Walking Data and Motion Model

Using the Zurich festival dataset we estimated the key walking parameters and used them in the simulation in a motion model. Specifically, we obtained data for average walking speed as well as average walking/standing times, together with their error terms. Our approach takes these values as averages over various crowd densities experienced in the dataset, referring to the fact that in the interesting range (3'000 - 4'000 pedestrians using our smartphone app) these parameters do not change dramatically.

Shape files represent streets by their segments, and our motion model consists of navigation along a network defined by these segments. Pedestrians share a common average walking speed but have individual walking and staying times taken from two normal distributions, characterized by the average walking/standing times and their standard errors, respectively. Pedestrians are heading towards an adjacent segment and upon reaching it select a new adjacent destination segment, never turning back in a single selection step.

Sample walking data is presented in Figure 2. Our motion model is based on a "mean field" approximation of the pedestrian population, hence we use the averages for walking speed 0,5 m/s, standing length 55s (std 100) and time between standing 40s (std 65), respectively. We take these as fixed external parameters in our simulations to test the internal parameters of the mode switching algorithm. With the motion model realized, we find that about half of all agents are immobile at any given time, about 20% never move at all and another 20% never stand still (those values seem intuitive for anyone having visited a street festival).

parameter	value	range
population N	1,000	fixed
device range r	50 m	fixed
access point mode uptime U	180 s	fixed
uptime std error std_U	35 s	fixed
max # clients per access point k	5	fixed
# clients required per access point n	1	1-4
device mobility weight DM_{weight}	1.0	0-1.0
device mobility threshold DM_{thresh}	0	0-0.6
Last Seen weight LS_{weight}	0	0-0.6
Last Seen threshold LS_{thresh}	0	0-60

TABLE II: Simulation parameters tested.

B. Simulation Runs and Parameter Sweeps

Here we present key results directly bearing upon open questions left about parameters in the empirical tests. In the simulation runs, we have used the parameters given in Table 2. The model is consistently scaled for all values as 1 patch = 7 meters and 1 tick = 1 sec. A useful visualization renders 1 patch as 4x4 pixels. The area tested was 1'500x1'500 meters, i.e. 2'250'000 square meters.

We are here presenting 2 cross-sections of analysis, first keeping the algorithmic parameters related to DM and LS fixed at $< 1.0, 0, 0, 0 >$ as in Table 2, and looking for the effects of n. In the second study we set n to the best value found and perform multi-dimensional sweeps for the 4 DM and LS parameters (prior tests justify this strategy and have indeed been suggesting optimization of n first.). The simultaneous analysis of all parameters of Table 2 would require advanced design of experiment tools such as adaptive sampling and is therefore beyond the focus of this paper.

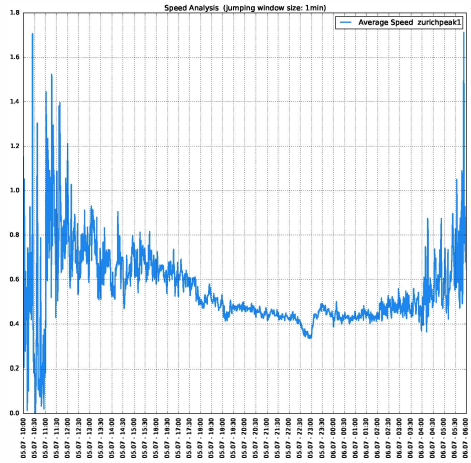
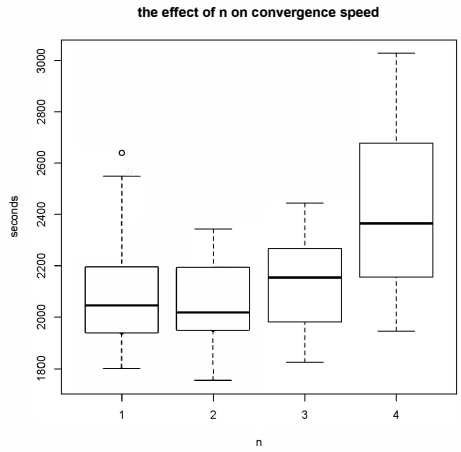


Fig. 2: Walking speed during one of the Zurich festival days.

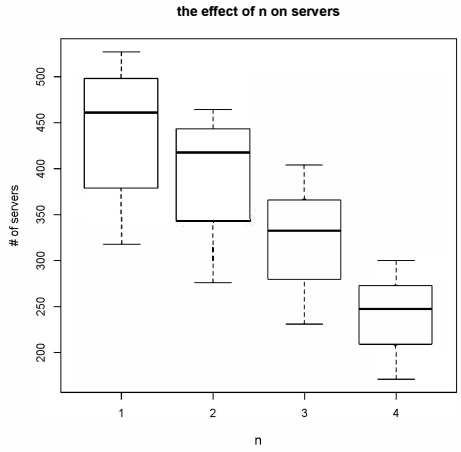
Among the output variables studied, of utmost importance are the ones related to efficiency, i.e. convergence time t (the time that **a single initial message** originating from a **a single device** reaches all devices) and the percentage of devices in access point mode. The first study shows a tradeoff between these two variables as a function of n using 10 random runs with identical parameters, as shown on Figure 3.

The number of access points used stands clearly in a strong negative linear relation with n (also tested but not shown using correlation package "car" in R). The effect of n on convergence speed is less profound however, and t has a flat minimum at $n = 1$ or 2. If access point efficiency is also taken into account this indicates $n = 2$ as best value. To appreciate these results, we observe that at this value (using ca. 40% of agents as access points) the message reaches all 1'000 individual users moving over more than 2 million square meters in a time of about 33 minutes; the same would be achievable by only 20-25% of servers in ca. 40 minutes. While these numbers may seem high at first sight, we refer back to the choice that we did not test U values smaller than 3 minutes on average. Clearly that value is a bottleneck in scenarios with low mobility. A dramatic increase of speed can be expected using smaller values (and a faster establishment of connections with the next generation of WiFi chips). Furthermore, please note that this simulation takes only one single device as a starting point for the propagation of the message whereas in reality, there would be more people moving into the affected area and thus more devices injecting the message into the AdHoc network. For reasons of reproducibility however, we decided for having only one device as initial information source for the scope of this simulation.

Using $n = 2$, our further simulation runs were intended to clarify the relations of DM_{weight}, DM_{thresh}, LS_{weight} and LS_{thresh}. In the intervals given in Table 2, the output variable t varied between 1'800 and 4'100 seconds, with all t values below 2'000 reached at $DM_{thresh} = 0$ at any value of the other 3 parameters. On the other end, high t values are characterized by high values of DM_{thresh}, but also by high values of the other parameters. Motivated by these findings,



(a) Convergence time t versus n.



(b) The number of access points versus n.

Fig. 3: Convergence speed and number of access points against n, minimum number of clients required per access points.

we look into the 3D sections of the resulting space (Figures 4 and 5). Convergence speed versus DM_{weight} and LS_{weight}, or versus LS_{weight} and LS_{thresh} do not give any interesting insight, reinforcing the recognition that LS is not offering the most relevant contribution. The 3D plots using the threshold variables (Fig. 4) and the plot of convergence speed against DM_{weight} and DM_{thresh} (Fig. 5) have a clearer message. Lower values of DM_{thresh} combined with higher values of DM_{weight} give best results largely independent from LS_{weight} and LS_{thresh} values.

Finally for this subsection, we ask about the impact of t on N, the number of agents. Results are shown in Figure 6. It can be seen that above 1'000 agents the time for a message to reach all agents remains constant, which justifies our choice of $N = 1'000$ for the main part of this study.

As a summary, we may conclude that LS is suppressed by the effects of high mobility values DM – using a realistic motion model this is expected indeed, as many agents are immobile or have low mobility, designating them as weak candidates for quick information propagation.

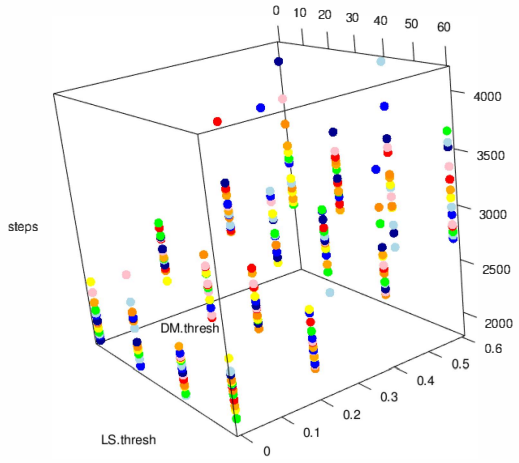


Fig. 4: DM_{thresh} and LS_{thresh} affecting convergence speed

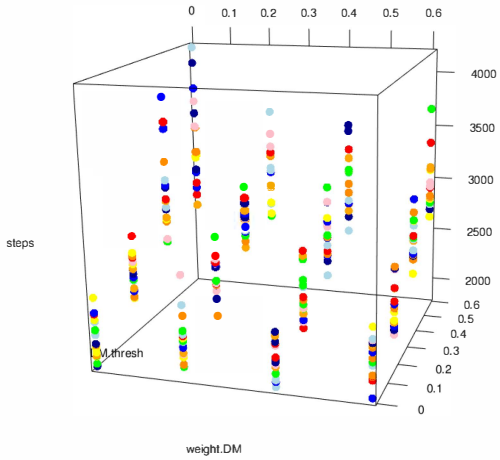


Fig. 5: DM_{thresh} and DM_{weight} affecting convergence speed

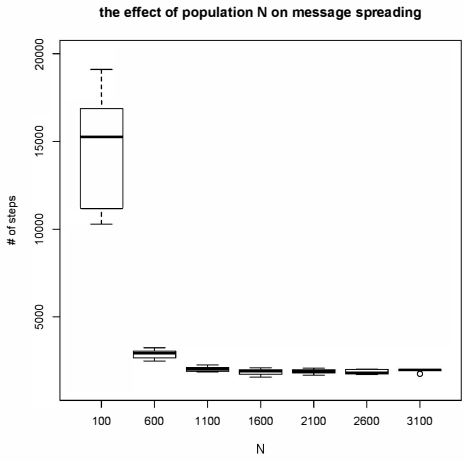


Fig. 6: The effect of population size N on the spreading of messages (i.e. convergence time t).

C. Comparison with Random Switching

Degrees of mobility have been tested and evaluated in the previous subsection with the conclusion that the amount of mobility is a decisive factor in the speed of signal propagation. Next, we also performed direct comparisons of our switching algorithm with a random switching model using the same optimal parameters. This also ensures fairness of comparison as it keeps the ratio of access points at around the same low level. It was found that our algorithm outperforms random switching by 10 percent or more (see Figure 7), with asymmetric error terms. Future refinements of the algorithm are likely to lead to further increases.

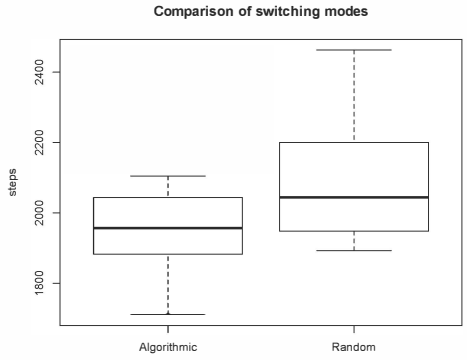


Fig. 7: Comparison of switching modes.

VII. Conclusion and Future Work

The core conclusion of this paper is that, under realistic circumstances, the use of smartphone hotspot functionality based Ad-Hoc networks for information distribution within large urban crowds is a viable approach. However, the mode switching strategy is critical and should carefully take into account what is known about the user and crowd dynamics.

In the future we want to do more research on defining suitable strategies for handling those members of the crowd which are "at the edge" of the region affected by a temporary network outage. Our current system assumes that those people will eventually walk into the affected region, thereby "carrying" messages that were received previously with them. If we found a suitable way for their devices to actively know that they could send data into the affected area without the users having to walk into it, we could greatly reduce the convergence time t.

VIII. Acknowledgements

This work is supported by the RESCUER project, funded by the European Commission (Grant: 614154), by the Brazilian National Council for Scientific and Technological Development CNPq/MCTI #(490084/2013-3), as well as the European Community's Seventh Framework Program (FP7/2007-2013) under grant agreement #600854 "Smart Society - hybrid and diversity-aware collective adaptive systems: where people meet machines to build smarter societies", and by the CoCoRec (Collaborative Context Recognition in Dynamic, Multimodal Smart Environments) project supported by the German Fedaral Ministry of Education and Research.

References

[1] M. Wirz, T. Franke, D. Roggen, E. Mitleton-Kelly, P. Lukowicz, and G. Troster, "Inferring crowd conditions from pedestrians' location traces for real-time crowd monitoring during city-scale mass gatherings," in *Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE), 2012 IEEE 21st International Workshop on*. IEEE, 2012, pp. 367–372.

[2] M. Wirz, T. Franke, E. Mitleton-Kelly, D. Roggen, P. Lukowicz, and G. Tröster, "Coenosense: A framework for real-time detection and visualization of collective behaviors in human crowds by tracking mobile devices," in *Proceedings of the European Conference on Complex Systems 2012*. Springer, 2013, pp. 353–361.

[3] T. Franke, P. Lukowicz, M. Wirz, and E. Mitleton-Kelly, "Participatory sensing and crowd management in public spaces," in *Proceeding of the 11th annual international conference on Mobile systems, applications, and services*. ACM, 2013, pp. 485–486.

[4] U. Blanke, G. Troster, T. Franke, and P. Lukowicz, "Capturing crowd dynamics at large scale events using participatory gps-localization," in *Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP), 2014 IEEE Ninth International Conference on*, April 2014, pp. 1–7.

[5] M. Abolhasan, T. Wysocki, and E. Dutkiewicz, "A review of routing protocols for mobile ad hoc networks," *Ad hoc networks*, vol. 2, no. 1, pp. 1–22, 2004.

[6] D. Camps-Mur, A. Garcia-Saavedra, and P. Serrano, "Device-to-device communications with wi-fi direct: overview and experimentation," *Wireless Communications, IEEE*, vol. 20, no. 3, 2013.

[7] F. Legendre, T. Hossmann, F. Sutton, and B. Plattner, "30 years of wireless ad hoc networking research: What about humanitarian and disaster relief solutions? what are we still missing?" in *Proceedings of the 1st International Conference on Wireless Technologies for Humanitarian Relief (ACWR), Amritapuri, Kollam, Kerala, India*, 2011, pp. 217–217.

[8] S. Trifunovic, B. Distl, D. Schatzmann, and F. Legendre, "Wifi-opp: ad-hoc-less opportunistic networking," in *Proceedings of the 6th ACM workshop on Challenged networks*. ACM, 2011, pp. 37–42.

[9] S. Trifunovic, M. Kurant, K. A. Hummel, and F. Legendre, "Wlan-opp: Ad-hoc-less opportunistic networking on smartphones," *Ad Hoc Networks*, 2014.

[10] N. Fet, M. Handte, and P. J. Marrón, "A model for wlan signal attenuation of the human body," in *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*. ACM, 2013, pp. 499–508.

[11] D. J. Dubois, Y. Bando, K. Watanabe, and H. Holtzman, "Lightweight self-organizing reconfiguration of opportunistic infrastructure-mode wifi networks," in *Self-Adaptive and Self-Organizing Systems (SASO), 2013 IEEE 7th International Conference on*. IEEE, 2013, pp. 247–256.

[12] E. Talipov, J. Yin, Y. Chon, and H. Cha, "A context-rich and extensible framework for spontaneous smartphone networking," *Computer Communications*, vol. 37, pp. 25–39, 2014.

[13] M. Raj, K. Kant, and S. K. Das, "E-darwin: Energy aware disaster recovery network using wifi tethering," in *Computer Communication and Networks (ICCCN), 2014 23rd International Conference on*. IEEE, 2014, pp. 1–8.

[14] S. Trifunovic, A. Picu, T. Hossmann, and K. A. Hummel, "Adaptive role switching for fair and efficient battery usage in device-to-device communication," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 18, no. 1, pp. 25–36, 2014.

[15] K. Kaemarungsi and P. Krishnamurthy, "Properties of indoor received signal strength for wlan location fingerprinting," in *Mobile and Ubiquitous Systems: Networking and Services, 2004. MOBIQUITOUS 2004. The First Annual International Conference on*. IEEE, 2004, pp. 14–23.

[16] H. C. Lukaski, "Methods for the assessment of human body composition: traditional and new." *The American journal of clinical nutrition*, vol. 46, no. 4, pp. 537–556, 1987.

[17] K. Kaemarungsi, "Distribution of wlan received signal strength indication for indoor location determination," in *Wireless Pervasive Computing, 2006 1st International Symposium on*. IEEE, 2006, pp. 6–pp.

[18] T. Kärkkäinen, M. Pitkänen, and J. Ott, "Enabling ad-hoc-style communication in public wlan hot-spots," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 17, no. 1, pp. 4–13, 2013.