This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/ACCESS.2017.2778347, IEEE Access

Using Firefighter Mobility Traces to Understand Ad-Hoc Networks in Wildfires

Sergio Cabrero, *Member, IEEE*, Xabiel G. Pañeda, *Senior Member, IEEE*, David Melendi, Roberto García, Thomas Plagemann

***Abstract*—Ad-hoc networks have long been studied as an ideal technology to provide communications in emergency operations when network infrastructures are not available. Nevertheless, this area has not yet delivered enough mature technologies or working prototypes. We suspect that, among other issues, this is due to a lack of understanding of this application domain that forces researchers to make too many assumptions. One of those concerns is the mobility of network nodes. This paper describes, analyzes and simulates the mobility traces of a fire department during 30 wildfires. The analysis shows interesting insights into the communication range and the type of network in these scenarios. For instance, multi-hop routes are unlikely, so the network behaves like a Delay-Tolerant Network. In addition, the simulation results present a clear image of the network performance under different circumstances that can be used to design applications. We found that the network capacity is low due to the sparse network connectivity. Moreover, the buffer size has a much bigger impact on data delivery than the delay-tolerant routing protocol selected. We think that these are valuable insights and that the traces constitute an important asset for the research community.**

I. Introduction

Communication networks are an important tool for emergency services. Ad-hoc networks are often proposed as an alternative when the access to a network infrastructure is difficult or impossible. Typical use cases include natural disasters where the conventional network infrastructure is destroyed, urban catastrophes where the network is collapsed due to the presence of large crowds, or remote areas where no infrastructure is deployed. Wildfires typically fall into the last category. For example, it is illustrative to compare the location of wildfires[1] with the 3G/4G network coverage map[2] in the region of Asturias (Spain). The overlap of these maps visualizes that most wildfires occur in areas where network coverage is low or inexistent. The main reason is that wildfires take place in areas with low population and that are difficult to reach. Obviously, the economical return of investment of setting up mobile networks in such areas is low. Nonetheless, due to the vital role of communications in emergencies and the frequency of wildfires, the social benefit is beyond doubt.

Despite the obvious interest and the promises of the technology, research and deployment of ad-hoc networks in emergencies is not trivial. There is a large body of previous work that has attempted to study different aspects of this topic from applications for emergency services, to applications to the general public, to theoretical models and frameworks. Nevertheless, the field is still not delivering mature technology or prototypes. We believe that the main reason for this is the amount of assumptions that researchers have to make, mainly due to the lack of context and data sources coming from the emergency and rescue realm. Specifically, there is a lack of mobility traces in this area. This is something that is crucial to first understand the potential of ad-hoc networks and then to evaluate the proposed systems under realistic constraints. This paper makes a contribution to cover this gap by presenting and analyzing real mobility traces captured during 30 wildfires.

The period from October 2011 to September 2012 presented an intense wildfire activity. Figure 1 summarizes all wildfires that burnt more than 100 ha in the region of Asturias (Spain). It visualizes their starting date, the burnt surface and the hours required for their extinction. This paper studies mobility traces of the regional fire department (Bomberos de Asturias / 112) during these wildfires. We analyze and simulate these traces assuming that network devices following this mobility form an ad-hoc network. The results show that in each of the wildfires the network resembles more a Delay-Tolerant Network (DTN), than a Mobile Ad-hoc Network (MANET). In addition, we evaluate the capacity of these networks to deliver different traffic patterns with different system configurations. Interestingly, the simulation results show that the delay-tolerant routing protocol selected has little impact on the results, but the buffer size has a relevant influence. For the shake of repeatability and to help other researchers in this area, we have published the GPS traces and part of the code used in this paper on GitHub[3].

The remainder of the paper is organized as follows. The next section presents background and related work. Section III describes the process followed to extract the mobility traces of wildfires from a bigger dataset—the Asturies-ER dataset [1]. Section IV analyzes the traces using network science metrics. Section V discusses the results of executing over 25,000 simulation runs with these traces. Finally, the main conclusions of this work are presented in Section VI.

Sergio Cabrero is with the Centrum Wiskunde en Informatica, Amsterdam, The Netherlands, e-mail: s.cabrero@cwi.nl

Xabiel G. Pañeda, David Melendi and Roberto García are with the Universidad de Oviedo, Xixón, Spain

Thomas Plagemann is with the Universitetet i Oslo, Oslo, Norway

[1]España En Llamas, http://espanaenllamas.es/visita-guiada/ (last visited November 22nd, 2017)

[2]OpenSignal, https://opensignal.com/networks/espa%C3%B1a/movistar-cobertura?z=8&minLat=42.15&maxLat=44.55&minLng=-8.72&maxLng=-3.00&s=&t=2-3-4&id=2144 (last visited November 22nd, 2017)

[3]See https://github.com/sergiocabrero/asturies-wildfires

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/ACCESS.2017.2778347, IEEE Access

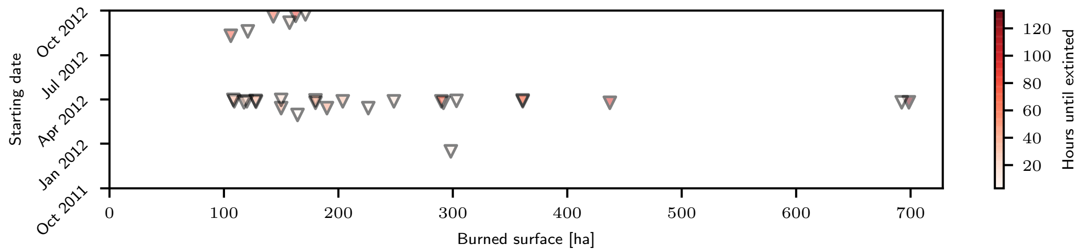

Fig. 1. Wildfires in Asturias (Spain) from October 2011 to October 2012

II. Background

The importance of communications during emergencies and in disaster relief scenarios is reflected in the amount of research and industrial effort in the area. Standards like TETRA are developed specifically for these situations with the aim of addressing their specific needs. In addition, there are ongoing discussions on how to make current LTE and future 5G networks suitable for disaster response [2], for example by increasing network resilience or with the support of device-to-device (D2D) communication [3]. Since relying on a network infrastructure during an emergency is too risky, many solutions aim to work without one. Thus, ad-hoc networks are proposed as a natural alternative.

Most protocols and applications that work on the Internet are not ready to work in an ad-hoc network. While on the Internet end-to-end connectivity between hosts is assumed, this is not the case in ad-hoc networks. In ad-hoc networks, devices often work as both hosts and routers, and end-to-end connectivity depends on their position and mobility. Therefore, network partitions, disruptions, and changes in the network topology are likely events. The different protocols in the stack need to cater for these issues, which affect aspects such as the way routes are calculated, the reliability of the network, or the delay experienced by applications. Location and mobility of network devices are the main factors that influence how frequent, how relevant or how difficult these problems are. They also influence the category or name given to ad-hoc networks: DTN, MANET, Vehicular Ad-Hoc Network (VANET), or Opportunistic Network. In principle, each of them has specific properties that differentiate it from the others, but sometimes this distinction is just by name. For instance, a VANET could be considered just a MANET where network devices are carried by cars, and a DTN could be just a very sparse MANET. This paper uses the terms MANET when talking about multi-hop routes, DTN when the store-carry-forward paradigm is needed, and ad-hoc network to refer to any of them. Next we explain multi-hop routes and the store-carry-forward paradigm.

Routing protocols are regarded as a key component in the performance of ad-hoc networks, so they have attracted a lot of attention. Researchers have focused on solving two problems. The first one is discovering devices in communication range and building multi-hop routes. This is a problem typically associated with MANETs. Two families of protocols are predominant: proactive and reactive, with Optimized Link State Routing (OLSR) [4] and Ad-hoc On-Demand Distance Vector Routing (AODV) [5], respectively, as the most frequently used. The second problem is communicating devices that cannot be connected using a single or multi-hop route. This is a problem associated to DTNs and requires the device sending a message—or other devices on its behalf—to store and forward the message until its destination is reached. This is known as the store-carry-forward paradigm [6]. DTN routing protocols define rules to know when to store, when to forward, and when to do both by creating copies of the data in the network. To do so, they take different approaches from random ones [7], to looking for patterns in mobility [8], to designing for specific scenarios including emergencies [9].

A plethora of applications have been proposed to aid emergency services and victims using these networking paradigms. To name a few: the application in [10] generates maps from disaster areas using a DTN; YouSoS [11] is another application that distributes multimedia content with the aid of victims; MOMENTUM [12] aims to distribute video merging MANET and DTN paradigms; and the work presented in [13] assists in the process of triage. These are just a few examples of the many that exist. Furthermore, some applications proposed for other ad-hoc network application domains could also be applied to emergency situations.

The problem is that most research in this area suffers from a lack of realistic knowledge of the mobility underneath the network. There are a few relevant mobility models, such as [14] or [15], and researchers have used them and general purpose mobility models to evaluate many different system aspects, especially routing strategies, see [16], [17], or [18]. Given the impact that mobility has on the performance of ad-hoc networks, this situation is far from optimal. Even if some traces from real mobility are available, they are uncommon and they are even more scarce in the emergency and rescue application domain. Not having enough mobility traces from emergencies produces two main issues. First, evaluation of proposals needs to be done over mobility that may be unrealistic. Second, there is not enough knowledge of the properties of the network, such as how dense the network is or how often devices connect to each other. Eventually, these issues hinder the process of creating new applications and the trust in the scientific validity of the results. This paper aims to push the state-of-the-art by making a new set of traces available in the specific context of wildfires, and by providing a thorough analysis of the ad-hoc networks in these situations.

III. Dataset & Method

The dataset in this paper—called the Asturies-Wildfires dataset—is the result of filtering a year of GPS traces from the fire department in the region of Asturias (Spain), with the data of wildfires occurred there from October 2011 to September 2012. Mobility traces from 30 different incidents were identified, each of them corresponding to a wildfire. Next we describe the process followed to prepare them for analysis.

The original mobility dataset (Asturies-ER [1]) was extracted from the Geographical Information System (GIS) database of a regional fire department in Spain: Bomberos de Asturias/112[4]. This system stores the location reported by GPS devices installed in most vehicles used by firefighters. The GPS devices report their location every 30 seconds when movement is detected using a 3G/4G network. If there is no coverage, the devices have the capacity to store a few of these positions and transmit them when possible. Nevertheless, the system is not perfect, so some of these positions may be lost in the process. After an intensive collaboration process with the fire department, a year of these traces— between October 2011 and September 2012—was made available to us for research purposes. In [19], we executed a first analysis of these traces that helped us to better understand the type of ad-hoc network that would be created following this mobility.

The motivation behind filtering the original traces is being able to analyze them with a better understanding of the context in which they happen. Different emergencies may have different properties, which affect mobility and eventually affect the type of ad-hoc network they produce. However, obtaining this context is not easy, since having access to reports and data from emergency services is generally a complicated task. The project *España en LLamas*[5] aims to raise awareness over wildfires in Spain. After questioning the right governmental entities, this project has gathered data such as the location—a GPS coordinate—, the day that the fire started, the time it took to extinguish it or the total surface burnt. These data points offered an interesting opportunity to provide some context to the original mobility dataset.

There were 30 wildfires reported between October 2011 and September 2012 that burnt more than 100 ha, see Figure 1. For each of them, we filtered the mobility traces using their location, their starting day, and the time it took to extinguish them. We defined two criteria to filter the traces with a combination of time and location. Since we did not have a specific starting time, we used the following procedure to establish the time period in which the wildfire occurred. We took the wildfire duration, converted it to days, and rounded it up. For example, a wildfire that took 3 hours to extinguish is rounded to a day, and a wildfire extinguished in 35 hours is rounded to two days. Then, we combined this duration with the starting day. To filter location, we used the GPS coordinate stated by the report from *España en LLamas* and defined a 10 Km circular area around it. All devices that were found in that circular area during the period of time determined for the wildfire were considered as taking part in the wildfire extinction. The maximum surface burnt in our data is: 2012 ha, which is an area equivalent to a circle of 2530 meters radius. Thus, we consider that a 10 Km radius is more than enough to include all the units aiding in the wildfire extinction. The final mobility traces for a wildfire include all the positions reported by these devices, both in and out of the circular area, during the period of time estimated for the wildfire.

[4]Bomberos de Asturias/112, http://www.112asturias.es

[5]España En Llamas, http://espanaenllamas.es/

The result of the previous step is a set of files containing a list of samples, each of them composed by a timestamp, an identifier of the device, and a GPS coordinate with latitude and longitude. To reduce the complexity of their analysis and simulation, we transformed this format into a more regular one. We established 30 second intervals—30 seconds was the default reporting time by the GPS devices—between the beginning and the end of each wildfire, and for each of them we assigned a position to every device included in the scenario. When interpolation was needed, we used the previous known position of the device. If there was not a previous position, we used the first position known. Figure 2 summarizes the properties of the resultant traces. As we see there is a heterogeneous number of nodes and samples. Moreover, some scenarios needed more interpolated traces, while others were more complete. Note that two of the wildfires only include one GPS device, hence they will not generate any ad-hoc network in our following analyses.

We study each wildfire using two different methods: an analysis based on well-known network science metrics [20], and simulation. In both of them, we assume that each of the GPS devices in the traces can be associated with a network device—we also often use the term (network) node. For the network science analysis, we need to define a criterium to create links between network nodes. Thus, two nodes are connected if their distance is shorter than a given communication range. This relationship is referred to as link or contact, and allows us to build a network graph (a topology) every 30 seconds. The following section discusses different metrics over these graphs. For the simulations, we introduce the mobility traces in The ONE [21] simulator. Different network configurations and traffic generation patterns are evaluated over each wildfire. We focus our discussion in metrics associated with message delivery.

IV. Network Analysis

A. Communication range

The communication range of devices in an ad-hoc network determines which network links can be formed and the properties of these links. It is determined in part by the network technology used, in part by how it is implemented with specific antennas and transmission power levels, and in part by the environmental conditions. Longer communication ranges create more opportunities to connect with devices further away, but they also increase chances of collision, power consumption, and propagation times. Understanding the effect of the communication range in the network is interesting to plan a real deployment and to define an appropriate network technology.

We look into the effect of a specific communication range on the wildfires. We measured the distance between nodes

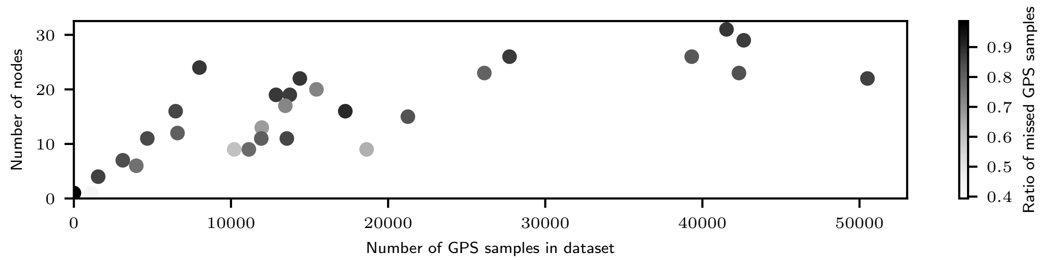

Fig. 2. Summary of Asturies-Wildfires dataset statistics

in the mobility traces and determined the contacts that a specific communication range produces in each scenario. In our analysis, two nodes are in contact—in other words, a network link between them can be established if the distance between them is shorter than the communication range. The GPS mobility traces are interpolated with 30 second intervals, so the contacts are also calculated using that interval. Then, we counted all the contacts that a given range produced in each scenario. Since the absolute value of this quantity is not important, because each wildfire has a different number of devices and length, we divided the total amount of contacts by the theoretical maximum, i.e., the number of devices in the wildfires squared multiplied by the number of 30 second intervals. This gives the ratio of contacts, see Equation 1.

$$Ratio\ of\ contacts = \frac{\#contacts}{\#devices^2 * \frac{wildfire\ duration}{30}} \quad (1)$$

We calculated the ratio of contacts for every scenario between 0 meters and 1000 meters taking steps of 20 meters. The line in Figure 3 corresponds to the mean value over the wildfires and the area around it represents the variance. Every contact that can be established with a given communication range can be established with a longer one, so the mean is always increasing. However, two different trends are observed. The growth of the ratio of contacts is bigger for shorter ranges, and becomes smaller between 150 and 200 meters. In addition, the variance increases with larger communication ranges. This is an interesting trend to be exploited, since at some point the benefit of making the network more connected by increasing the range of devices will not compensate other aspects, such as a higher power consumption or a lower bitrate. Interestingly, a communication range of this magnitude is well aligned with current 802.11 standards, such as 802.11p for vehicular networks.

In the following sections, we use 50, 200, and 1000 meters as representative values of the communication range. These ranges are representative, because they create different networks—ratio of contacts—, they require different technologies, and they will deliver different network performances.

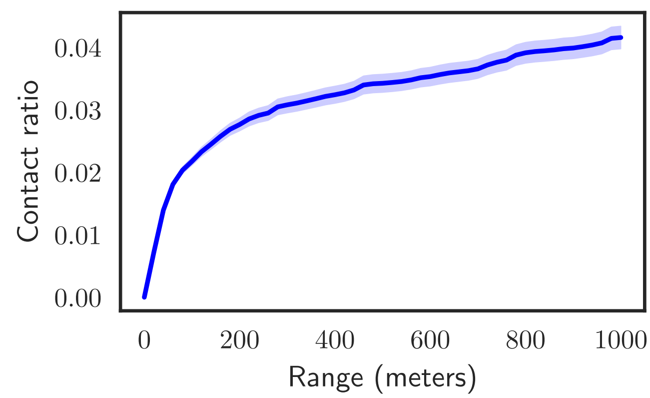

Fig. 3. Ratio of possible contacts compared with the communication range

B. Network density & node popularity

In an ad-hoc network, the links that devices establish according to their position determines the network topology. If a device can be reached in a single transmission—i.e. in a single hop—it is called a neighbor. The number of neighbors is also called degree in the fields of social network analysis or network science. This metric is helpful to understand two properties of the network: its density and the popularity of nodes. Network density is the ratio between the number of links in the network and the maximum number of links possible, which is the number of nodes squared. Thus, if most nodes in the network are neighbors to each other—i.e. node degrees are relatively high in comparison with the total number of nodes—, the network is considered dense; otherwise, it is sparse. In dense networks, sending data to any other node should be easier than in sparse networks, since there are many connections. At the same time, nodes in dense networks are more likely to suffer congestion, because many nodes compete to access the same wireless medium, so the number of collisions increases. Due to the implications of network density in network performance, understanding it is interesting to determine the protocols that better suit real deployments.

We inspected all network topologies formed by the mobility traces and calculated node degree for every node. Figure 4 shows a histogram that tells the relative frequency, in a logarithmic scale over the y-axis, of having a given number of neighbors, on the x-axis. We used ranges 50, 200 and 1000 meters. With a 1000 meters communication range, nodes have more than a 38% chances of having just one neighbor, and in less than 20% of the occasions node degrees are over 3. The shorter the range, the less likely it is to see

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/ACCESS.2017.2778347, IEEE Access

high node degrees, 6 being the maximum value found for 50 meters. This indicates that we are looking at networks that are mostly sparse, although exceptions are seen, e.g., for long communication ranges where large clusters of up to 19 nodes are formed.

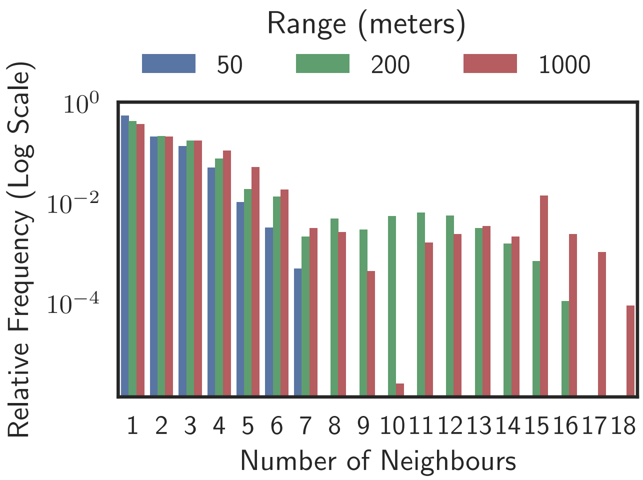

Fig. 4. Number of neighbors

In many scenarios, network degree is not a uniform property. On the contrary, some nodes have more neighbors than others. The nodes with more neighbors are more popular. Following our previous argumentation, nodes that are popular are likely to suffer more congestion. They are also more likely to forward packets on behalf of others either in real-time over a multi-hop route or storing them when using the store-carry-forward paradigm. Being popular can hinder their performance, and drain their battery—if battery is a limitation. Nonetheless, popular nodes are also the best place to disseminate information in the network. Identifying them is useful, for instance to design routing algorithms.

We look for popular nodes in each of the wildfires. For that purpose, we calculate the mean number of neighbors of each node during the wildfire. Figure 5 represents violin plots with the distribution of these means for three different communication ranges. In line with our previous analysis, most nodes have a mean of one or close to one neighbors. However, there are some nodes with a much higher mean, especially when the range increases. Furthermore, the plots resemble a long-tail distribution, which present interesting properties in different types of networks [20]. It is beyond the scope of this paper to further investigate the effect of these popular nodes on the network, but we believe a deeper study of this property could reveal interesting insights to design routing protocols or content distribution policies.

C. Multi-hop routes

Some ad-hoc networks can form multi-hop routes in which nodes are used to forward traffic in real time. This is the task of MANET routing protocols such as OLSR [4] or AODV [5]. Using these protocols introduces an overhead in the network, as building multi-hop routes requires more signaling than just

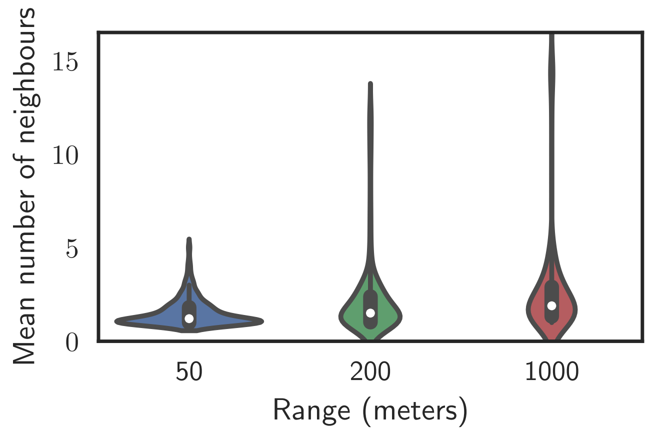

Fig. 5. Mean popularity of nodes in each wildfire

detecting nodes in range. In this section, we explore how likely multi-hop routes are in our scenarios to see if their use is justified and what the benefit would be.

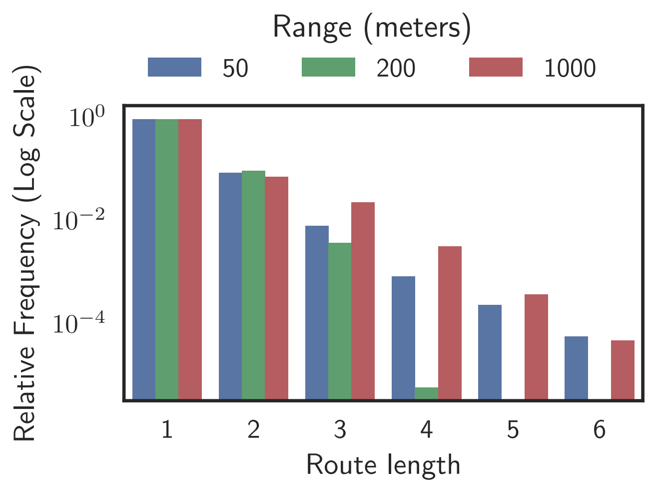

Fig. 6. Frequency of occurrence of a specific route length

We looked into every network topology formed by the mobility traces calculating the shortest path between every two nodes. Then, we counted the number of occurrences of every shortest path length. Figure 6 illustrates the relative frequency of finding a route of a given length in the wildfires. The x-axis shows the lengths in hops. No route with more than 6 hops was found. The y-axis shows the frequency in a logarithmic scale. For any communication range, less than 10% of the routes are multi-hop, and less than 3% of them have 3 hops or more. If a 200 meter range is used, there are no routes with more than 4 hops, but this not the case for both 50 meters and 1000 meters. Overall, these results imply that there is little incentive in using a routing protocol to build multi-hop routes.

D. Connectivity over time

Another interesting property of networks is their connectivity. In an ideal situation, all nodes would be connected

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/ACCESS.2017.2778347, IEEE Access

to each other either by direct links or by multi-hop routes. However, ad-hoc networks are normally partitioned in different groups of nodes that can connect to each other, but not to other groups. These are called network partitions. Using the store-carry-forward mechanism, it is possible to send data from one partition to another, if a device is moving between them. Hence, it is possible to connect nodes that are not connected in real-time. In this context, an interesting question is how partitioned the networks in our traces are. However, analyzing all the possible dynamics is a complex and computationally intensive task. Thus, in this section we will just scratch the surface and let the results from our extensive simulations in Section V to give more insights.

We built aggregated network topologies to understand how nodes connect over periods of time longer than the 30 second interval used originally. In an aggregated topology, two nodes have a link if they have a link at any moment during that period. First, we look at the aggregated topology for the whole scenario in each of the wildfires, again using 50, 200 and 1000 meter communication ranges. Many of them converge to a single network partition, but there are instances in which this does not happen. The mean values are: 4.1 partitions for 50 meters, 2.4 partitions for 200 meters, and 1.5 partitions for 1000 meters. From these results, we can ensure that some nodes in some wildfires would have no chance of exchanging data. So, if communication among these nodes is needed, alternative mechanisms must be found for real deployments.

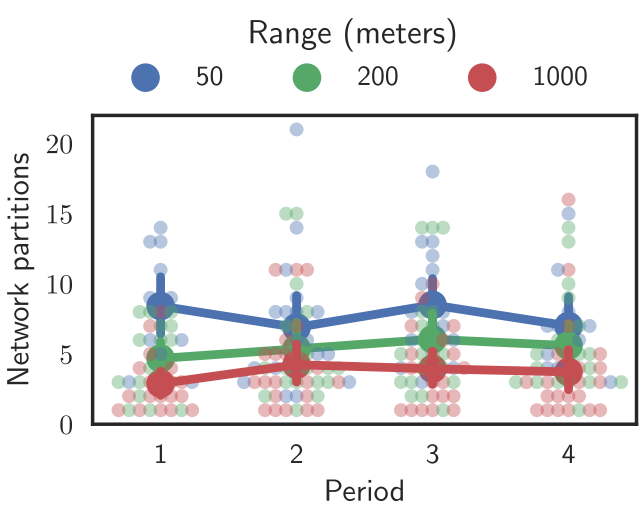

Fig. 7. Aggregated network partitions during each quarter of wildfire

We also want to understand if the number of partitions changes over time and if looking at a specific period is similar to looking at the whole scenario. For that purpose, we divided each scenario into four periods of equal duration and analyzed the aggregated topologies in each of these periods. Figure 7 represents the period of time on the x-axis and the number of partitions found in that period on the y-axis. The solid dots connected with lines represent the mean value for all scenarios and each range, the attenuated dots represent all the values. From this figure, we conclude that the number of partitions in each quarter is very similar, and it is not far from the number of partitions obtained for the aggregated topologies for each wildfire.

E. Coverage

Until now, we have assumed that the network is only composed by the nodes in the traces, but they could also be part of a larger ad-hoc network. For example, they could gather data from sensors in the area or allow communication between devices carried by firefighters. Since these devices are carried by vehicles, they are less constrained in terms of power, size, storage, or computation. Therefore, they could form a mobile infrastructure for the rest of mobile devices in the area. If used for this purpose, it is interesting to study the area they cover.

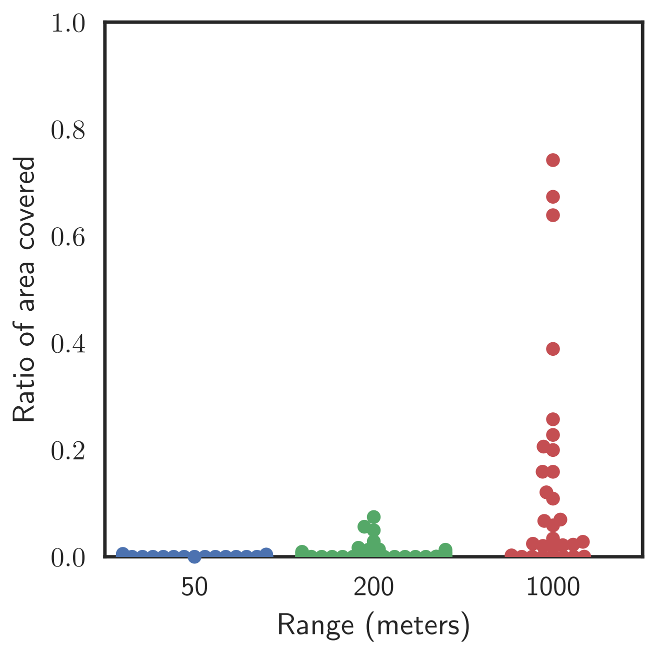

Fig. 8. Coverage of the 1 Km circle around the incident location

We took a circular area of 1 Km radius around the location reported for each wildfire. Then, we calculated a mean value of the ratio of area that would be covered by the devices using three communication ranges 50 meters, 200 meters, and 1000 meters. The results in Figure 8 show that there are a few cases, and only with the longest range, in which the ratio of area covered is significant. One reason for this may be that wildfires happen in places where access by road is difficult, and therefore, firefighters are deployed closer to the wildfire while vehicles are parked and moving at a distance greater than 1 Km.

V. Network Simulation

A. Experimental setup

The simulation results presented here are produced by the ONE Simulator [21], which simulates a DTN of network devices that follow the movements of the GPS devices during wildfires. As stated in the previous section, the amount of

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/ACCESS.2017.2778347, IEEE Access

multi-hop routes in the scenarios is marginal. For that reason, simulating them as DTNs is likely to produce similar results to simulating them as MANETs, i.e., using a multi-hop routing protocol. ONE implements the store-carry-forward paradigm and allows the configuration of a large set of parameters. Our focus is on understanding the capacity of these networks to exchange data under different conditions. Thus, we study how the results change when the configuration of the network changes, what we call system factors, and when the traffic generated by the nodes changes, what we call workload factors.

We have 96 different experiment configurations. Each of them was executed ten times with each scenario, resulting in more than 25,000 simulation runs in total. Each configuration is the result of combining one out of four system settings with one out of twenty-four traffic generation settings. The network settings are the result of combining two different routing protocols: Epidemic [7] and Prophet [8], with two different buffer sizes: 7500 KB and 100 MB. The routing protocols were chosen as relevant representatives of two different approaches of routing. We study the effects of a small buffer size, which is the same used in [18], and a much larger one. In all simulations, nodes are configured to simulate a WiFi interface that varies rate with distance. The traffic generation patterns are a product of combining four different message generation intervals: 5, 10, 30, and 60 seconds with six different message sizes: 1, 10, 128, 256, and 512 KB, and 1 MB. Again, many of these settings are taken from [18]. We only consider homogeneous traffic generation, this is that all the nodes follow the same pattern, e.g., a 10 KB message every 30 seconds, and all nodes are equally probable destinations. Nonetheless, we believe that heterogeneous traffic patterns may be very interesting for future work.

Next, we discuss the effects of mobility, and the effects of system and workload configuration.

B. Effects of mobility

In this section, we examine the results of the simulations without considering how they were configured. This helps us understand what the overall performance of the network in all scenarios is and how different scenarios perform. Three metrics are used: delivery probability, message latency, and network bitrate. Delivery probability is the number of messages delivered during the simulation divided by the number of messages created, which depends on the traffic generation settings explained previously. Message latency measures the time elapsed between a message creation and its delivery. Finally, network bitrate is the number of bytes delivered—which is the number of messages delivered multiplied by their size— divided by the scenario duration. These metrics provide interesting insights into how network performance is globally and how different from each other the scenarios are.

Figure 9 represents delivery probability on the x-axis and mean message latency in each simulation run on the y-axis. A histogram for each variable is represented on the axes, the main plot is a heat map where the darkness of each bin is proportional to the number of runs in it. This visualization[6] summarizes the results for these two variables for all simulation runs. Most simulations fall into the low latency and low delivery probability category. This means that in most situations the networks do not provide high chances of connectivity, but when they do messages are delivered relatively fast. Note that fast can still be very slow for what is common on the Internet—the range of latency is so wide that the first bin contains messages delivered in less than 5 hours. Despite the low delivery probability in most cases, there are some exceptions, i.e., in some scenarios around 80% of the messages are delivered.

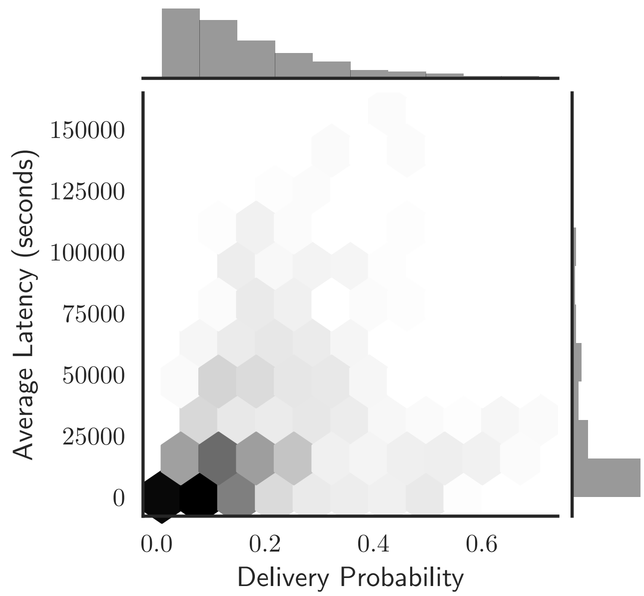

Fig. 9. Delivery Probability and Latency for all simulations

Figure 10 represents the network bitrate obtained in each simulation run. Each row corresponds to a wildfire, the x-axis is the bitrate in megabits per second. There are 960 dots for each wildfire—combining configurations and runs—so many of them are plotted on top of each other. Three insights can be extracted from this figure. First, bitrates are below 1 Mbps for the whole network. This is not surprising, as we already know that connectivity is challenging. Second, there is heterogeneity in the performance among wildfires. Some of them achieve higher bitrates than others, although all of them cover their full range. There is not a single scenario that does not result in a low bitrate for some configurations. Third, heterogeneity is also present within wildfires. Thus, mobility is not the unique influential factor in the performance of the network. Hence, we explore the impact of system and workload factors next.

C. Effects of system and workload

Figure 11 summarizes the impact of using a different routing protocol or a different buffer size on the delivery probability. It contains two violin plots where the distribution of each buffer/routing protocol combination is represented[7]. The top

[6]See http://seaborn.pydata.org/generated/seaborn.jointplot.html for details.
[7]See http://seaborn.pydata.org/generated/seaborn.violinplot.html for details.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/ACCESS.2017.2778347, IEEE Access

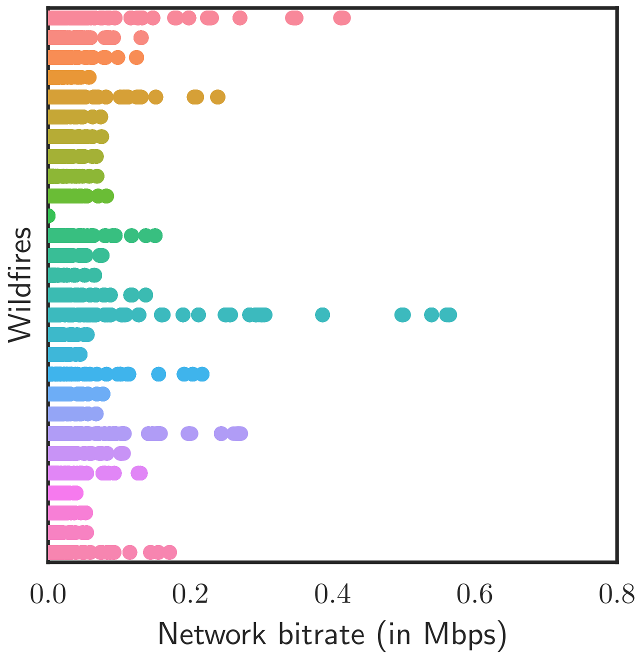

Fig. 10. Network bitrates obtained on each wildfire

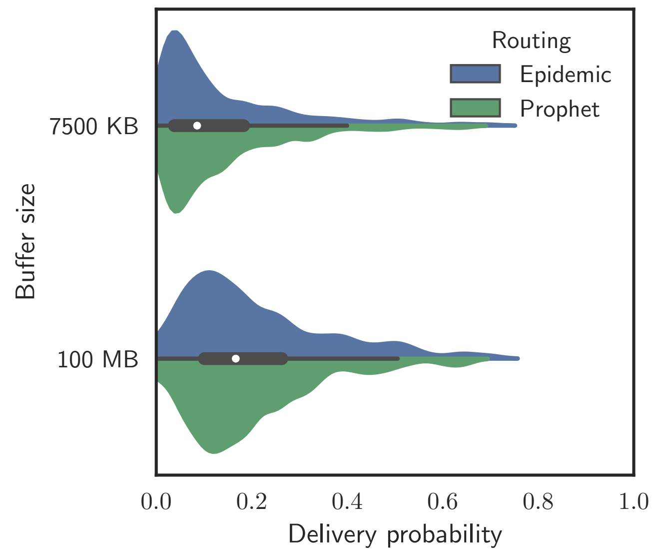

Fig. 11. Effect of system factors on the delivery probability

plot illustrates the smaller buffer size (7500 KB). This violin plot also has a top half with the distribution of the simulation runs using Epidemic routing and a bottom half for Prophet. The bottom violin plot illustrates the same for the larger buffer size (100 MB). This figure shows a similar performance for both routing protocols, but an interesting difference between buffer sizes. By comparing the mean delivery probability for each condition, this difference becomes more apparent. The mean delivery probability for all Prophet runs is 0.164, while for Epidemic is 0.167. Meanwhile, the mean delivery probability for a 7500 KB buffer is 0.131, and for a 100 MB is 0.20. Therefore, the buffer size plays a much more important role in our experiments than the routing protocol. This is a relevant insight, because it may be worth revisiting the effort spent on evaluating routing protocols for emergency scenarios, while far less effort is spent on analyzing other system factors.

Figure 12 illustrates the distribution of delivery probability for all combinations of message interval and message size. Each of them is represented with a violin plot. In general, the more traffic generated, the lower the mean delivery probability. Messages of 1 and 10 KB are delivered with much higher frequency than larger messages. A similar trend is observed with message generation intervals: longer intervals produce higher delivery. These effects, which are expected in a conventional network, would not be present if mobility was the only factor determining network performance. This trend is similar to that found by [18] in their simulations. However, the actual numbers are very different, which supports our claim on the importance of using mobility traces to understand the constraints of real deployments.

Finally, Table I breaks down the delivery probability for all system and workload factors. So, we can analyze how the different factors interact. Buffer size is a limitation, especially when messages are large and their generation interval is long. However, when messages are generated every 5 seconds, the network throughput seems to play a bigger role, so delivery probability is low for both small and large buffers. These performance details are important insights when designing applications, because the traffic pattern generated may require different decisions for the network configuration.

VI. Conclusion

This paper revisited the idea of using ad-hoc networks to support emergencies. Unlike most previous research, real mobility traces of wildfires were used. This allows us to better understand this application scenario. We also want to encourage other researchers to use these traces for this research—we published them on GitHub[8]—and to collaborate with emergency services in making more traces available. We believe that this is the path to produce solid progress in this area and overcome the lack of real deployments.

Our analysis shows that the network during a wildfire is closer to a DTN than to a MANET, if we understand MANET as a network that uses multi-hop routes. The number of multi-hop routes found is relatively low, so the benefits of using a MANET routing protocol are questionable. In addition, the network is partitioned, so the store-carry-forward paradigm is essential. Furthermore, we have demonstrated that the size of the buffer in nodes plays a crucial role in delivery, while the role of the delay-tolerant routing protocol in our simulations is minimal. This opens an interesting discussion about their role in emergencies. In the metrics analyzed, none of the protocols used showed better performance delivering messages than any other. It may be that another protocol is able to improve delivery, but it is also possible that most delivery is just carried out using store and forward, and that the role of the routing protocol in the network can be neglected. More studies with real mobility could provide a clearer view on this.

[8]See https://github.com/sergiocabrero/asturies-wildfires

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/ACCESS.2017.2778347, IEEE Access

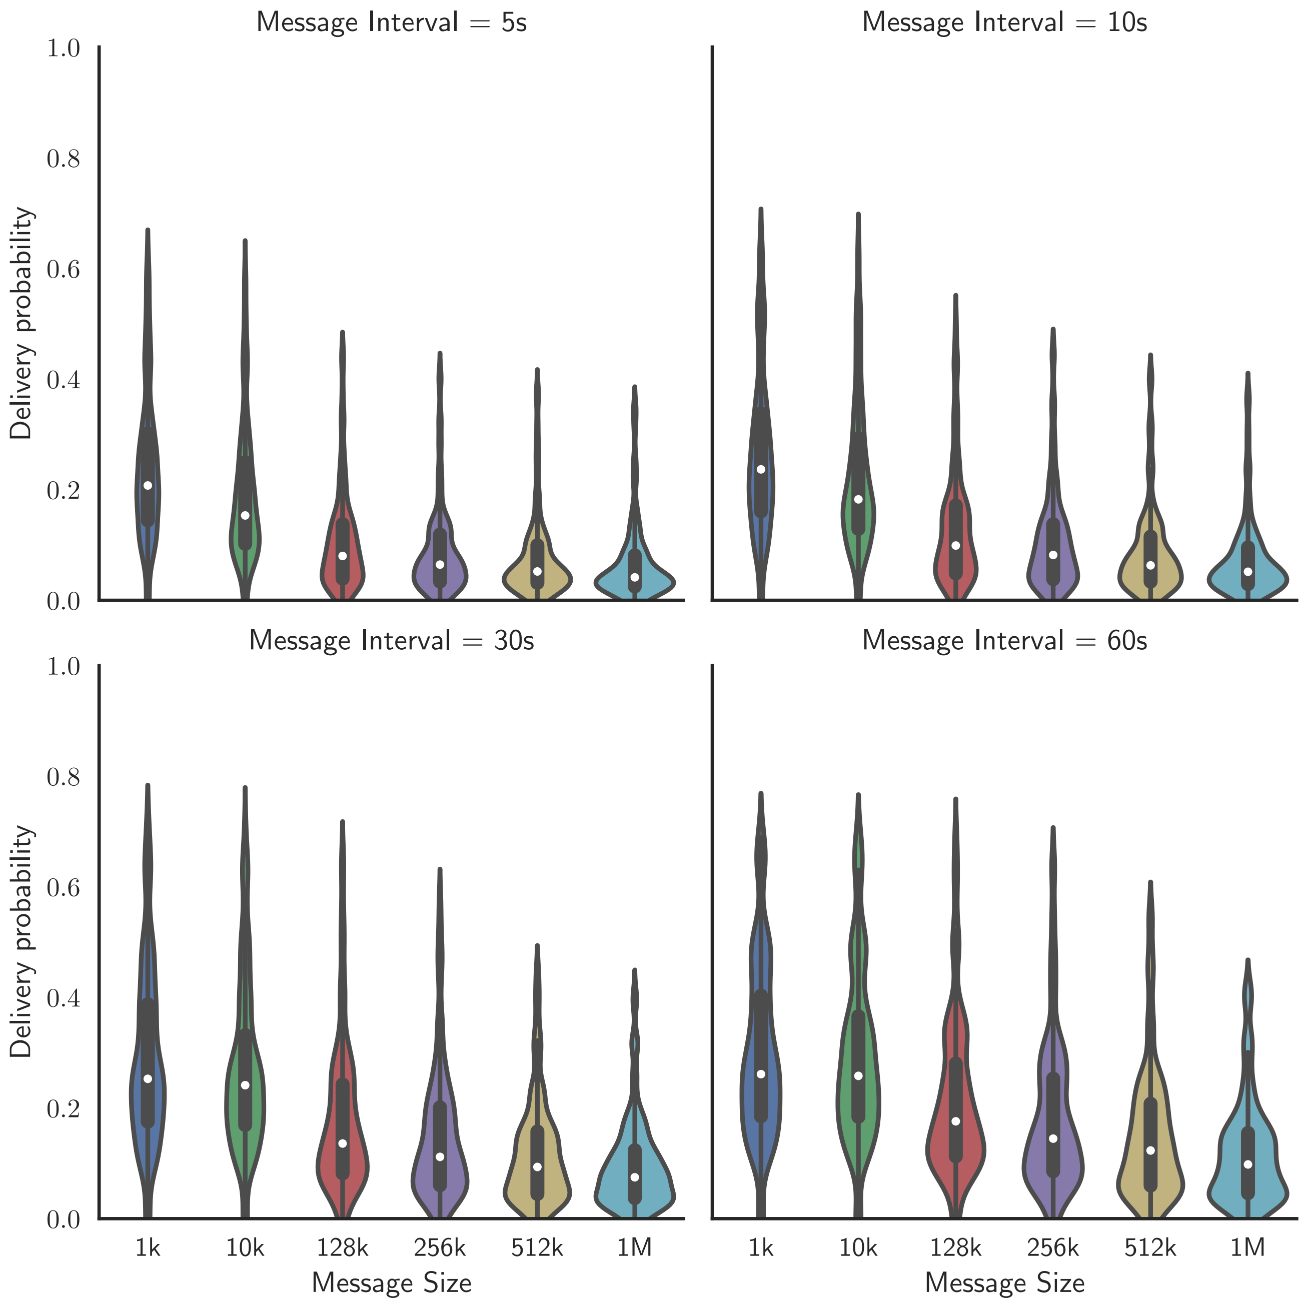

Fig. 12. Effect of workload in the delivery probability

The low delivery rates shown by the network is the major concern extracted from our results. On the one hand, it is questionable that a network with such a low capacity is useful. On the other hand, if an ad-hoc network is the only communication alternative, using it is worth a try. Moreover, knowing its limitations is the best way to improve it and to optimize its performance. Understanding mobility constraints is key to design applications that are ready to get the most out of network capacity, our previous work in disruption-tolerant video adaptation is an example of this [22]. In addition, connectivity during wildfires could also be improved in different ways. Drones are a promising alternative to assist emergency services, they could also be used to extend network connectivity. It is unlikely that drones will offer a complete infrastructure, but instead they can be useful connecting isolated partitions.

Although the technology has existed for several years now, the use of ad-hoc networking in different application domains, not only emergencies, is not yet massive. In this discussion, we align with part of the community, see [23], stating that the previous efforts in this area will now start paying dividends with emerging technologies. We have already mentioned drones, but also the proliferation of sensors and the Internet Of Things paradigm are likely to need from ad-hoc networking. Public

TABLE I
MEAN MESSAGE DELIVERY PROBABILITY BY SYSTEM AND WORKLOAD FACTORS

		Buffer			
		100 MB		7500 KB	
		Routing			
Interval	*Size*	Epidemic	Prophet	Epidemic	Prophet
5s	1k	0.26	0.23	0.24	0.23
	10k	0.25	0.23	0.13	0.15
	128k	0.13	0.15	0.06	0.07
	256k	0.10	0.12	0.06	0.07
	512k	0.08	0.09	0.05	0.06
	1M	0.07	0.08	0.05	0.06
10s	1k	0.28	0.25	0.28	0.24
	10k	0.28	0.24	0.18	0.18
	128k	0.18	0.19	0.07	0.09
	256k	0.14	0.15	0.06	0.07
	512k	0.11	0.12	0.05	0.07
	1M	0.08	0.09	0.05	0.06
30s	1k	0.31	0.27	0.31	0.27
	10k	0.31	0.27	0.25	0.24
	128k	0.25	0.24	0.10	0.12
	256k	0.21	0.21	0.08	0.09
	512k	0.16	0.17	0.06	0.08
	1M	0.12	0.14	0.05	0.07
60s	1k	0.33	0.28	0.33	0.28
	10k	0.33	0.28	0.29	0.27
	128k	0.30	0.27	0.13	0.15
	256k	0.26	0.25	0.09	0.12
	512k	0.21	0.22	0.07	0.09
	1M	0.17	0.18	0.05	0.08

Safety Networks will have to support heterogeneous devices and traffic requirements in situations where infrastructures are not always available. Future work in this area will have to design adaptable protocols and applications that are able to stand different forms of connectivity and different degrees of capacity. Nonetheless, for this to be possible, the research community must also produce and share different forms of data from real-life scenarios, i.e., mobility traces, so this work is built over solid foundations.

ACKNOWLEDGMENT

The authors would like to thank Bomberos de Asturias/112 for the GPS traces and España en Llamas for the wildfire historic data.

REFERENCES

[1] S. Cabrero, R. García, X. G. Pañeda, and D. Melendi, "CRAWDAD dataset oviedo/asturies-er (v. 2016-08-08)," Downloaded from http://crawdad.org/oviedo/asturies-er/20160808, Aug. 2016.

[2] T. Doumi, M. F. Dolan, S. Tatesh, A. Casati, G. Tsirtsis, K. Anchan, and D. Flore, "Lte for public safety networks," *IEEE Communications Magazine*, vol. 51, no. 2, pp. 106–112, February 2013.

[3] M. Usman, A. A. Gebremariam, U. Raza, and F. Granelli, "A software-defined device-to-device communication architecture for public safety applications in 5g networks," *IEEE Access*, vol. 3, pp. 1649–1654, 2015.

[4] T. Clausen and P. Jacquet, "Optimized link state routing protocol (olsr)," Internet Requests for Comments, RFC Editor, RFC 3626, October 2003, http://www.rfc-editor.org/rfc/rfc3626.txt. [Online]. Available: http://www.rfc-editor.org/rfc/rfc3626.txt

[5] C. Perkins, E. Belding-Royer, and S. Das, "Ad hoc on-demand distance vector (aodv) routing," Internet Requests for Comments, RFC Editor, RFC 3561, July 2003, http://www.rfc-editor.org/rfc/rfc3561.txt. [Online]. Available: http://www.rfc-editor.org/rfc/rfc3561.txt

[6] Z. Zhang, "Routing in intermittently connected mobile ad hoc networks and delay tolerant networks: overview and challenges," *IEEE Communications Surveys Tutorials*, vol. 8, no. 1, pp. 24–37, First 2006.

[7] A. Vahdat, D. Becker *et al.*, "Epidemic routing for partially connected ad hoc networks," 2000.

[8] A. Lindgren, A. Doria, and O. Schelén, "Probabilistic routing in intermittently connected networks," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 7, no. 3, pp. 19–20, Jul. 2003. [Online]. Available: http://doi.acm.org/10.1145/961268.961272

[9] M. Lindeberg, J. E. Haavet, S. Cabrero, V. Goebel, and T. Plagemann, "Message lost or message taken - on message ferry selection in dtns -," in *2012 IFIP Wireless Days*, Nov 2012, pp. 1–7.

[10] E. M. Trono, Y. Arakawa, M. Tamai, and K. Yasumoto, "Dtn mapex: Disaster area mapping through distributed computing over a delay tolerant network," in *2015 Eighth International Conference on Mobile Computing and Ubiquitous Networking (ICMU)*, Jan 2015, pp. 179–184.

[11] S. Jelassi and G. Rubino, "Yousos: A victim-centered multimedia sos emergency framework," in *2015 2nd International Conference on Information and Communication Technologies for Disaster Management (ICT-DM)*, Nov 2015, pp. 232–239.

[12] S. Cabrero, X. G. Pañeda, T. Plagemann, V. Goebel, and M. Siekkinen, "Overlay solution for multimedia data over sparse manets," in *Proceedings of the 2009 International Conference on Wireless Communications and Mobile Computing: Connecting the World Wirelessly*, ser. IWCMC '09. New York, NY, USA: ACM, 2009, pp. 1056–1061. [Online]. Available: http://doi.acm.org/10.1145/1582379.1582611

[13] R. Mart, S. Robles, A. Martín-Campillo, and J. Cucurull, "Providing early resource allocation during emergencies: The mobile triage tag," *Journal of Network and Computer Applications*, vol. 32, no. 6, pp. 1167 – 1182, 2009. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1084804509000769

[14] N. Aschenbruck, E. Gerhards-Padilla, M. Gerharz, M. Frank, and P. Martini, "Modelling mobility in disaster area scenarios," in *Proceedings of the 10th ACM Symposium on Modeling, Analysis, and Simulation of Wireless and Mobile Systems*, ser. MSWiM '07. New York, NY, USA: ACM, 2007, pp. 4–12. [Online]. Available: http://doi.acm.org/10.1145/1298126.1298131

[15] Y. Huang, W. He, K. Nahrstedt, and W. C. Lee, "Corps: Event-driven mobility model for first responders in incident scene," in *MILCOM 2008 - 2008 IEEE Military Communications Conference*, Nov 2008, pp. 1–7.

[16] E. Rosas, N. Hidalgo, V. Gil-Costa, C. Bonacic, M. Marin, H. Senger, L. Arantes, C. Marcondes, and O. Marin, "Survey on simulation for mobile ad-hoc communication for disaster scenarios," *Journal of Computer Science and Technology*, vol. 31, no. 2, pp. 326–349, 2016. [Online]. Available: http://dx.doi.org/10.1007/s11390-016-1630-x

[17] L. E. Quispe and L. M. Galan, "Behavior of ad hoc routing protocols, analyzed for emergency and rescue scenarios, on a real urban area," *Expert Systems with Applications*, vol. 41, no. 5, pp. 2565 – 2573, 2014. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0957417413008099

[18] A. Martín-Campillo, J. Crowcroft, E. Yoneki, and R. Mart, "Evaluating opportunistic networks in disaster scenarios," *Journal of Network and Computer Applications*, vol. 36, no. 2, pp. 870 – 880, 2013. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1084804512002275

[19] S. Cabrero, R. García, X. G. Pañeda, and D. Melendi, "Understanding opportunistic networking for emergency services: Analysis of one year of gps traces," in *Proceedings of the 10th ACM MobiCom Workshop on Challenged Networks*, ser. CHANTS '15. New York, NY, USA: ACM, 2015, pp. 31–36. [Online]. Available: http://doi.acm.org/10.1145/2799371.2799381

[20] A.-L. Barabási, "Linked: The new science of networks," 2003.

[21] A. Keränen, J. Ott, and T. Kärkkäinen, "The one simulator for dtn protocol evaluation," in *Proceedings of the 2Nd International Conference on Simulation Tools and Techniques*, ser. Simutools '09. ICST, Brussels, Belgium, Belgium: ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering),

2009, pp. 55:1–55:10. [Online]. Available: http://dx.doi.org/10.4108/ICST.SIMUTOOLS2009.5674

[22] S. Cabrero, X. Pañeda, R. García, D. Melendi, and T. Plagemann, "Dynamic temporal scalability: Video adaptation in sparse mobile ad-hoc networks," in *2012 IEEE 8th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, Oct 2012, pp. 349–356.

[23] M. Conti and S. Giordano, "Mobile ad hoc networking: milestones, challenges, and new research directions," *IEEE Communications Magazine*, vol. 52, no. 1, pp. 85–96, January 2014.