

Understanding Opportunistic Networking for Emergency Services

Analysis of One Year of GPS Traces

Sergio Cabrero
Centrum Wiskunde & Informatica
Science Park 123
Amsterdam, The Netherlands
s.cabrero@cwi.nl

Roberto García, Xabiel G. Pañeda,
David Melendi
University of Oviedo
Campus de Xixón s/n
Xixón, Spain
{garciaroberto, xabiel,
melendi}@uniovi.es

ABSTRACT

Opportunistic networking can help emergency services in both their daily operation and disaster relief. This idea has been extensively explored in previous research, but most studies are based on little knowledge of real mobility. In order to support future research, this paper analyses one year of GPS traces from a fire department. The results reveal the characteristics of hypothetical opportunistic networks formed by devices following this mobility considering different communication ranges. We found that the networks analysed are heterogeneous in many dimensions. They are also sparse and partitioned, but delay-tolerant routes connecting these partitions exist. To ease the discovery of these routes, we reveal in the connections between nodes. These findings can be applied in the design and deployment of solutions from the physical to the application layer.

Categories and Subject Descriptors

H.4 [Information Systems Applications]: Miscellaneous

General Terms

Theory

Keywords

Mobility analysis, Delay Tolerant Networks, Network Science, Emergency and Rescue

1. INTRODUCTION

Communication networks are an essential tool for emergency services. Nowadays, they are used for coordination, information gathering, alerting population and more. In the near future, the proliferation of new technologies, such

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CHANTS'15, September 11, 2015, Paris, France.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-3543-0/15/09 ...\$15.00.

DOI: <http://dx.doi.org/10.1145/2799371.2799381>.

as sensors, are going to increase even more their importance. However, the use of existent network infrastructures, e.g. commercial 3G/4G networks or specialised TETRA networks, is not always possible nor the best solution. Emergency services often act in remote locations, where the infrastructure may not exist, or after natural disasters, which may destroy it. Furthermore, costs can be saved by not using a commercial network for services that do not rely on continuous connectivity. For that reason, Mobile Ad-hoc Networks (MANETs) are an attractive alternative, not only in disaster relief, but also in the daily operation of emergency services. MANETs are deployed using wireless protocols, such as 802.11. Devices, also known as nodes, establish network links with others within their communications range. They act as both hosts and routers, forwarding traffic on behalf of others. Sometimes, when nodes can not be connected through a multihop route, they can still leverage Delay-Tolerant Networking (DTN) mechanisms to exchange data. MANETs and DTNs can be combined to get the most of the network. In general, nodes use opportunities given by their location to communicate with other nodes. Thus, these networks are generally called opportunistic networks. Understanding node mobility is key to engineer realistic solutions for this type of networks. The research community is aware of this fact and has worked on it. Mobility models have been proposed, for example in the area of tactical networks [1], which includes emergencies [2]. However, researchers confront the scarcity of real mobility traces to support these models. This trend is fortunately changing and more mobility and network traces are being gathered in different application domains¹.

This paper describes the results of analysing one year of mobility traces of an emergency service. We have worked together with the regional fire department of Asturias (Spain) to analyse their mobility without compromising privacy. The goal is to provide the properties of hypothetical opportunistic networks given their real movement. As far as we know, nobody has tackled the analysis of such a big mobility dataset in the context of emergencies. Our results reveal several insights that can be used for network engineering, protocol design and future research in opportunistic networking for emergency services.

¹See the CRAWDAD database - crawdad.org

The remainder of the paper is organised as follows. Next, we describe the mobility traces and the method to process them. Then, we present the most relevant results from applying our method to analyse mobility. Section 4 discusses the implications of the most relevant results. Finally, we formulate some conclusions and potential future work.

2. MOBILITY TRACES & METHOD

Our original data source is one year of GPS traces extracted from the Geographical Information System (GIS) of a regional Fire Department (Bomberos de Asturias / 112). The traces were generated by devices embedded mainly in cars and trucks, but also in a helicopter and a few personal radios. They represent their movement in emergency and rescue operation, but also in their daily routines. A total of 229 devices reported 19,462,339 locations. A new location is reported with an interval of approximately 30 seconds when the GPS device detects movement.

GPS traces allow us to hypothesise about the potential of opportunistic networks. Due to the size of the dataset, we make a set of assumptions to simplify its analysis. We assume that there is a network node located in the same positions registered by each GPS device. At any given moment, the position of a node is the last GPS position registered. Links between nodes are estimated calculating the distance between them. In our analysis, all nodes have the same communication range and a link is established if the distance between two nodes is shorter than it. We consider three ranges: 10, 50 and 200 meters; which can be associated with radio technologies such as Bluetooth, WiFi and WiMAX. Therefore, we obtain three hypothetic opportunistic networks from the original mobility dataset. These assumptions simplify our calculations at the cost of slightly decreasing the realism of the results obtained. First, the real position of a node may be an intermediate points between two reported positions. Second, the establishment of a link between nodes depends on many other factors apart from the distance, such as obstacles that may attenuate the signal.

Given the communication range and the GPS location of every node, we are able to estimate the existence of network links. We carry out two types of analyses. On the one hand, we analyse the dynamics of each link individually, i.e. when each link exists or not. We define a contact as the period of time when the link is uninterruptedly established. We define a break as the time between consecutive contacts. Our analysis examines the duration of contacts and breaks for each pair of nodes. They are fundamental to determine network properties, such as its capacity or delay. On the other hand, we analyse the dynamic topology of each network. The topology is defined by all the nodes and the links between them. Since nodes move, topology changes over time. To discover delay-tolerant properties of the networks, we consider snapshots of the network topology using different time windows: 1 hour, 1 day and 1 year. We calculate consecutive network topologies that link two nodes if there is a link between them at any moment within the given time window. For example, the time window of 1 day builds one network topology for each day. Each topology contains links between all nodes that have had a contact during that day, independently of its duration. Therefore, each of these windows generates a set of snapshots, which reveal different network properties. The 1 year window pro-

Table 1: Contact duration metrics

Range (meters)	10	50	200
# Links	3,755	11,057	19,291
# Contacts	399,129	2,440,652	1,458,481
Mean contact (seconds)	398	3,886	15,640
Aggregated contacts (hours)	431,340	2,398,270	5,686,176
Mean break (seconds)	8,634	55,490	212,500

vides one network topology representing the big picture of the relationships between nodes in our dataset. The 1 day window provides 365 network topologies, revealing the network properties for applications that could support delays lower than a day. The 1 hour window provides 8,740 network topologies, revealing properties for applications that are more sensitive to delay and closer to real time. Then, network science metrics [3] are used to analyse each of these snapshots. We look into partitions, clustering coefficients and different centrality metrics. These metrics reveal network properties, problems and protocol requirements, as it has been showcased by previous related work [6].

3. RESULTS

3.1 Contacts & Breaks

In this section, we analyse the distribution of several metrics: contact & break duration for all nodes (Figures 1 and 2), mean contact & break duration for each pair of nodes (Figures 3 and 4), and aggregated contact duration for each pair of nodes (Figure 5). The x-axis corresponds to the duration in seconds using a logarithmic scale. To ease readability, we add markers equivalent to long periods of time in seconds. We also provide a summary of relevant values for each communication range in Table 1.




Figure 1: Contact duration distribution for all links

In Figure 1, we analyse the duration of all the contacts found in our traces. It provides a good overview of how long contacts between nodes last. Most contacts last more than a 1 minute and less than 1 hour using a 10 meters range. For 50 meters and 200 meters, contact duration increases

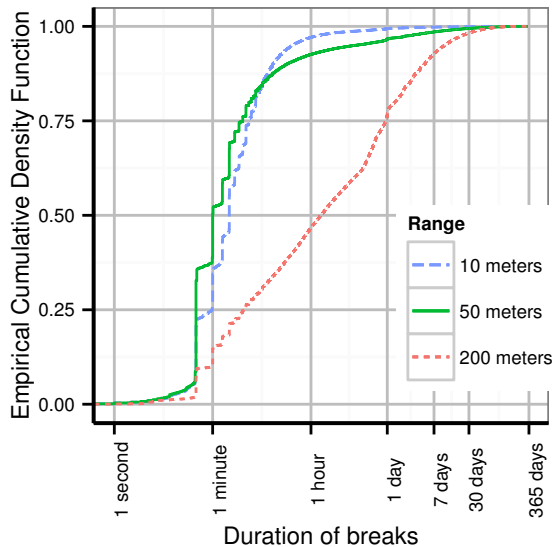


Figure 2: Break duration distribution for all links

significantly. As expected, the larger the range, the longer the contacts. Contacts longer than 1 day are unfrequent, although existent. Indeed, the longest contact detected lasts for 292 days, which most likely corresponds to vehicles that have been parked closely and barely used. Contacts shorter than 1 minute are difficult to detect due to the frequency in which GPS traces are collected (30 seconds), because it limits the maximum resolution for this analysis. Figure 2 represents the distribution of the duration of breaks. It is remarkable how the distribution for the 200 meter range differs from the others. Whereas going for 10 meters to a 50 meters range has a small impact in the duration of breaks, increasing the range to 200 meters decreases strongly the existence of short breaks.

Table 1 provides numerical results of the analysis of all contacts. We have extracted some relevant metrics to understand the relationships between nodes. We calculated the number of different links detected for each of the ranges. We also calculated the total number of contacts, considering all links. Whereas the number of links obviously increases with the range, the number of contacts is the highest for the 50 meters range. Thus, a significant amount of breaks (close to a million) that would occur in an ad-hoc network with 50 meters range could be avoided by increasing it to 200 meters. This is consequent with the result obtained analysing the duration of breaks, since we have already observed that a 200 meters range removes short contacts.

Table 1 provides the mean duration of contacts and breaks, as well as the aggregated durations of all contacts. These metrics provide a good insight of the influence of the range in the link stability -how frequently and for how long breaks occur- and also the network capacity -how much time nodes have to exchange information. Given the amount of time that nodes could be in contact, it seems like there are plenty of opportunities to exchange information without using any network infrastructure. Figures 3, 5 and 4 illustrate the distribution of these metrics, but individually for each pair of nodes (or link). Examining links separately reveals the heterogeneity of the network. Links present strong differences among them. Therefore, these networks are far from being

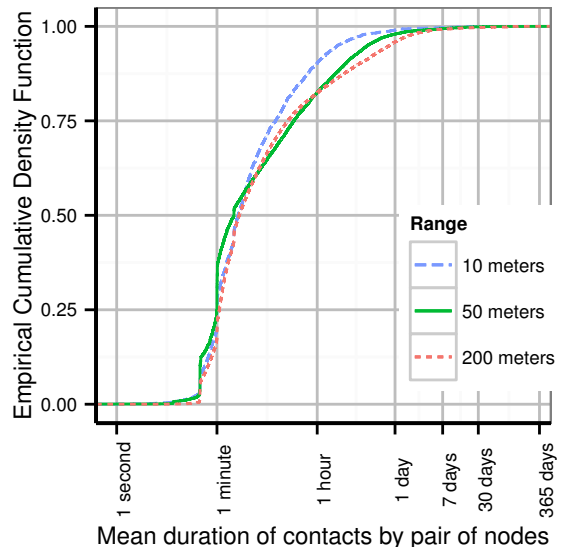


Figure 3: Mean contact duration distribution for all links

uniform, which would condition the way in which information can be propagated.

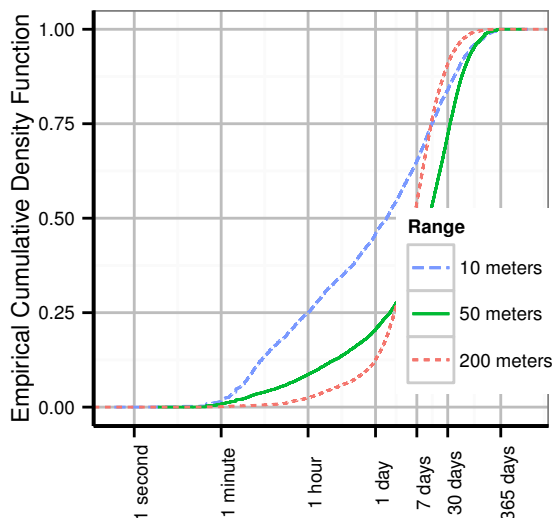
3.2 Network topologies

This section describes the results obtained from calculating snapshots of the network topology considering different communication ranges (10 meters, 50 meters and 200 meters) and time windows (1 hour, 1 day and 1 year). First, we present the results of analysing network partitions. Second, we look into clustering coefficients and centrality metrics. We have summarised the average values of these metrics in Table 2. To obtain each value, metrics are averaged from all the values obtained for each time window and range. Furthermore, we illustrate the Normalised Degree distributions of the topologies analysed in Figure 3.2.

Network partitions are groups of isolated nodes that can only communicate among them. Partitions are important from the networking point of view, because they give an idea on how nodes are connected. Two nodes are not able to exchange information within the given time window if they are in different partitions. We have measured the number and size of partitions in the network topologies generated by each time window and range. The general trend is that the number of partitions decreases with the length of the time window and with the range. This is sound with our expectations, because, in a longer period of time, nodes are more likely to contact with others, build new links and eventually merge in fewer (and bigger) partitions. In addition, a longer communication range increases the possibilities of establishing links. On the contrary, and also as expected, the average size of the partitions grows with the range and the time window length. The biggest diameter metric indicates the longest multihop route present in the network (often in the biggest partition). In other words, it is the maximum number of nodes that a packet would have to traverse if routing was optimal. For all the network topologies examined, the average of the longest route is below 9 hops, which can have key implications for the design of routing and forwarding policies. There are two relevant take-aways from analysing partitions. First, our results indicate that when

Table 2: Mean values of network metrics

Time window	1 year			1 day			1 hour		
Range (meters)	10	50	200	10	50	200	10	50	200
# Partitions	1	1	1	14.35	4.64	1.07	12.47	10.25	3.40
Partitions size	228	229	229	7.83	32.99	169.80	4.32	6.08	24.36
Biggest diameter	4	3	2	8.18	7.86	3.99	2.77	3.75	5.14
Clustering coefficient	0.44	0.71	0.88	0.56	0.40	0.31	0.48	0.28	0.09
Degree centrality	49.01	125.20	183.60	3.65	6.78	18.84	2.26	2.51	3.43
Betweenness centrality	95.81	51.40	22.20	45.03	154.50	101.40	1.22	7.80	92.05
Normalised closeness centrality	0.55	0.70	0.85	0.01	0.10	0.47	0.02	0.02	0.13



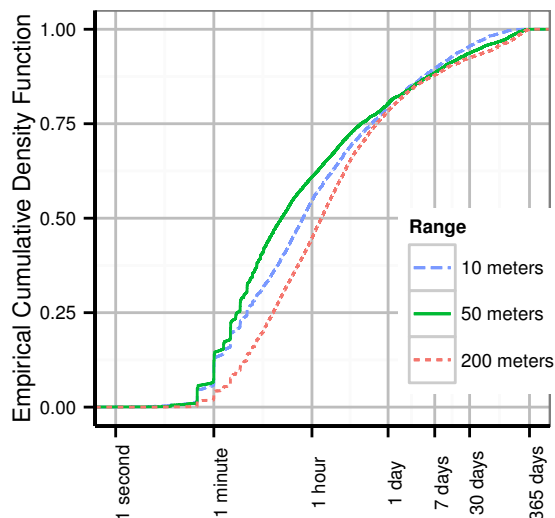
Mean duration of breaks by pair of nodes

Figure 4: Mean break duration distribution for all links

using a shorter time window, a larger number of nodes are isolated. Thus, the network topology is sparse in real time. Second, when the time window is big enough (e.g. 1 year) almost all nodes merge in a big partition. This means that communication between any two nodes is possible, although possibly with long delays.

The *clustering coefficient*, or transitivity, of a network is the probability that two nodes connected to a third one (its neighbours) are also connected. It is a measure of the connectivity between neighbours: a high clustering coefficient indicates that the neighbours of a node are very likely to be connected as well. This metric gives a basic idea about network density and about groups of nodes (clusters) that are highly interconnected. Clusters are interesting because they can be leveraged in the design of distributed systems. In addition, the clustering coefficient also gives insights on how strongly connected is the network, which is important for network resilience. For example, a network suffers less when a node drains its battery if its neighbours are well connected among them. It would be expected that a longer range and time window would produce higher clustering coefficients, however, this is not always the case, as observed in Table 2.

Centrality metrics measure the importance or popularity of nodes in the network. There are several centrality metrics, but we consider three that are interesting to analyse ad-hoc networks according to Katsaros et al. [7]: degree, betweenness and closeness. Degree centrality is the number



Aggregated duration of contacts by pair of nodes

Figure 5: Aggregated contact duration distribution for all links

of neighbours of a node. As expected, it increases with the range and the time window. Understanding degree distributions is key to understand the structure of a network. Figure 3.2 represents these distributions for all the network topologies generated with all ranges and time windows. We can observe that the trend is similar for different ranges using the same time window, but not for different time windows (especially for 1 year). These distribution also show the heterogeneity of the network: some nodes have a few links and some have many. Popular nodes, the ones with many links, are referred to as hubs. In our observations for the 1 year time window, there are hubs with a relatively high degree, e.g. 221. *Betweenness centrality* is the number of shortest paths that traverse a node. This metric can reveal the existence of bottlenecks in the network. A node with high betweenness is likely to forward more packets on behalf of others and, as a consequence, may become congested or drain its battery faster. *Closeness centrality* is the inverse of the distance in hops from a node to the rest of the nodes in the network. This metric is relevant for information dissemination in a network. If we wanted to send a message to all nodes in a partition, the most efficient way would be to use the node with the highest closeness. The most interesting result is for the 1 year window, where the average of the observations are relatively high. This means that most nodes are close to each other in the long term and that there would be many candidates to disseminate information.

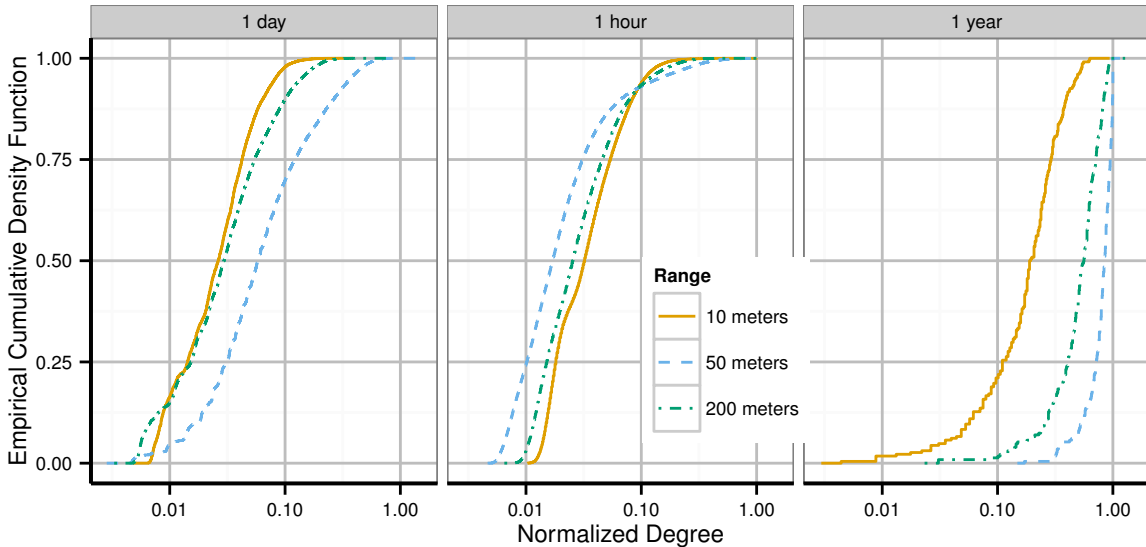


Figure 6: Normalised degree distribution for all ranges and time windows

4. DISCUSSION

This section discusses some implications of the described results in the design and deployment of solutions from the physical to the application layer. The stability shown by contacts is relevant for the protocols in the lower layers. Fast connections and disconnections of nodes are infrequent. Therefore, there is time enough for state-of-the-art link layer protocols, e.g. 802.11, to manage links between nodes. If contacts were too short, these protocols would have to support very fast link establishment and soon link break detection. Shared medium is also problematic in wireless ad-hoc networks. There are two main issues: collisions produced by too many nodes sharing the medium and the well-known hidden node problem. These issues are only relevant if nodes compete for the medium at the same time. The analysis of partitions revealed that in real time the network would be sparse, with many isolated nodes, 2 or 3 nodes partitions and short multihop routes. Therefore, the network is unlikely to suffer heavily from any shared medium problem.

Although the network is sparse in real time, nodes connect when observing longer timespans. This demonstrates a common assumption when designing systems for MANETs in emergencies: the network is sparse and partitioned, but there are nodes that can be used as data ferries. The store-carry-forward paradigm [11] can be applied to design delay-tolerant applications, in which a key problem is how to find ferry nodes. Delay-tolerant routing protocols, e.g. PROPHET [8] or dLife[9], aim to solve this issue. They typically use information from past contacts to predict future contacts. To understand if this approach is valid in our scenarios, we look for existent patterns in the contacts. In specific, we aim to predict future frequency of contact from past frequency of contact. The *frequency of contact* for a pair of nodes can be calculated dividing the aggregated contact duration by the total time analysed. For example, if one minute of mobility is analysed and two nodes are connected for 30 seconds, the frequency of contact is 0.5.

Lets define a variable t that takes as value all the seconds in our traces. Then, for each t we calculate the frequency of contact before t ($F_{c,bt}$) and the frequency of contact after t

($F_{c,at}$). So, for a pair of nodes, $F_{c,bt}$ is the aggregated contact duration before t divided by the seconds elapsed from the beginning of the traces to t . On the other hand, $F_{c,at}$ is the aggregated contact duration after t divided by the seconds left from t to the end of the traces. We now suppose that every time a contact ends, the nodes predict that $F_{c,at}$ is equal to $F_{c,bt}$. This is a simple prediction that can be implemented in real systems. Since we know $F_{c,at}$, the error of making these predictions in our traces is $F_{c,at} - F_{c,bt}$. Figure 7 shows the Probability Density Function for the prediction errors. The error means are 0.01, 0.12, and 0.05 for 10, 50 and 200 meters respectively. Standard deviations are small. Thus, capacity estimation errors are around 0, which may be assumable in some real systems. These results indicate that repeating patterns are present in the contacts. Therefore, PROPHET-like protocols are adequate for this application domain. To increase packet delivery probability and overcome prediction errors, several ferries may be used. This is possible in these networks because they are well connected and several routes between nodes exist.

A MANET routing protocol discovers multihop routes in real time. These protocols are sometimes criticised due to their bad scalability. However, this would not be a problem, because multihop routes are likely to be short and the partitions small, which implies a small routing table. Short routes positively affect communications reliability as well. It is well known that the probability of losing a packet increases with the number of hops that it has to traverse.

Centrality metrics indicate that hubs are common in the network and that nodes are highly interconnected. There are many popular nodes, which is positive for network resilience. If a node failed, alternative nodes would be able to rebuild connectivity. These results are also relevant for network congestion as bottlenecks are unlikely according to observed betweenness. Finally, a problem in MANETs is the underperformance of TCP. Basically, TCP confuses disruptions and congestion [5]. Thus, if disruptions are frequent, TCP obtains a very low throughput. In our analysis we have observed short contacts, but also very long ones, in which it should be possible to use TCP. Therefore, TCP use may be reconsidered in some situations.

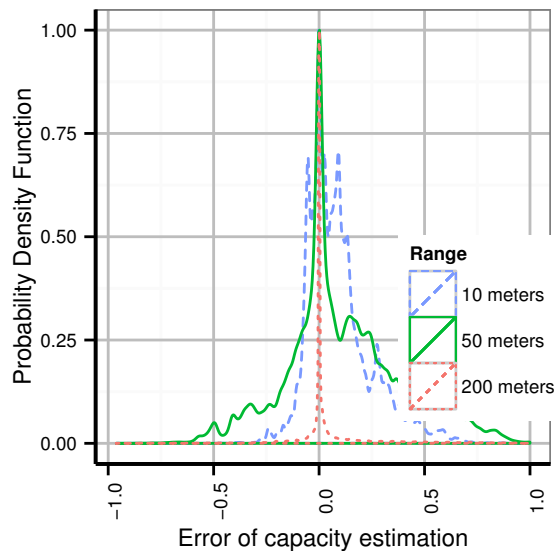


Figure 7: Capacity estimation error distribution

5. CONCLUSIONS

This paper proposes the use of GPS traces to discover the properties of opportunistic networks for emergency services. Our approach differs from others that use connectivity traces, e.g. [10]. GPS traces introduce a higher level of abstraction that could reduce the realism of the results. Nonetheless, they also introduce more freedom in the analysis, such as in the usage of different communication ranges. As a result, the traces reveal interesting properties of an opportunistic network for an emergency service. The results and discussion of this paper can help other researchers in the design of experiments and protocols with a realistic approach. Although mobility traces can not be made public due to privacy issues, the contacts will be available in the CROWDAD database to drive simulations or apply further analysis.

Our analysis has revealed three important properties in the networks analysed. First, they are heterogeneous in several dimensions. Degree distributions are not uniform. Hence, nodes in the network can have a very different number of neighbours. In addition, every link behaves differently, producing different durations of contacts and breaks. Second, mobility creates sparse MANETs in real-time, but with the possibility of delay-tolerant communication. At a given moment, nodes are likely to be isolated or in a small partition. However, their movement opens the possibility to connect with others and use store-carry-forward to transport information. Third, the error of predicting future link capacity with past link capacity is small. The direct implication is the design of efficient routing strategies. However, other applications can be envisioned from capacity estimation, such as adaptive video transport [4].

Future work will be directed to deepen the analysis of these mobility traces. By using information from the emergency services, it will be possible to extract mobility from specific situations, e.g. wildfires. The analysis of these scenarios could reveal specific properties, e.g. linked to the type of emergency operation. Furthermore, models for some parameters will be studied to enable the generation of networking scenarios with realistic characteristics. The analysis

of mobility will be complemented with the evaluation of existent delay-tolerant routing protocols, as well as the design of new solutions for emergency services.

6. ACKNOWLEDGMENTS

We would like to thank Bomberos de Asturias and 112 Asturias for providing the GPS traces and the information for their analysis.

7. REFERENCES

- [1] N. Aschenbruck, E. Gerhards-Padilla, and P. Martini. A survey on mobility models for performance analysis in tactical mobile networks. *Journal of Telecommunications and Information Technology*, pages 54–61, 2008.
- [2] N. Aschenbruck, A. Munjal, and T. Camp. Trace-based mobility modeling for multi-hop wireless networks. *Computer Communications*, 34(6):704 – 714, 2011.
- [3] A.-L. Barabasi. *Linked the new science of networks*. Perseus Pub., Cambridge, Mass., 2002.
- [4] S. Cabrero, X. G. Pañeda, R. García, D. Melendi, and T. Plagemann. Dynamic temporal scalability: video adaptation in sparse mobile ad-hoc networks. In *Wireless and Mobile Computing, Networking and Communications (WiMob), 2012 IEEE 8th International Conference on*, pages 349–356. IEEE, 2012.
- [5] G. Holland and N. Vaidya. Analysis of tcp performance over mobile ad hoc networks. *Wireless Networks*, 8(2/3):275–288, 2002.
- [6] P. Hui, J. Crowcroft, and E. Yoneki. Bubble rap: Social-based forwarding in delay-tolerant networks. *Mobile Computing, IEEE Transactions on*, 10(11):1576–1589, 2011.
- [7] D. Katsaros, N. Dimokas, and L. Tassiulas. Social network analysis concepts in the design of wireless ad hoc network protocols. *Network, IEEE*, 24(6):23–29, November 2010.
- [8] A. Lindgren, A. Doria, and O. Schelén. Probabilistic routing in intermittently connected networks. *ACM SIGMOBILE mobile computing and communications review*, 7(3):19–20, 2003.
- [9] W. Moreira, P. Mendes, and S. Sargento. Opportunistic routing based on daily routines. In *World of wireless, mobile and multimedia networks (WoWMoM), 2012 IEEE international symposium on*, pages 1–6. IEEE, 2012.
- [10] P.-U. Tournoux, J. Leguay, F. Benbadis, V. Conan, M. Dias de Amorim, and J. Whitbeck. The accordion phenomenon: Analysis, characterization, and impact on dtn routing. In *INFOCOM 2009, IEEE*, pages 1116–1124. IEEE, 2009.
- [11] W. Zhao, M. Ammar, and E. Zegura. A message ferrying approach for data delivery in sparse mobile ad hoc networks. In *International symposium on Mobile ad hoc networking and computing (MobiHoc)*, pages 187–198. ACM, 2004.