Quantifying and Minimizing the Impact of Disasters on Wireless Communications

Francesco Malandrino
Politecnico di Torino
Torino, Italy

Carla Fabiana Chiasserini
Politecnico di Torino
Torino, Italy

ABSTRACT

Communication networks are often disrupted during emergency situations, which can make relief operations slower, riskier, and less effective. In this context, our paper has the twofold purpose of (i) *quantifying* how significant such disruptions are, and (ii) evaluating the best strategies to *minimize* them. For both tasks, we leverage a set of real-world, crowd-sourced traces, including information on the position and connectivity (cellular and Wi-Fi) of thousands of smartphone users across the United States. Through our analysis, we find that enabling device-to-device communications can be extremely beneficial in dense urban centers, while including Wi-Fi access points in the safety network is the best strategy in suburban areas.

CCS CONCEPTS

• **Networks** → *Network performance analysis*; *Network reliability*;

KEYWORDS

Disaster relief, wireless networks, device-to-device communications, real-world traces.

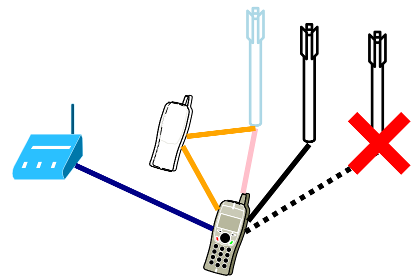

Figure 1: Connectivity options of a mobile user whose cell has become unavailable due to a disaster (dotted black line), in counter-clockwise order: connecting to a different cell of the same operator (black line); connecting to a cell of a different operator (blue line); use a device-to-device link to reach the Internet (orange lines); using a Wi-Fi access point (pink line).

1 INTRODUCTION

Smartphones and smartphone apps are ever present throughout all our lives, and tragic moments such as emergencies and disasters are no exception. To the potential victims, Facebook provides a *Safety Check* feature [1], allowing users that are detected to be close to an emergency situation to mark themselves as "safe". As far as responders are concerned, LTE is widely [2, 3] regarded to as the natural successor of special-purpose technologies such as TETRA or Project 25, providing broadband capacity in mission-critical situations.

In this context, the extent to which communication networks are impaired by disasters has a critical impact on the success of disaster relief efforts. If the victims of a disaster have no network access, they cannot inform the responders about their position and situation. Similarly, if the responders have to use special-purpose technologies only, rescue operations become slower, harder, and potentially more dangerous.

The first objective of this paper is therefore to *assess* the potential impact of disasters on wireless communication networks, as a function of factors like the size of the area affected by the disaster itself or the considered wireless technology. A second, more important, objective is to investigate possible strategies to *minimize* such an impact. As summarized in Fig. 1, there can be several connectivity options available to a mobile user whose cell becomes unavailable; it is thus important to assess their effectiveness and the extent to which they can be combined.

In order to achieve our goals, we leverage a set of real-world, large-scale, crowd-sourced traces, including the position of the users *and* the connectivity options available to them (e.g., Wi-Fi and LTE). This allows us, for example, to establish a link between the size of the area affected by a disaster, the number of users that remain isolated as a consequence, and the effectiveness of alternative connectivity strategies. In particular, this analysis can provide useful guidelines on the benefits of adopting *ad hoc* devices like Sonnet [4], which, exploiting device-to-device (D2D) communications, allow smartphones to connect to other user devices even when they are out of cellular coverage.

Our analysis is mainly related to two broad research topics, namely, characterizing the victims of a disaster and enhancing the connectivity options available to them. Among the first body of works, [5, 6] study the problem of localizing the victims of a disaster, using help messages [5] or social network posts [6], while [7] and [8] model, respectively, the generated traffic and the mobility of both victims and responders. With regard to enhancing the connectivity options during a disaster, [9, 10] survey and compare the available hardware options, while [11, 12] focus on scheduling and routing, and [10, 13, 14] discuss possible network architectures.

The remainder of the paper is organized as follows. Sec. 2 describes the traces we use for our analysis. Then, Sec. 3 presents our model for the disasters themselves and the possible remedy strategies. Finally, Sec. 4 summarizes our numerical results and Sec. 5 concludes the paper and sketches directions for future work.

Figure 2: Area covered by the San Francisco trace. Dots correspond to locations with active users.

Figure 3: Subset of the area covered by the Los Angeles trace that we consider, whose size matches the area covered by the San Francisco trace. Dots correspond to locations with active users.

2 REAL-WORLD TRACES

The data traces we use come from an Android app called WeFi [15], providing its users with up-to-date, location-specific information on available Wi-Fi access points and their features (encryption, speed, reliability...). In return, users share with WeFi information about their location and activity, including the apps they use and the available connections at their location. In our paper, we use two datasets, collected in the U.S. cities of Los Angeles and San Francisco. Tab. 1 summarizes main features of the datasets, while Fig. 2 and Fig. 3 respectively show the areas they cover. Notice that we only consider a subset of the area covered by the Los Angeles trace, so as to be able to directly compare it with the San Francisco trace.

The traces are made of several *records*. For each user, a new record is generated every time one of this event occurs: (i) a one-hour period passes, (ii) the user moves to a new location, (iii) the app active on the smartphone changes, (iv) the user connects to a different cell or network (e.g., Wi-Fi). Each record contains time and location information, an anonymized user identifier, the mobile operator and ID of the cell the user is connected to (if any), and the anonymized BSSID of the Wi-Fi AP the user is connected to (if any).

Table 1: WeFi datasets

	Los Angeles	San Francisco
Time of collection	Oct. 2014	Mar. 2015
Covered area [km^2]	46×73	14×11
Total traffic [TB]	35.61	9.18
Number of records	81 million	60 million
Unique users	64,386	14,018
Unique cells	36,09	14,728

These traces are especially useful for our analysis, for three main reasons. First, we can compare disaster relief situations in two major U.S. cities, with substantially different urban structures. Second, the traces span different mobile operators *and* Wi-Fi access points (APs). This allows us to understand the benefits of making different network technologies interact, and to which extent they can complement each other. Finally, we have information on the individual cells composing the cellular network of each mobile operator covering a given geographical area. This makes it possible to study scenarios where some cells of a mobile operator are active while others are not, as well as to assess the benefits of mobile operators cooperating to provide service in disaster areas.

3 NETWORK AND DISASTER MODELS

This section presents the graph-based model we use to describe mobile networks. Additionally, it discusses how the network model can describe both disasters (i.e., their impact on network connectivity) and remedy strategies (i.e., the additional connectivity options they bring).

3.1 Network model

Any network snapshot can be described through a directed graph, with vertices representing cells, Wi-Fi APs, and users. A special vertex called `Internet` is also added, in order to track which users have Internet access. The following edges are created:

- from `Internet` to all cells;
- from `Internet` to all Wi-Fi APs;
- from each cell to all users of the same network operator that are within the cell coverage area[1];
- from each Wi-Fi AP to all users that are observed, in the trace, to download data from it.

An example of the network graph is depicted in Fig. 4. Given the graph, a user has Internet connectivity whenever there is a path on the graph that connects the vertex corresponding to the user with the `Internet` node. Thanks to this property, we can study the effect of disasters by removing some nodes and/or edges from the graph, and observing how the set of users connected to the Internet

[1] We assume [16] the coverage area of a cell to correspond to the convex hull of all locations from which users have reported being covered by the cell itself.

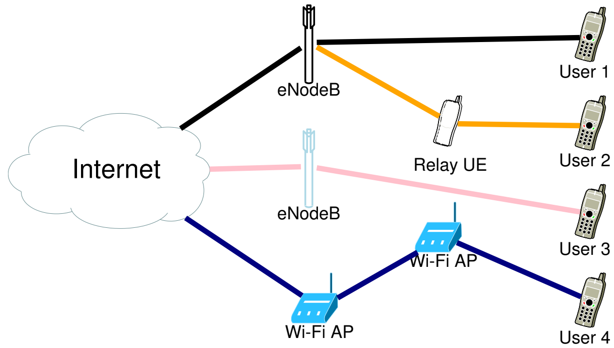

Figure 4: Example graph. Users 1–4 connect to the Internet through a cell belonging to their operator, roaming, D2D, and Wi-Fi respectively.

shrinks. Similarly, we can model remedy strategies by adding edges to the graph and then studying how this enlarges the number of users connected to the Internet.

3.2 Disasters

Disasters are described through three properties, namely:

- their *focus* location F, e.g., the epicenter of an earthquake;
- the *radius* r of the area around the focus that is affected by the disaster;
- their *severity* s, i.e., the fraction of cells and Wi-Fi APs within the disaster area that are disabled.

Given the disaster properties, the network graph is changed as follows. First, a fraction s of the cells within the disaster area, i.e., whose location[2] is less than r away from F, are removed from the graph, along with all incoming and outgoing edges. The other cells that are within the disaster area are left in the graph, but their connection with the `Internet` node is removed; this models the IOPS emergency mode in LTE, that allows eNBs to provide local IP connectivity via a local EPC.

Similarly, we remove the link between `Internet` and a fraction s of Wi-Fi APs within the disaster area, but not the APs themselves or their links. This corresponds to the case when wired Internet connectivity is disabled, but the equipment itself is unharmed.

Users in the graph can now be in one of three states. If they are out of the disaster area, they are *unaffected* by the disaster itself. If they are within the disaster area and are unreachable from `Internet`, we call them *isolated*. Finally, users within the disaster area for which there is still a path on the graph connecting them with `Internet`, are *reconnected*. Clearly, it is our purpose to have as few isolated users as possible, and as many nodes within the disaster area as possible in the "reconnected" state.

3.3 Remedy strategies

In the following, we describe several remedy strategies, that can be adopted in order to move users from the "isolated" to the "reconnected" state. They vary in complexity, i.e., hardware and/or

[2]The location of the base station serving a cell is assumed [16] to correspond to the barycenter of the cell coverage area.

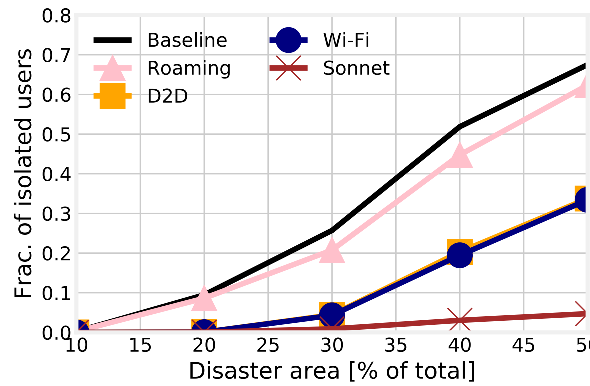

Figure 5: San Francisco trace: fraction of isolated users as a function of the disaster area size, for different sets of remedy strategies. Remedy strategies are applied in cascade, in increasing order of complexity, e.g., the "D2D" line refers to the case when both roaming and D2D are enabled.

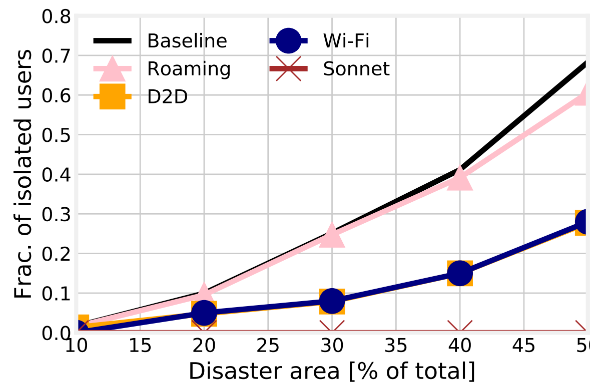

Figure 6: Los Angeles trace: fraction of isolated users as a function of the disaster area size, for different sets of remedy strategies. Remedy strategies are applied in cascade, in increasing order of complexity, e.g., the "D2D" line refers to the case when both roaming and D2D are enabled.

software changes required for their implementation, and – of course – effectiveness. In increasing order of complexity, we consider: data roaming; cellular D2D data transfer; open Wi-Fi APs; special-purpose *ad hoc* devices.

Data roaming. This strategy simply consists in granting all users within the disaster area access to any cell, regardless its operator. It is a direct extension of how emergency numbers (e.g., 999 in the U.S.) can be dialed using any network. In the graph, it corresponds to adding new edges from every cell to all users within its coverage area. Its implementation would require a change to the configuration of cellular networks.

Cellular D2D data transfers. D2D communication between users' smartphones is a very popular and promising way to improve the coverage and capacity of cellular networks, in both normal and emergency conditions [17]. Its implementation is foreseen by current 3GPP standards, and is expected to be commonplace

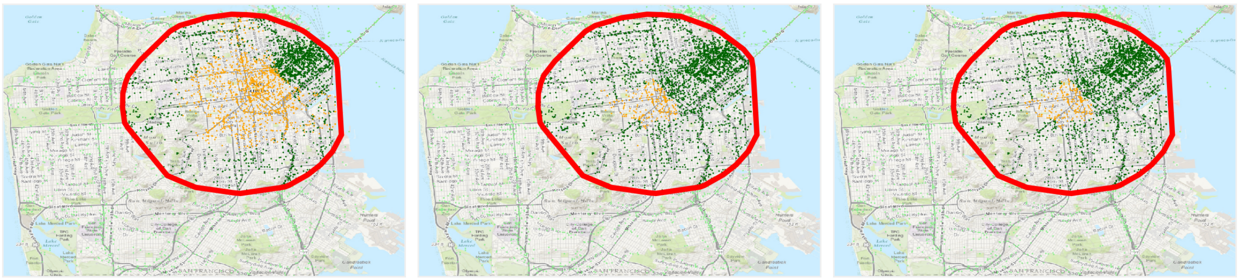

(a) (b) (c)

Figure 7: San Francisco, scenario when the disaster area (enclosed by the red line) is 30% of the total topology. Location of unaffected (pale green), reconnected (dark green), isolated (orange) users, under the roaming (a), D2D (b), AHD (c) remedy strategies.

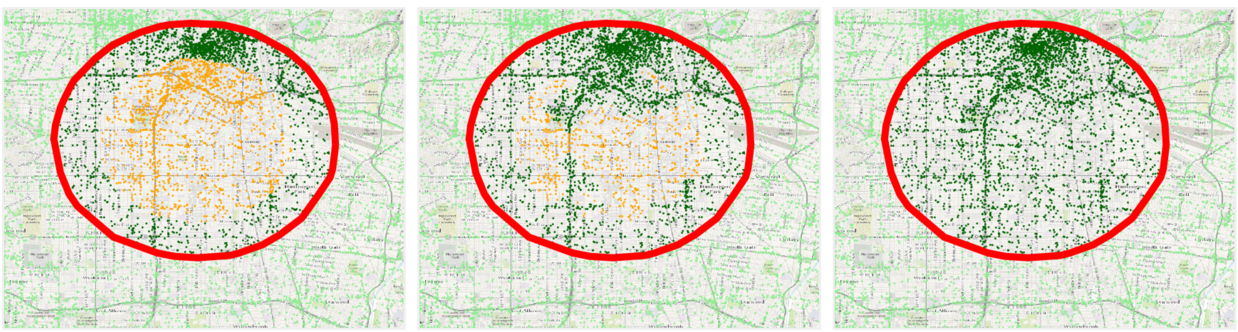

(a) (b) (c)

Figure 8: Los Angeles, scenario when the disaster area (enclosed by the red line) is 30% of the total topology. Location of unaffected (pale green), reconnected (dark green), isolated (orange) users, under the roaming (a), D2D (b), AHD (c) remedy strategies.

among smartphones in the near future [18]. In our model, assuming symmetric links, it corresponds to adding an edge between any two users within D2D range of each other.

Open APs. This strategy envisions granting users access to any Wi-Fi AP within their proximity, i.e., disabling security-layer function in the event of a disaster. In the graph, it corresponds to adding edges from every AP to all users within its coverage area, as well as between APs within coverage of each other. Its implementation would require updating the firmware of recent APs, and replacing older ones.

Special-purpose ad-hoc devices (AHD). Several devices have been proposed that enhance the connectivity of smartphones during emergency situations, or simply if the user is abroad and needs to save on roaming fees. An example is Sonnet [4], a device that offers Wi-Fi connectivity to the user smartphone within its range, and can communicate with other devices up to 5 km away. In our model, this strategy corresponds to adding a given number of AHD nodes to the graph (as many as the number of users that own an AHD). Each AHD node is then connected through a AHD-smartphone edge to the corresponding user node (which should be closer than the typical Wi-Fi range), and through an AHD-to-AHD edge with another AHD closer than 5 km. Implementing this strategy clearly requires a widespread adoption of ad-hoc devices.

4 NUMERICAL RESULTS

We evaluate the impact of disasters and the effectiveness of the above remedy strategies in scenarios where:

- we consider a time snapshot of the topologies, taken at noon;
- the focus location F of disaster is the location that minimizes the average distance from all users, thus putting us in the most challenging possible scenario;
- the severity of the disaster is always $s = 1$, i.e., all cells within the disaster area are disabled;
- the disaster radius r is adjusted in such a way that the disaster area varies between 10% and 50% of the total topology;
- 5% of users, randomly chosen, are equipped with an ad-hoc device.

Furthermore, we set the range for D2D links and Wi-Fi APs to 100 and 150 meters, respectively. In all our experiments, we *cascade* remedy strategies, in increasing order of complexity: as an example the "D2D" lines in the plots refer to the case when both roaming and D2D are enabled.

Fig. 5 and Fig. 6 summarize the most important metric we are interested in, i.e., the fraction of users that are isolated from the network. As we can expect, such a quantity grows as the area hit by disaster becomes larger and, when no remedy strategies are in place, exceeds 60% for both cities (black lines).

The different remedy strategies have similar impact in the two cities. Allowing roaming, for example (pink lines in the plots), brings limited improvement in San Francisco and almost none in Los Angeles; this is due to the higher density of cells in the first city, which makes it more likely that portions of the disaster area are covered by base stations that are themselves unharmed. D2D transfers (orange lines in the plots), on the other hand, substantially decrease the fraction of isolated users in both cases. Adding Wi-Fi (blue lines) again has negligible effects; intuitively, people and Wi-Fi APs tend to be located in the same areas, so adding both to the network does not increase its coverage – though it does improve its robustness. Finally, enabling AHDs (brown lines) brings the fraction of isolated users to almost zero both cases, owing to their long-distance communication capabilities.

We can get a better understanding of how different relief strategies considering the location of reconnected users. In Fig. 7 and Fig. 8, the area affected by the disaster is enclosed by a thick red line, and dots correspond to user locations. All users outside the disaster area are in the "unaffected" state, and are colored in pale green. Users within the disaster area, instead, are colored in orange if they are isolated, and in dark green if they are reconnected.

Fig. 7(a) and Fig. 8(a) depict the effects of the "roaming" strategy: in both cities, this strategy affects those users that are closer to the border of the disaster area. Furthermore, the location of reconnected users is roughly symmetric with respect to the disaster focus: as one might expect, the probability of being covered by an unaffected base station mostly depends on how close one is to the edge of the disaster area.

More interestingly, looking at the effect of D2D, depicted in Fig. 7(b) and Fig. 8(b), we note that the location of reconnected users is *not* symmetric with respect to the disaster focus. Indeed, most of the users reconnected by D2D are located in the most densely populated areas, e.g., Chinatown, and SoMa in San Francisco.

5 CONCLUSION AND FUTURE WORK

We set the twofold goal of estimating the extent to which disasters can disrupt wireless communications, and evaluating the remedy strategies to reduce it. To this end, we leveraged two real-word, large-scale, crowd-sourced traces collected in the U.S. cities of San Francisco and Los Angeles. We then combined them with a graph-based model accounting for both the effects of disasters and the results of remedy strategies.

Our results show that a combination of device-to-device transfers between mobile users and mesh-like networks composed of Wi-Fi access points can substantially decrease the number of users that remain isolated from the network. Furthermore, the relative effectiveness of different strategies strongly depends upon the features of the scenario under consideration, e.g., its user density.

One prominent direction for future work is the characterization of disasters. We plan to develop different models for different types of disasters, e.g., floods and earthquakes, in order to account for the different impact they have on people and infrastructure.

REFERENCES

[1] Facebook. Safety Check. https://www.facebook.com/about/safetycheck.
[2] Nokia. LTE to the rescue in public safety networks. https://blog.networks.nokia.com/public-safety/2015/01/14/improving-public-safety-with-lte.
[3] R. Ferrus, O. Sallent, G. Baldini, and L. Goratti, "Lte: the technology driver for future public safety communications," *IEEE Communications Magazine*, 2013.
[4] Sonnet Labs. Sonnet: revolutionary communications device for the adventurer in all of us. https://www.sonnetlabs.com/.
[5] N. Suzuki, J. L. F. Zamora, S. Kashihara, and S. Yamaguchi, "Using sos message propagation to estimate the location of immobilized persons," in *ACM MobiCom*, 2012.
[6] R. C. Jain, "Disasters response using social life networks," in *ACM WWW*, 2013.
[7] K. Cho, C. Pelsser, R. Bush, and Y. Won, "The japan earthquake: The impact on traffic and routing observed by a local isp," in *ACM CoNEXT SWID Workshop*, 2011.
[8] S. C. Nelson, A. F. Harris, III, and R. Kravets, "Event-driven, role-based mobility in disaster recovery networks," in *ACM CHANTS*, 2007.
[9] K. Miranda, A. Molinaro, and T. Razafindralambo, "A survey on rapidly deployable solutions for post-disaster networks," *IEEE Communications Magazine*, 2016.
[10] M. Herlich and S. Yamada, "Comparing strategies to construct local disaster recovery networks," in *IEEE AINA*, 2016.
[11] V. S. H. Huynh and M. Radenkovic, "A novel cross-layer framework for large scale emergency communications," in *IEEE IWCMC*, 2017.
[12] Q. T. Minh, Y. Shibata, C. Borcea, and S. Yamada, "On-site configuration of disaster recovery access networks made easy," *Ad Hoc Netw.*
[13] M. Raj, K. Kant, and S. K. Das, "E-darwin: Energy aware disaster recovery network using wifi tethering," in *IEEE ICCCN*, 2014.
[14] T. Ngo, H. Nishiyama, N. Kato, S. Kotabe, and H. Tohjo, "A novel graph-based topology control cooperative algorithm for maximizing throughput of disaster recovery networks," in *IEEE VTC Spring*, 2016.
[15] WeFi. WeFi Mobile App. https://www.wefi.com/.
[16] F. Malandrino, C.-F. Chiasserini, and S. Kirkpatrick, "Cellular Network Traces Towards 5G: Usage, Analysis and Generation," *IEEE Transactions on Mobile Computing*, 2017.
[17] F. Malandrino, Z. Limani, C. Casetti, and C. F. Chiasserini, "Interference-aware downlink and uplink resource allocation in hetnets with d2d support," *IEEE Transactions on Wireless Communications*, vol. 14, no. 5, pp. 2729–2741, May 2015.
[18] 3GPP. 3GPP TS 23.303 V13.4.0 (2016) Rel. 13 – Proximity-based services (ProSe). https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=840.

ACKNOWLEDGMENTS

This work was supported by the European Commission through the I-REACT project (grant agreement no. 700256) and by the Fondazione CRT under the initiative *La Ricerca dei Talenti*, through the MIMOSE project.