Building Dissent Networks: Towards Effective Countermeasures against Large-Scale Communications Blackouts

Shaddi Hasan
UC Berkeley
shaddi@cs.berkeley.edu

Yahel Ben-David
UC Berkeley and De Novo Group
yahel@cs.berkeley.edu

Giulia Fanti
UC Berkeley
gfanti@eecs.berkeley.edu

Eric Brewer
UC Berkeley and Google
brewer@cs.berkeley.edu

Scott Shenker
UC Berkeley and ICSI
shenker@cs.berkeley.edu

Abstract

Large-scale communications blackouts, such as those carried out by Egypt and Libya in 2011 and Syria in 2012 and 2013, have motivated a series of projects that aim to enable citizens to communicate even in the face of such heavy-handed censorship efforts. A common theme across these proposals has been the use of wireless mesh networks. We argue that such networks are poorly equipped to serve as a meaningful countermeasure against large-scale blackouts due to their intrinsically poor scaling properties. We further argue that projects in this space must consider user safety as first design priority and thus far have failed to preserve user anonymity and to rely only on innocuous hardware. From these two insights, we frame a definition of *dissent networks* to capture the essential requirements for blackout circumvention solutions.

1 Introduction

In the wake of the 2011 Arab Spring, international attention focused on the role that the Internet and social media services such as Facebook and Twitter can play in supporting popular uprisings against repressive regimes. At the same time, the actions of these regimes demonstrated the fragility of the infrastructure that connects people to these services, as well as their willingness to use the full power of the state to engage in large-scale censorship of the Internet and other communication networks. In response, researchers and technically-minded activists around the world have started projects which aim to build censorship-resistant communication networks. Their goals vary, ranging from building an alternative Internet infrastructure outside the control of corporate or government interests to building emergency communications infrastructures for times of crisis. Yet for the most part, they all share a common goal—building networks that can survive serious disruption to existing communications infrastructure while ensuring free expression among their users.

We concur that this is a worthy goal. However, we believe that much of the work in this space suffers a disconnect from reality which stems from a lack of a clearly defined set of properties for such networks. To this end, we propose and define "dissent networking", discuss the desired properties, and address the suitability of proposed technologies and solutions (or lack thereof). Dissent networks aim to allow free expression even in the face of censorship and communications blackouts. Dissent networks are:

- *Resilient against communications blackouts.* Should be challenging for any entity to disable.
- *Resistant to monitoring and tracking of users.* Both who is using the network and any sensitive messages they send should be secret.
- *Able to be built from innocuous components.* Should only require readily available hardware, and the possession and use of required hardware shouldn't be illegal or suspicious.
- *Able to run at meaningful scales.* Should be more effective at disseminating information than people with megaphones; more broadly, given a level of service, should be able to run at non-trivial scales.

The most common proposal to meet the needs of this space are wireless mesh networks. We argue that traditional mesh networks face an inherent tension between their ability to fulfill the first three facets of our definition, which they can do at small scale, and the last, which they can only do by compromising on one (or more) of the first three. As a result, we do not believe that current proposals for such networks constitute an effective countermeasure to Internet censorships or blackouts. We emphasize that we do not aim to dismiss wireless mesh networks out of hand, but instead focus our criticism on common assumptions in proposed mesh-based solutions and present design-level approaches for getting around

these shortcomings.

The core contributions of this work are a taxonomy of wireless mesh networks and their relationship to dissent networking, and a set of requirements for effective countermeasures to blackout circumvention.

2 Related Work

Many systems for blackout circumvention have been proposed recently. The Commotion Wireless project [24] is building a customized firmware to enable WiFi access points and other devices to form mesh networks, with a focus on ease-of-deployment. Serval has developed a WiFi mesh mobile telephony system [15]. The Free Networking Foundation [4] aims to support the development of community-owned censorship-resistant networks. These projects all leverage WiFi-based mesh networks to varying degrees and each carry the explicit goal of building censorship-resistant networks.

Beyond these projects that aim to build independent network infrastructure, several others focus on circumventing other forms of Internet censorship. Tor is an overlay network for secure and anonymous communications on the Internet using a peer-to-peer network of onion routers [10]. Ultrasurf [5] and Freegate [1] likewise enable secure and anonymous Internet access, though these rely on centralized proxy servers. VPNs and proxy servers are also commonly used to bypass censorship. While these projects fill a similar need to the one we discuss in this paper, they all assume the existence of some underlying form of connectivity and thus provide no resistance to blackouts.

3 What is a mesh network?

The basic idea of a wireless mesh network is relatively universal: multiple devices ("nodes") each communicate directly with their neighbors, and messages from one node to another are forwarded through the mesh via intermediate nodes. This contrasts with "infrastructure" wireless networks, such as cellular phone networks, where client devices (e.g., cell phone) communicate with a master device (e.g., a cell phone tower), which is connected via a separate, independent link (often a wired one) to the rest of the provider's network. While "infrastructure" networks are best thought of as a hierarchical tree, mesh networks are often thought of as a well-connected graph. Akyildiz et al. [7] provide an overview of the space.

Beyond this basic definition, however, mesh networks can take a variety of forms. Consider the following two definitions of mesh networking:

> [A] mesh wireless network offers the ability of users to connect directly to each other and facilitate a distributed network infrastructure that provides multiple paths for communication to the network and does not require a centrally-located towers [. . .]They can bypass obstacles, [. . .]have no single point of failure, and are easily expandable.
>
> (Commotion Wireless, user of Serval)

> Mesh networking [...] creates a self-healing network that is resilient to cable and switch failures. [...] By using Cisco Meraki mesh, organizations can extend the wireless network to areas that are difficult or expensive to connect via Ethernet cabling. (Cisco Meraki)

The first definition highlights the promise of decentralization provided by mesh networks; the second describes mesh networks as an easy way to expand wired enterprise networks (indeed, note a key selling point used by the latter is their *centralized* control platform!). These examples illustrate a range of (sometimes conflicting) attributes that characterize mesh networks. In general, mesh networks fall across a design space defined by three main tradeoffs.

Planned vs. Organic growth. The growth of a planned mesh network is intentionally designed and laid out. Such a network may use antennas which require careful alignment, implement strict policies regarding which devices can and cannot join the mesh, or rely on careful management of radio spectrum use. In contrast, organically-grown mesh networks grow without a particular goal for their topology without needing to coordinate the placement of new nodes. These networks typically utilize non-directional, low-gain antennas and rely on automated routing protocols to allow the mesh network to grow without explicit human involvement.

Centralized vs. Decentralized management. The human organization that operates a mesh network can be centralized or decentralized. In an organizationally-centralized network, a single person or group is responsible for a network's operation and management. In a decentralized network, multiple independent groups cooperate in some way to build a single mesh network, and no one entity has control over an entire network.

Static vs. Mobile topology. In a static mesh network, the mesh nodes are fixed and immobile. Typically, static networks utilize dedicated wireless routers as mesh nodes. Mobile mesh networks use mobile devices as mesh nodes, such as smartphones or laptops. In general, the dynamically changing conditions of a mobile mesh network make them harder to manage than a static mesh network. Note that we refer to the mobility of the *infrastructure* from which a network is built; the fact that a mobile device can connect to a network

Project	Characteristics
Freifunk [2]	Planned, Centralized, Static
Meraki [3]	Organic, Centralized, Static
Serval [15]	Organic, Decentralized, Mobile
Freedom Tower [4]	Organic, Decentralized, Static

Table 1: Example mesh networking projects and their space in our taxonomy.

(such as a smartphone using a WiFi access point) does not make the network a mobile one.

Table 1 shows how various major mesh networking systems fit into this taxonomy. This taxonomy highlights a key tension for projects that wish to use mesh networks to overcome censorship. To successfully resist communications blackouts, a networking technology should grow organically, be mobile, and employ decentralized management—widely available radio direction finding equipment can identify the location of mesh nodes, and any centralized management system represents a single point of failure for the whole network. Unfortunately, as we'll see in §4, building mesh networks that scale and function efficiently is challenging without being planned, static, and centrally managed. The incompatibility of these two goals places serious constraints on the viability of wireless mesh networks as an effective blackout circumvention tool.

4 Scaling Mesh Networks

The capacity scaling of wireless mesh networks has been well-studied in the literature. Gupta and Kumar's foundational result [19] proved that the per-node capacity of a multihop wireless network approaches zero as the number of nodes increases. Li et al. provided experimental validation of this result for 802.11-based networks [21]. This point bears repeating—under reasonable and practical assumptions, the capacity of a mesh network provably tends to zero as it grows. Both of these results, however, are primarily theoretical, and make strong assumptions about properties of the network such as link rates, external interference, coverage radius, and node layout. While we emphasize these results are nonetheless quite general (the Gupta/Kumar result, for example, holds for arbitrary networks), an intuitive understanding of how and why mesh networks scale is useful for practical situations.

4.1 Capacity of Mesh Networks

Channel contention is the primary factor that prevents per-node capacity in mesh networks from scaling. Mesh nodes carry traffic on behalf of other nodes in the network; critically, each node can transmit and receive from multiple other nodes. Mesh networks typically use omnidirectional antennas ("omnis") to support communication regardless of the relative orientation of nodes. Antennas are passive devices that concentrate RF energy; omnis have radiation patterns resembling spheres or disks. Other radiation patterns are possible using directional antennas, but again these can only focus a node's energy over a smaller area; these are less useful for mesh networks since they limit the degree of each node. Using omnis is a design decision to prioritize unplanned deployment over efficiency: most of the energy transmitted by each node is wasted by being radiated away from the recipient.

Yet poor efficiency is not the real problem. Note that the radiation pattern of an antenna applies to both what it transmits and what it receives, and rather than just two nodes, consider a regular lattice of nodes that is evenly spaced. For simplicity, we assume that each node has a fixed radius over which it can successfully transmit and receive messages, and that nodes are spaced by less than this radius.[1] When node A transmits to node B, none of A's neighboring nodes can receive any transmissions due to collisions. To these nodes, A acts as source of interference at node A's location, no different than government jamming equipment. This highlights a key property of nodes with omnis—not only do they cause interference in all directions when they transmit, they are susceptible to interference from *any* direction. If nodes use a carrier-sense MAC protocol such as 802.11, the problem is more insidious—even if one of A's neighbors wanted to transmit to a node outside A's transmission radius, it must wait until A's transmission ended.

The problem is further compounded by the fact that most commonly available radio equipment used for mesh networks has only a single transceiver, which is a half-duplex device. While multi-radio equipment is available today, laptops, mobile phones,[2] and consumer-grade access points rarely have more than one. The multiradio equipment that is available is specifically designed for mesh networks; we argue that such purpose-built hardware makes targeting dissidents easier (§5.1). Finally, although we assume that nodes themselves generate the network's traffic, nodes can also serve as access points for devices like laptops or phones; this (suboptimal) design can lead to increased contention.

Channel contention carries two implications. First, mesh networks have poor performance because of time wasted waiting for opportunities to send traffic. Second,

[1] This is essentially the model used by Li et al., though here we assume the transmission and reception radius are equal. Of course, real-world RF behavior is much more complex.

[2] While phones and laptops often *do* have multiple radios (e.g., WiFi, Bluetooth, and cellular) typically only the WiFi radio is used for mesh due to support for "ad-hoc mode" and legal constraints.

mesh networks have highly variable performance [6] since the scale of contention varies significantly based on workload (along with environmental factors that affect radio propagation).

4.2 Application Support

Despite these challenges, meshes can provide a useful degree of service—*if* applications running on them can tolerate their unique shortcomings. For example, so-called "smart meters" use mesh networks to report customers' usage to their utility companies; messages are forwarded across the network to "gateway" nodes connected to the Internet. This is an application particularly well-suited for mesh networks. First, it is highly *delay tolerant*—as long as the utility company receives its billing data within a few minutes or even hours the data is still useful. Secondly, it requires *little bandwidth*—even with low absolute efficiency the mesh is still able to meet the application's performance requirements.

In contrast, web traffic is a workload that performs poorly in these high-contention environments. Consider the basic task of a user sending a TCP request and receiving a response over a multi-hop mesh network as depicted in Fig. 1. We assume "oracle routing" that determines the optimal path for all traffic, though doing so in practice is challenging. TCP requires the network to send bidirectional traffic: packets from A to B will generate acknowledgements upon receipt. This is a problem for typical mesh networks that use single-radio nodes which share the same channel. When A in Fig. 1 transmits to B, each packet must be received by nodes 1-4 and can only be re-transmitted when the channel is free, halving effective bandwidth at each point. Each time a node in the path transmits a packet, none of its neighboring nodes may transmit or receive, lest they create collisions. Synchronizing transmissions is a challenging problem; WiFi-based networks utilize a mechanism known as RTS/CTS to announce their intention to transmit. While this mechanism reduces collisions, it increases the amount of time the channel is idle: each RTS/CTS exchange between nodes requires at least two transmissions before sending actual data. Every ACK that B sends back to A undergoes the same process. The end result is that each wireless hop substantially decreases effective bandwidth and increases latency and loss, even in this simple case. Multiple pairs of communicating nodes exacerbate the problem.

Mobile nodes enable capacity to scale linearly under certain assumptions [18] but introduce new opportunities for loss and delay (e.g. nodes not being in range of each other). Highly variable latency and loss due to collisions are standard conditions in a mesh network, and since these violate assumptions of TCP congestion control mesh networks tend to be ill-suited for TCP-based applications. Mesh networks present a challenging environment for voice traffic (which requires low jitter) for similar reasons. Directional antennas also improve capacity [27], though networks using them topologically resemble non-mesh networks and rely on purpose-built hardware, a problem addressed in §5.1.

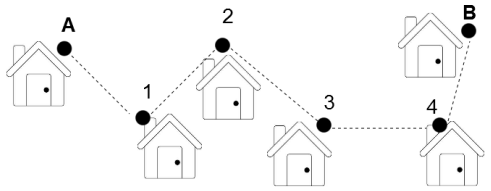

Figure 1: Multihop communication across a mesh network.

This discussion suggests two strategies for building mesh networks that scale. The first is to reduce the degree of channel contention in the mesh network by carefully planning how nodes can interfere with each other and where new nodes are added to the network. Such a network provides a high level of service, but wouldn't be a "dissent network". The second strategy is to accept the limitations of mesh networks and build applications that can work under those regimes. For example, applications that leverage delay-tolerant networking [12] principles can cope with such limitations [16, 20], as can very low bandwidth applications. Such applications could prove quite useful for dissent networks [13].

We finally note that our discussion ignores several key unsolved problems in scaling wireless mesh networks. Most notably, routing across ad-hoc mesh networks continues to be an area of active research and engineering effort. We've chosen to ignore this for two reasons. First, this paper focuses on real-world networks in real-world environments. Few mesh routing protocols have seen the level of sustained development and testing necessary to fairly judge their ability to function in such environments. Second, and more importantly, our criticism of mesh networks for blackout circumvention is an *architectural* one, and is orthogonal to the routing protocol used. Even with a "perfect" routing protocol, mesh networks cannot overcome the fundamental physics of radio from which their scaling properties derive.

5 Supporting Dissent

Our objective here, of course, is not simply to tear down wireless mesh networks. There are several examples of mesh networks that have scaled well and serve large numbers of users, such as Freifunk or the Athens Wireless Network. Yet the bar for dissent networking

is higher—such networks will be used in environments where even the act of using such a network puts the user at risk. Centralized and planned networks can't work in this environment, as they have a single point of failure, and static mesh nodes are easy targets for a government with even the most basic electronic surveillance equipment. Not only can a repressive government shut down a network by attacking the technology itself, it can also attack the organization and people behind it.

The goal of work in this space is to promote freedom of expression under oppressive regimes—in short, to support political dissent. At their core, censorship and suppression of communication are *non-technical* problems; while technical solutions may alleviate their direct impacts, the root issue is one of unjust governance. Technology doesn't produce political movements. A key idea from the technology for international development literature states that technology only amplifies human intent [26]. Put differently, technology plays a *multiplicative* role, not an additive one. Moreover, technology amplifies both positive and negative intentions [22]. Any anti-censorship tool can thus only build upon existing social movements and simultaneously carries the potential to amplify the efforts of repressive regimes (e.g., by providing another mechanism to track dissident activity).

This presents a pair of related challenges to dissent-oriented projects. First, such projects should leverage existing social trust networks. Doing so simultaneously builds upon pre-existing social infrastructure while using that infrastructure to reduce risk to users. Secondly, such projects should minimize the extent to which the systems they are developing could be used for harm. We emphasize two particular elements of this second challenge—the need to use "innocuous" hardware that doesn't raise suspicion and the need to provide anonymity (not pseudonymity) guarantees to users.

5.1 Innocuous Hardware

Projects that propose illegal or restricted hardware face challenges for sourcing equipment, may put activists and users at increased risk, and provide an easy excuse for government crackdowns. Import restrictions on radio equipment are enforced the world over; attempts to smuggle such gear could end not only in confiscation of the equipment, but even arrests and severe punishment. Exceptions are typically made for WiFi devices and similar gear that operates on unlicensed spectrum, yet some countries have not deregulated the use of such spectrum. While customs tend to overlook importation of innocuous equipment like laptops and smartphones, they reserve the right to confiscate equipment and harass citizens who attempt such import when the regime feels it might be used "inappropriately".

Additionally, using licensed spectrum makes it easy for a regime to identify and locate who is using it (particularly if the user has applied for a license). A network using such spectrum without a license provides an easy excuse for a government to terminate operation and dole out punishment to the network's operators. Moreover, illegally using licensed spectrum carries the risk of disrupting operations of legitimate license holders, who have an incentive to report such activity to authorities. This is a plausible outcome for activists who, for example, set up unlicensed cellular systems (even low-power ones) as has been proposed by some groups [24].

In countries which allow use of deregulated spectrum, setting up rooftop WiFi antennas may not be outright illegal, yet may be a cause for government scrutiny and harassment. Due to the conspicuous nature of such equipment, its existence can raise suspicion from neighbours or agents and collaborators of the regime. The governments of such countries incentivize citizens to report any suspicious dissent activity and thereby turning them into informants and collaborators. Fear of persecution therefore renders widespread adoption of such rooftop technology unlikely, even if the actual gear and spectrum use is within the boundaries of the law.[3]

Given the above, we believe the use of unconventional, purpose-built, or otherwise uncommon equipment is likely to be shut down quickly, limiting the impact of projects using such hardware. Worse, such equipment could put users at risk or aid an oppressive regime in tracking users. We conclude that the only devices that can be used in a meaningful dissent network are innocuous, ubiquitous devices such as smartphones and preexisting indoor WiFi access points.

5.2 The Need for Anonymity

Adversaries can analyze communication content and patterns to identify dissidents. While tools like encryption ensure communication security, dissent networks also require *privacy*. While security protects communications from eavesdroppers, privacy aims to limit the information revealed by legitimate communications. Such communications may involve malevolent agents, so it is critical to avoid leaking condemning information. We wish to prevent persecution of individuals based on their involvement in such a network.

We contend that the only truly safe solution to this problem is anonymity. Users should ideally be unlinkable with their true names, but this may be impossible in practice due to surveillance. A potentially sufficient alternative is deniability—users should be able to make a plausible case for innocence. This requires two levels

[3] One of the authors has been arrested and interrogated on suspicion of espionage for installing perfectly legal rooftop WiFi antennas.

of protection: 1) The network should be useful for non-dissent purposes, so usage is not incriminating in and of itself. 2) Activity on the network should be impossible or difficult to trace. To achieve the latter point, systems may need to satisfy a number of notions of anonymity:
Author anonymity: It is impossible to link a message with its author.
Reader anonymity: It is impossible to link a document with its readers.
Document anonymity: Servers do not know which documents they are storing.
Query anonymity: A server does not know what client request it is filling.

These varieties of anonymity are defined in [9], and many dissent networking solutions address subsets thereof via pseudonymity—i.e., users are associated with network identities disjoint from their true identities. However, pseudonymity is not safe enough for dissent networking, since attributing profile information to individuals facilitates identification. It has been shown repeatedly that personal information in social networks can be correlated with external information to deanonymize users [17]. Pseudonymity can also be implicit, enabling similar threats. For instance, fixed-infrastructure networks can lead to localization and deanonymization of users [14, 25]. In the allegedly anonymous Bitcoin network, researchers learned information about individual users by observing transaction patterns [23]. Decentralized mesh networks are more robust to traffic analysis because interaction records are difficult to trace, so the main concern is avoiding explicit pseudonymity.

Ideally, a dissent networking solution should have all the above anonymity properties, but the network must still function as a communication tool. Theoretical results from other domains have demonstrated fundamental tradeoffs between privacy and system utility [11], suggesting that similar tradeoffs may exist for communication networks. For instance, some networks rely on user trust graphs (useful for deanonymization) to defend against sybil attacks. Indeed, adversaries enjoy advantages on truly anonymous networks due to such networks potential lack of reputation or accountability, such as the ability to send false messages, impersonate other users, or execute sybil attacks. The relation between privacy and communication efficiency is an important research question, but strong privacy should nonetheless remain a conscious design goal for this space.

6 Moving Forward

This work takes a critical view of proposed blackout circumvention systems; we acknowledge we offer few explicit solutions. We nonetheless believe that there is good work to be done in this space.

Dissent-oriented mesh networks can improve by leveraging mobility, directional antennas, and limitation-tolerant applications, while providing strong anonymity. There are several examples of work that partially meets these requirements for a successful dissent network. For instance, the Dissent and TOR projects incorporate notions of deniability and anonymity into the system functionality [8, 10]. Projects like Commotion and Serval exploit mobility and delay-tolerance in a mobile mesh setting, while avoiding exotic hardware [24, 15]. Ideally, systems should aim to address *all* the requirements; this is attempted in [13], though the practical scalability of such a solution is yet unproven. Along these lines, we hope the community will consider "communications" broadly while pushing to build workable dissent networks. Though mesh networks will likely encounter scalability problems for applications like telephony or even point-to-point communications, other models (e.g. one-to-many communication) have yet to be explored.

7 Conclusion

Developing effective countermeasures to communications blackouts involves requirements beyond what most existing projects have set out to meet. Mesh networks, the most commonly proposed solution, suffer a fundamental tension between scale and safety for use under a repressive regime. Such networks can reach meaningful scales by adopting centralized management, planned growth, and a static topology, making them more susceptible to government interference. Networks can retain a decentralized nature at the cost of lower quality of service, requiring applications tailored to their limitations.

More than this, we feel that prior work has not paid enough attention to the fact that building alternative network infrastructure is itself a subversive act. Those who build such systems do so with the full awareness that the design choices they make can have grave consequences for their users. At the same time, given the public resources that have been directed to this space in lieu of other forms of support for promoting dissent, these blackout circumvention systems should be able to scale to meaningful sizes—beyond just demonstration deployments. We believe that our definition of dissent networking captures these two goals, and that projects that attempt to meet our definition will produce more effective countermeasures to communications blackouts.

8 Acknowledgements

We thank Kurtis Heimerl for helpful discussions and our shepherd Rodger Dingledine and the anonymous reviewers for their insightful feedback. This work was supported in part by an NSF Graduate Research Fellowship.

References

[1] Freegate. www.dit-inc.us/freegate.

[2] Freifunk Wireless Network. http://start.freifunk.net/.

[3] Meraki. http://www.meraki.com/.

[4] The Free Networking Foundation. http://thefnf.org.

[5] Ultrasurf. https://ultrasurf.us/.

[6] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-Level Measurements from an 802.11b Mesh Network. *ACM SIGCOMM Computer Communication Review*, 34(4):121–132, 2004.

[7] I. F. Akyildiz, X. Wang, and W. Wang. Wireless Mesh Networks: A Survey. *Computer Networks*, 47(4):445–487, 2005.

[8] H. Corrigan-Gibbs and B. Ford. DISSENT: Accountable Anonymous Group Messaging. In *Proceedings of the 17th ACM Conference on Computer and Communications Security*, pages 340–350. ACM, 2010.

[9] R. Dingledine, M. J. Freedman, and D. Molnar. The Free Haven Project: Distributed Anonymous Storage Service. In *Designing Privacy Enhancing Technologies*, pages 67–95. Springer, 2001.

[10] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The Second-Generation Onion Router. Technical report, DTIC Document, 2004.

[11] J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Privacy aware learning. *arXiv preprint arXiv:1210.2085*, 2012.

[12] K. Fall. A Delay-Tolerant Network Architecture for Challenged Internets. In *Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, pages 27–34. ACM, 2003.

[13] G. Fanti, Y. B. David, S. Benthall, E. Brewer, and S. Shenker. Rangzen: Circumventing Government-Imposed Communication Blackouts. Technical Report UCB/EECS-2013-128, EECS Department, University of California, Berkeley, Jul 2013.

[14] J. Freudiger. *When Whereabouts is No Longer Thereabouts: Location Privacy in Wireless Networks*. PhD thesis, École Polytechnique Fédérale de Lausanne, 2011.

[15] P. Gardner-Stephen. The Serval Project: Practical Wireless Ad-hoc Mobile Telecommunications, 2011.

[16] M. Garetto, P. Giaccone, and E. Leonardi. Capacity Scaling in Delay Tolerant Networks with Heterogeneous Mobile Nodes. In *Proceedings of the 8th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 41–50. ACM, 2007.

[17] O. Goga, H. Lei, S. Parthasarathi, G. Friedland, R. Sommer, and R. Teixeira. Exploiting Innocuous Activity for Correlating Users Across Sites. In *World Wide Web Conference (WWW)*, May 2013.

[18] M. Grossglauser and D. N. Tse. Mobility Increases the Capacity of Ad Hoc Wireless Networks. *IEEE/ACM Transactions On Networking*, 10(4):477–486, 2002.

[19] P. Gupta and P. R. Kumar. The Capacity of Wireless Networks. *IEEE Transactions on Information Theory*, 46(2):388–404, 2000.

[20] U. Lee, S. Y. Oh, K.-W. Lee, and M. Gerla. Scaling Properties of Delay Tolerant Networks with Correlated Motion Patterns. In *Proceedings of the 4th ACM Workshop on Challenged Networks*, pages 19–26. ACM, 2009.

[21] J. Li, C. Blake, D. S. De Couto, H. I. Lee, and R. Morris. Capacity of Ad Hoc Wireless Networks. In *Proceedings of the 7th Annual International Conference on Mobile Computing and Networking*, pages 61–69. ACM, 2001.

[22] E. Morozov. *The Net Delusion: The Dark Side of Internet Freedom*. PublicAffairs, 2012.

[23] F. Reid and M. Harrigan. An Analysis of Anonymity in the Bitcoin System. In *Security and Privacy in Social Networks*, pages 197–223. Springer, 2013.

[24] A. Reynolds, J. King, S. Meinrath, and T. Gideon. The Commotion Wireless project. In *Proceedings of the 6th ACM Workshop on Challenged Networks*, pages 1–2. ACM, 2011.

[25] M. Srivatsa and M. Hicks. Deanonymizing mobility traces: Using social network as a side-channel. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, pages 628–637. ACM, 2012.

[26] K. Toyama. Technology as Amplifier in International Development. In *Proceedings of the 2011 iConference*, pages 75–82. ACM, 2011.

[27] S. Yi, Y. Pei, and S. Kalyanaraman. On the Capacity Improvement of Ad Hoc Wireless Networks using Directional Antennas. In *Proceedings of the 4th ACM International Symposium on Mobile Ad Hoc Networking & Computing*, pages 108–116. ACM, 2003.