Designing Networks for Large-Scale Blackout Circumvention

Shaddi Hasan

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2013-230
http://www.eecs.berkeley.edu/Pubs/TechRpts/2013/EECS-2013-230.html

December 19, 2013

Acknowledgement

This thesis is based in large part on work I conducted jointly with Yahel Ben-David, Giulia Fanti, Scott Shenker, and my advisor, Eric Brewer; I am deeply grateful for their contributions and feedback. I thank Kurtis Heimerl, Paul Pearce, Kashif Ali, Javier Rosa, and Matt Podolsky for helpful discussions on this topic, as well as Roger Dingledine for his insightful feedback on an earlier version of this work. I also thank members of the mesh networking community for helpful discussions and their candor in speaking with me about their projects, particularly the folks at Commotion Wireless and Isaac Wilder of the Free Network Foundation. I have great respect for their dedication to the cause of Internet freedom and for the inspiration their projects have given to people worldwide.

Designing Networks for Large-Scale Blackout Circumvention

by Shaddi Hasan

Research Project

Submitted to the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, in partial satisfaction of the requirements for the degree of **Master of Science, Plan II**.

Approval for the Report and Comprehensive Examination:

Committee:

Professor Eric Brewer
Research Advisor

(Date)

* * * * * * *

Professor Scott Shenker
Second Reader

(Date)

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 safety of both users and network operators as a first-order design priority. From these two insights, we frame a definition of *dissent networks* to capture the essential requirements for blackout circumvention solutions.

Acknowledgements

This thesis is based in large part on work I conducted jointly with Yahel Ben-David, Giulia Fanti, Scott Shenker, and my advisor, Eric Brewer; I am deeply grateful for their contributions and feedback. I thank Kurtis Heimerl, Paul Pearce, Kashif Ali, Javier Rosa, and Matt Podolsky for helpful discussions on this topic, as well as Roger Dingledine for his insightful feedback on an earlier version of this work. I also thank members of the mesh networking community for helpful discussions and their candor in speaking with me about their projects, particularly the folks at Commotion Wireless and Isaac Wilder of the Free Network Foundation. I have great respect for their dedication to the cause of Internet freedom and for the inspiration their projects have given to people worldwide.

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 that 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.

This is a worthy goal. However, we believe that much of the work in this space suffers a disconnect from reality that stem from a lack of a clearly defined set of properties for such networks. We further believe that having a crisp definition and set of desired properties for this type of network will help distill the problems with current proposals. To this end, we propose and define "dissent networking", discuss the desired properties, and address the suitability of proposed technologies and solutions. 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.
- *Safe for network operators and users in dissent scenarios:* Should be minimally dangerous to build, operate, and use.
- *Able to run at meaningful scales:* Should be more effective at disseminating information than people with megaphones; more broadly, should be able to run at non-trivial scales.

Achieving any of these properties is difficult, and solving all three with the same system represents a "grand challenge" for the censorship circumvention community. A true dissent network would fundamentally change the balance of

power between repressive regimes and dissidents in terms of access to communications.

As yet, no such system exists. One enduringly popular concept within this space is that of the "wireless mesh network". For some, mesh networks have become something of a panacea for censorship circumvention, oppressive governments, corrupt incumbent telecom providers, and all manner of other bad actors standing in the way of freedom of expression. Mesh networks promise to build decentralized, ubiquitous, resilient, and community-owned network infrastructure. This vision is admittedly compelling, particularly for those working against Internet censorship. Proponents of mesh networks claim that such a network would have enabled revolutionaries in Egypt to continue communicating with each other and the outside world even after the Mubarak regime took the draconian step of shutting down all Internet infrastructure in the nation.

As always, if it sounds too good to be true, it probably is. We argue that traditional mesh networks face an inherent tension between their ability to fulfill the first two 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 two. 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. This is the core contribution of this work. We present a taxonomy of existing wireless mesh networks in the context of dissent networks, and a set of requirements for effective countermeasures to communications blackouts. Through a critique of existing, but flawed, proposed "dissent networks", we can identify design objectives for future systems.

2 Related Work

Many systems for blackout circumvention have been proposed recently.

The Commotion Wireless project [50] 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 [28]. The Free Networking Foundation [8] aims to support the development of community-owned censorship-resistant networks. Rangzen [25] is a privacy-preserving mobile mesh network that leverages social ties for making routing de-

Network	Location	Size (nodes)
Guifi.net	Spain	~20,000
Athens Wireless Metropolitan Network	Greece	2396
Vienna Funkfeuer	Austria	~750
Wlan Slovenija	Slovenia	304
Hamburg Friefunk	Germany	300

Table 1: Some notable operational "mesh" networks. Network size is self-reported by each network, as of November 2013. Note that not all nodes in a network may be in the same mesh "cloud"; some groups of nodes may be isolated.

cisions. These projects all leverage WiFi-based mesh networks to varying degrees and each carry the explicit goal of building censorship-resistant networks.

Several operational "mesh" networks also exist. Guifi.net in Spain, the Athens Wireless Network in Greece, Freifunk in Germany, and Funkfeuer in Austria are examples of large networks. Not all of these networks use mesh routing protocols—the Guifi network uses a combination of OSPF and BGP, for example. Many smaller networks also exist, such as the Kansas City Freedom Network [5] and the Red Hook Wifi Network [40] (operated in partnership with the Free Networking Foundation and Commotion Wireless, respectively). Table 1 highlights a few particularly large mesh networks; this list is by no means exhaustive, of course.

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 [22]. Ultrasurf [9] and Freegate [2] 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. Hyperboria [4] is an overlay network consisting of nodes running CJDNS [1], essentially a distributed virtual private network. Although 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.

Finally, mesh networks and privacy have received an extensive treatment in the literature. Akyildiz et al. [13] provide an overview of the space. Wu et al. [54] consider privacy in mesh networks, specifically towards providing confidentiality

of traffic content and patterns of communication. Zhou et al. [56] consider the threats faced by mesh networks and offers solutions for protecting against denial of service attacks and establishing a reliable PKI in an ad hoc network. A variety of routing protocols for mesh networks have been proposed; among the most popular in deployed networks include OLSR [10, 16], AODV [43], and Babel [15].

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. Although "infrastructure" networks are best thought of as a hierarchical tree, mesh networks are often thought of as a well-connected graph. The literature on mesh networks (also known as adhoc wireless networks) is extensive; a survey of the many variations on the basic theme described above that have been proposed is beyond the scope of this work. Akyildiz et al. [13] provide an overview of the space, to which we direct interested readers.

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

> Mesh networks afford an alternative to [the] centralized "hub-and-spoke" WLAN model: rather than relying on the ISP for Internet connectivity, mesh technologies can produce ad hoc networks that allow distributed nodes to act as the senders, receivers, and conduits of information. In the mesh model of networking, "each user has the capability to receive and send information and to relay information on behalf of other connected computers." (Berkman Center)

> [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)

> A mesh network is one where any device can be connected to one or more other neighbor devices in an unstructured (ad-hoc) manner. Mesh networks are robust and simple to configure because the software determines the routing of data automatically in real-time based on sensing the network topology. Traditional mesh networks are limited in scale because they rely on single radio, wireless-only connections and omni-directional antennas. By using directed wireless links and wired transfers whenever possible, the Fabfi system is optimized for building very large-scale static (as opposed to mobile) mesh networks. (FabFi)

> 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 that 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.

Stationary vs. Mobile topology. In a stationary mesh network, the mesh nodes are fixed and immobile. Typically, stationary networks utilize dedicated

Project	Characteristics
Freifunk [3]	Planned, Centralized, Static
Meraki [7]	Organic, Centralized, Static
Serval [28]	Organic, Decentralized, Mobile
Freedom Tower [8]	Organic, Decentralized, Static
Guifi [39]	Planned, Decentralized, Static

Table 2: Example mesh networking projects and their space in our taxonomy.

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

Table 2 shows how various major mesh networking systems fit into this taxonomy. There's no "right" way to build a mesh network; as Table 2 shows there are examples of systems that choose a wide variety of points within this design space. Yet 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 Section 4, building mesh networks that scale and function efficiently is challenging without being planned, stationary, 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 [33] 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 [36]. This point bears repeating—under reasonable and prac-

tical 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. Although 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; channel contention is the factor that prevents mesh networks from scaling. 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, as in Figure 1. 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

[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; e.g., nodes can receive (both signal and noise) from nodes further away than they can transmit, and the radiation pattern of an omnidirectional antenna is not a perfect sphere.

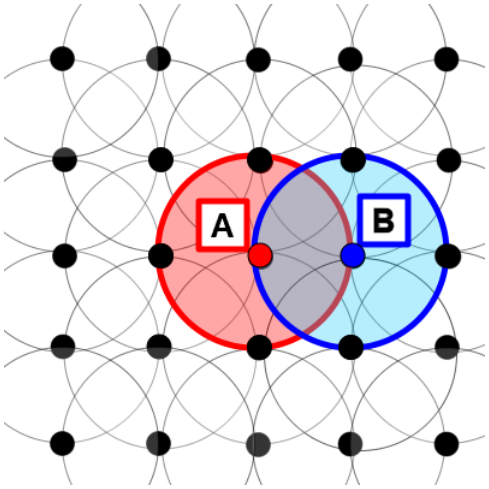

Figure 1: Regular lattice of mesh nodes with omnidirectional antennas. Each node is able to transmit to and receive from each of its neighbors.

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 only has a single half-duplex transceiver. Although 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; not only does obtaining such hardware substantially increase cost and logistical difficulty, we argue that such purpose-built hardware makes targeting dissidents easier (Section 5.1). Finally, we note that we assume nodes themselves generate the network's traffic. In some designs, nodes also serve as access points for client devices such as phones and laptops, presenting an additional source of channel contention. We don't consider this case further as it is a suboptimal design.

Channel contention carries two implications. First, mesh networks suffer a decrease in per-node performance as they grow because of time wasted waiting for opportunities to send traffic. Second, mesh networks have highly variable performance [11] since the scale of contention varies significantly based on workload (along with environmental factors that affect radio propagation).

[2]Although 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.

Approach	Drawback
Directional antennas	Decreased resilience due to fewer redundant paths
Mobile nodes	Significantly increased routing complexity
Multi-radio nodes	Increased per-node costs, specialized hardware
Delay tolerant applications	Poor user experience, requires new applications
Smaller networks	Limits reach of network
Fewer users	Limits reach of network
Network planning	Slows network growth
Centralized management	Introduces single point of failure

Table 3: Drawbacks for various techniques for scaling mesh networks.

4.2 Designing for Scale

Despite these challenges, meshes can provide a useful degree of service—*if* network designers and implementers make careful decisions about how to mitigate channel contention. The techniques listed in Table 3 can each mitigate channel contention in wireless mesh networks. Large mesh networks such as Freifunk and Guifi, for example, rely on directional antennas and partitioning their networks into smaller chunks in order to operate at the scale of thousands of nodes. However, these approaches represent a tradeoff space among reduced channel contention and one or more of the design imperatives for dissent networks: resilience, scalability, use of common components, and resistance to tracking.

Application Support. One solution is to specifically design applications that will run on mesh networks to 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. Delay tolerance is particularly helpful in an environment with variable contention (and hence delay). Secondly, it requires *little bandwidth*—even with low absolute efficiency the mesh is still able to meet the application's performance requirements. Moreover, smart metering makes economic sense for utility companies since it allows them to stop manually reading meters, the savings from which outweighs the cost of the required mesh infrastructure.

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

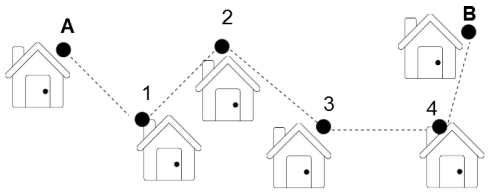

Figure 2: Multihop communication across a mesh network.

and receiving a response over a multi-hop mesh network as depicted in Figure 2. We assume "oracle routing" that determines the optimal path for all traffic, though doing so in practice is challenging and imparts non-trivial overhead onto the network. TCP requires the network to send bidirectional traffic: packets from A to B will generate acknowledgements upon receipt. In a wired Ethernet network, such a task would cause no issues, since acknowledgements over the reverse path would not interfere with the transmission of packets in the forward path. Wireless mesh networks present two key challenges: 1) nodes are half-duplex and 2) the wireless channel is shared by all nodes. The bidirectional nature of TCP is a problem for mesh networks that use single-radio nodes that share the same channel (or even a limited set of channels). When A in Figure 2 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. Although 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, further increasing contention for the channel. 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.

Thus, while innovation at the application layer can mitigate the poor performance of mesh networks, it comes at a high cost—namely, losing the ability to leverage existing applications and Internet-based tools like Twitter and Facebook. Bootstrapping new communications tools is difficult given the inherent chicken-and-egg problem of attracting users to such a service (without other users to interact with, the service is of little value). Existing Internet-based tools are already widely used, avoiding this issue, and have the benefit of capturing community

structure, allowing dissidents to take advantage of their existing organizational and social connections. Developing the technology for new applications, then, is only part of the problem.

Mobility and Directionality. Mobile nodes enable capacity to scale linearly under certain assumptions [32] 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. In mesh networks, highly variable latency and loss due to collisions is a standard operating condition. Mesh networks present a challenging environment for voice traffic (which requires low jitter) for similar reasons. Alternatively, directional antennas can also solve this problem [55], though using such antennas hurts the resilience of the mesh network since nodes are able to communicate with fewer of their neighbors and limits mobility due to reduced or uneven coverage. Also, while directional antennas aren't necessarily rare or expensive, they are purpose-built hardware, a problem we return to in section 5.1.

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 [24] principles can cope with such limitations [29, 35], as can very low bandwidth applications.

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 thesis 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 Guifi.net, 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 stationary 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 its core, censorship is a *non-technical* problem; while technical solutions may alleviate its direct impact, 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 [52]. Put differently, technology plays a *multiplicative* role, not an additive one. Moreover, technology amplifies both positive and negative intentions [37]. 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 Keeping Operators Safe

Although any act of censorship circumvention is risky, dissent networks should not present undue risks to network operators. The main role of network operators in dissent networks is overseeing the construction and maintenance of network

Location	Date	Duration
Egypt [19, 20]	Jan.-Feb. 2011	5 days
Libya [20]	19 Feb. 2011	7 hours
Libya [20]	20 Feb. 2011	7 hours
Libya [31]	March - July 2011	171 days
Syria [49]	June 2011	1 day
Syria [48]	Dec. 2012	1 day
Syria (partial)	Jan 2013	1 day
Syria [47]	May 2013	1 day
Burma [46]	Aug. 2013	1 hour
Sudan [45]	Sept. 2013	1 day

Table 4: Large-scale Internet blackouts and their durations since 2011.

infrastructure. Given the challenges faced by mesh networks outlined in Section 4, one might consider using more exotic technologies to build dissent networks, such as satellite terminals. Despite their practical limitations, wireless mesh networks such as those proposed by Commotion have one key benefit: in general, they can be built using commodity equipment that is widely available. This property is of course advantageous from the perspective of ease of deployment, but we argue that it is also vital to ensure user safety.

Designers of would-be dissent networks must consider how their users will source the equipment needed to build their system. Recall that the intended goal of dissent networks is to allow activists to build alternative infrastructure networks that can survive government shutdowns. With the notable exception of the Libyan shutdown of 2011[3], nation-scale Internet blackouts to date have been of relatively limited duration. For example, Egypt's 2011 blackout lasted three days. As of this writing the most recent similar event was carried out by Syria in May 2013; it lasted only a matter of hours preceding a significant government military operation against rebel forces [17]. While there's no guarantee that future large-scale blackout events will be as short, the massive collateral damage they inflict provides a strong incentive for government actors to keep them as short as possible.

Thus, designers of effective dissent networks need to design around having requisite equipment pre-staged, if not actually operational, before any blackout

[3]This blackout coincided with a full-scale civil war, and Libya had relatively low Internet penetration to begin with (7% at the time of the blackout). These factors could in part explain the unusually long duration of this blackout.

event occurs; building an alternative communications system in a matter of hours or even days in the type of chaotic environment that has surrounded previous blackout events is a tall order for dissidents. This in turn implies that activists must have the necessary equipment in operation in the days and weeks leading up to a blackout event for their dissent network to be effective. As a result, necessary hardware should be safe for an activist to posses for an extended period of time. 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.

One such challenge is that of getting equipment into a country. Import restrictions on radio equipment are common and strongly enforced worldwide; most countries require radio equipment to undergo an approval process to ensure compliance with radio regulations. Although smuggling equipment can be a viable option for small volumes of equipment, being caught could lead not only to confiscation of the equipment, but even fines, arrests, or other severe punishments. This suggests that dissent networks should limit their equipment needs to that which is readily available in whatever countries the network intends to operate.

Another challenge is choice of spectrum bands used by the network's equipment. Operating on licensed bands enables a regime to more easily identify and locate dissidents' radio equipment. Legal use of such spectrum by definition requires registering with a government authority. Operating in these bands without a license is risky: doing so 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 likely outcome for activists who, for example, set up unlicensed GSM[4] cellular networks (even low-power ones) as has been proposed by groups such as the Commotion Wireless project [50]. Even commodity WiFi devices and similar equipment that operate in unlicensed spectrum may be illegal depending on the country in which they are being used—regulations on WiFi, for example, vary widely by country and some maintain regulations limiting WiFi use to indoor areas, limited frequency bands, or low power levels [6].

Unfortunately, even when equipment can be legally operated in a country, dissidents won't necessarily be able to do so free of risk. In countries which allow use of deregulated spectrum, setting up rooftop WiFi antennas may not be out-

[4]More seriously, GSM networks have well-known [38] security flaws that allow adversaries to eavesdrop on communications and track the identities of those using the network.

right illegal but can be a cause for increased scrutiny and harassment. Due to the conspicuous nature of such equipment, its existence can raise suspicion from neighbors or agents and collaborators of the regime. This is a real threat: the founder of AirJaldi, a large wireless network in northern India, was once arrested under suspicion of espionage as a result of setting up rooftop WiFi equipment, despite the fact that both the equipment and the way it was being used was legal and fully compliant with communications regulations in India at the time [14].

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 successful dissent networking designs should only rely on ubiquitous devices that are inconspicuous, legal to operate, and can be plausibly used for non-dissent purposes—in short, *innocuous hardware*. Smartphones and residential WiFi access points easily fall into this category, whereas the software-defined radios required to run popular standalone GSM networking software [41, 42] or satellite Internet terminals do not. Similarly, while small omnidirectional antennas or parabolic dish antennas are unlikely to raise suspicion, sector antennas are unusual enough that they may not be innocuous in some dissent situations.

5.2 Keeping Users Safe

Effective dissident communications networks are a prime target for government adversaries for tracking usage, gathering intelligence, and spreading misinformation. Just as designers of dissent networks need to consider the safety of network operators when making decisions about hardware requirements for their systems, so too should they design their system to avoid jeopardizing the safety of network users. Unfortunately the strategies for keeping users safe in mesh networks tend to be poorly considered or to offer minimal guarantees. We begin with a brief analysis of the user safety properties of two major blackout circumvention tools—Serval and Commotion.

Serval. The Serval project claims to provide a (poorly defined) level of confidentiality and authenticity, but in order to achieve those assumes the existence of a public key infrastructure among nodes in its mobile mesh network [27]. Nodes in Serval are mobile phones that serve a dual role as both infrastructure nodes, relaying traffic for other nodes in the mesh, as well as end host devices. Each node has a public/private keypair; all communication to a node is encrypted with the destination nodes' public key. Because of its decentralized nature, verifying

the identity of nodes is left to an out-of-band process, namely in-person verification of keys. This model is deeply flawed in numerous ways. First, it essentially assumes the existence of a decentralized public key infrastructure for each Serval network in order to provide authentication of communication. Building a PKI is challenging under the best of circumstances, and since Serval falls back to informing the user their communications are untrusted in the absence of such a PKI the practical effectiveness of their scheme is unclear (a close analogue is that of SSL warnings in web browsers, which are frequently ignored by users [12]). Secondly, it provides no method for key rotation—a lost or confiscated mobile device could be used by an adversary to impersonate its owner. Finally, it does not provide perfect forward secrecy. Because Serval relies on potentially untrusted nodes to relay traffic across the network, we argue that providing such a guarantee is vital since it provides captured dissidents a means of deniability.

Commotion. The Commotion project makes similarly vague references to providing "secure" communications but provides no mechanisms for such communications. Although Commotion networks support link-layer encryption, to support organic growth their entire network shares a single key that must be known by all users of the network. It also provides no protection against "rogue" nodes joining the mesh, relying on users to implement their own end-to-end encryption. These unclear guarantees and flawed security models put users at risk. Rather than solving the very difficult problems around providing safety guarantees to users of their networks, the project includes a "warning label" to inform potential users of its confidentiality, integrity, availability, and anonymity limitations. One key issue with their approach, however, is the fact that the contents of the warning label conflict with the project's (perhaps aspirational) marketing as a solution for "secure" communications. Technically unsophisticated users who take such marketing at face value put themselves at risk. Moreover, the warning label appears primarily on the download page for the project's software; if a user receives the software by other means (such as from another activist, or through a mirror) they may not be exposed to the warning label.

Despite these shortcomings, the *need* for communication security—particularly authenticity and confidentiality—seems to be well-understood in the mesh networking community, and while leading systems such as those highlighted above have flawed security models this is an area of active development and research. Unfortunately, protecting user identities is a far less common goal; as far as we are aware the only significant mesh networking project that aims to provide this is Rangzen [25]. Adversaries can analyze communication content and patterns to identify dissidents. Although tools like encryption ensure communication secu-

rity, dissent networks also require *privacy*. While encryption 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, and thus a strategy for protect.

5.2.1 Anonymity

Anonymity is an obvious way to protect user safety: simply hiding users' identities can prevent adversaries from threatening them. Users should ideally be unlinkable with their true names, but this may be impossible in practice due to surveillance. Indeed, a variety of degrees of anonymity exist in the literature, which systems should aim to satisfy precisely:
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 [21], 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 [30]. Pseudonymity can also be implicit, enabling similar threats. For instance, fixed-infrastructure networks can lead to localization and deanonymization of users [26, 51]. In the allegedly anonymous Bitcoin network, researchers learned information about individual users by observing transaction patterns [44]. Decentralized mesh networks are more robust to traffic analysis because interaction records are difficult to trace, so the main concern is avoiding explicit pseudonymity. Theoretical results from other domains have demonstrated fundamental tradeoffs between privacy and system utility [23], suggesting that similar tradeoffs may exist for communication networks.

Anonymity also offers adversaries advantages due to potential lack of reputation or accountability, such as the ability to send false messages, impersonate other users, or execute sybil attacks. Anonymous systems implicitly rely on their users to not unintentionally reveal their identity, a major risk for potentially technically unsophisticated dissidents.

In a network that completely discards the notion of authorship, an adversary could (in principle) correlate information within the network to attribute communications to a single author, and subsequently deanonymize those authors. Some networks rely on graphs of user trust to prioritize communications and provide resistance to false content dissemination (e.g. sybil attacks). Although robustness to such attacks is important, the structure of these social graphs can be correlated with external information to deanonymize users [53], raising important performance questions: how robust can a decentralized network be against sybil or denial of service attacks without knowledge of global graph structure? However, the feasibility of such an attack is dependent on multiple assumptions about the uniqueness of individual communication patterns and network size.

Thus, protecting user safety through anonymity, while an important goal, will be difficult in practice. Designers of dissent networks need to be aware of the subtleties and risks inherent in any such scheme. We argue their objective should be to make deanonymization as difficult as possible, with clearly stated and well-understood limitations to the anonymity guarantees their systems provide.

5.2.2 Deniability

Another approach for keeping users safe is that of deniability, i.e., making it difficult for a regime to use mere use or possession of a dissent networking technology as an effective way to target dissidents. Cell phones are an example of a technology that could be used for dissent networking that have strong deniability properties. In particular, cell phones (1) are ubiquitous and (2) have clear non-dissent uses (i.e., as a basic communications device). Of course, we refer only to the cell phone itself here, not cell phones using the cell phone network: the *contents* of a user's cell phone activity can be used to gather intelligence against dissidents.

The Tor Project [22] provides a better example of a system that supports dissent while providing some degree of deniability for its users. First, having Tor on a computer doesn't necessarily implicate the owner of that computer in being involved with dissident activity; there are a variety of innocuous reasons why a person might use Tor (such as to access content not available in their country, or to protect their identities from websites they visit. Second, if we imagine an adversary that can monitor network traffic, such as a state-run ISP, a person running a Tor middle relay node along with using Tor for their own traffic has plausible deniability that encrypted traffic originating from their device is simply relay traffic. Even if such an adversary were to successfully decrypt some portion of that users traffic, because Tor provides perfect forward secrecy they would not (easily)

be able to decrypt future messages, keeping the user safe. Thus, monitoring Tor usage doesn't provide a reliable signal to a monitoring authority that its originator is involved with dissident activity, and monitoring Tor traffic is difficult to begin with. Since Tor is an overlay network it is susceptible to the type of communications blackouts we are concerned with, but the deniability properties it provides can inform the design of future dissent networks.

5.2.3 Transparency

A more radical approach to providing user safety is transparency. Implicit in this strategy is a reliance on *non-technical* mechanisms to protect user saftey, even after an adversary has identified particular dissidents. For example, Chris Hedges argues that one of the strengths of the Occupy movement in the United States was its complete transparency, which hindered its adversaries' ability to sow seeds of doubt about the motivations and intentions of the movement [34]. Similarly, dissident news organizations have relied on transparency to ensure the safety of their staff, knowing their government would suffer serious backlash should they attempt to interfere with their operations or harass their employees.[5]

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—even if the right answers have yet to be found. 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 [18, 22]. Projects like Commotion and Serval exploit mobility and delay-tolerance in a mobile mesh setting, while avoiding exotic hardware [50, 28]. Short-range content transmission via Bluetooth or WiFi is standard in modern smartphones; applications building on this and leveraging opportunistic interactions between phones could provide a rich, safe means of communicating among dissidents [25]. Ideally, systems should aim to address *all* the requirements; Rangzen attempts this,

[5]This anecdote comes from personal communications with the founder of a dissident news organization in sub-Saharan Africa.

though its practicality and scalability is unproven.

Moving beyond 802.11 WiFi hardware, on which most projects to date have focused, presents both risks (Section 5.1) and opportunities. Radio stations are key sources of information during crises; exposing the RDS data stream in the FM standard to users and applications could provide a means of delivering low-bitrate broadcast information more efficiently than a voice broadcast. Amateur radio has a long history of operating in emergency situations; lessons from that community, both technical and operational, could prove instructive. Along these lines, we hope the community will consider "communications" broadly while designing practical dissent networks. Though mesh networks will likely encounter scalability problems for applications like telephony or point-to-point communications, other models (e.g. one-to-many communication) have yet to be explored. Although we do not know what effective blackout circumvention systems will look like, we strongly believe they will meet our definition of dissent networks.

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. They can achieve meaningful scale by adopting centralized management, planned growth, and stationary topologies, making them susceptible to government interference, or they can retain a decentralized nature at the cost of lower quality of service, requiring applications tailored to their limitations. This tradeoff stems from fundamental properties of wireless mesh networks, particularly the impact of channel contention.

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 should 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.

References

[1] CJDNS. https://github.com/cjdelisle/cjdns. Retrieved Nov. 2013.

[2] Freegate. www.dit-inc.us/freegate. Retrieved Apr. 2013.

[3] Freifunk Wireless Network. http://start.freifunk.net/. Retrieved Apr. 2013.

[4] Hyperboria. http://hyperboria.net/. Retrieved Nov. 2013.

[5] Kansas City Freedom Network. http://www.kcfreedom.net/. Retrieved Nov. 2013.

[6] Linux Wireless Regdb. http://wireless.kernel.org/en/developers/Regulatory/Database. Retrieved Oct. 2013.

[7] Meraki. http://www.meraki.com/. Retrieved Apr. 2013.

[8] The Free Networking Foundation. http://thefnf.org. Retrieved Apr. 2013.

[9] Ultrasurf. https://ultrasurf.us/. Retrieved Apr. 2013.

[10] C. Adjih, T. Clausen, P. Jacquet, A. Laouiti, P. Muhlethaler, and D. Raffo. Securing the OLSR Protocol. In *Proceedings of Med-Hoc-Net*, pages 25–27, 2003.

[11] 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.

[12] D. Akhawe and A. P. Felt. Alice in Warningland: A Large-Scale Field Study of Browser Security Warning Effectiveness. In *Proceedings of the 22th USENIX Security Symposium*, 2013.

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

[14] Y. Ben-David. Personal communication.

[15] J. Chroboczek. The Babel Routing Protocol. Internet Draft, Internet Engineering Task Force, Apr 2011.

[16] T. Clausen, P. Jacquet, C. Adjih, A. Laouiti, P. Minet, P. Muhlethaler, A. Qayyum, L. Viennot, et al. Optimized Link State Routing Protocol (OLSR). 2003.

[17] K. Coleman. Syria Gets Its Internet Back. May 2013.

[18] 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.

[19] J. Cowie. Egypt Leaves the Internet. `http://www.renesys.com/wp-content/uploads/2013/05/nanog-51-Egypt-Gone.pdf`, 2011.

[20] A. Dainotti, C. Squarcella, E. Aben, K. C. Claffy, M. Chiesa, M. Russo, and A. Pescapé. Analysis of Country-Wide Internet Outages Caused by Censorship. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement*, pages 1–18. ACM, 2011.

[21] 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.

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

[23] J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Privacy Aware Learning. *arXiv preprint arXiv:1210.2085*, 2012.

[24] 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.

[25] 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.

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

[27] P. Gardner-Stephen. A Framework for a Robust Architecture Offering Pervasive Security in Isolated and Infrastructure-Deprived Networks. Technical report, Serval Project.

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

[29] 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.

[30] 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.

[31] Google. Libya 2011 Internet Outage. `http://www.google.com/transparencyreport/traffic/disruptions/36/`. Retrieved Oct. 2013.

[32] 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.

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

[34] C. Hedges. Occupy Draws Strength From the Powerless. *Truthdig*, 2012.

[35] 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.

[36] 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.

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

[38] K. Nohl and C. Paget. GSM – SRSLY? In *26th Chaos Communications Congress*, 2009.

[39] M. Oliver, J. Zuidweg, and M. Batikas. Wireless commons against the digital divide. In *Technology and Society (ISTAS), 2010 IEEE International Symposium on*, pages 457–465. IEEE, 2010.

[40] Open Technology Institute. Case Study: Red Hook Initiative Wifi and Tidepools. Technical report, New America Foundation, Feb 2013.

[41] OpenBSC. `http://openbsc.osmocom.org/`. Retrieved Nov. 2013.

[42] OpenBTS. `http://openbts.org`. Retrieved Apr. 2013.

[43] C. Perkins, E. Belding-Royer, and S. Das. Ad Hoc On-Demand Distance Vector (AODV) Routing. RFC 3561 (Experimental), July 2003.

[44] 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.

[45] Renesys. Internet Blackout in Sudan. `http://www.renesys.com/2013/09/internet-blackout-sudan/`. Retrieved Oct. 2013.

[46] Renesys. Myanmar Internet Disruptions. `http://www.renesys.com/2013/08/myanmar-internet/`. Retrieved Oct. 2013.

[47] Renesys. Syrian Internet... Fragility? `http://www.renesys.com/2013/05/syrian-internet-fragility/`. Retrieved Oct. 2013.

[48] Renesys. Syrian Internet Off the Air. `http://www.renesys.com/2012/11/syria-off-the-air/`. Retrieved Oct. 2013.

[49] Renesys. Syrian Internet Shutdown. `http://www.renesys.com/2011/06/syrian-internet-shutdown/`. Retrieved Oct. 2013.

[50] 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.

[51] M. Srivatsa and M. Hicks. Deanonymizing Mobility Traces: Using Social Networks as a Side-Channel. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, pages 628–637. ACM, 2012.

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

[53] B. Viswanath, A. Post, K. P. Gummadi, and A. Mislove. An Analysis of Social Network-Based Sybil Defenses. In *ACM SIGCOMM Computer Communication Review*, volume 40, pages 363–374. ACM, 2010.

[54] T. Wu, Y. Xue, and Y. Cui. Preserving Traffic Privacy in Wireless Mesh Networks. In *Proceedings of the 2006 International Symposium on on World of Wireless, Mobile and Multimedia Networks*, pages 459–461. IEEE Computer Society, 2006.

[55] 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.

[56] L. Zhou and Z. J. Haas. Securing Ad Hoc Networks. *Network, IEEE*, 13(6):24–30, 1999.