Evading Cellular Data Monitoring with Human Movement Networks

Adam J. Aviv Micah Sherr Matt Blaze Jonathan M. Smith
University of Pennsylvania
{aviv,msherr,blaze,jms}@cis.upenn.edu

Abstract

Cellular networks are centrally administered, enabling service providers and their governments to conduct system-wide monitoring and censorship of mobile communication. This paper presents HUMANETS, a fully decentralized, smartphone-to-smartphone (and hence human-to-human) message passing scheme that permits unmonitored message communication even when all cellular traffic is inspected.

HUMANET message routing protocols exploit human mobility patterns to significantly increase communication efficiency while limiting the exposure of messages to mobile service providers. Initial results from trace-driven simulation show that 85% of messages reach their intended destinations while using orders of magnitude less network capacity than naïve epidemic flooding techniques.

1 Introduction

Mobile telephony networks are highly centralized, making network-wide eavesdropping, filtering and blocking feasible. This is a consequence of their architecture – cellular networks provide wireless access to a switched telecommunications system where routing and billing are performed. Cellular service providers and their governments can exploit this architecture to control *who* may use cellular services and – in the case of data communication – *what* content they may access. Governments imposing restrictive policies on cellular carriers may monitor, censor, or deactivate internal communication networks at will, as was the case in Iran [3] following the country's tumultuous elections in June 2009.

This paper presents our initial design of unmonitored and fully decentralized out-of-band communication protocols for mobile smartphone devices. Our new communication paradigm, *Human-to-human Mobile Ad hoc Networks* (HUMANETS), provides communication by exchanging messages over a "sneakernet" of mobile phones. HUMANETS thereby avoid centralized cellular systems and their controls, at a cost of increased latency.

HUMANETS' decentralized architecture presents a number of interesting research challenges. Messages must be efficiently routed toward their intended receivers, without relying on any fixed infrastructure. Moreover, the system should be robust against malicious insiders who attempt to either intercept communication or disrupt the network. This paper proposes efficient distributed routing protocols for HUMANETS, and outlines defenses against adversaries who infiltrate the network.

To provide scalability and efficient routing, HUMANETS leverage features of newer smartphone devices – in particular, the ability to discern the operator's physical location. By taking advantage of the particularities of human movement, predictive profiles can be constructed for each HUMANET node (smartphone) and serve as a basis for local routing decisions. A sender can then use information about where a receiver is (or will likely be) to address a message which the HUMANET will greedily route toward the addressed destination. Messages are passed from phone to phone through ad hoc WiFi networking, bypassing the cellular system.

Our phone-to-phone messaging passing scheme is similar to (but not identical to) more traditional mobile ad hoc networks. However, many classic geographic and position based [5, 16, 17, 19, 26] routing routines are ill-suited for completely decentralized message passing, as they often require a route discovery phase (assuming quiescent nodes and/or static neighbor sets), knowledge of the positions of all other participants, or the mapping of a grid overlay. Additionally, epidemic [25] and gossip [14] routing protocols exhibit poor scalability, and do not support multiple simultaneous senders and receivers efficiently.

Our preliminary results are encouraging. In a trace driven simulation based on actual movement within a metropolitan area, HUMANET-based smartphone routing successfully delivered 85% of the messages to their intended destinations. Moreover, 75% of the messages reach their destination within a day. In comparison to epidemic flooding, our techniques incur fixed storage costs, requiring only a small fixed-sized subset of the network to carry message copies. Our routing algorithm is highly scalable, permitting many more concurrent messages in the system than allowed by flooding and gossiping techniques.

2 HUMANETS

HUMANETS are intended to be used as fully decentralized networks with no fixed infrastructure, and can be composed of participants who frequently change location at will and without coordination. Our routing protocol for HUMANETS aims to achieve efficient routing in such highly dynamic and unstructured networks

This work was supported by: NoBot, Networks Opposing Botnets, ONR N00014-09-1-0770; Application-Aware Anonymity (A3) for the Masses, CNS-08-31376; and, Foundational and Systems Support for Qunatitative Trust Management, ONR MURI N00014-07-0907

by exploiting *human mobility patterns*. In particular, HUMANET routing depends on a specific property of human movement, which we call the *Return-to-Home Principle*. This hypothesizes that a person is likely to return to places that he frequented in the past (*e.g.*, home, work, favorite coffee shop, etc.). We believe that sufficiently many network users will have a number of such locations, which we call *homes*, and HUMANETS exploit this fact to provide efficient routing.

At a high level, HUMANETS operate by forwarding messages toward receivers' homes. A message *carrier* – a smartphone that stores a copy of a message – will transfer the message to another phone if the operator of the latter phone tends to frequent the same (or nearby) homes as the receiver. Hence, messages are relayed via a heuristic-based protocol: rather than route toward the receiver's current location (which may not be readily known and can change), messages are relayed toward probable future locations of the receiver.

HUMANETS support three messaging primitives: *Unicast* messaging targets a specific receiver, and requires that the sender possess some knowledge about the receiver's homes. Alternatively, the sender may send messages to any (*anycast*) or all (*multicast*) phones in a targeted area. Unicast messaging permits point-to-point communication, whereas anycast and multicast disseminate information (e.g., "come to the rally at 4pm at the castle") to physically proximate individuals.

2.1 Location Profiles

To permit intelligent routing decisions, each phone maintains a *location profile* that compactly defines the geographic areas (homes) in which its owner tends to locate. Smartphones will periodically record their physical locations using GPS or E911 services, and construct profiles from collected locations (*e.g.*, overnight as the phone is charging).

Profile construction begins by applying k-means clustering to partition recorded physical locations into k clusters. GPS and E911 may produce spurious locations, so a bounded region is formed for each cluster by first dividing the cluster radially about its centroid into equal-sized partitions (Fig. 1, *left*). Next, within each division, the location that exhibits the median distance, measured with respect to the cluster's origin, is chosen as a polygon vertex (Fig. 1, *right*). The resulting polygon formed becomes a *home*. To represent long-term movement patterns (*e.g.*, a weekly commute between cities), the location profile consists of h such home areas. Newly computed homes are added to the location profile if they intersect with existing home areas; non-intersecting homes are eventually expired. We provide additional details regarding home formation and profile home selection strategies in a technical report [4].

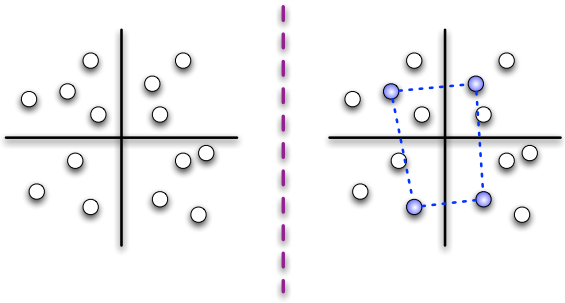

Figure 1: The construction of homes from recorded physical coordinates. All points reside within the same cluster. *Left:* The points within the cluster are divided into four quadrants. *Right:* The median point in each quadrant is chosen as a vertex of a home.

2.2 HUMANET Routing Protocol

To route messages toward a geographic location, we rely on the Return-to-Home Principle: a person (and the phone she carries) is likely to visit places that she has frequented in the past. Conceptually, our routing algorithm transfers a message to a nearby neighbor if the neighbor is more likely to visit a place frequented by the intended receiver. To increase the likelihood of message delivery, multiple message copies are inserted into the network. The sender will address each copy to a separate geographic region that she believes the receiver is likely to visit. However, it is worth emphasizing that the number of message copies remains fixed in HUMANETS. As explained below, messages are transferred (*i.e.*, handed off) between phones.

Message Format. Message payloads are prefixed with several fields used for HUMANET routing (Fig. 2)[1]:

- `type`: unicast, anycast, or multicast delivery.
- `destination`: geographic polygon that the message is to be routed toward.
- `destination id`: (unicast only) a pre-shared secret identifier known only to the sender and the receiver.
- `message id`: a randomly chosen identifier. To ensure acyclic routing, phones store the `id`'s of previously received messages, and refuse to accept messages that they have already carried.
- `message timeout`: expiration time of the message (after which it can be discarded).

Greedy Routing. To forward messages, HUMANET-enabled smartphones participate in a synchronous message passing protocol. With some fixed periodicity, phones engage in an *exchange round* by joining an ad hoc WiFi network with a known SSID. Each phone X broadcasts the tuple $(nonce_X, profile_X)$, where $nonce_X$

[1]The message format is designed to fit in one MTU.

Size (bytes)	1	20	20	20	4	1400
Field	type	message id	destination	destination id	message timeout	payload

Figure 2: HUMANET message format.

is a large random number and $profile_X$ is a serialized representation of X's location profile. (We discuss the privacy implications of broadcasting location profiles in Section 4.)

During the exchange round, a carrier (a phone that possesses a message) listens to the broadcasts. For each broadcast $profile_X$, the carrier counts the number of intersections between the polygons defined by the homes in $profile_X$ and the region described by the message's `destination` field. The carrier then *transfers* its stored message to the smartphone that has the highest number of such intersections, provided that the count is greater than the number of intersections between the carrier's profile and `destination`.[2]

To transfer the message to phone X', the carrier will broadcast the tuple $(nonce_{X'}, msg)$, where $nonce_{X'}$ is the random number broadcast by phone X' and *msg* is the message (including the headers shown in Fig. 2). Recall that X' may reject the transfer if it has already seen a message with the same identifier. In such cases, the carrier can either attempt to transfer the message to another phone or defer until the next exchange round.

Preventing Sinkholes. HUMANET routing does not guarantee that a message will take the optimal path. For example, a message may be transferred to a phone whose operator had previously shared multiple home regions with the receiver, but the phone may then exhibit a new movement pattern that is dissimilar from the receiver's. To prevent messages from becoming *sinkholed*, a carrier sets a *local timeout* threshold whenever it receives a message. If the carrier does not transfer the message before the timeout expires, it will then transfer the message to the next HUMANET-enabled smartphone that it encounters, regardless of that phone's location profile. This effectively "resets" the protocol.

The Last Mile, Literally. The greedy routing algorithm forwards messages to phones that have frequented the geographic area defined by the messages' `destination` fields. For anycast messages, once the message arrives in the targeted area, it is considered delivered. Multicast and unicast messages use an additional *flooding stage* in which messages are spread epidemically but only within the area specified by the `destination` field. A carrier who enters the targeted area will transmit a copy to each smartphone that it encounters in that area. Hence, there may be many message copies for multicast and unicast messaging, but such copies are restricted to a confined physical space.

[2] In the case that multiple phones have the highest count, a phone is chosen among them at random.

3 Preliminary Evaluation

HUMANETS implement a greedy routing algorithm that relies on the Return-to-Home Principle, with messages addressed to a targeted geographic area. Message carriers transfer messages only if a better candidate is encountered – *i.e.*, a phone that has visited the targeted area more frequently.

The reliability of our heuristic-based approach depends on several factors: (1) the correctness of the Return-to-Home Principle; (2) the contact rates of HUMANET participants; and (3) the patterns of human movement. In the next section, we demonstrate the feasibility of our approach via trace-driven simulation.

We evaluate HUMANETS using publicly available traces of actual human movement [1, 10, 21], as existing human movement simulation models (*e.g.*, Levy Walks [22]) do not incorporate the Return-to-Home Principle, a fundamental tenet of HUMANET's routing protocols. (We are currently pursuing the development of synthetic mobility models that more accurately capture humans' routine movement patterns.)

A significant challenge to evaluating the performance of the HUMANET protocols is the lack of suitable large-scale datasets of fine-grained human movements over metropolitan (and larger) areas. However, one publicly available dataset, widely used for human movement simulation, is suitable here. For our preliminary evaluation, we used the *Cabspotting Dataset* [21] which contains GPS coordinates and time-stamps collected over 20 days from 536 licensed taxicabs operating in the San Francisco area. Compared with other publicly available human movement traces, the Cabspotting Dataset provides finer-grained movement history and covers a longer period of time. Although the movements of taxis may not be perfectly representative of that of the general population of all phone users, our simulation can be best interpreted as representing a HUMANET network of taxi drivers' smartphones, possibly applicable to other "fast moving" populations. We are currently seeking similarly fine-grained datasets for other user populations.

3.1 Testing the Return-to-Home Principle

To evaluate the correctness of the Return-to-Home Principle, we measure how often a phone returns to a home defined in its location profile. A profile consists of a maximum of 15 homes, 5 generated using the previous day's data (*i.e.*, $k = 5$) and 10 homes from previous

profiles (see [4] for more details on this selection criteria). Using the Cabspotting dataset, we found that 65% of GPS coordinates fell within the previous day's profile and that phones reside in their profiles' homes for 65% of the day. Even in the worst performing case, 39% of recorded locations were inside home regions and 45% of the day was spent within these homes. Similar results were obtained using additional human mobility datasets [1, 10]. As we show in Section 3.2, the Return-to-Home Principle is sufficiently accurate to enable highly effective routing.

3.2 Performance Under Simulation

We constructed a discrete-event simulator to evaluate the performance of HUMANET routing. Movement data from the Cabspotting traces were used as input to our simulator. Our evaluation focuses on unicast messaging, with senders and receivers chosen randomly from the taxicabs. Due to the limited size of the Cabspotting dataset (536 cabs), we did not implement the flooding stage of the HUMANET protocol. Instead, we conservatively consider a message to be successfully delivered only if it is directly received by its intended target.

We compare the performance of HUMANETS against three alternative phone-to-phone message passing techniques: *epidemic flooding*, *probabilistic epidemic flooding*, and *probabilistic random walk*. In contrast to our profile based approach, epidemic flooding sends a copy of a message to each phone that comes in contact with a carrier. In probabilistic epidemic flooding, a carrier transfers a message to an encountered phone with some fixed probability.

Both flooding techniques result in multiple copies of a single message. In contrast, the probabilistic random walk approach *transfers* messages between phones, with some fixed probability. The random walk technique therefore functions similarly to the HUMANET routing protocol, but the former does not exploit location information. This allows us to directly measure how effectively our location-based heuristics aid routing performance.

Parameters. Each simulation consists of 300 runs in which there is a single sender and (intended) receiver, both chosen uniformly at random. For probabilistic epidemic flooding and probabilistic random walk, we reduce the spread rate by an order of magnitude by using a modest transfer probability of 0.05 (the exploration of the effects of various transfer probabilities is left as future research). For HUMANET routing, profiles consist of at most 15 homes and are recomputed daily based on the previous day's movement history. Messages are addressed using the cab's current profile at the time of the initial send, and one message copy is inserted per home in the profile (*i.e.*, at most 15).

Local and message timeouts are enforced for all routing protocols. When a local timeout occurs, a phone will not accept another copy of the message. No messages are transferred after the message timeout. The local and message timeouts are 10 hours and three days, respectively.

Metrics. We measure the effectiveness of the techniques using three performance metrics: *success rate* is the fraction of sent messages that reach the receiver before the message timeout expires; *latency* is the time interval between when a message is sent and when it is received; and *network load* is the number of duplicate message copies in the network during the message's lifespan (*i.e.*, the time before the message timeout).

Simulation Results. As shown in Fig. 3 (*left*), HUMANET's routing protocol produces the highest success rates of the tested techniques. 85% of messages were successfully delivered using HUMANETS, compared to 76.3% for epidemic flooding[3]. Probabilistic random walk exhibited the lowest success rate – only 28% of the messages were delivered, highlighting the benefits of routing based on location profiles.

Fig. 3 (*center*) plots the cumulative distribution of latencies achieved by the four routing techniques *for the cases in which messages reach their destinations*. When messages are successfully delivered, epidemic flooding and probabilistic epidemic flooding deliver messages with less latency than the random walk and HUMANET techniques. The respective median latencies for epidemic and probabilistic epidemic routing are 277 and 676 minutes, compared with 870 minutes (14.5 hours) for HUMANET routing. Probabilistic random walk performs the worst, with a median latency of 1849 minutes.

The cumulative distribution of network load resulting from the four routing strategies is shown in log scale in Fig. 3 (*right*). Although epidemic techniques deliver messages faster, they incur a significantly higher storage cost. Neither HUMANETS nor probabilistic random walk ever impose a network load of more than the number of homes in the receiver's profile. In comparison, epidemic and probabilistic epidemic routing both incur significant network load, and in all simulations, a majority (more than 60% for probabilistic and 80% for epidemic routing) of *all phones in the network* carried the message at some point. Such loads would render the network unusable for even a small number of simultaneous sender and receiver pairs.

Summary. Our preliminary results show that HUMANETS deliver 85% of messages to their receivers.

[3]The performance of epidemic flooding was hindered by the pace of the "infection" – messages were transferred too quickly, causing several local timeouts early in the simulation. Since phones that have already received a message do not accept future copies of that message, this resulted in lower than expected success rates.

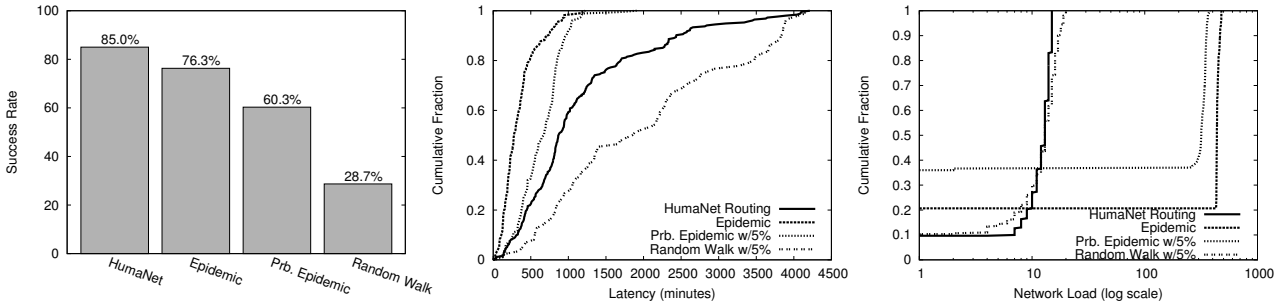

Figure 3: Success rates (*left*), and CDFs of latency (*center*) and network load (*right*) for HUMANET compared with epidemic, probabilistic epidemic, and random walk routing.

Although HUMANET routing imposes a modest latency cost, it is far more scalable than epidemic flooding approaches. The simulation suggests location-based routing is a feasible approach for surveillance-resistant messaging in metropolitan areas.

4 Challenges

Although our initial results are encouraging, there are several challenges that must be met to construct robust and surveillance-resistant HUMANETS. Addressing these challenges is the subject of our ongoing work.

Reliability HUMANETS exploit the predictive nature of human movement. Although our preliminary evaluation shows that most messages are delivered in a city-sized HUMANET, the heuristic routing protocol does not *guarantee* delivery. Like Internet (IP) routing, HUMANETS provide best-effort network-layer messaging. We are investigating techniques for applying reliability protocols on top of HUMANETS, borrowing approaches from delay tolerant networking. We are also examining techniques that compensate for expected message losses (*e.g.*, via redundant message copies or erasure codes). Such strategies do not require bidirectional communication, but are not resilient to loss bursts.

Location Privacy. HUMANET participants periodically broadcast their location profiles. However, the HUMANET message format (Fig. 2) does not contain any attributes that uniquely identify the phone or its operator. However, given knowledge of a person's routine movements, it may of course be possible to link an intercepted profile with a person. We are investigating metrics for quantifying and minimizing privacy disclosures. One potential approach is to allow users to define private geographic areas that are excluded from their profiles.

Anonymity. HUMANETS resist surveillance from cellular service providers by providing an "out-of-band" channel for users to communicate short messages. However, since HUMANETS rely on unprotected WiFi broadcasts, eavesdroppers can monitor exchange rounds to glean information about participants. In particular, the `message timeout` field leaks information that may be useful to determine whether a carrier that broadcast a message is likely to have originated it: longer timeouts indicate that the message is fresher, and consequently, suggests that the carrier may be the initial sender. The sender may increase her anonymity by fuzzing `message timeout` values, although eavesdroppers can subtract such noise if they have knowledge of the sender's chosen probability distribution. A potential solution that we are exploring is the application of a *k-anonymity* scheme in which exchange rounds are conducted only in crowds of k or more participants. However, defending fully decentralized k-anonymity schemes against Sybil attacks (for example, in which the $k-1$ neighbors may be controlled by the adversary) is an open problem.

Unicast messaging provides receiver-anonymity since the `destination id` is assumed to be a pre-shared secret known only to the sender and receiver. (The receiver's likely whereabouts are not kept private by our scheme, as such information is necessarily disclosed to route packets.) However, if an eavesdropper monitors the receiver after he receives a message, then he risks revealing his role as the receiver if he no longer forwards the message. A simple countermeasure is for receivers to continue participating in the protocol, letting all messages be transferred until their `message timeout` expires.

Routing Attacks. HUMANETS are fully distributed peer-to-peer (p2p) networks and are vulnerable to the same classes of routing attacks that afflict Internet-based p2p systems. For example, malicious carriers may trivially cause denial-of-service by refusing to forward their stored messages. We are investigation whether mitigation techniques developed for use on the Internet (*e.g.*, the sending of redundant message copies via separate paths [7]), may also be used to defend HUMANETS.

5 Related Work

HUMANETS draw on previous work in the areas of geographic and position based routing (such as [5, 16, 17]), but a lengthy survey is beyond the scope of this work and infeasible due to space considerations (see [26, 19] for a survey of these techniques). However, many of these strategies often assume fixed neighbor sets, require the centralization of position information, or use epidemic [14, 25] principles that do not scale. In contrast, HUMANETS are completely decentralized and incur fixed bandwidth overheads. HUMANETS also build upon previous work that profiles node movements and contact histories to make local routing decisions [6, 13, 15, 18], but these prior techniques do not provide sender anonymity.

Human contact networks [8, 12] have been studied in the context of cellular botnets [11, 23]. Interestingly, a HUMANET constitutes an ideal platform for mobile botnets. Not only is the botmaster protected with sender anonymity, the ability to route to a particular location could be used for highly coordinated, targeted attacks against cellular infrastructure [24].

Finally, classic sender-anonymity schemes, such as Tor [9], have been ported to cellular devices [2]. However, service providers can easily obtain the addresses of anonymizing relays and block access to the anonymity system[4]. Unlike Tor, HUMANETS are completely decentralized and rely on phone-to-phone message passing, making it very difficult to enforce centralized filtering policies.

6 Conclusion

In many countries, wireless cellular services are the dominant communication medium. The architecture of these systems is highly centralized, enabling intrusive monitoring and filtering. This paper proposes an alternative means of communicating using the same smartphones that were designed for the centralized cellular systems. Human-to-human Mobile Ad hoc Networks (HUMANETS) leverage the ubiquity of smartphone devices by transmitting messages out-of-band (directly between proximate smartphone handsets) rather then via (easily monitored) cellular networks.

By taking advantage of the location-awareness of smartphone devices and common regularities in human movement, messages can be routed reasonably efficiently and reliably toward their intended recipients. Preliminary trace-driven simulation results are encouraging: 85% of HUMANET messages are eventually delivered, and 75% of those messages are received within 24 hours. Additionally, HUMANETS are highly scalable, incurring only small fixed network overhead for each message.

This work represents an initial step toward constructing robust anonymity networks for smartphones. The relationship many people today have with their computing devices, and in particular their mobile phones, is a rich area for future research. In addition to obtaining new data sources and performing active experiments, we intend to investigate and exploit the potential of smartphones (and autonomous personal devices generally) as a platform for new kinds of anonymous communication.

[4]For example, China has been blocking access to Tor relays since September 2009 [20].

References

[1] IEEE VAST 2008 Challenge. http://www.cs.umd.edu/hcil/VASTchallenge08/.

[2] Tor on android. https://www.torproject.org/docs/android.html.

[3] Poll results prompt Iran protests. Al Jazeera (English), June 14, 2009.

[4] A. J. Aviv, M. Sherr, M. Blaze, and J. M. Smith. Moving targets: Geographic routed human movement networks. Technical Report MS-CIS-10-12, University of Pennsylvania, March 2010. Available at http://repository.upenn.edu/cis_reports/926/.

[5] S. Basagni, I. Chlamtac, V. R. Syrotiuk, and B. A. Wordward. Distance routing effect algorithm for mobility (DREAM). In *MobiCom*, 1998.

[6] C. Becker and G. Shiele. New mechnaism for routing in ad hoc networks. In *4th Plenary Cabernet wksp.*, October 2001.

[7] M. Castro, P. Druschel, A. J. Ganesh, A. I. T. Rowstron, and D. S. Wallach. Secure Routing for Structured Peer-to-Peer Overlay Networks. In *OSDI*, 2002.

[8] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott. Impact of human mobility on opportunistic forwarding algorithms. *IEEE Transactions on Mobile Computing*, 6(6):606–620, 2007.

[9] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The Second-Generation Onion Router. In *USENIX Security*, 2004.

[10] N. Eagle and A. (Sandy) Pentland. Reality mining: sensing complex social systems. *Personal Ubiquitous Comput.*, 10(4):255–268, 2006.

[11] C. Fleizach, M. Liljenstam, P. Johansson, G. M. Voelker, and A. Mehes. Can you infect me now?: malware propagation in mobile phone networks. In *WORM*, 2007.

[12] N. Glance, D. Snowdon, and J.-L. Meunier. Pollen: using people as a communication medium. *Computer Networks*, 35(4):429–442, 2001.

[13] M. Grossglauser and M. Vetterli. Locating Nodes with EASE: Last Encounter Routing in Ad Hoc Networks Through Mobility Diffusion. In *InfoCom*, 2003.

[14] Z. J. Haas, J. Y. Halpern, and L. Li. Gossip-based ad hoc routing. *IEEE/ACM Trans. Netw.*, 14(3):479–491, 2006.

[15] P. Juang, H. Oki, Y. Wang, M. Martonosi, L. S. Peh, and D. Rubenstein. Energy-efficient computing for wildlife tracking: design tradeoffs and early experiences with zebranet. In *ASPLOS-X*, pages 96–107, 2002.

[16] B. Karp and H. T. Kung. GPSR: greedy perimeter stateless routing for wireless networks. In *MobiCom '00*, 2000.

[17] Y. B. Ko and N. H. Vaidya. Location-aided routing (LAR) in mobile ad hoc networks. In *MobiCom*, 1998.

[18] A. Lindgren, A. Doria, and O. Schelen. Probablisitic routing in intermittently connected networks. In *SAPIR (LNCS 3126)*, pages 239–254, 2004.

[19] M. Mauve, J. Widmer, and H. Hartenstein. A survey on position-based routing in mobile ad-hoc networks. *IEEE Network*, 15:30–39, 2001.

[20] phobos. Tor partially blocked in China, September 2009. Available at https://blog.torproject.org/blog/tor-partially-blocked-china.

[21] M. Piorkowski, N. Sarafijanovoc-Djukic, and M. Grossglauser. A Parsimonious Model of Mobile Partitioned Networks with Clustering. In *COMSNETS*, 2009.

[22] I. Rhee, M. Shin, S. Hong, K. Lee, and S. Chong. On the levy-walk nature of human-mobility. In *IEEE Infocom*, 2008.

[23] K. Singh, S. Sangal, N. Jain, P. Traynor, and W. Lee. Evaluating bluetooth as a medium for botnet command and control. In *DIMVA*, 2010.

[24] P. Traynor, M. Lin, M. Ongtang, V. Rao, T. Jaegar, P. McDaniel, and T. L. Porta. On cellular botnets: Measuring the impact of malicious devices on a cellular network core. In *CCS*, 2009.

[25] A. Vahdat and D. Becker. Epidemic routing for partially-connected ad hoc networks. Technical report, Duke University, 2000.

[26] Z. Zhang. Routing in intermittently connected mobile ad hoc networks: Overview and challenges. *IEEE Comunications Surveys*, 8(1), 2006.