

Using Predictable Mobility Patterns to Support Scalable and Secure MANETs of Handheld Devices

David R. Bild[†] Yue Liu[†] Robert P. Dick[†] Z. Morley Mao[†] Dan S. Wallach[‡]
[†]University of Michigan [‡]Rice University
{drbild,liuyue,zmao}@umich.edu, dickrp@eecs.umich.edu dwallach@cs.rice.edu

ABSTRACT

Mobile ad-hoc networks of wireless devices (MANETs) hold the promise of providing network services without traditional infrastructures that could fall victim to manipulation and censorship. Unfortunately, current MANET systems suffer significant scalability problems, effectively precluding their use for general-purpose networking. We suggest tailoring MANETs to particular classes of application and leveraging application-specific properties to increase scalability. This paper describes the design of a scalable and secure MANET for text-based personal communication. Our design is based on geographic routing and uses human motion and communication patterns to facilitate location tracking and distribution, thereby increasing scalability above that of traditional geographic routing. We provide location privacy by transplanting mix-net techniques into MANETs.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design— Distributed networks; C.2.0 [Computer-Communication Networks]: General— Security and protection

General Terms

Design, Security

1. INTRODUCTION

Wireless mobile ad hoc networks (MANETs) composed of volunteer, mobile devices offer some advantages over traditional infrastructure networks because their nonhierarchical nature eliminates critical points of failure that can be exploited by attackers to reduce reliability and enable censorship, surveillance, and other forms of undesirable interference. Attacks upon communication systems are easier when most network traffic is routed through backbone networks owned by a few ISPs or a state [17]. MANETs have the potential to significantly increase the cost of large-scale censorship or shutdowns. Unfortunately, communication and computation capacities of individual nodes limit scalability [18] and have, thus far, undermined general-purpose use. However, use in specific applications remains a possibility. In particular, while MANET bandwidths and end-to-end latencies may be insufficient to support voice conversations or video, they may support valuable services like text messaging.

This work was supported in part by the National Science Foundation under awards TC-0964545 and CNS-0347941.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiArch '11, June 28, 2011, Bethesda, Maryland, USA.

Copyright 2011 ACM 978-1-4503-0740-6/11/06 ...\$10.00.

1.1 MANETs May Offer A More Robust Supplement to the Internet

The Net interprets censorship as damage and routes around it. — *John Gilmore, 1993*

Although the Internet has been heralded for being robust to censorship, ongoing events in the Middle East, North Africa, Asia, and elsewhere falsify this belief; governments can exploit the hierarchical nature of the Internet to censor news as well as limit and monitor communication. In an extreme example, Egypt completely disabled Internet access for several days in February 2011 by forcing their five major ISPs to withdraw Border Gateway Protocol routes [14]. In Tunisia, where bandwidth is leased from the government [13], Internet access is heavily filtered. Many websites (e.g., YouTube) are blocked [13]. Others (e.g., Facebook and Twitter) are modified to steal login credentials [2]. Emails and attachments are filtered and scrubbed [13]. In all these cases, the choke-points inherent to the Internet's hierarchical structure help facilitate the censorship.

In contrast, mobile ad hoc networks composed of volunteer, wireless devices (e.g., smartphones and laptops) have the potential to be more resistant to corruption. Due to their nonhierarchical, ad hoc structures, censoring communication requires controlling many of the nodes in the network. When these nodes are handheld devices owned by private individuals numbering in the tens of thousands or more, acquiring such control is vastly more difficult and expensive than adding filtering software to a few backbone routers. Although MANETs will not help for long-distance or transocean communication, they have the potential to provide secure and uncensored communication within contiguously populated local regions, which may be sufficient to support communication among friends and family members.

1.2 MANET Architectures Should Exploit Application-Specific Properties

An ideal robust supplement to the Internet would support all types of traffic. Unfortunately, poor MANET scalability precludes their use for general-purpose networking. Thus, instead of seeking a general MANET architecture, we argue that MANET architectures must be tailored towards specific application-classes.

This poor scalability stems from two primary properties. (1) The traffic forwarded by each node increases with network size, reducing throughput for originating traffic [18]. (2) The traffic required to maintain routing state for the mobile nodes increases with network size, reducing available bandwidth [12]. Simulations indicate that current MANETs scale to only a few thousand nodes, with low per-node throughput (<5 kbps) [3].

We argue that these limitations imply that useful MANET architectures must be tailored to specific application-classes. First, the throughput and latency induced by the required network size must be acceptable. Second, properties of the application should be leveraged to design more efficient routing methods. In this work, we use predicted human motion patterns to support a MANET for text-based personal communication (e.g., text messaging), a low-bandwidth and latency-tolerant application.

1.3 MANET Architecture for Text-Based Personal Communication Applications

Text-based personal communication among friends and family members is both useful to many people (as evinced by the popularity of text and instant messaging) and particularly suited to a town-sized MANET, as indicated by the following two properties. (1) The required per-node throughput is low (<500 bps) and relatively high latency is acceptable (1–5 sec). (2) People frequently communicate with relatively small groups of contacts in close geographic proximity [6], implying a short average link length, which improves scaling properties. Furthermore, properties of human motion patterns can be leveraged to provide efficient routing.

A MANET architecture supporting text-based personal communications should satisfy the following requirements.

- **Scalability.** A useful personal communication network must cover a region of non-trivial area (e.g., a small town or a university campus), providing reliable delivery for all participants (e.g., a few thousand nodes) without imposing much computation or battery energy overhead on participating nodes.
- **Confidentiality.** The network should guarantee end-to-end message confidentiality. Packets should therefore be protected from eavesdropping and traffic analysis as they are relayed through arbitrary nodes untrusted by the source and destination.
- **Location Privacy,** defined as “the ability to prevent other parties from learning ones’ current or past location” [4]. Persistent identifiers must not be linkable to node locations.
- **Social Network Privacy.** A person’s social network, i.e., the set of network peers he communicates with, should be protected. No one (except the sender and receiver themselves) should be able to determine both the sender and receiver of any packet (by real identity, network identity, or location).

In this paper, we present the design of a location-centric MANET architecture supporting text-based personal communication within town-sized regions. Properties of human mobility patterns motivate a novel routing method, *location profile-aided geographic routing*. Geographic routing [12] is at the core of its scalability: next-hop selection requires only local knowledge within one-hop neighborhoods. However, to address a message the sender needs to know the destination locations, which are traditionally provided by distributed location services [7] that scale poorly and do not easily support confidentiality and privacy. We observed that (1) humans have highly predictable motion patterns, spending the majority of time in a few locations [11] and (2) the frequency of change in mobility patterns is on the order of months and years. We propose to model location patterns as *location profiles* (e.g., location–probability pairs), distributing them face-to-face, instead of real locations via the network, to reduce overheads (see Section 2). Direct visibility of location profiles is often unacceptable, so we embed the pre-shared location profiles in encrypted reply blocks [5], thus preserving location privacy by hiding the destination from the sender (see Section 3). The reply blocks also provide sender–receiver unlinkability and public key encryption provides confidentiality (keys are shared along with the location profiles, so PKI is not necessary).

Note that our primary goal is providing a censorship-resistant communication system for day-to-day use, when human motion is highly routine and predictable. Our primary target is not Internet shutdowns in an active protest or revolution scenario (à la Egypt in February 2011) where movements may be highly varied and non-routine. However, our system still enables communication in these scenarios, with the scalability dependent on the extent that locations are predictable (e.g., when protesters are at home). Supporting communication during protests is a secondary goal. In general, political organizations use censorship to distort the views of the masses with the goal of causing them to take potentially self-destructive acts that benefit the political organization. If long-term,

every-day access to information is thus distorted, there is no reason to believe that successful protests and revolutions by those with the resulting censorship-based world views will bring the intended results. Our primary goal is therefore supporting communication among friends and family members.

We make the following primary contributions:

- We propose leveraging the predictability of human motion to reduce routing costs in MANETs comprising handheld devices.
- We develop a reply block-based scheme to add location privacy to geographic-based routing.
- We describe a location-centric MANET architecture that provides scalable and secure text-based personal communication that resists censorship and shutdown.

The rest of the paper presents a detailed description and justification for this architecture. Section 2 presents location profile-aided routing. Section 3 proposes the location reply blocks used to address the security and privacy issues induced by location profiling. With these two fundamental components developed, Section 4 describes the full scalable and secure location-centric network architecture.

2. LOCATION PROFILE AIDED ROUTING

Geographic routing techniques [12] scale well because only local knowledge (the destination node’s location, available in the packet header, and the neighboring nodes’ locations) is needed to determine the next hop. The cost for maintaining valid routes does not increase with network size, unlike proactive and reactive techniques that flood routing control packets across the network [1]. Scalability is not guaranteed, however; senders must know the current locations of the destination nodes. Indeed, the cost of maintaining valid routes is replaced by the cost for determining individual nodes’ locations.

Traditionally, location tracking is performed by distributed location services [7]. Each node appraises a subset of the network of its current location. Other nodes may then query this subset for current location information. However, such services have two major problems. First, the update and query costs are high, limiting scalability. Second, providing sender–receiver anonymity is difficult because location queries are sent to third-parties. Such security features might potentially be added, but the greater complexity would increase the risks of security holes due to design flaws or implementation bugs.

Instead, we turn to two properties of the text-based personal communication applications to motivate a more-efficient, albeit less-general, distribution mechanism. First, human-carried nodes have predictable motion patterns, implying that the location information necessary for geographic routing can be exchanged a priori, requiring only infrequent updates. More specifically, most humans spend the majority of their time in a small number of locations and thus have simple, often-repeated motion patterns [11], suggesting that node locations can be predicted with simple models. Second, a significant fraction of most people’s communication is with a relatively small number of contacts in close geographic proximity [6], indicating that supporting a useful fraction of messages requires each node to know locations for only a small number of contacts.

Based on these two observations, we propose replacing current locations with automatically-developed predictive models of node mobility, which we call *location profiles*. The first observation indicates that profiles change less frequently than locations, reducing distribution cost. The second indicates that profiles can be exchanged directly between contacts, eliminating the dependence on third-party location servers and simplifying implementation of security features. The profiles can be shared initially face-to-face and then updated via the network.

Directly sharing location profiles violates location privacy, so instead we propose embedding profiles in reply blocks, as described in Section 3. For clarity, the remainder of this section is written as

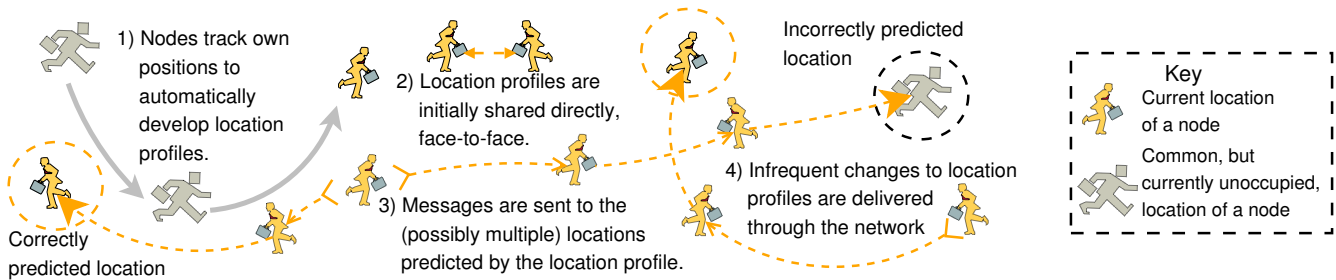


Figure 1: Illustration of the main steps in location profile routing.

though location profiles are exchanged directly. However, the ideas described are fully compatible with the location-privacy-preserving reply-block addition proposed in Section 3.

2.1 Evidence of Beneficial Human Mobility Patterns

Gonzalez et al. [11] used the trajectories of 100,000 cellphone users during a six-month period to show that humans generally have simple, repeated motion patterns, spending most of their time in a few locations. To be useful for geographic routing, these common locations should have size similar to the wireless transmission range (~ 100 m for 802.11b), but the coarse granularity of the cellphone-based location data (several kilometers) prevents such precise characterization. More fine-grained characterization is needed.

Using thirty-five fine-grained GPS traces from a study of university students following their daily routine [15], we characterized the common location size. Trace durations range from 1.7 hr to 21.7 hr, averaging 10.2 hr. These durations are too short to analyze temporal predictability, so we assume the common locations change infrequently, as found by Gonzalez et al. [11].

To characterize the size and count of the students' common locations, we calculated the percentage of the locations in each trace covered by circular regions of varying radius and count. On average, a single circular region of 100 m radius covered 72% of each trace, indicating that common location behavior exhibited in the cellphone traces also exists at finer spacial granularity.

Figure 2a show the coverage percentage for multiple 100 m radius regions. Three regions are enough to cover over 90% of a trace, on average. Increasing the count further only slightly improves the coverage. With five regions, Figure 2b illustrates the coverage percentage for various region radii. With a 10 m radius, five regions cover only an average 82% of a trace. The percentage increases sharply to 92% with a 50 m radius, but slowly after that point. These results confirm that the common location behavior is exhibited at the scale of 802.11b transmission ranges, indicating the practicality of developing simple location profiles. Specifically, on average the students spend most of their time (93%) in 4 common locations of small size (100 m radius). To successfully reach a receiver, it is only necessary to transmit to an average of 1.2 locations.

2.2 Details of Location Profile-Aided Routing

The main steps of location profile-aided routing for personal communication are shown in Figure 1. Nodes continuously monitor their positions to build location profiles (step 1), which are then shared with contacts directly (step 2). A message to a contact is addressed to the location(s) predicted by the corresponding location profile (step 3). When motion patterns change (e.g., a user starts a new job or moves to a new home), updated profiles are shared with contacts via the network (step 4). Routing fails if a receiver is not in any of the predicted locations. Possible solutions include delay tolerant networking, but are not the focus of this work.

Location profile-aided routing has four components: *location profiles* (used in step 1) describing nodes' motion patterns, a *distribution method* (used in steps 2 and 4) for sharing location profiles,

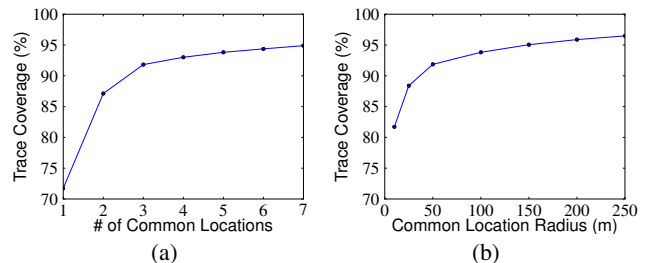


Figure 2: Characterizations of common locations.

and an *addressing policy* (used in step 3) specifying a precedence over the predicted locations. Details and design trade-offs follow.

Location Profiles: Motions patterns can be modeled in many ways. A location profile is any model or function P mapping a time to a set of location–confidence tuples, with higher confidence indicating greater belief in the node occupying the corresponding location at the given time:

$$P : \text{time} \mapsto \{(loc_1, conf_1), \dots, (loc_n, conf_n)\}$$

In our example model (< 1 kB per profile), a node's most-common locations (circular regions of fixed radius, possibly several 802.11b hops) are associated with times-of-day and days-of-week. More sophisticated predictive models that improve latency and overhead are possible, but even this simple example model is surprisingly effective. Full location profiles are sensitive, so in Section 3 we describe a method to use them for routing that hides profile contents, providing location privacy.

Profile Distribution Method: Human mobility and communication patterns enable a simple profile distribution mechanism. Location profiles can be first exchanged via face-to-face contact, just as phone numbers are often initially exchanged by family members and friends. Subsequent updates can be made in two ways: (1) pushing directly to contacts through the network or (2) pulling from a third-party location profile server (e.g., distributed location services). The push mechanism works when each user has a small number of contacts, most in close geographic proximity, who communicate more frequently than profiles change and eliminates location-server complexity. The pull mechanism is suitable when a large number of contacts communicate less frequently than profiles change. Distributed location services could also benefit from human mobility patterns by replacing real-time locations with location profiles, but providing anonymity and privacy would remain cumbersome.

Addressing Policy: The addressing policy translates the location–confidence tuples of the profile to a message delivery strategy specifying when and where packets will be sent. Only one of the locations can be correct, so the order and method in which they are tried influences the network throughput and latency trade-off (and may influence the privacy protocols). Directly routing to all predicted locations concurrently minimizes latency. Sequentially routing over a Steiner tree minimizes the throughput and energy cost.

Fallback Method: Location profile based routing fails when nodes are in unpredictable locations (roughly 7% of the time for the traces described in Section 2.1). For personal communication, this

may be tolerable (like cellphone dead-zones), although non-ideal. Space constraints prevent detailed discussion, but possible solutions include delayed delivery (buffering messages until the recipient returns to a predictable location) and rendezvous delivery (messages are routed to a predetermined rendezvous point, which the receiving node apprises of current forwarding instructions).

3. PRIVACY AND ANONYMITY

MANETs are open to untrusted observation and participation, inducing several security concerns, e.g., location and social network privacy. Furthermore, our proposed routing scheme at first appears to require that users trust their contacts enough to share location profiles, selectively giving up location privacy. Although this might be acceptable in some applications (e.g., when one’s contacts already know the motion patterns), often it is undesirable. We propose a reply-block- and pseudonym-based scheme that enables location profile-aided routing to operate without exposing location profiles (or identifying information), even to contacts. In this section, we define the desired security properties and describe our solution.

3.1 Attack Model

We assume that the attackers, in addition to participating, can observe all links in the network and collaborate using side-channels. They may have storage and processing capabilities exceeding those of a typical handheld device, allowing for traffic analysis of accumulated observations, and may triangulate the position of transmitters. We do restrict their number, assuming that economics dictates that conforming nodes will generally outnumber attacking nodes.

We do not consider attacks using information from outside of our protocols, e.g., taking photos of the human carrying a node for later identification. Similarly, we assume the other protocol layers, e.g., physical and application, are secure (as defined in Section 3.2). For example, wireless transmissions should not contain identifying analog “fingerprints” that would allow a node to be tracked. Of course, full system security requires that all layers have these security properties, but such provision is orthogonal across layers, so this work focuses on the network layer. Finally, we assume the majority of nodes obey our protocols, thus resisting routing attacks. We plan to quantify this resistance in future work.

3.2 Desired Anonymity and Privacy Properties

The trust concerns in MANETs are often addressed by listing specific requirements for privacy (the confidentiality of information) and anonymity (the confidentiality of the relationship between an identity and its information, i.e., attributes or actions). We believe this approach has two primary flaws. First, it focuses attention on the security provided, when the security *not provided* is of greater importance and interest. Second, it suggests a false separation between the attributes of an entity and its identity. In reality, the attributes themselves often allow identification (e.g., the Netflix dataset fiasco [16]), so separating them from a “traditional” identity (e.g., a name or social security number) is false protection. Further, predicting which attributes could, in the hands of a clever-enough attacker, allow identification is difficult. Therefore, we adopt a methodology that puts focus on the security not provided and endeavors to provide complete anonymity, removing the need to attempt to accurately distinguish identifying and non-identifying attributes.

We focus attention on the unprovided security by starting from an unrealistically strong, but easy-to-define, security goal, and relaxing it by describing specific security properties that it implies but we cannot yet provide. These relaxations have two sources. Type 1 relaxations are inherent to the underlying implementation technology (e.g., with wireless communication technology, the location of the transmitter of a packet is always linkable to the packet). These cannot be considered flaws of our protocols and must be accepted.

Type 2 relaxations are those induced by our protocols (e.g., we employ per-location pseudonyms to prevent tracking a node across space, but it remains possible to track a node in one location, across time). These are clearly flaws of our protocols and are opportunities for future improvement.

We term our unrealistically strong starting point *complete anonymity* and use it to address the false separation of identity and attributes. Put simply, complete anonymity requires that each observable attribute in the network (i.e., the act of transmission and each data attribute within) be unlinkable to the other attributes from the same entity (i.e., node). More precisely, in a network comprising n nodes, an observer should have belief $\frac{1}{n}$ for each node that a given attribute originated from that node. Equivalently, for any two attributes, an observer will have an equal belief in their originators being the same node or different nodes. This strong unlinkability requirement prevents the inference of identifying information. For example, network participation is anonymous because an identifier (a set of data attributes) is unlinkable to transmission (an action).

In MANETS, we can decompose complete anonymity into six unlinkability relationships over three attributes: actions (e.g., packet transmission), traditional identifiers (e.g., MAC address, name, pseudonym), and locations. The following list summarizes these six relationships, describing the Type 1 and 2 relaxations:

action–location: In a MANET, transmission location is obviously visible, so this link is an allowed Type 1 relaxation. Actions must still be unlinkable to past or future locations of their entity.

action–identifier: Our protocols (given in the following subsection) use visible per-location pseudonyms, resulting in a Type 2 relaxation: each action is linkable to exactly one pseudonym. Actions must still be unlinkable to other identifier types.

action–action: Action–location linkability induces a slight Type 1 relaxation: two actions linkable to the same location can be linked. However, two actions at different locations must still be unlinkable.

identifier–location: An identifier should not be linkable to the location of its entity. Again, our solution using per-location pseudonyms will violate this slightly, resulting in a Type 2 relaxation: each pseudonym is linkable to exactly one location of its entity.

identifier–identifier: Two identifiers for the same entity should not be linkable. For example, a pseudonym must be unlinkable to other identities (e.g., real name) and multiple pseudonyms for an entity must be unlinkable. For network transmissions, this means that the identifiers of the sender and receiver cannot be linked, providing *social network privacy*. For personal communication, contacts know each other, resulting in one Type 1 relaxation: communicating contacts can link each other’s identifiers. Pseudonyms induce one Type 2 relaxation: the pseudonyms of the sender and receiver for a one-hop (i.e., forwarding) transmission are linkable.

location–location: The current, past, and future locations of a node must be unlinkable. We allow one Type 2 relaxation: the allowed identifier–location link for per-location pseudonyms implies that past and future locations can be linked, but only when they are the same location. Critically, this provides *location privacy*, with the exception that the existence of a node at a single location may be tracked across time.

3.3 Unlinkability via Reply Blocks and Pseudonyms

This section presents our reply block- and pseudonym-based solution to provide the desired unlinkability. A formal argument that the properties are satisfied requires complete enumeration of all types of actions, identifiers, and locations and lengthy analysis. Such detail is beyond the scope and space constraints of this paper, so instead the solutions are presented with high-level arguments for their correctness. Roughly, the following arguments derive from the premise that two attributes are unlinkable if (1) they are never both available in the same context and (2) transitive application of known relationships cannot be used to link them.




Figure 3: Message flow for ordinary and multi-server reply blocks.

Geographic routing lends itself to our unlinkability requirements, because messages are addressed to locations, not identifiers. Identifiers are not visible in packet headers and thus the three identifier relationships are implicitly unlinkable by third parties. The sender and receiver themselves do know each other’s identifiers, so we use reply blocks, a variant of Chaum’s *mix-nets*, to disassociate information available at the sender (receiver’s identifier) from the information available at the receiver (receiver’s location and receiver’s actions) and vice versa, explicitly protecting the identifier relationships¹. A reply block is a routing instruction that guides a message from a sender through a mix-chain leading to the receiver. A mix-chain is composed of *mix-servers*, each of which disassociates the incoming and outgoing messages by reordering them and changing their appearance via layered decryption. Thus, observers (including the sender itself) cannot track the original message; at any point, only the previous and next mix-servers are known. We give detailed descriptions of applying reply block techniques in MANETs, including how senders choose the mix-servers composing a chain, in the remaining parts of this subsection.

Action-action links are also protected. This linking would require transitive application of other relationships: action A linked to X and action B linked to X implies A is linked to B. Aside from the allowed Type 1 exception when X is a location, no such X exists; the action-identifier and identifier-location relationships are unlinkable.

Location-location links are also protected. The location caches shared with a contact are encapsulated in reply blocks, so the actual locations are not revealed to the contact. Further, as with the action-action link, transitive linking of locations is not possible: the mix-chain dissociates locations from other attributes.

Location-based addressing has one significant problem. The predicted locations are inherently imprecise, so messages must be addressed to relatively large regions (several 802.11b hops in radius) and then flooded, wasting significant bandwidth and energy. To address this, we introduce pseudonyms as secondary addresses. Messages are addressed to both a location and a pseudonym (both encapsulated in the reply block), with the location used for initial routing and the pseudonym used in the destination region. Different pseudonyms are used in each location, preventing the pseudonyms from transitively linking other attributes. However, they still violate the strictest requirements, resulting in the previously mentioned Type 2 relaxations. Three of these, action-pseudonym, pseudonym-location, and, for one-hop sender-receiver links, pseudonym-pseudonym, are acceptable because the pseudonyms map one-to-one to an already visible attribute, location, and contain no additional useful information. The fourth though, is unfortunate. Pseudonyms persist across time and can be used to link the times when a node is in the same location (a type of location-location link). We are investigating possible remedies. An obvious possibility is frequently changing pseudonyms.

3.3.1 Reply Block Operation and Management

The chain in Figure 3 illustrates the use of an ordinary reply block, specifying a two-server mix-chain. Each transmission depends on those before, posing a deliverability problem. Each mix-server pro-

¹The usual caveats for mix-chains apply. Linking is possible if all nodes in the chain collaborate and global traffic analysis can potentially reveal message flows in some special circumstances.

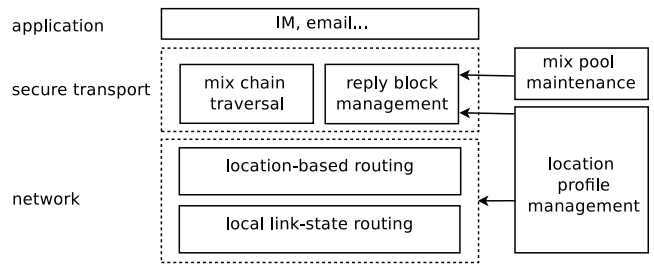


Figure 4: Main components of the location-centric network, with arrows representing service relationships.

vides a single common location, so with non-negligible probability the server will be unreachable at the time of attempted contact. We solve this problem by specifying multiple mix servers at each layer (also in Figure 3), increasing the probability of successful delivery. Each layer of the reply block is encrypted to three servers, who each remove an encryption layer and each forward the packet to the next three mix-nodes. Each server remembers the previous-next hop association. The receiver sends a message back through the fastest chain to complete, marking it as available for subsequent packets.

Reply blocks are location profiles anonymized by mix-chains, so managing them includes two main tasks: location profile management and mix-pool management. Each device needs to track its motions and keep its location profile up to date. Additionally, the mix-servers used in one’s reply blocks also need to be valid. When there are significant changes in a device’s location profile, or there are too many unreachable mix-servers in a reply block to permit any valid route, the reply blocks need to be updated accordingly.

3.3.2 Mix-Server Pool Management

Mix-server selection is important because if all mix-servers in a chain collaborate on an attack, the sender and receiver can be linked. Two selection requirements need to be satisfied. (1) Servers should have high probability of protocol compliance, reducing the chance that all servers in a chain improperly collaborate to trace a message. (2) They must be directly reachable by locations, instead of reply blocks, to prevent an infinite chain of reply blocks. For traditional Internet mix-chains, services are chosen from semi-trusted published lists, as with Tor [9]. However, this method is not suitable for MANETs; no semi-trusted authority who could publish such a list exists. A new method for choosing mix-servers is needed.

We assume that physical attacker density is limited by economic constraints, and thus propose that each node individually maintain a pool of mix-servers chosen randomly from the various one-hop neighbors it encounters. This density assumption could be violated by Sybil attacks [10], in which one device pretends to be many, so existing techniques leveraging signal strength measurements are used to detect Sybil identities during pool population [8]. As a node moves in the network, it asks one-hop neighbors to act as future potential mix-servers. Willing neighbors respond with a single (common location, pseudonym) address and an associated contact probability. Entire profiles are not shared to preserve location privacy. The requester saves the information from non-Sybil neighbors in its mix-server pool for future usage.

4. LOCATION-CENTRIC NETWORK

Now that the two most important pieces—location profile-aided routing (see Section 2) and reply block-based privacy (see Section 3)—have been described, we present the architecture of our location-centric network for secure personal communication. System scalability relies on location-profile-aided routing, into which we incorporate confidentiality and privacy mechanisms. As illustrated in Figure 4, the system comprises three layers, (1) application, (2) secure transport, and (3) network. The target application is low-bandwidth and delay-tolerant text-based communication, e.g., email

and text messaging. The secure transport layer provides confidential and anonymous host-to-host delivery using mix-chains. The reply blocks constructed by a host and shared during face-to-face contact act as the transport layer addresses. The network layer delivers messages between mix-nodes using geographic routing. A network address is a two-tuple containing a pseudonym and location. Keys for encryption are exchanged face-to-face between contacts, so no PKI is required. In this section, we will describe the network and secure transport layers in more detail.

Network. Geographic routing (e.g., GPSR [12]) is the backbone of the network, providing routing scalability. Location profiles are exchanged face-to-face, providing location-distribution scalability, normally the Achilles' heel of geographic routing. A node's movement within a small region prevents addressing destinations by precise coordinates, so we propose using geographic routing for coarse delivery and reactive routing near the receiver. Thus, a receiver is addressed by both a destination region and a pseudonym. When a message reaches its destination region, the intermediate node at the boundary transitions from geographic routing to local link-state routing. If a route is known, the message is delivered along it. Otherwise, a route discovery message is broadcast to discover one. If the node is unreachable, the message is dropped.

Secure transport. The transport layer provides host-to-host secure communication channels. A channel is a mix-chain between the sender and the receiver, constructed according to the receiver's reply block, that provides the desired location privacy and sender-receiver unlinkability. It is constructed according to the receiver's reply blocks. End-to-end encryption provides confidentiality.

We now describe the operation of the transport layer, responsible for delivering messages from the application layer to the destination node. To deliver a message, the sender first determines whether a channel to the destination is already available. If so, the message is sent via the channel. Otherwise, the sender sends a setup message using the receiver reply block with the highest contact probability. If the sender does not receive a response within a constrained time, it concludes that the receiver is not at the corresponding location of that reply block and repeats this process for the remaining reply blocks, until a response is received or all the reply blocks have been used. Our preliminary analysis indicates that, on average, receivers will be contacted via a reply block 93% of the time. When a response is received, the sender marks the channel as valid, sets a timeout for it, and messages are delivered thorough this channel. The receiver can respond via this channel as well, although the routing is not, in general, symmetric. Messages are encrypted with a session key established during the channel setup process.

Overhead. Energy consumption is a significant concern, especially since much of the work is forwarding others' traffic and does not directly benefit the user paying the cost. Current 802.11 ad hoc technology is inefficient, depleting cellphone batteries in several hours (the power save mode is only for AP networks). Implementing a periodic sleep option for ad hoc mode will be necessary. Even with reasonable battery life, some selfish users might refuse to forward traffic for others, but we believe they will be in the minority. Most people derive some satisfaction from helping others, particularly at low cost (e.g., charging ones' phone each night instead of every other). An application feature displaying statistics about the number of conversations relayed could encourage such altruism.

5. FUTURE WORK

Performance-cost trade-offs (e.g., latency vs. number of concurrent connections, anonymity vs. mix-chain length, etc.) influence efficiency (e.g., power, hops, latency, etc.) and thus require further investigation. **Human mobility patterns:** Predictable motion patterns can substantially improve routing efficiency for networks of human-carried devices, but deriving full benefit requires future study and characterization. **Security:** The reply block mechanism must be formalized and verified. Further, the dependence on uncontrol-

lable parameters (e.g., temporal and spatial distributions of available mix-servers) must be analyzed. **Applications:** Adapting application semantics to the native properties of MANETs (e.g., user-controlled selective application-layer forwarding) may yield further gains.

6. REFERENCES

- [1] M. Abolhasan, T. Wysocki, and E. Dutkiewicz. A review of routing protocols for mobile ad hoc networks. *Ad Hoc Networks*, 2(1):1–22, Jan. 2004.
- [2] N. Anderson. Tweeting tyrants out of Tunisia: Global Internet at its best. *Wired.com*, Jan. 14 2011. <http://www.wired.com/threatlevel/2011/01/tunisia>.
- [3] R. Barr, Z. J. Haas, and R. van Renesse. Scalable wireless ad hoc network simulation. In J. Wu, editor, *Handbook on Theoretical and Algorithmic Aspects of Sensor, Ad Hoc Wireless, and Peer-to-Peer Networks*, chapter 19, pages 297–311. CRC Press, 2005.
- [4] A. Beresford and F. Stajano. Location privacy in pervasive computing. *IEEE Pervasive Computing*, 2:46–55, Jan. 2003.
- [5] D. L. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, Feb. 1981.
- [6] C. Cortes and D. Pregibon. Signature-based methods for data streams. *Data Mining and Knowledge Discovery*, 5(3):167–182, July 2001.
- [7] S. M. Das, H. Pucha, and Y. C. Hu. Performance comparison of scalable location services for geographic ad hoc routing. In *Proc. Int. Conf. Computer Communications*, pages 1228–1239, Mar. 2005.
- [8] M. Demirbas and Y. Song. An RSSI-based scheme for Sybil attack detection in wireless sensor networks. In *Proc. Int. Symp. on a World of Wireless, Mobile, and Multimedia*, pages 564–570, June 2006.
- [9] R. Dingleline, N. Mathewson, and P. Syverson. Tor: the second-generation onion router. In *Proc. USENIX Security Symp.*, page 21, Aug. 2004.
- [10] J. Douceur. The Sybil attack. In *Proc. Int. Wkshp. Peer-to-Peer Systems*, pages 251–260, Mar. 2002.
- [11] M. C. Gonzalez, C. A. Hidalgo, and A.-L. Barabasi. Understanding individual human mobility patterns. *Nature*, 453:778–782, June 2008.
- [12] B. Karp and H. Kung. GPSR: Greedy perimeter stateless routing for wireless networks. In *Proc. Int. Conf. Mobile Computing and Networking*, pages 243–254, Aug. 2000.
- [13] OpenNet Initiative. Internet filtering in Tunisia, 2009. <http://opennet.net/research/profiles/tunisia>.
- [14] R. Pike. More unrest in the Middle East results in Internet disruptions. *TechieInsider.com*, Feb. 19 2011. <http://www.webcitation.org/query?url=www.techieinsider.com/news/6485&date=2011-03-17>.
- [15] I. Rhee, M. Shin, S. Hong, K. Lee, and S. Chong. On the Levy-walk nature of human mobility. In *Proc. Int. Conf. Computer Communications*, pages 924–932, Apr. 2008.
- [16] B. Schneier. Why 'anonymous data' sometimes isn't. *Wired.com*, Dec. 13 2007. http://www.webcitation.org/query?url=www.wired.com/politics/security/commentary/securitymatters/2007/12/securitymatters_1213&date=2011-04-27.
- [17] J. Wu, Y. Zhang, Z. M. Mao, and K. Shin. Internet routing resilience to failures: Analysis and implications. In *Proc. Int. Conf. Emerging Networking Experiments & Technologies*, pages 1–12, Dec. 2007.
- [18] F. Xue and P. Kumar. *Scaling Laws for Ad Hoc Wireless Networks: An Information Theoretic Approach*. NOW Publishers, 2006.