



 Latest updates: <https://dl.acm.org/doi/10.1145/3719027.3765133>

RESEARCH-ARTICLE

Amigo: Secure Group Mesh Messaging in Realistic Protest Settings

DAVID INYANGSON, Johns Hopkins University, Baltimore, MD, United States

SARAH RADWAY, Harvard University, Cambridge, MA, United States

TUSHAR M. JOIS, City College of New York, New York, NY, United States

NELLY FAZIO, City College of New York, New York, NY, United States

JAMES MICKENS, Harvard University, Cambridge, MA, United States

Open Access Support provided by:

Harvard University

City College of New York

Johns Hopkins University



PDF Download
3719027.3765133.pdf
01 April 2026
Total Citations: 0
Total Downloads: 2286



Published: 22 November 2025

Citation in BibTeX format

CCS '25: ACM SIGSAC Conference on
Computer and Communications Security
October 13 - 17, 2025
Taipei, Taiwan

Conference Sponsors:
SIGSAC

Amigo: Secure Group Mesh Messaging in Realistic Protest Settings

David Inyangson*
Johns Hopkins University
Baltimore, Maryland, USA
dinyang1@jhu.edu

Sarah Radway*
Harvard University
Cambridge, Massachusetts, USA
sradway@g.harvard.edu

Tushar M. Jois
City College of New York
New York, New York, USA
tjois@ccny.cuny.edu

Nelly Fazio
City College of New York
New York, New York, USA
nfazio@ccny.cuny.edu

James Mickens
Harvard University
Cambridge, Massachusetts, USA
mickens@g.harvard.edu

Abstract

During large-scale protests, a repressive government will often disable the Internet to thwart communication between protesters. Smartphone mesh networks, which route messages over short-range, possibly ephemeral, radio connections between nearby phones, allow protesters to communicate without relying on centralized Internet infrastructure. Unfortunately, prior work on providing secure communication in Internet shutdown settings fails to adequately consider protester needs. Previous attempts fail to support efficient *private group* communication (a crucial requirement for protests), and evaluate their solutions in network environments which fail to accurately capture link churn, physical spectrum contention, and the mobility models found in realistic protest settings. In this paper, we introduce Amigo, a novel mesh messaging system which supports group communication through a decentralized approach to continuous key agreement, and forwards messages using a novel routing protocol. Amigo is uniquely designed to handle the challenges of ad-hoc routing scenarios, where dynamic network topologies and node mobility make achieving key agreement non-trivial. Our extensive simulations reveal the poor scalability of prior approaches, the benefits of Amigo's protest-specific optimizations, and the challenges that still must be solved to scale secure mesh networks to protests with thousands of participants.

CCS Concepts

• Security and privacy → Security protocols.

Keywords

mesh messaging, censorship circumvention, anonymity, routing

ACM Reference Format:

David Inyangson, Sarah Radway, Tushar M. Jois, Nelly Fazio, and James Mickens. 2025. Amigo: Secure Group Mesh Messaging in Realistic Protest Settings. In *Proceedings of the 2025 ACM SIGSAC Conference on Computer*

*Both authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution 4.0 International License. *CCS '25, Taipei, Taiwan*

© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1525-9/2025/10
<https://doi.org/10.1145/3719027.3765133>

and *Communications Security (CCS '25)*, October 13–17, 2025, Taipei, Taiwan. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3719027.3765133>

1 Introduction

Protests are an important tool in the fight against authoritarian governments. By occupying a physical space, ordinary citizens can express their grievances and bring visibility to their concerns. In countries where political freedoms are curtailed, protests may be the *only* mechanism for the public to express their grievances.

When protests become movements comprising thousands of people, communication becomes vital for organizing protest actions and keeping participants safe. In a modern protest, the demonstrators often rely on smartphone-based messaging platforms like WhatsApp or Telegram to communicate; using these apps, protesters learn about rally points and inform each other about the location of the police or the military [10]. Authoritarian governments are thus highly motivated to intercept or alter the messages that protesters exchange. Using intercepted messages, the government can identify protest leaders and shift police deployments to disrupt nascent rally points; using altered messages, the government can send spoofed messages, ostensibly from protest leaders, that direct protesters towards harm. These concerns are not theoretical. For example, law enforcement – even in relatively liberal governments – possess *cell-site simulators* which create fake cellphone towers [2]. By inducing protester phones to pair with a simulator, law enforcement can launch a variety of attacks, e.g., downgrading the phone's cellular connection to the unencrypted 2G protocol, allowing the government to intercept messages to and from that phone [33].

A repressive government may also try to degrade or completely disrupt app-based communication during a protest. For example, during the 2019 Hong Kong anti-ELAB protests, the government (or its partisans) launched a denial-of-service attack against Telegram servers [3]. In 2019, Iran responded to protests involving economic discontent by enacting a large-scale Internet shutdown [14]. Similarly, in Myanmar, Internet shutdowns are frequently imposed to prevent civilians from documenting human rights violations or accessing information about military operations [6]. Overall, there were as many as 283 Internet shutdowns across 39 counties in 2023 alone – the worst year for Internet shutdowns on record [7].

Thus, protesters need a robust way to exchange authenticated, encrypted messages, even if the government has disabled Internet connectivity or tampered with cellular routing infrastructure.

A promising approach is to exchange encrypted messages via an *ad-hoc mesh network* built from point-to-point radio links (e.g., Bluetooth between smartphones). In a mesh network, there is no centralized routing infrastructure; instead, user devices exchange messages over (potentially ephemeral) short-range links created via physical proximity as users move in space. Multi-hop message forwarding over these point-to-point links enable information exchange between users too far away to establish direct connections.

Because mesh networks do not rely on centralized infrastructure, there is no single point of failure for a regime to attack. Unfortunately, classic mesh protocols [8, 9, 73] do not safeguard user privacy, revealing message senders and recipients to even passive observers. These protocols are ill-suited for protest scenarios where users wish to hide their identities from the government [10].

The research community has introduced mesh networks that attempt to hide user activity [22, 35, 52, 57–61, 63, 68]. However, previous systems cannot practically handle large-scale protests for two major reasons. First, prior systems lack efficient (or any) support for private groups, despite the widespread use of group-based messaging in protest settings [10]. Second, prior systems use routing protocols which are scalable in benign network conditions, but can experience persistent congestion collapse for the mobility patterns and smartphone radio technologies found in realistic protest settings; the result of the persistent network congestion is a very low rate for application-level message delivery.

To address these challenges, we introduce *Amigo*, a new system for efficient, secure mesh communication in protest settings. *Amigo*'s high-level novelty is that it is designed *full stack*, accounting for both security challenges at the cryptographic layer and performance/reliability challenges at Layers 1 and 2 of the network stack. To protect messages exchanged by protesters, *Amigo* leverages a new protocol for continuous group key agreement (CGKA), efficiently providing security properties specifically required by protesters (e.g., post-compromise security, and fast removal of a member from a group). Additionally, *Amigo* introduces a new mesh routing protocol which, compared to earlier protocols, enjoys higher delivery rates for mobility patterns often seen in real protests. Our comprehensive simulations, which leverage those realistic mobility models and high-fidelity representations of low-level network phenomena like radio collisions between nearby phones, demonstrate that *Amigo*'s protest-specific optimizations provide secure, protester-desired functionality with better routing performance than earlier work. However, our evaluation results also reveal fundamental scalability challenges for all known approaches to secure mesh networking (including *Amigo*); a fundamental contribution of the paper is illuminating those practical challenges (previously unknown due to the lower-fidelity simulations of prior work), and suggesting several directions for future work.

Contributions. In summary, we contribute the following:

- We introduce a new protocol for decentralized group key agreement in mesh networks. The protocol tolerates an unreliable mesh that may drop or reorder messages, providing important security properties missing from prior schemes. It is also efficient with respect to network consumption and CPU usage – critical for a protocol that runs atop resource-constrained phones.
- We introduce a new routing approach, based on clique-level forwarding, tailored to protest-based mobility patterns.
- To evaluate *Amigo*, we introduce new mobility models which capture how real-life protesters organize in physical space, both during stable times and when external shocks arise (e.g., the unexpected appearance of the police). Our evaluation experiments also capture important Layer 1 and Layer 2 network phenomena that were ignored in prior work on secure mesh networking.
- We test *Amigo*'s performance when *Amigo* uses our clique-based routing approach and routing approaches from prior work, demonstrating (1) the benefits of *Amigo*'s protest-specific optimizations at the cryptographic layer and the routing layer, and (2) the scalability and robustness challenges (previously unknown) that still remain unsolved, given the challenging network environment created by protest-based mobility patterns and the nuances of real-life short-range communication technologies. Our results suggest mesh routing protocols previously thought to be practical will often function poorly in realistic protest environments due to a self-reinforcing cycle of network congestion. We discuss the implications of our results on future research, and will open-source all of our code upon publication as artifacts.

In the full version of our paper [38], we provide additional details on our protocols as well as supplementary results.

2 Motivation and Background

A large-scale protest involves hundreds or thousands of people gathered in (and moving throughout) a shared physical space. The success of a such a protest increasingly depends on the ability of protesters to organize via smartphone-based communication [10, 13, 48, 65]. Without the ability to effectively communicate, critical information regarding group movement, law enforcement activity, and collective decision-making propagates slowly among protesters. The consequences for the protesters are dire: physical assault, imprisonment, or even death [28].

As a case study, consider the widespread 2019–2020 protests in Hong Kong. The protests erupted in response to the government's attempt to pass the Extradition Law Amendment Bill (ELAB). The bill would have allowed Hong Kong citizens to be extradited to China; many Hong Kong citizens feared that the bill would further erode Hong Kong's legal independence from China [49, 50]. Using smartphone apps like WhatsApp, Telegram, and Signal, protesters decided where and when to hold demonstrations. Once protesters arrived on site, these apps allowed protesters to respond in real time to events like the arrival of police [10].

The ability of citizens to efficiently organize *during* the protest itself was a marked departure from static, pre-smartphone-era demonstration strategies. In the anti-ELAB protests, a column of marchers might quickly decide to split into groups, each taking a path down a different street and abruptly changing directions in response to newly-placed barricades; using apps like Telegram, the independent groups could synchronize their long-term movement trajectories and later reform into a single human chain [3].

Prior work has explored smartphone-based *mesh networks* as a way for protesters to communicate without centralized Internet access. Mesh networks leverage short-range radio technologies like Bluetooth or Wi-Fi Direct, routing messages between a sender

and a receiver across potentially ephemeral radio connections between mobile nodes. In the protest context, *end-to-end encryption* of signed/authenticated messages ensures that network eavesdroppers cannot inspect cleartext or alter messages without detection. An ideal mesh network would also limit the impact of an attacker who discovers the keys used by protesters at time t . *Forward secrecy* ensures the attacker cannot break the confidentiality of messages sent before t . *Post-compromise security* means that, after the key leakage at t , the attacker cannot break the confidentiality of messages sent *after* t . Broadly, forward and post-compromise security are implemented by having participants rotate their keys periodically.

Protester needs. Albrecht et al. performed wide-ranging interviews with 11 anti-ELAB protestors from Hong Kong [10]. Below, we highlight the relevant findings from that study with respect to how communication apps for protesters should be designed:

- (1) **Protest communication occurs primarily in group chats.** Groups can range from a handful of members to thousands.
- (2) **Groups must be able to securely add members as a protest unfolds.** Some protesters may hear about a demonstration offline, e.g., word-of-mouth. Upon arrival, such a person may wish to join a group for protest coordination. So, protest-focused messaging requires the ability to add new members on-the-fly.
- (3) **Groups must be able to remove members.** Protesters are resigned to the fact that law enforcement may be able to infiltrate groups (especially larger ones). Protesters also understand that law enforcement may coerce a loyal group member to nonetheless relinquish their device and their login credentials. Thus, revoking a user’s group membership is critical functionality.
- (4) **Communications should be possible during an Internet shutdown.** Protest-focused communication apps should still function if centralized Internet infrastructure has been degraded. During the Telegram blackout in the anti-ELAB protests, some protesters attempted to use the mesh networking app Bridgefy [26] to exchange messages, but it failed to handle the load presented by the large numbers of protesters [10].
- (5) **Communication should be anonymous.** The threat of retaliation from law enforcement requires that participation in an online protest group should not reveal the real-life identity of the sender or the receiver. The peer-to-peer nature of mesh networking does not automatically provide anonymity. For example, security researchers have found that Bridgefy [26] protects neither message confidentiality nor message authenticity, and allows outsiders to deanonymize users [11].

3 System Design

At a high level, Amigo’s goal is to provide a mesh network that meets the protester needs described above while being reliable, low-latency, secure, and resource-efficient. *Reliability* means that Amigo successfully forwards the vast majority of messages to their intended recipients, despite the stochastic nature of pairwise node connectivity. *Low-latency* forwarding means that Amigo delivers messages as quickly as possible, given the prevailing network conditions. *Secure* forwarding ensures that messages are end-to-end encrypted, integrity-protected, and enjoy both forward security and post-compromise security; additionally, key agreement among

group members should not leak cryptographic information to participants outside of the group. *Resource efficiency* is important because Amigo targets smartphone deployments in which network bandwidth, CPU cycles, memory space, and available power are limited compared to traditional server-based environments. To the best of our knowledge, Amigo meets or outperforms all prior work (§8) along these evaluation metrics. However and importantly, our evaluation experiments (§7) illuminate fundamental performance and robustness challenges for secure mesh networks—challenges that must be solved for these systems to reliably scale to protests involving thousands of participants (§9).

Architecture. Figure 1 depicts Amigo’s high-level architecture. Users interact with Amigo via the **application layer**, creating and destroying groups, adding or removing members to groups, and exchanging messages with group members. Users directly generate and read text messages or picture messages; behind the scenes, those messages are protected with keys negotiated by Amigo’s **session layer**. The session layer implements an efficient CGKA scheme which we describe in detail in Section 4. Amigo focuses on small to medium-sized groups with <200 members; such close-knit groups are perceived by protesters as more trustworthy and less susceptible to infiltration by the government [10].

The session layer sends and receives user-generated messages and CGKA-generated messages via the **routing layer**. As we describe in Section 5, Amigo is compatible with three routing schemes from prior work on mesh networking. Amigo also introduces a new approach, dynamic clique routing, which is tailored for the unique connectivity challenges of protest scenarios. As shown in Figure 1, the basic idea behind all of the routing schemes is that, when two nodes wander into communication range, they exchange locally-buffered messages that were generated locally or received from other peers in the past; the routing schemes differ in which messages are exchanged with which peers.

At the bottom of the Amigo stack is the **link layer**. Similar to prior work on secure meshes [51, 57, 58], Amigo leverages off-the-shelf point-to-point radio technologies like Bluetooth and Wi-Fi Direct. Using off-the-shelf hardware for Layers 1 and 2 of the network stack makes Amigo easily deployable, but traditional short-range radio protocols lack understanding of end-to-end congestion metrics, and therefore may unwittingly overload the network and trigger precipitous declines in application-level delivery rates.

Threat model. As explained by Albrecht et al. [10], the core goals of an authoritarian regime during a protest are to:

- (1) identify message senders and group members,
- (2) read the cleartext content of in-transit messages,
- (3) tamper with in-transit messages, and
- (4) compromise protester devices after a protest to decrypt messages that protesters have already received, or inspect/alter messages that protesters will send in the future.

We assume that the adversary has multiple vantage points in the network, but cannot view all traffic emanating from all nodes. We also assume that the adversary, through malicious or lawful means, can physically access a group member’s device [76]; after doing so, the adversary gains access to the messages and the key material associated with all groups on that device. Knowledge of the keys

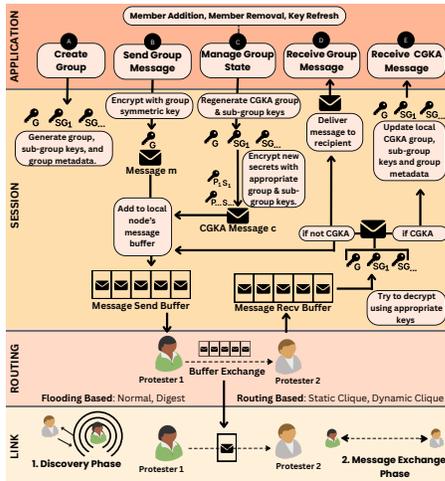


Figure 1: An overview of Amigo and its layers.

in particular allows the attacker to monitor group communication indefinitely until the attacker is removed from the group.

Given such an adversary, Amigo has three primary security goals. First, Amigo wants to provide messages with confidentiality and integrity (including forward security and post-compromise security). Second, Amigo wants to provide *outsider anonymity*, such that if an eavesdropper receives a message intended for a group that does not include the eavesdropper, the eavesdropper cannot learn the identities of the group members. We discuss these two first security goals in Section 4, evaluate Amigo’s ability to provide these goals in Section 7, and provide a proof in the full version of this paper [38].

Amigo’s third security goal is to defend against denial of service attacks in which malicious nodes intentionally drop messages. Amigo’s defenses leverage prior work in which nodes prefer to route messages via trusted peers [51, 58]. However, Amigo’s core novelty lies in the way that Amigo protects group messaging, so we do not focus on DoS protection, e.g., circumventing signal jamming attacks, in the remainder of this paper.

Non-goals. In designing Amigo, we explicitly sought to avoid network behavior that would trigger pathologies in Layer 1 or Layer 2 dynamics. Indeed, a primary contribution of this paper is the demonstration that secure ad-hoc meshes can appear performant when Layer 1 and Layer 2 dynamics are ignored, but may actually experience debilitating congestion collapse when these dynamics are considered (§7). This decision left some attacker behaviors outside of our threat model. For instance, Amigo does not prevent a multi-vantage-point attacker from determining when a user sends an encrypted message because, to the best of our knowledge, state-of-the-art defenses against such attacks require the insertion of dummy messages to thwart correlation analyses [57]; our simulations (which incorporate Layer 1 and 2 effects) show that secure meshes can experience network collapse *even if those meshes do not inject dummy messages*. In contrast, DoS protections that leverage trusted routing peers [51, 58] add minimal network load.

4 A CGKA Protocol for Mesh Messaging

In this section we discuss Amigo’s CGKA which is designed to provide the properties protesters need for group communication.

Intuition. Protesters need secure group communication, which intuitively calls for a key agreement protocol that addresses the nuances of mesh messaging during a large-scale protest. For example, the unreliable nature of smartphone mesh links requires our protocol to minimize the number of messages required for consistent key agreement. Another nuance is that in a protest setting, parties do not require *full* anonymity. Protesters vet other members before inviting them into group communication and do not need to be anonymous within their groups [10]. But, they do require *outsider anonymity* [34]; group messages sent within the mesh should not include any sender or recipient identifiers to those who are outside the intended group – a natural and appropriate relaxation. Notions of *forward secrecy* and *post-compromise security* are also necessary to limit the impact of a compromise (e.g., due to arrest), and to return the group to a safe state afterwards.

CGKA protocols are promising here [21, 29, 45]. These protocols can provide forward secrecy and post-compromise security while also efficiently structuring group members (i.e., through a binary tree) to enable updates of group key material and membership with *sublinear communication complexity*, unlike pairwise approaches (e.g., Signal [1]).

Adapting existing CGKA protocols to a large-scale protest is not trivial, however. Traditional CGKA protocols [21, 29, 37] are oriented towards centralized infrastructure, requiring all members to apply protocol messages in the same sequence to maintain consistent state. But, we cannot assume this constraint holds in a mesh during an Internet shutdown.

Another issue is that existing CGKAs protocols do not consider the unique threats of the protest setting. For instance, while CGKAs refresh cryptographic state, the time intervals between these states, or *epochs*, vary. Epochs are not tied to messages but rather group state changes, e.g., a new member. The frequency of these key rotations may be inadequate given the dynamics of the protest environment; forward secrecy and post-compromise security *per message* may be more desirable.

Due to the aforementioned challenges, CGKA protocols remain largely unexplored in the protest use case. In this paper, we leverage the advantages of CGKA protocols by appropriately adapting them to the Internet-shutdown protest setting. We propose a CGKA that is bandwidth-efficient, dynamic, outsider anonymous, and enables consistent state for group members regardless of protocol message order. We also orient our CGKA to specifically handle compromise during a protest scenario (e.g., through efficient removals) and provide stronger (per message) and standard (per epoch) notions of forward secrecy and post-compromise security.

Comparison to MLS. As a concrete example of the challenges in porting existing CGKAs to the setting of Amigo, let us consider Messaging Layer Security (MLS) [37], perhaps the most prominent prior CGKA protocol. MLS is an IETF-standardized asynchronous group key agreement protocol that provides forward secrecy and post-compromise security for a wide range of group sizes, ranging from two to thousands. In a typical MLS configuration, a centralized provider is responsible for two services: authentication and

delivery. The authentication service binds public key material to application-specific identifiers and enables clients to receive and verify these bindings. The delivery service is responsible for distributing messages and key material, but can also help reconcile group states and enforce a global ordering on messages [37].

While sufficient for traditional group messaging scenarios, MLS fails to adequately meet the protester specific needs we highlight in Section 2. One reason for this is that MLS is primarily intended for end-to-end encrypted messaging *over the Internet*, taking advantage of centralized infrastructure (i.e., servers) to realize the authentication and delivery services. The mesh setting we consider is fundamentally different, involving (potentially ephemeral) short-range wireless links, dynamic network topology, and resource-constrained devices. Thus, the reliable centralized services required by MLS for authentication or delivery are not feasible. Without these services, the MLS specification is unclear about important details, such as group synchronization, authentication, and secure message dissemination; these are left to implementations to determine [37].

Additionally, MLS does not sufficiently defend against the mesh adversary (§3). MLS considers the Internet Threat Model [36], which does not consider an adversary who can shut down Internet connectivity entirely. Furthermore, its threat model does not sufficiently consider threats arising from metadata analysis, which an adversary can use to break the outsider anonymity of groups. In fact, the core MLS specification does not attempt to protect certain group metadata like unique group IDs and epoch numbers [37]. As a result, a compromised delivery service or eavesdropper may use this metadata to enumerate unique groups, message frequencies, and message types, leading to inferences on group patterns, membership churn, and other protest dynamics.

Furthermore, MLS does not enforce secure configurations for the issuance of key packages, transport security, ratchet tree distribution, and other design decisions [37], which is intended to allow for customization. In our setting, however, relaxed security and privacy configurations can easily lead to identifying group members, behaviors, and dynamics, putting protesters at risk.

4.1 CGKA Definition

We first introduce the definition and syntax of a CGKA, adapted from Klein et al. [45], which we will use throughout this section.

Definition 4.1 (Asynchronous Continuous Group Key Agreement). An asynchronous continuous group key agreement protocol is a tuple of algorithms $\text{CGKA} = (\text{Init}, \text{Add}, \text{Rem}, \text{Upd}, \text{Proc})$ with:

$(\gamma) \leftarrow \text{Init}(N, ID_0, 1^\lambda)$: A founding member, ID_0 , generates their initial group state γ . This initializes the set of members $G = \{ID_0\}$ and generates a group state γ with a maximum capacity of N participants, where λ is the security parameter.

$(\gamma', S) \leftarrow \text{Add}(\gamma, pk_{new}, ID_{new})$: A member of a group can invite a new party, ID_{new} , to join by first sending the current group state in a welcome message, W . After receiving and processing W , the new member inserts themselves into the group G and sends a protocol message S to each member in G .

$(\gamma', S) \leftarrow \text{Rem}(\gamma, ID_{rem})$: A member removes another, ID_{rem} , by sending a protocol message, S , intended for $G \setminus \{ID_{rem}\}$.

$(\gamma', S) \leftarrow \text{Upd}(\gamma, ID_i)$: A member can refresh key material in γ by sending the corresponding protocol message S to G .

$(\gamma') \leftarrow \text{Proc}(\gamma, S)$: Upon receiving a protocol message S , a member modifies their local group state, γ .

Security game. We use our threat model (§3) to inform our game-based security definition, which can be found in our full version [38].

4.2 Meeting Protester Needs

Enabling group synchronization over a mesh. As noted previously, traditional CGKA protocols require mechanisms for ensuring protocol messages arrive in a consistent total order for all recipients, which is not realistic over a mesh. Amigo applies the idea proposed by Weidner [70]; instead of changes to state that require overwriting key material held by nodes in the ratchet tree (e.g., in TreeKEM [21]), state changes *combine* existing and newly generated material, an operation denoted by \star . So, members can apply protocol messages regardless of the order. In this way, Amigo can achieve *eventual consistency* of the CGKA state over a mesh as long as all protocol messages are received.

We build this \star primitive in Amigo’s CGKA via elliptic curves, where a public key pk is derived by multiplying a (secret) scalar sk with the curve’s generator point P . Thus, we can instantiate \star as: $(sk_{x\star y} = sk_x + sk_y, pk_{x\star y} = (sk_x + sk_y) \cdot P)$.

Stronger security guarantees. Notions of forward secrecy and post-compromise security for traditional CGKAs [21, 29] relate to key rotation per epoch rather than per message. Protester safety may sometimes require stronger guarantees. As a result, we utilize updatable public key encryption (UPKE). Informally, a UPKE scheme is a tuple of algorithms $\text{UPKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ where encrypting a message outputs a new public key alongside a ciphertext, and decryption of a ciphertext outputs the message and associated new secret key. These operations can be used to refresh cryptographic state per message by applying the new key material on the old using our aforementioned \star operation. We define our UPKE and security model formally in the full version of this paper [38].

4.3 The Amigo CGKA

With CGKA, \star , and UPKE as building blocks, we now show how we can achieve secure group mesh messaging in the large-scale protest setting. Our CGKA implements the Amigo session layer (Figure 1). The core data structure of Amigo’s CGKA is a *ratchet tree*, a binary tree where each node holds asymmetric key material [21, 29]. The leaf nodes are associated with group members, while intermediary nodes up to and including the root allow for (public key) communications to subsets or the entire group respectively. Group members should possess all public keys within the ratchet tree, but only the secret key material of their leaf node and its direct ancestors. As opposed to pairwise approaches where key updates and messages may require linear encryptions, a CGKA protocol utilizing a ratchet tree reduces key updates to logarithmic encryptions and messaging to even a single encryption. Ratchet trees have found practical use as a result (e.g., Messaging Layer Security [37]).

Amigo is designed and optimized for group communication. Our CGKA enables parties to form groups on the fly, make changes to group membership, and communicate efficiently within those groups. Members can belong to a single group, or multiple distinct groups each holding a unique group state. This means parties can

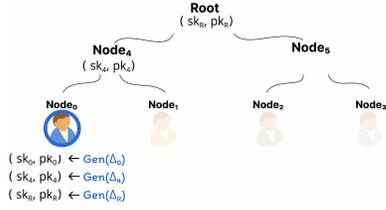


Figure 2: Initializing a group in Amigo. Only member ID_0 is in the group, occupying $Node_0$.

be members in the same groups without being aware of this fact, a level of privacy useful in a protest setting. Every group G has its own group state γ , consisting of the following fields:

- $\gamma_{members}$: List of members' group-specific identifiers¹
- γ_{pk} : Vector of all known public keys in the group
- γ_{sk} : Vector of all known secrets in the group
- γ_{tree} : Binary tree mapping between identifiers and keys

State operations transition the group state held by members over periods denoted as epochs, θ . The CGKA operations (Definition 4.1) enable dynamism, forward secrecy, and post-compromise security by rotating keys held in the ratchet tree via the generation of random seed values Δ_i , corresponding to a unique node in the ratchet tree.² A party initiating a state operation generates these seeds and uses them to create and send protocol message S to the entire group. Upon processing S , members have applied the necessary key updates, evolving their group state from γ to γ' , and entering into a new epoch. S contains the following fields:

- S_{id} : Identity of initiator
- S_{op} : Type of operation ($op \in \{\text{Add, Rem, Upd}\}$)
- S_{mod} : Identity of the node to be added or removed³
- S_{seeds} : Vector of random seeds, where each seed is encrypted under its associated node's public key
- S_{keys} : Vector of public keys derived from S_{seeds}

We now describe how Amigo instantiates the CGKA algorithms specified in Definition 4.1:

Init: As shown in Figure 2, the founding group member, ID_0 , generates a ratchet tree of a specified size N . This ratchet tree is initially empty. ID_0 inserts themselves into the left-most leaf in the ratchet tree by deriving separate, randomly sampled Δ_i for themselves and each ancestor node along its direct path to the root. From each Δ_i , a keypair is generated and stored at the appropriate node: $(sk_i, pk_i) \leftarrow \text{Gen}(\Delta_i)$, where Gen is the generation algorithm of the underlying UPKE scheme (available in [38]). Thus, initial group state γ_0 for ID_0 contains the key material they generated, the ratchet tree mapping, and their group specific identifier.

Add: If ID_i wants to add ID_j to the group, ID_j needs the necessary group state to become a member. Thus, ID_i reserves the left-most open leaf node in the ratchet tree for ID_j and sends a *welcome*

¹These identifiers, e.g., a pseudorandom string, are used in processing changes to group cryptographic state. They are group-specific, so an individual in multiple groups has a different identifier per group. An adversary cannot use one identifier to track members across groups.

²For **Add** and **Upd**, these nodes include the leaf and its direct path up to and including the root. For **Rem**, these are the minimum subset of nodes such that every member in the group, excluding the removed party, can decrypt an element in S_{seeds} .

³Only used when $S_{op} = \text{Add}$ or $S_{op} = \text{Rem}$.

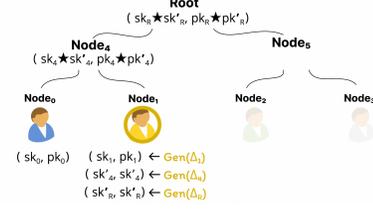


Figure 3: A new member, now occupying $Node_1$, joins the group thus updating the ratchet tree.

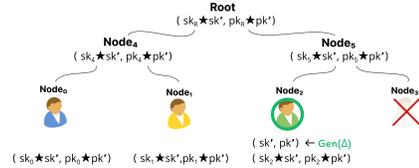


Figure 4: The member occupying $Node_2$ removes the member occupying $Node_3$ from the group.

message, W , encrypted to ID_j 's public key pk_j . This communication can be bootstrapped via an in-person interaction (e.g., via a QR code), in line with how groups are formed on the front lines during large-scale protests [10]. W consists of group state γ_i , where $\gamma_i.secret_keys$ only includes the secret keys associated to common ancestors in the path from ID_j to the root.

Upon receiving W , ID_j randomly samples Δ_i for themselves and each ancestor node along their direct path to the root, generating the corresponding key pairs $(sk_i, pk_i) \leftarrow \text{Gen}(\Delta_i)$, and using this to update each node's key material on their path to the root. If key material for an ancestor node already exists, it is combined using our \star operator. This is shown in Figure 3. Finally, ID_j sends S to the rest of the group, where: $S_{id} = ID_i$, S_{op} denotes **Add**, $S_{mod} = ID_j$, S_{seeds} contains the relevant encrypted seeds, and S_{keys} contains all known public keys. Group members can process S to update their local γ and enter a new epoch.⁴

Rem: Essential to the protest setting is efficient removal from a group, as mass arrests or police infiltration could compromise the keys of group members. A member, ID_i , removes ID_j by generating Δ , deriving a keypair $(sk', pk') \leftarrow \text{Gen}(\Delta)$ and using this to update every node in the ratchet tree, rather than just the nodes on their direct path to the root. The remove message is generated by recursing the tree to find the minimum number of subgroups to encrypt to such that everyone except ID_j can decrypt an element in S_{seeds} . For this message, each encrypted Δ in S_{seeds} may be the same; all remaining members will use this seed to update all nodes within the ratchet tree. Finally, remaining members clear the state of the removed member, $S_{mod} = ID_j$, by deleting the material held by that leaf, illustrated in Figure 4.

⁴In Amigo, addition is most effective when members join left-to-right, i.e., the most recently joined member adds the new user; upon joining they have all the needed secrets (if they exist) for their ancestor nodes. If this order cannot be enforced, at a minimum a newly joined member will have the group secret, which they can use to communicate with the group and query nodes for the relevant secrets they are missing. We sketch a mechanism for learning missing keys in the full version of this paper [38].

Upd: Consistent and frequent updates lead to fresh key material for the updating nodes and their path to the root, ensuring stronger notions of forward secrecy and post-compromise security for themselves and the group. A member ID_i initiates **Upd** by first randomly sampling Δ_i for themselves and their direct ancestors. The key material derived from each Δ_i is used to update the relevant node's key material and generate an update message intended for the group, similar to the process for addition. After processing this update, the updating member and the entire group enter a new epoch with a fresh group state.

Proc: When a member initiates a state operation, other members of the group need to process the corresponding protocol message S . When a member receives this protocol message, they first determine $S_{.op}$ and apply the operation to their local state γ , effectively moving it forward to γ' and entering a new epoch θ' . This entails updating (either combining or doing an initial write) the appropriate nodes in the ratchet tree using $S_{.seeds}$, and $S_{.keys}$. Depending on $S_{.op}$, this may also entail changes to $\gamma_{.members}$ and $\gamma_{.tree}$.

Security proof. We adapt our proof from that of Klein et al. [45] within the framework introduced by Jafarholi et al. [39]. We first prove security in the *selective* setting, in which the adversary is required to corrupt all the users at the beginning of the security experiment, even before receiving the public keys of the system. More realistically, however, an adversary could *adaptively* choose to corrupt nodes over a period of time, based on some strategy, so we also consider this setting. Details are available in the full version of our paper [38].

4.4 Sending and Receiving Messages

Amigo's CGKA effectively manages group cryptographic state via state operations, so adding secure messaging on top is straightforward and done at the application layer. Members can chat with each other via the operations **send-message** and **read-message**. A member runs **send-message** to generate A , where $A_{.payload}$ contains the actual message content; upon reception, intended recipients run **read-message** on A .

To provide *outsider anonymity*, every message sent within the group, whether from the CGKA (S) or otherwise (A) is encrypted under the group symmetric key (held by the root) and does not contain any identifying metadata. A party who is not an intended recipient has no information about the originator or destination of the message. As a result, parties utilize *trial decryption* to determine whether a message is for them; members will attempt decryption, first with the group symmetric key, before further processing.

To encrypt a message to the entire group, a group member derives a symmetric key from the root node's key material. The member concatenates these keys, hashes them, and passes this as input to a KDF where: $key_i \leftarrow KDF(H(pk_i || sk_i))$. The underlying authenticated encryption algorithm we use is AES-GCM.

Better security through UPKE. Encrypting solely under the group symmetric key provides forward secrecy and post-compromise security at an epoch granularity, i.e., until a state operation is processed. Protesters, however, may require more frequent key rotation, e.g., in protests where the threat of compromise is higher. We provide the option for this increased level of protection by incorporating UPKE in our CGKA. At the expense of efficiency, group

members can encrypt messages via the UPKE keypairs in the CGKA ratchet tree. Thus, forward secrecy and post-compromise security can be achieved per message within the current epoch.

When using UPKE for messaging, $A_{.payload}$ is a UPKE ciphertext, which contains the new public key pk' resulting from UPKE encryption. Intended recipients know where to apply pk' and the corresponding secret: the node whose key material successfully decrypted $A_{.payload}$. Our key combination operator, \star , can be used to combine the newly generated material with the old, effectively rotating keys per message. But, excluded recipients will not know where to apply pk' in this way. We discuss different strategies to address synchronizing public key state in our full version [38].

Handling mesh conditions. Although our CGKA allows us to provide secure, concurrent group communication and state management in a decentralized mesh setting, there are subtle but important deployment challenges due to the lack of a centralized delivery service. For instance, handling concurrent group additions may require a specific point of control for group admission, e.g., a group administrator, as used by protesters [10]. We discuss how Amigo handles these and other challenges in our full version [38].

5 Routing in Mesh Messaging

Any message must be *routed* over the mesh. In Amigo routing, a *message* represents (1) user-generated content or (2) cryptographic material for the CGKA protocol. Each message is tagged with a time-to-live (TTL). A message's destination is either a group or an individual user (a size-1 group).

Each Amigo node has a *message buffer* which stores a single copy of each unique message the node has recently received. When a node wants to send a message, the node inserts the message into the buffer, evicting the oldest message in the buffer if necessary to make space; given the choice between a CGKA message and a non-CGKA message, Amigo prefers to evict the latter, as CGKA messages are important in synchronizing group states. A node periodically decrements the TTL field of all messages in the buffer, evicting messages when their TTLs reach 0.

When two nodes discover each other through physical proximity, the nodes use a *routing protocol* to determine whether and how to share their local message buffers. Amigo can use four such routing protocols – three from prior work in mesh networking, and a fourth, dynamic clique routing, which Amigo introduces, tailored to the mobility patterns of protests. We now describe these four in detail.

Epidemic flooding. In this approach, each node divides time into 60-second intervals, transmitting a *beacon* at a random time within each interval to avoid Layer 2 collisions. If a peer receives such a beacon, the peer establishes a Layer 2 channel with the beacon originator, and the two nodes exchange their complete message buffers. Prior work has used epidemic flooding as a baseline for evaluating more sophisticated mesh routing protocols [22, 51, 57, 58]. Epidemic flooding aggressively probes the network for routes between senders and receivers, but risks network congestion due to the indiscriminate nature of message forwarding.

Digest flooding. To prevent two peers from exchanging messages that are already present on both nodes, each peer can first exchange Bloom filters [23] which summarize the locally-buffered messages; a peer X only sends a particular message to peer Y if Y 's Bloom filter

indicates that Y does not already store the message. Perry et al. [57] first propose this optimization to epidemic flooding, referring to a Bloom filter as a *message digest*. Details about Amigo’s Bloom filter implementation can be found in the full version of this paper [38].

Static clique routing. Perry et al. [57] sketched an optimized form of digest routing based on *cliques*. The basic idea is that each node belongs to a clique, with one node in each clique serving as the *clique leader*.⁵ When node X wants to send a message to node Y , X sends the message to X ’s clique leader; the clique leader then exchanges messages with other clique leaders via digest routing, with the hope that the message will eventually reach Y ’s leader, who will then transmit that message to its clique members (including Y). Users form a static clique, and choose a clique leader, prior to the protest. A node’s leader is static and does not change in response to the spatial arrangement of nodes or other properties of the mesh.

Intuitively, static clique routing should decrease overall network pressure compared to flooding, since only clique leaders perform bulk message exchanges. This approach might result in worse end-to-end delivery rates if clique leaders do not encounter each other often, or if clique members are not near their leaders.

Dynamic clique routing. To mitigate the problems mentioned above, Amigo introduces *dynamic clique routing*, in which both a node’s clique and clique leader can change over time.

To determine a node’s dynamic clique, Amigo divides the 2D plane into adjacent grid squares, of 3m width. A node uses GPS to identify its current grid square; all members within the same grid reside in the same dynamic clique. Amigo divides time into *rounds*, where each round requires the selection of a clique leader. The round length balances maximizing channel utilization for message delivery and ensuring swift recovery when bootstrapping is unsuccessful. For most mobility models, cliques remained largely stable over a five-minute period; therefore, Amigo uses a five-minute round length. The first minute of a round is reserved for clique leader election; the remaining time is used to exchange messages.

Amigo’s leader elections are based on Android’s implementation [15]. Nodes choose a “willingness” value (0-100), indicating their readiness to serve as a clique leader. The node in each given region with the highest “willingness” value is elected as leader. Each node picks a random offset within the election period of the round, and generates a beacon containing the node’s approximate location, the node’s Layer 2 identifier, and the leader “willingness” value; the random offset decreases the likelihood of Layer 2 collisions.

A node also listens for beacons from other peers. At the end of the election period, a node selects the beacon-generating peer with the highest leader willingness in their region to be its clique leader. If a node hears no other beacons, it becomes its own clique leader.

Under perfect conditions, all nodes in a clique will appoint the same leader—the node with the highest “willingness” value; however, if beacons are not consistently received, nodes may not elect the correct node as a leader. If so, a node might believe itself to be the leader based on its received beacons, even though another node was actually elected leader by the others. Thus, this “false leader” node would not receive any messages for exchange, and

⁵We note that routing-layer *cliques* are separate from the application-layer *groups* used for messaging. Due to the protections from our CGKA (§4.1), even if an adversary learn cliques, they learn about clique members’ groups or communications.

essentially sits out the rest of the round. This node can then try again by participating in the leader election for the next round.

The “willingness” value is vulnerable to attack; malicious nodes could send high “willingness” values to be elected clique leaders, and subsequently disrupt traffic. While we do not address this problem in Amigo, we note that for added robustness, nodes could instead include random values in beacons to determine the leader, using methods such as distributed randomness beacons [27].

Empirically, a protest may alternate between moments of tight, highly-ordered spatial density (e.g., when protesters form a human chain), interspersed with less dense spatial arrangements and/or more chaotic mobility patterns (e.g., when protesters randomly scatter due to the sudden arrival of law enforcement). We expect dynamic clique routing will better adapt to these patterns than static clique approaches.

Our dynamic routing protocol self-organizes using location data, but this use does not leak additional information to a global eavesdropper. Such an eavesdropper already knows the location of each node because eavesdropping implies physical proximity to a victim node, and an eavesdropper knows its own location.

Routing privacy. While our CGKA provides the privacy properties required for mesh messaging, our routing protocols also have certain privacy implications. We discuss further in our full version [38].

6 Modeling Protester Dynamics

To rigorously evaluate the performance of a secure mesh network, we need realistic models of user and network dynamics.

6.1 Mobility Models

A mobility model describes how protesters move through a physical space. The model determines how the spatial density of protesters varies over time, and how long protesters linger within the communication radii of each other’s short-range radios. Accurate mobility models are crucial for evaluating mesh networks because different mobility models have a dramatic impact on when (and whether) full or partial spanning trees for routing will emerge.

Unfortunately, mobility models from prior work are not well-suited for analyzing protest dynamics. For example, Perry et al. [57] considered spatial densities of one person every 2.5-15 ft², with each person located at the centroid of a grid square, and with a communication radius of 10 m. This communication radius would mean that each protester could directly contact 400 peers. A protest’s high density area might have as little as 0.2 m² per protester, however; if protesters were arranged in a grid, then 1,600 protesters would be within the communication range of any particular individual [67]. Furthermore, in a large-scale protest, multiple kinds of mobility patterns are likely to occur through space and time, as different protesters responds to different stimuli (e.g., the arrival of police or more protesters) in different parts of the protest region.

To build more realistic mobility models, we draw insights from a crowd-sourced guide to protesting, written in 2019 by Hong Kong anti-ELAB demonstrators [16]. The document provides practical advice about the planning and operational execution of a protest, describing various spatial arrangements that protesters can use.

Based on this information, we create several mobility models rooted in protest scenarios, which we describe below.

Marches. A march involves a group walking together in the same direction. In our march mobility model, we assign a starting area for the march, randomly distribute nodes within the starting area, and have all nodes move in the same direction at approximately 1.3 m/s [53]. The spatial density is about 2 m² per protester.

Human chains. A human chain entails a line of stationary people who pass materials from one end of the line to the other. Typically, the line begins in a storage area, and ends in front of the protest [16]. Unlike a march, there is no mobility in a human chain, and the spatial density is sparser, with nodes in a line about 2 m apart.

In our human chain mobility model, we pick two points within the simulated space, and generate a pathway between them. This pathway is not a direct line – instead, we add randomness to account for obstructions between the two points and general stochasticity in the locations of people along the line. We implement this randomness by selecting 0-10 additional random points between the two endpoints, and applying subtle curves to the paths originally generated with linear interpolation. Finally, we assign each protester to a location on the chain, spacing the protesters as above.

Gatherings. In the Hong Kong protest guide, the authors state that protests can also take the form of “gatherings with a specific focus.” Examples of such gatherings are musical performances and communal art projects [16]. These kinds of protests involve crowds in which individual participants generally stay within the congregation area, but can migrate within that area at walking speeds.

In our gather mobility model, we assign an area (roughly a city block in size) where an event will take place. We randomly assign nodes to occupy the space, and select a subset of nodes to move randomly within the space (as random waypoints [40]).

Blockades. A blockade occurs when protesters prevent entrance to a given space, e.g., anti-ELAB protesters blockaded the Cross-Harbor Tunnel to stop police from entering a university [56].

We implement two blockade-based mobility models. The first represents an “object blockade” in which demonstrators place objects along the blockade frontier. For example, Hong Kong protesters scattered road blocks and then set them ablaze. In our object blockade mobility model, protesters move at roughly 1.4 m/s between areas where materials are stored and the protest frontier. At each location, protesters briefly pause to pick up or drop off material.

Our “human blockade” model captures a scenario in which protesters have gathered near a blockade frontier, and are being monitored by law enforcement. In this model, clumps of police officers move towards the frontier, as do the protesters. If a police officer and a protester get within half a meter of each other for 30 seconds, one of them will retreat to the nearest clump of like-minded individuals; this behavior represents a participant being injured or physically overwhelmed by a member of the opposing group.

Additional models. For comparison, we also implement the static and random waypoint mobility models, which are standard in the literature. For more details on these models, see our full version [38].

6.2 Physical Layer Simulations

Prior work [51, 57, 58] uses custom network simulators to evaluate mesh routing protocols. Unfortunately, these works do not model physical-layer events like message collisions between two peers within communication range who transmit data simultaneously. Because prior simulations did not capture these Layer 1 effects, they overestimated the available bandwidth at Layers 2+.

To avoid this problem, we evaluate Amigo using ns-3, a discrete-event network simulator which *does* model important low-level characteristics. We extended ns-3 to support Wi-Fi Direct, a Layer 1+2 technology that, like Bluetooth, allows devices to communicate in a peer-to-peer fashion. We built our Wi-Fi Direct implementation atop ns-3’s preexisting support for ad-hoc Wi-Fi; implementation details are in our full version [38]. Our implementation uses 802.11ax/Wi-Fi 6, with a maximum data rate of 143.4 Mbps for advertising (MCS Index 11), and 263.3 Mbps for data exchange (MCS Index 6). This network setup is feasible on modern smartphones and representative of Wi-Fi Direct, which can provide up to 250 Mbps [71]. We limit signal propagation to 10 m, as in previous work [57].

We found ns-3 hit scaling limits as we increased the simulated protest size. For example, to complete a single run of a 250 node simulation using a static mobility model and digest routing, we required ~10 GB of memory and one week of wall clock time, even on a high-end cloud instance (GCP c3d-highmem-180). Nonetheless, as we demonstrate in Section 7, even modestly-sized protests can trigger crippling network congestion in standard approaches for secure mesh routing (and sometimes in the Amigo-specific dynamic clique approach), suggesting important areas for future study (§9).

6.3 Traffic Model

Our traffic model expresses the rate at which application-level behavior generates application messages and CGKA messages (§4).

For a given simulation lasting an hour of simulated time, we generate 20 CGKA messages, each sent from a random group member and at a random time in the first half of the simulation; the latter constraint ensures the message has sufficient time to propagate to the other group members before the simulation ends. This is roughly equivalent to each group adding a member, removing a member, and performing a key refresh during a protest.

Non-CGKA traffic corresponds to standard text messages sent to a group. We implement a basic traffic model for non-CGKA messages: each node in the simulation sends a message every minute. We choose these values based on the protest manual, which states chats see an excess of 100 messages/min during an event’s peak [16].

To generate the senders and receivers for messages, we create group chats of 25 members, based on nodes’ locations: we assume 0.8 of a clique is nearby the node, and select the other 0.2 randomly. We imagine many members of a close-knit group chat may enter the protest together, and a fraction may attend separately, or remain home or nearby to perform wellness checks. Mobility models may cause group members to move physically apart from one another during the simulation; for example, to carry supplies, talk to someone new, etc. We assume people within a group chat send messages to each other. Each message is 250 bytes and has a TTL of 5,000 seconds—slightly less than the length of an average action [16].

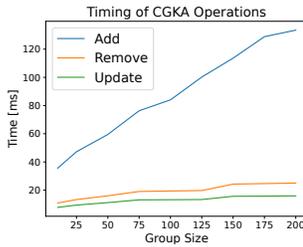


Figure 5: Timing of Amigo CGKA operations.

7 Results

We now discuss our extensive end-to-end benchmarks evaluating Amigo’s feasibility in realistic protest scenarios.

7.1 CGKA Microbenchmarks

Amigo is intended for deployment atop resource-constrained devices. To evaluate the resource consumption of Amigo’s CGKA, we ran microbenchmarks on both a 2020 Moto E phone (1.8 GHz octa-core, 2GB RAM) and a Raspberry Pi 4 with a 1.8 GHz quad-core ARMv8 SoC and 8 GB RAM. The smartphone emulates low-end performance while the Raspberry Pi offers similar specifications to medium/high-end modern smartphones. We used the Criterion [5] framework for the benchmarks, running each for a minimum of 1000 iterations.

Are CGKA operations fast? As discussed in Section 2, updates and removals have to be fast; when a group member is compromised, e.g., because their phone is confiscated by authorities, all other group members are at risk. Figure 5 shows, in a 200 person group, Amigo on a low end budget phone can perform removal and updates in under 26 ms. Addition is much slower, but adding a new member to a group is not as critical as removing a compromised one; furthermore, we expect additions will be infrequent. At a group size of 200, on our Android phone we could execute roughly 39 member removals, 62 key refreshes, or 7 member additions in one minute. Performance was significantly better on the Raspberry Pi. See the full version of this paper for detailed results [38].

What is the CGKA-induced network load? For non-CGKA messages, Amigo is more efficient than pairwise group approaches [1] where the same message is sent encrypted multiple times, once per each symmetric key negotiated with an individual group member. For state operations, Amigo generates one message for each removal or update. Member addition requires two messages: a welcome message to bootstrap in the new member, and an addition message that brings the group into a new epoch with updated group state. See our full version for more detailed results [38].

Is Amigo battery efficient? Our full version contains a complete discussion [38]. For now, we note Amigo’s energy consumption is reasonable and will not unduly drain a phone’s battery.

7.2 Network Congestion and Stability

How does buffer size impact network saturation? Figure 6 shows the packet volume generated by various buffer sizes, with all other aspects of the simulation held constant (e.g., traffic model,

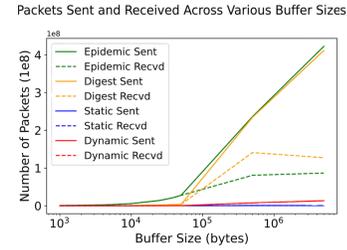


Figure 6: Packets sent and received by routing protocol, for 100 nodes in a static mobility model; average of 3 runs.

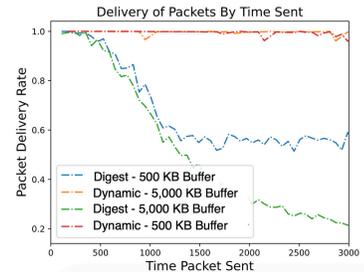


Figure 7: Average delivery rates for packets sent at time t for our digest routing and dynamic routing, for 100 nodes in a static mobility model.

mobility model, etc.). Table 1 breaks down the various packet types sent, and their end-to-end delivery rates, given a 5,000 KB buffer.

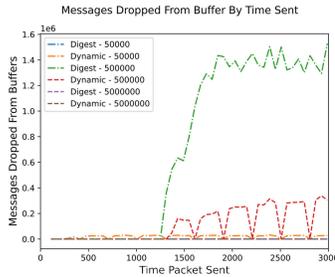
As expected, flooding generates significantly more traffic than other approaches; as buffer size increases, the difference grows (Figure 6). The bulk of the traffic arises from messages and digests, not from Layer 2 management packets (i.e., beacons, probe requests, and probe responses).

Notably, end-to-end delivery rates for epidemic routing and digest routing are very low (e.g., ~30% for digest routing). The reason is that both schemes gradually congest the network and then make no attempt to reduce their sending rate, even once message buffers have become steady-state full. When buffers become full and sending rates do not back off, new messages are constantly added to buffers, reducing the effectiveness of the digest optimization. As a result, nodes often exchange complete or near-complete copies of their local buffers, resulting in larger Layer 2 messages and more Layer 1 collisions (triggering retransmissions of the associated Layer 2 messages). The net result is, as time progresses, the expected delivery rate (Figure 7) for packets sent at time $t + 1$ trend monotonically lower than the expected delivery rate for packets sent at time t .

The relatively poor performance of digest routing is surprising, given previous work [57] introduced digest-based routing as an optimization. However, by modeling realistic Layer 1 and Layer 2 effects, we see with epidemic routing, only message packets collide with themselves, whereas with digest routing, message packets *and digest packets* collide, exacerbating the effects of network saturation. Thus, as shown in Table 1, digest routing’s performance is not superlatively better than that of vanilla epidemic routing.

Table 1: Average packets sent of various type across various routing mechanisms; results averaged across 3 runs of static mobility models for 5,000 KB buffer size. Packet delivery rate in parentheses.

	Epidemic Flooding	Digest Flooding	Static Clique Routing	Dynamic Clique Routing
Messages	422,623,852 (0.21)	411,590,675 (0.31)	2,030,682 (1.0)	13,522,625 (0.98)
Digests	-	24,087,630 (0.95)	472,836 (1.0)	776,972 (1.0)
Willingness Beacons	-	-	-	113,859
Message Beacons	-	-	-	4,800
Link Layer Beacons	6,000	6,000	6,000	1,200
Probe Requests	575,236 (1.0)	574,816 (1.0)	11,258 (1.0)	114,359 (1.0)
Probe Responses	574,064 (1.0)	573,609 (1.0)	11,258 (1.0)	113,875 (1.0)
Total	423,779,152	436,832,727	2,532,033	14,647,688

**Figure 8: Messages evicted in digest and dynamic mechanisms; 100 nodes in gather mobility model, bucketed into 60s intervals.**

Note that the network saturation in Table 1 represents a scenario in which each Layer 1 collision domain contains 100 protesters. A high-density protest could have as little as 0.2 m^2 per protester [25], meaning 1,600 protesters would be in a single collision domain and would suffer from an even more intense congestion collapse.

How does buffer size impact end-to-end delivery rates? All of Amigo’s routing schemes use per-node buffers to hold the messages nodes exchange with each other. Given a fixed mobility pattern and traffic matrix, decreasing the buffer sizes will result in fewer messages exchanged when nodes encounter each other, resulting in less network congestion; however, due to contention for buffer space, older messages (i.e., messages with lower TTLs) in a buffer will be evicted more aggressively to make room for new messages, potentially reducing how often messages are shared (and thus potentially reducing end-to-end message delivery rates). Conversely, increasing the buffer size will give messages more opportunities to be transmitted, potentially increasing message delivery rates; however, the resulting increase in traffic volume may induce network congestion that *decreases* end-to-end delivery rates. The tensions between increased congestion and increased delivery rates materialize in different ways for different protest environments.

Figure 8 is a time-series view of the number of messages evicted from node buffers during simulations using the gather mobility model with dynamic routing or digest routing. For 5,000 KB buffers, no evictions occur. For 500 KB, buffer evictions begin to increase at approximately $t = 1200$. Eviction rates remain high, with a brief hiatus for dynamic routing during advertising periods.

This eviction behavior is relevant to Figure 9, which shows the delivery rates for packets sent at various time intervals. We see

**Figure 9: Message delivery for the gather mobility model based on initial send times, bucketed into 60s intervals.**

delivery rates are initially high across all routing mechanisms for buffer sizes of 500 KB and 5,000 KB, but at approximately $t = 1200$ delivery rates become consistently poor. The reason is, once buffer evictions start at $t = 1200$, the effective number of transmission chances a message gets will decrease, hurting the message’s likelihood of being end-to-end delivered. For the simulation settings in Figure 9, the network has not become fully utilized before buffer evictions start occurring; this means, for the gather mobility pattern examined in Figure 9, a buffer of 500 KB is unable to fully utilize the available network bandwidth, resulting in unnecessary degradation of end-to-end delivery rates. The overall takeaway is that, even for Amigo’s protest-optimized dynamic routing, the lack of on-the-fly adjustment to buffer dynamics and network congestion leads to non-optimal routing performance. We discuss further in Section 9.

7.3 Protester Mobility and Network Partitioning

How does protester mobility impact message delivery? Table 2 explores the number of packets sent and message delivery rates for the routing protocols and various mobility patterns when using a 5,000 KB buffer. Simulations ran for 3,600 seconds, but Table 2 only considers messages sent before $t = 3000$ so as to not mark as dropped a message that was in-transit at the end of the simulation and would have eventually been delivered if the simulation had continued. Table 2 does include the “warm-up” period at the beginning of the simulation, because such a warm-up would also occur at the beginning of a real protest.

As shown in Table 2, static clique routing was the worst performing approach across all mobility patterns. The reasons were that (1) a clique leader often had trouble encountering its clique members or

Table 2: Total message packets sent for all routing mechanisms and mobility models; results averaged across 3 runs for 5,000 KB buffer size. Message delivery rate in parentheses. Checkmarks if CGKA group consistent at the end of the simulation.

	Epidemic Flooding	Digest Flooding	Static Clique Routing	Dynamic Clique Routing
static	422,613,855 (1.000) ✓	411,579,195 (1.000) ✓	596,074 (0.705) ✗	13,555,092 (1.000) ✓
random	425,916,153 (1.000) ✓	420,206,611 (1.000) ✓	2,073,922 (0.980) ✗	8,626,361 (0.963) ✗
gather	416,506,046 (1.000) ✓	405,078,143 (1.000) ✓	1,964,185 (0.941) ✗	12,956,503 (1.000) ✓
chain	34,706,437 (1.000) ✓	34,732,147 (1.000) ✓	36,241 (0.017) ✗	14,378,903 (0.999) ✓
march	50,609,963 (0.867) ✗	50,325,119 (0.867) ✗	29,145 (0.026) ✗	6,208,141 (0.726) ✗
blockade1	11,073,410 (0.271) ✗	11,079,296 (0.271) ✗	18,214 (0.008) ✗	87,310 (0.025) ✗
blockade2	435,542 (0.033) ✗	434,633 (0.033) ✗	8,286 (0.005) ✗	12,380 (0.006) ✗

other clique leaders, and (2) peers had no way to attach to a better-suited leader. Ignoring the blockade mobility patterns for a moment, epidemic routing and digest routing outperformed dynamic clique routing because those algorithms did not trigger congestion collapse for the mobility patterns and node counts explored in the simulations; as a result, the more aggressive forwarding behavior of these algorithms (compared to that of dynamic clique routing) led to better end-to-end delivery rates.

Static clique routing’s performance for the chain and march patterns were essentially the same. For the other three routing protocols, the march model was more challenging than the chain model; the reason being, while the chain model is static, the march model is dynamic, making message transmission and clique dynamics harder to maintain.

Table 2 shows that all four routing protocols struggled with blockade mobility patterns. In blockade1, protesters retrieved materials to form a barrier made of physical objects, whereas in blockade2, protesters generally approach a police line, and then spread out due to police intimidation. Both models had low spatial density, with nodes often ten meters apart or more, reducing chances for nodes to hear beacons (or exchange messages more generally).

Overall, we note that our new dynamic clique routing protocol achieves 95% or higher delivery rate on multiple mobility models (including some of the protester-oriented ones) while sending far fewer messages over the course of the simulation. This differences ranges from $\sim 2\times$ to $\sim 50\times$, depending on the mobility model evaluated. In other words, our simulations show that dynamic clique routing can achieve similar levels of message delivery at a fraction of the messages sent – important for smartphones, which are constrained by battery size and CPU capability, as well as for overall network stability, as discussed in Section 7.2.

Can protester groups maintain key agreement? The strong protections of Amigo’s CGKA protocol are moot in practice if Amigo nodes struggle to maintain CGKA key state due to an unreliable routing layer. Amigo’s CGKA protocol tolerates out-of-order messages (§4.2), but it does not tolerate a group member *never* receiving a CGKA message—in this scenario, the member will be unable to continue in the group. As discussed in Section 5, Amigo prioritizes the delivery of CGKA messages, but is this policy sufficient to allow the CGKA protocol to succeed in practice?

Table 2 shows that dynamic clique routing achieves comparable message delivery rates with far fewer messages, and achieves CGKA convergence for most of the same cases as the flooding protocols, but not for the random mobility model.

To understand Amigo’s support for highly dynamic group states, we increase the number of group updates that occur throughout the simulation, modifying our traffic model to include $5\times$ more group update (CGKA) packets. In our experimental setup, this equates to each group experiencing a change in group state every minute. Under these conditions, total message and group state update delivery rates are consistent with those of less dynamic groups. Our complete results are in the full version of this paper [38].

Amigo’s CGKA can achieve eventual consistency if all CGKA packets are delivered (§4.2); flooding protocols are highly redundant, so they are more likely to achieve this state (i.e., perfect CGKA message delivery) with enough time. But, dynamic clique routing still has a relatively high delivery rate. In the event a member does miss a CGKA message, a straightforward solution is to be re-added to the group by bootstrapping key material with another group member. We believe decentralized CGKAs that are robust to low message delivery rates is an interesting direction for future work.

8 Related Work

Below we discuss related work in the context of mesh messaging, routing, and key agreement for protest scenarios.

Mesh messaging in protest scenarios. Table 3 summarizes prior work in mesh messaging and compares it to Amigo.

Bridgefy [26] is a mesh network supporting pairwise messaging via flooding. Though Bridgefy claims to provide anonymity and end-to-end encryption via the Signal protocol, significant vulnerabilities were discovered which put these claims into question [11, 12].

Moby [58] is a flooding-based mesh network that supports pairwise messaging. Moby uses Signal’s symmetric ratchet [4] to provide end-to-end encryption with forward secrecy. Anonymity is provided through HMAC validation trials.

ASMesh [22] augments Signal’s Double Ratchet algorithm for stronger properties and routes messages using flooding.

Perry et al. [57] use flooding to provide both end-to-end encryption and anonymity via a CPA-secure signcryption scheme. This scheme supports group communication by having a group creator generate a single secret key and share it with each member of the group, but does not provide additional security. To protect anonymity from traffic analysis attacks, parties constantly send a message at a fixed interval, using “dummy” messages if needed.

Rangzen [51] propagates messages via epidemic flooding, giving preference (i.e., more local buffer space) to “trusted” messages. Trust is determined by examining how many mutual friends users share

Table 3: Related work in mesh messaging. E2E: end-to-end encryption; FS: forward secrecy; PCS: post-compromise security; Anon: anonymity; Pair: pairwise communication; Groups: Group communication.

Scheme	E2E	FS	PCS	Anon	Pair	Groups
Bridgefy [26]	✗	✗	✗	✗	✓	✗
Moby [58]	✓	✓	✗	✓	✓	✗
ASMesh [22]	✓	✓	✓	✓	✓	✗
Perry et al. [57]	✓	✗	✗	✓	✓	✓
Rangzen [51]	✗	N/A	N/A	✓	✗	✗
Anix [41]	✗	N/A	N/A	✓	✗	✗
Amigo (this)	✓	✓	✓	✓	✓	✓

in a social graph maintained out-of-band. Rangzen does not end-to-end encrypt messages because Rangzen is intended to spread “micro-blogs” (akin to tweets) that are destined for all users.

Anix [41] is also a “micro-blog” platform over flooding, introducing notions of trust via ephemeral but linkable pseudonyms. Parties can remotely allow their messages to be connected back to them by “trusted users” while remaining anonymous to everyone else.

Mobile ad-hoc networks. In mobile ad-hoc networks (MANETs), mobile devices communicate with each other in an infrastructure-less setting. There is rich literature focusing on *multicast* MANET protocols: settings where senders can communicate with multiple nodes using a single destination address (e.g., [30, 32]). We refer readers to Badarneh and Kadoch [18] for a full survey. While specific details differ, multicast protocols generally provide functionalities for joining/leaving multicast routing groups and maintaining routing state. Multicast protocols aim to establish *forwarding paths* between a source and receivers that allow messages to be efficiently propagated through the network.

Most MANET routing protocols do not consider anonymous and private route construction/maintenance. Multicast groups are typically addressed by persistent identifiers [18], which our adversary (§3) can collect (passively or actively) to identify groups and infer their behavior. Also, routing information sent in the clear and cached routing tables can leak topological network information to an adversary [18, 47]. Previous literature has shown that this topological view of the network can be used to reveal the approximate physical locations of nodes [54, 64]. While there has been work on securing multicast MANET protocols from attacks [30, 32, 62], they do not adequately discuss anonymity and privacy in these settings.

There is a promising line of work exploring anonymous route constructions for MANETs [20, 24, 43, 46, 66, 69, 72, 74]. These works typically provide properties like sender/recipient unlinkability, route untraceability, and location privacy. These protocols are most effective when there exist relatively stable forwarding paths that remain valid during multicast sessions. However, our simulations suggest that such topological stability is far from the norm. Our results under protest-specific movement and traffic patterns (§7) suggest that handling congestion and high mobility is a challenge in smartphone mesh messaging approaches, and one that must first be addressed before anonymous MANET protocols can be effective in the environments we consider.⁶

⁶We hope our high-fidelity simulations (§6) can guide the design of future anonymous, multicast MANET protocols in smartphone mesh networks during Internet shutdowns.

Group key agreement in wireless networks. A line of work [19, 31, 44, 75] focuses on providing mechanisms to establish and rotate group keys, achieve dynamic group membership, and provide data confidentiality in wireless mesh and ad-hoc networks. These works, however, leverage assumptions and models that differ significantly from those underlying Amigo. Dong et al. [31] provide secure group communication by taking advantage of stationary wireless routers that form a network backbone, infrastructure Amigo cannot assume is present. Zhu et al. [75] provide group rekeying via a keyserver node, a single point of failure where a compromise could put all members of a group at risk indefinitely. Balachandran et al. [19] incur issues with scalability in our setting, requiring a number of messages linear in the group size to refresh keys. Kaya et al. [44] treat group member revocations as rare, requiring reasonably more computation and bandwidth resources. Protester safety, however, requires removals to be fast and efficient (§3).

9 Discussion

Secure communication protocols require participants to create and destroy cryptographic keys. Messages associated with key management are exchanged over the same network that handles regular message traffic. Thus, the reliability of the network impacts the reliability of key management. In traditional (non-ad-hoc) networks, a Layer 4 security protocol like TLS can rely on the Layer 3 TCP protocol to provide dependable, in-order message delivery and avoid network congestion. However, the ephemeral, unpredictable nature of node connectivity in a mesh network makes it hard for Layer 3+ protocols to understand (and react to) network dynamics like congestion in real time. Prior work on secure mesh routing has focused on the details of key management, and made simplifying assumptions about the behavior of the network layer (§6.2). However, our results in Section 7 demonstrate that low-level network behavior has a critical impact on the performance and correctness of a secure mesh network. For example, Layer 1 phenomena like collisions and multi-path interference will frequently be triggered in real-life protest situations involving densely-packed people and physical obstacles like buildings. Furthermore, Layer 2 network partitions in the peer-to-peer spanning tree can also be triggered by realistic movement patterns. The limited channel count of real-life peer-to-peer radio technologies like Wi-Fi Direct also limits the scalability of network bandwidth, reducing the effectiveness of the mesh network. These factors suggest secure mesh networking protocols which appear feasible when deployed atop well-behaved networks may in fact perform badly in realistic scenarios. Optimizing cryptographic protocols is not enough; we must also complement improved cryptographic schemes with network-layer innovations.

Future work. Our results (§7) suggest important future research:

- Congestion detection and avoidance for mesh networks.
- Key management protocols more tolerant to packet loss.
- Adaptive routing protocols which modify forwarding in response to dynamic estimates of global properties.
- User-facing feedback mechanisms that allow users to understand current network dynamics and possibly adapt user-level behavior (e.g., by not sending video messages when high network congestion is detected).

These features will be key to providing sufficient bandwidth to support stronger security and anonymity properties [42, 57].

Properly evaluating these approaches for larger protests will also require fundamental optimization work involving network simulators like ns-3. ns-3 currently uses standard C++ data structures for priority queues [55], but our experience is that stock data structures prevent ns-3 from scaling to simulation protests beyond a few hundred nodes – ns-3 either runs out of memory and crashes, or cannot finish a simulation despite running for days of wall-clock time. Optimizing ns-3's use of memory and compute, e.g., Devastator-style techniques [17], is crucial for understanding how secure mesh networks behave in practice.

10 Conclusion

We present Amigo, a mesh messaging system for protesters that combines a tailored CGKA with a new clique-based routing protocol. Our simulations show substantial improvements over prior work, while highlighting the need for further advances at lower layers of the protocol stack to ensure strong privacy and efficient routing.

Acknowledgments

The authors would like to thank Jonathan Rozen, Mathias Jud, and the anonymous reviewers from Real World Cryptography (RWC) 2025 for their feedback, as well as Rosario Gennaro for contributing cloud computing resources to our evaluation.

This work is supported by National Science Foundation (NSF) awards 1955172 and 2451597, an NSF Graduate Research Fellowship, a Sui Foundation Academic Research Award, a PSC-CUNY award, and awards from Google as a part of the Cyber NYC program. This work also used Jetstream2 computing resources through allocations from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, which is supported by NSF awards 2138259, 2138286, 2138307, 2137603, and 2138296.

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the sponsors. Any mention of specific companies or products does not imply any endorsement by the authors, by their employers, or by the sponsors.

References

- [1] 2014. Signal: Private Group Messaging. <https://signal.org/blog/private-groups/>.
- [2] 2015. Department of Justice Policy Guidance: Use of Cell-Site Simulator Technology. <https://www.justice.gov/opa/file/767321/dl>.
- [3] 2019. BeWater: Seven Tactics That Are Winning Hong Kong's Democracy Revolution. <https://www.newstatesman.com/politics/2019/08/be-water-seven-tactics-that-are-winning-hong-kongs-democracy-revolution-2>.
- [4] 2020. The Double Ratchet Algorithm. <https://signal.org/docs/specifications/doublersatchet/#symmetric-key-ratchet>.
- [5] 2021. Criterion.rs Documentation. https://bheisler.github.io/criterion.rs/book/criterion_rs.html.
- [6] Access Now. 2024. Myanmar's iron curtain: internet shutdowns and repression in 2023. <https://www.accessnow.org/press-release/myanmar-keepit-on-internet-shutdowns-2023-en/>. Online; accessed 1 Sept 2024.
- [7] Access Now. 2024. Shrinking democracy, growing violence: Internet shutdowns in 2023. <https://www.accessnow.org/wp-content/uploads/2024/05/2023-KIO-Report.pdf>. Online; accessed 1 Sept 2024.
- [8] Ian F Akyildiz and Xudong Wang. 2009. *Wireless mesh networks*. John Wiley & Sons.
- [9] Ian F Akyildiz, Xudong Wang, and Weilin Wang. 2005. Wireless mesh networks: a survey. *Computer networks* 47, 4 (2005), 445–487.
- [10] Martin R. Albrecht, Jorge Blasco, Rikke Bjerg Jensen, and Lenka Mareková. 2021. Collective Information Security in Large-Scale Urban Protests: the Case of Hong Kong. In *USENIX Security 2021*, Michael Bailey and Rachel Greenstadt (Eds.). USENIX Association, 3363–3380.
- [11] Martin R. Albrecht, Jorge Blasco, Rikke Bjerg Jensen, and Lenka Mareková. 2021. Mesh Messaging in Large-Scale Protests: Breaking Bridgefy. In *Topics in Cryptology – CT-RSA 2021: Cryptographers' Track at the RSA Conference 2021, Virtual Event, May 17–20, 2021, Proceedings* (San Francisco, CA, USA). Springer-Verlag, Berlin, Heidelberg, 375–398. doi:10.1007/978-3-030-75539-3_16
- [12] Martin R. Albrecht, Raphael Eikenberg, and Kenneth G. Paterson. 2022. Breaking Bridgefy, again: Adopting libsignal is not enough. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA, 269–286. <https://www.usenix.org/conference/usenixsecurity22/presentation/albrecht>
- [13] Nezar ALSayyad and Muna Guvenc. 2015. Virtual Uprisings: On the Interaction of New Social Media, Traditional Media Coverage and Urban Space during the 'Arab Spring'. *Urban Studies* 52, 11 (2015), 2018–2034. doi:10.1177/0042098013505881 arXiv:https://doi.org/10.1177/0042098013505881
- [14] Amnesty International. 2020. Iran: Internet deliberately shut down during November 2019 killings – new investigation. <https://www.amnesty.org/en/latest/press-release/2020/11/iran-internet-deliberately-shut-down-during-november-2019-killings-new-investigation/>. Online; accessed 1 Dec 2023.
- [15] Android Developers. 2024. WifiP2pConfig. <https://developer.android.com/reference/android/net/wifi/p2p/WifiP2pConfig>. Online; accessed Oct 2024.
- [16] Anonymous Authors. 2023. The HK19 Manual - Part 2B: How Tos. https://docs.google.com/document/d/1UOURUN37_gUqrDd4FYDFYXQAmYuXtsBAfADLo47Im-Kk/edit#heading=h.9n9f9eiq3xhs. Online; accessed 1 Dec 2023.
- [17] John Bachan, Jianlan Ye, Xuan Jiang, Tan Nguyen, Mahesh Natarajan, Maximilian Bremer, and Cy Chan. 2024. Devastator: A Scalable Parallel Discrete Event Simulation Framework for Modern C++. In *Proceedings of the 38th ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*. 35–46.
- [18] Osamah S. Badarneh and Michel Kadoch. 2009. Multicast routing protocols in mobile ad hoc networks: a comparative survey and taxonomy. *EURASIP J. Wirel. Commun. Netw.* 2009, Article 26 (Jan. 2009), 42 pages. doi:10.1155/2009/764047
- [19] R.K. Balachandran, B. Ramamurthy, Xukai Zou, and N.V. Vinodchandran. 2005. CRTDH: an efficient key agreement scheme for secure group communications in wireless ad hoc networks. In *IEEE International Conference on Communications, 2005. ICC 2005*, 2005, Vol. 2. 1123–1127 Vol. 2. doi:10.1109/ICC.2005.1494522
- [20] Lichun Bao. 2007. A new approach to anonymous multicast routing in ad hoc networks. In *2007 Second International Conference on Communications and Networking in China*. IEEE, 1004–1008.
- [21] Karthikeyan Bhargavan, Richard Barnes, and Eric Rescorla. 2018. *TreeKEM: Asynchronous Decentralized Key Management for Large Dynamic Groups A protocol proposal for Messaging Layer Security (MLS)*. Research Report. Inria Paris. <https://inria.hal.science/hal-02425247>
- [22] Alexander Bienstock, Paul Rösler, and Yi Tang. 2023. ASMesh: Anonymous and Secure Messaging in Mesh Networks Using Stronger, Anonymous Double Ratchet. In *ACM CCS 2023*. Weizhi Meng, Christian Damsgaard Jensen, Cas Cremers, and Engin Kirda (Eds.). ACM Press, 1–15. doi:10.1145/3576915.3616615
- [23] Burton H Bloom. 1970. Space/time trade-offs in hash coding with allowable errors. *Commun. ACM* 13, 7 (1970), 422–426.
- [24] A. Boukerche, K. El-Khatib, Li Xu, and L. Korba. 2004. SDAR: a secure distributed anonymous routing protocol for wireless and mobile ad hoc networks. In *29th Annual IEEE International Conference on Local Computer Networks*. 618–624. doi:10.1109/LCN.2004.109
- [25] Robert M. Brecht. [n. d.], Festival and Concert Production: Crowd Safety. <https://tseentertainment.com/festival-and-concert-production-crowd-safety/>. Online; accessed 1 Sept 2024.
- [26] Bridgefy. 2022. Bridgefy Messaging App - Offline Messaging. <https://bridgefy.me>. Online; accessed 1 Dec 2023.
- [27] Kevin Choi, Aathira Manoj, and Joseph Bonneau. 2023. Sok: Distributed randomness beacons. In *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE, 75–92.
- [28] Darin Christensen and Francisco Garfias. 2018. Can you hear me now? How communication technology affects protest and repression. *Quarterly journal of political science* 13, 1 (2018), 89.
- [29] Katriel Cohn-Gordon, Cas Cremers, Luke Garratt, Jon Millican, and Kevin Milner. 2018. On Ends-to-Ends Encryption: Asynchronous Group Messaging with Strong Security Guarantees. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security (Toronto, Canada) (CCS '18)*. Association for Computing Machinery, New York, NY, USA, 1802–1819. doi:10.1145/3243734.3243747
- [30] Reza Curtmola and Cristina Nita-Rotaru. 2007. BSMR: Byzantine-Resilient Secure Multicast Routing in Multi-hop Wireless Networks. In *2007 4th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks*. 263–272. doi:10.1109/SAHCN.2007.4292838
- [31] Jing Dong, Kurt Ackermann, and Cristina Nita-Rotaru. 2009. Secure group communication in wireless mesh networks. *Ad Hoc Networks* 7, 8 (2009), 1563–1576. doi:10.1016/j.adhoc.2009.03.004 Privacy and Security in Wireless Sensor and Ad Hoc Networks.

- [32] Jing Dong, Cristina Nita-Rotaru, and Reza Curtmola. 2011. Secure High-Throughput Multicast Routing in Wireless Mesh Networks. *IEEE Transactions on Mobile Computing* 10, 05 (May 2011), 653–668. doi:10.1109/TMC.2010.194
- [33] Electronic Frontier Foundation. 2023. CELL-SITE SIMULATORS/ IMSI CATCHERS. <https://sls.eff.org/technologies/cell-site-simulators-imsi-catchers>. Online; accessed 1 Sept 2024.
- [34] N. Fazio and I.M. Perera. 2012. Outsider-Anonymous Broadcast Encryption with Sublinear Ciphertexts. In *IACR Public Key Cryptography—PKC '12*. Springer, Heidelberg, 225–242. LNCS 7293.
- [35] Shaddi Hasan, Yahel Ben-David, Giulia Fanti, Eric Brewer, and Scott Shenker. 2013. Building Dissent Networks: Towards Effective Countermeasures against {Large-Scale} Communications Blackouts. In *3rd USENIX Workshop on Free and Open Communications on the Internet (FOCI 13)*.
- [36] IETF. 2003. *Guidelines for Writing RFC Text on Security Considerations*. Technical Report. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/html/rfc3552>
- [37] IETF. 2023. *Message Layer Security (MLS)*. Technical Report. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/rfc9420/>
- [38] David Inyangson, Sarah Radway, Tushar M. Jois, Nelly Fazio, and James Mickens. 2024. Amigo: Secure Group Mesh Messaging in Realistic Protest Settings. Cryptology ePrint Archive, Paper 2024/1872. <https://eprint.iacr.org/2024/1872>
- [39] Zahra Jafarholi, Chethan Kamath, Karen Klein, Ilan Komargodski, Krzysztof Pietrzak, and Daniel Wichs. 2017. Be Adaptive, Avoid Overcommitting. In *Advances in Cryptology – CRYPTO 2017*, Jonathan Katz and Hovav Shacham (Eds.). Springer International Publishing, Cham, 133–163.
- [40] D Johnson. 1996. Dynamic source routing in ad hoc wireless networks. *Mobile Computing/Kluwer Academic Publishers* (1996).
- [41] Sina Kamali and Diogo Barradas. 2025. Anix: Anonymous Blackout-Resistant Microblogging with Message Endorsing. In *2025 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, Los Alamitos, CA, USA, 15–15. doi:10.1109/SP61157.2025.00015
- [42] Sina Kamali and Diogo Barradas. 2025. Anix: Anonymous Blackout-Resistant Microblogging with Message Endorsing. (2025).
- [43] Jung-Chun Kao and Radu Marculescu. 2007. Energy-efficient anonymous multicast in mobile ad-hoc networks. In *2007 International Conference on Parallel and Distributed Systems*. 1–8. doi:10.1109/ICPADS.2007.4447716
- [44] T. Kaya, G. Lin, G. Noubir, and A. Yilmaz. 2003. Secure multicast groups on ad hoc networks. In *Proceedings of the 1st ACM Workshop on Security of Ad Hoc and Sensor Networks (Fairfax, Virginia) (SASN '03)*. Association for Computing Machinery, New York, NY, USA, 94–102. doi:10.1145/986858.986872
- [45] Karen Klein, Guillermo Pascual-Perez, Michael Walter, Chethan Kamath, Margarita Capretto, Miguel Cueto, Iliia Markov, Michelle Yeo, Joël Alwen, and Krzysztof Pietrzak. 2021. Keep the Dirt: Tainted TreeKEM, Adaptively and Actively Secure Continuous Group Key Agreement. In *2021 IEEE Symposium on Security and Privacy (SP)*. 268–284. doi:10.1109/SP40001.2021.00035
- [46] Jiejun Kong and Xiaoyan Hong. 2003. ANODR: anonymous on demand routing with untraceable routes for mobile ad-hoc networks. In *Proceedings of the 4th ACM International Symposium on Mobile Ad Hoc Networking & Computing (Annapolis, Maryland, USA) (MobiHoc '03)*. Association for Computing Machinery, New York, NY, USA, 291–302. doi:10.1145/778415.778449
- [47] Jiejun Kong, Jun Liu, Xiaoyan Hong, Dapeng Wu, and Mario Gerla. 2007. *On Performance Cost of On-demand Anonymous Routing Protocols in Mobile Ad Hoc Networks*. Springer US, Boston, MA, 119–142. doi:10.1007/978-0-387-71058-7_6
- [48] Jeroen Van Laer and Peter Van Aelst. 2010. INTERNET AND SOCIAL MOVEMENT ACTION REPERTOIRES. *Information, Communication & Society* 13, 8 (2010), 1146–1171. doi:10.1080/136991181003628307 arXiv:<https://doi.org/10.1080/136991181003628307>
- [49] Francis Lee, Samson Yuen, Gary Tang, and Edmund Cheng. 2019. Hong Kong's Summer of Uprising: From Anti-Extradition to Anti-Authoritarian Protests. *China Review* 19 (11 2019), 1–32.
- [50] Francis L. F. Lee. 2020. Solidarity in the Anti-Extradition Bill movement in Hong Kong. *Critical Asian Studies* 52 (2020), 18 – 32. <https://api.semanticscholar.org/CorpusID:214191879>
- [51] Ada Lerner, Giulia C. Fanti, Yahel Ben-David, Jesus Garcia, Paul Schmitt, and Barath Raghavan. 2016. Rangzen: Anonymously Getting the Word Out in a Blackout. *ArXiv abs/1612.03371* (2016). <https://api.semanticscholar.org/CorpusID:2659855>
- [52] Xiangfang Li, Lijun Qian, and Joseph Kamto. 2009. Secure anonymous routing in wireless mesh networks. In *2009 International Conference on E-Business and Information System Security*. IEEE, 1–5.
- [53] Elaine M Murtagh, Jacqueline L Mair, Elroy Aguiar, Catrine Tudor-Locke, and Marie H Murphy. 2021. Outdoor walking speeds of apparently healthy adults: A systematic review and meta-analysis. *Sports Medicine* 51 (2021), 125–141.
- [54] D. Niculescu and B. Nath. 2001. Ad hoc positioning system (APS). In *GLOBECOM'01. IEEE Global Telecommunications Conference (Cat. No.01CH37270)*, Vol. 5. 2926–2931 vol.5. doi:10.1109/GLOCOM.2001.965964
- [55] ns 3. 2021. Events and Simulator. <https://www.nsnam.org/docs/manual/html/events.html#scheduler>. Online; accessed 1 Sept 2024.
- [56] Jessie Pang and Kate Lamb. 2019. Highway blockade reveals splits in Hong Kong protest movement. <https://www.reuters.com/article/world/highway-blockade-reveals-splits-in-hong-kong-protest-movement-idUSKBN1XP06R/>. Online; accessed 1 Dec 2023.
- [57] Neil Perry, Bruce Spang, Saba Eskandarian, and Dan Boneh. 2022. Strong Anonymity for Mesh Messaging. *arXiv preprint arXiv:2207.04145* (2022).
- [58] Amogh Pradeep, Hira Javaid, Ryan Williams, Antoine Rault, David Choffnes, Stevens Le Blond, and Bryan Alexander Ford. 2022. Moby: A blackout-resistant anonymity network for mobile devices. *Proceedings on Privacy Enhancing Technologies* 2022, 3 (2022), 247–267.
- [59] Abhinav Prakash, Dharma P Agrawa, and Yunli Chen. 2014. Network coding combined with onion routing for anonymous and secure communication in a wireless mesh network. *International Journal of Computer Networks & Communications (IJCNC)* 6, 6 (2014), 1–14.
- [60] Kui Ren and Wenjing Lou. 2008. A sophisticated privacy-enhanced yet accountable security framework for metropolitan wireless mesh networks. In *2008 The 28th International Conference on Distributed Computing Systems*. IEEE, 286–294.
- [61] Kui Ren, Shucheng Yu, Wenjing Lou, and Yanchao Zhang. 2009. PEACE: A novel privacy-enhanced yet accountable security framework for metropolitan wireless mesh networks. *IEEE Transactions on Parallel and Distributed Systems* 21, 2 (2009), 203–215.
- [62] S. Roy, V.G. Addada, S. Setia, and S. Jajodia. 2005. Securing MAODV: attacks and countermeasures. In *2005 Second Annual IEEE Communications Society Conference on Sensor and Ad Hoc Communications and Networks, 2005. IEEE SECON 2005*. 521–532. doi:10.1109/SAHCN.2005.1557104
- [63] Jaydip Sen. 2012. Secure and privacy-preserving authentication protocols for wireless mesh networks. In *Applied Cryptography and Network Security*. IntechOpen, 3–34.
- [64] Yi Shang, Wheeler Ruml, Ying Zhang, and Markus P. J. Fromherz. 2003. Localization from mere connectivity (*MobiHoc '03*). Association for Computing Machinery, New York, NY, USA, 201–212. doi:10.1145/778415.778439
- [65] Clay Shirky. 2011. The political power of social media: Technology, the public sphere, and political change. *Foreign Affairs* 90 (2011), 28–41. <https://api.semanticscholar.org/CorpusID:202307119>
- [66] Ronggong Song, Larry Korba, and George Yee. 2005. AnonDSR: efficient anonymous dynamic source routing for mobile ad-hoc networks. In *Proceedings of the 3rd ACM Workshop on Security of Ad Hoc and Sensor Networks (Alexandria, VA, USA) (SASN '05)*. Association for Computing Machinery, New York, NY, USA, 33–42. doi:10.1145/1102219.1102226
- [67] G. Keith Still. 2019. Standing Crowd Density. <https://www.gkstill.com/Support/crowd-density/CrowdDensity-1.html>. Online; accessed 1 Oct 2024.
- [68] Nazatul Haque Sultan and Nityananda Sarma. 2014. PAPAR: Pairing based authentication protocol with anonymous roaming for wireless mesh networks. In *2014 International Conference on Information Technology*. IEEE, 155–160.
- [69] Somayeh Taheri, Salke Hartung, and Dieter Hogrefe. 2015. Anonymous group-based routing in MANETs. *Journal of Information Security and Applications* 22 (2015), 87–98. doi:10.1016/j.jisa.2014.09.002 Special Issue on Security of Information and Networks.
- [70] Matthew Weidner. 2019. *Group Messaging for Secure Asynchronous Collaboration*. Master's thesis. University of Cambridge.
- [71] Wi-Fi Alliance. 2021. *Wi-Fi Direct Specification*. Wi-Fi Alliance. https://www.wi-fi.org/system/files/Wi-Fi_Direct_Specification_v1.9.pdf
- [72] Liu Yang, Markus Jakobsson, and Susanne Wetzel. 2006. Discount Anonymous On Demand Routing for Mobile Ad hoc Networks. In *2006 Securecomm and Workshops*. 1–10. doi:10.1109/SECCOMW.2006.359533
- [73] Yan Zhang, Jijun Luo, and Honglin Hu. 2006. *Wireless mesh networking: architectures, protocols and standards*. CRC Press.
- [74] Bo Zhu, Zhiguo Wan, M.S. Kankanhalli, Feng Bao, and R.H. Deng. 2004. Anonymous secure routing in mobile ad-hoc networks. In *29th Annual IEEE International Conference on Local Computer Networks*. 102–108. doi:10.1109/LCN.2004.21
- [75] S. Zhu, S. Setia, S. Xu, and S. Jajodia. 2004. GKMPAN: an efficient group rekeying scheme for secure multicast in ad-hoc networks. In *The First Annual International Conference on Mobile and Ubiquitous Systems: Networking and Services, 2004. MOBIQUITOUS 2004*. 42–51. doi:10.1109/MOBIQ.2004.1331709
- [76] Maximilian Zinkus, Tushar M. Jois, and Matthew Green. 2021. SoK: Cryptographic Confidentiality of Data on Mobile Devices. *Proceedings on Privacy Enhancing Technologies* 2022 (2021), 586 – 607. <https://api.semanticscholar.org/CorpusID:237605508>