Diogo Barradas

- **Assistant Professor**, Cheriton SCS @ University of Waterloo, CrySP
- **Associate Director**, Waterloo Cybersecurity and Privacy Institute (CPI)

- **Previously:** Post-doc @ Rice University (2021), Research Intern @ Cylab/CMU (2020), Research Intern @ UT Austin (2017), PhD researcher @ Técnico Lisbon (2016-2021)

Research Interests:

- “Privacy-preserving Networking”
  - Internet Censorship
  - Privacy-enhancing Technologies
  - Encrypted Traffic Analysis
  - Network Forensics

cs.uwaterloo.ca/~dbarrada

CrySP
Cryptography, Security, and Privacy Research Group

CYBER SECURITY AND PRIVACY INSTITUTE
UNIVERSITY OF WATERLOO

UNIVERSITY OF WATERLOO

Privacy-preserving Blackout-resistant Technologies

Diogo Barradas
diogo.barradas@uwaterloo.ca
University of Waterloo

LASIGE

FCUL, Lisbon, Portugal

January 14th, 2026

Internet Shutdowns (*a.k.a. blackouts*)

- Repressive governments often aim to control/restrict the flow of information
  - Network-level interference
  - Social media monitoring
  - Messaging filters
- Today, censors are choosing to instate region/country-wide Internet shutdowns
  - Lasting up to weeks in a row

Kashmiri journalists protest against internet blockade put by India's government in Srinagar on October 12, 2019. TAUSEEF MUSTAFA/AFP/AFP via Getty Images

https://www.cnn.com/2019/12/21/asia/internet-shutdowns-china-india-censorship-intl-hnk/index.html

An Ongoing Shutdown during Iranian Protests

NetBlocks 13 Jan 2026
@netblocks@mastodon.social

⚠️ Update: #Iran has now been offline for 120 hours.

Despite some phone calls now connecting, there is no secure way to communicate and the general public remain cut off from the outside world.

What footage makes it through shows extensive use of force against civilians 🔞

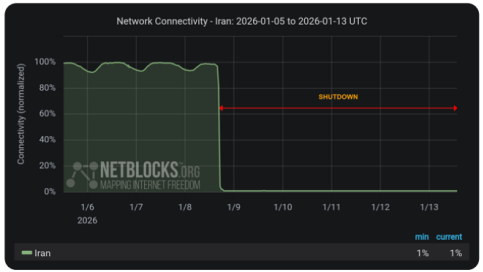

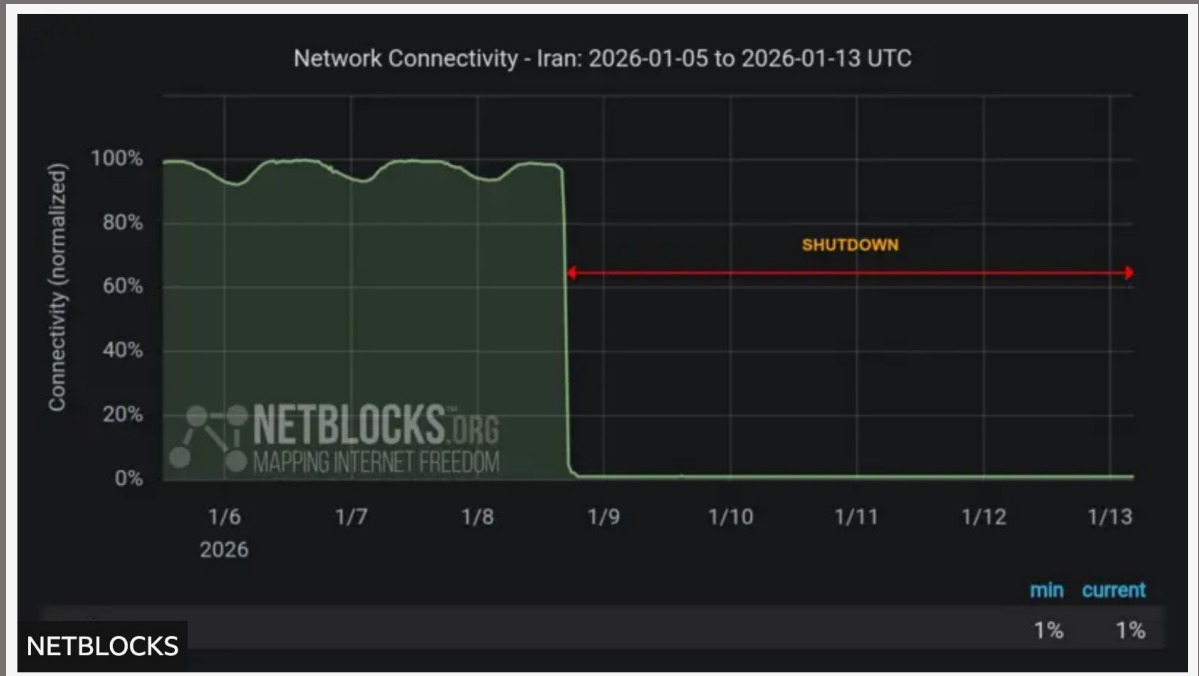

Shutdowns are on the Rise

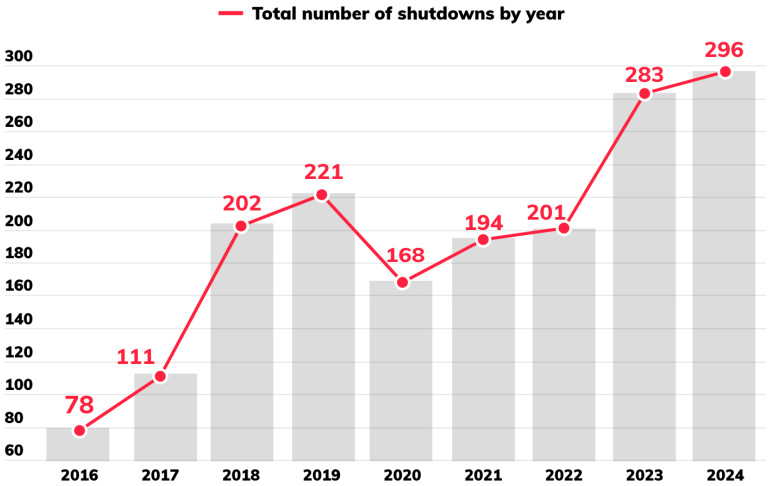

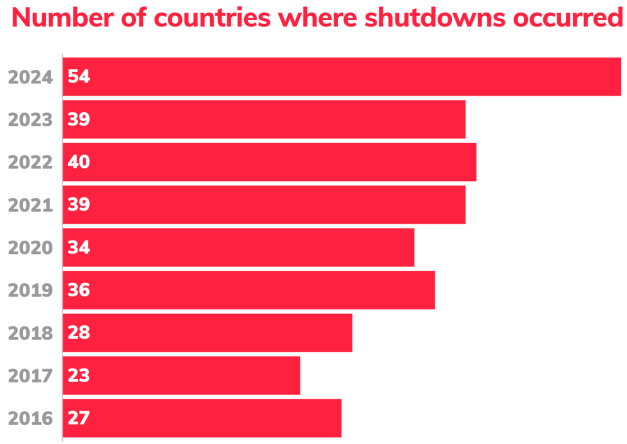

https://www.accessnow.org/internet-shutdowns-2024/

Shutdowns are on the Rise

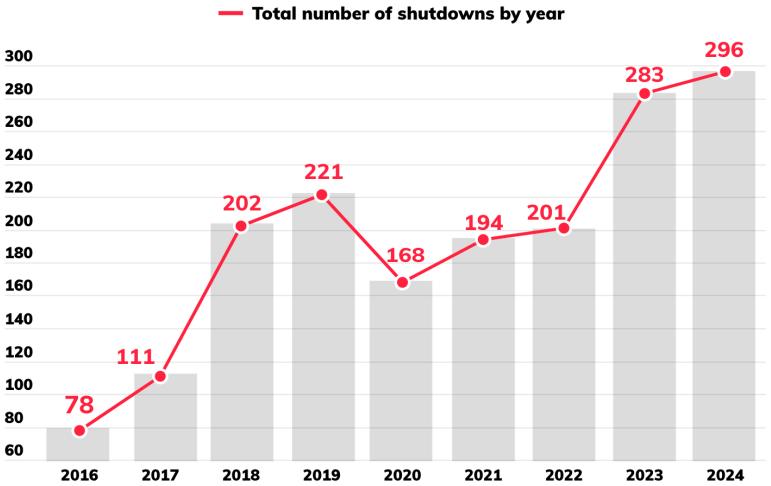

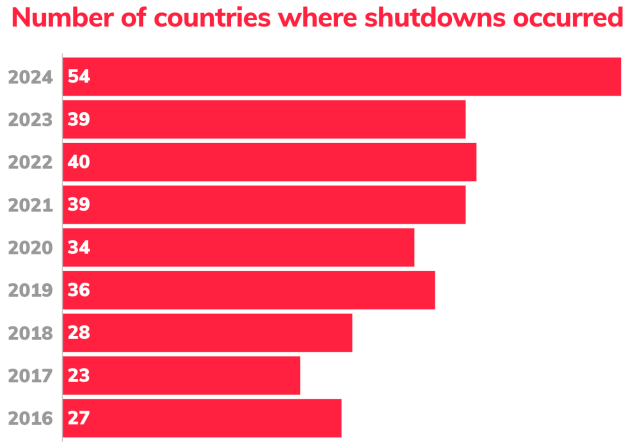

https://www.accessnow.org/internet-shutdowns-2024/

“How can we tackle these shutdowns?”

Blackout-resistant Messaging via Mobile Mesh Networks

- Allow for communication without Internet or cellular access
  - Rely on wireless capabilities (Bluetooth + WiFi Direct) of modern smartphones
  - Messages hop from phone to phone (gossip-based)

FireChat – the messaging app that's powering the Hong Kong protests

The internet is vulnerable to state intervention, but demonstrators have found a way around it

Pro-democracy supporters checking their phones during the protests in Hong Kong. Photograph: Anthony Kwan/Getty Images Photograph: Anthony Kwan/Getty Images

https://www.theguardian.com/world/2014/sep/29/firechat-messaging-app-powering-hong-kong-protests

Hong Kong protesters using Bluetooth Bridgefy app

3 September 2019

Share Save

Jane Wakefield
Technology reporter

Pro-democracy protesters in Hong Kong have been turning to a new app to communicate - one that does not use the internet and is therefore harder for the Chinese authorities to trace.

https://www.bbc.com/news/technology-49565587

Offline message app downloaded over million times after Myanmar coup

By Fanny Potkin and Jessie Pang

February 2, 2021 1:06 PM EST · Updated 4 years ago

Myanmar Army armored vehicles drive past a street after they seized power in a coup in Mandalay, Myanmar February 2, 2021. REUTERS/Stringer Purchase Licensing Rights

https://www.reuters.com/article/technology/offline-message-app-downloaded-over-million-times-after-myanmar-coup-idUSKBN2A22H0/

Desirable Properties for Mesh Messaging Apps

Flexible Communication Models

- One to one
- Some to some
- One to many (broadcast)

User Anonymity

- Sender and receiver
- Forward anonymity
- Post-compromise anonymity

Trust Systems

- Direct Trust
- Direct Trust Mediation
- Transitive Trust

Identity Revocation

- Soft revocation
- Hard revocation

The Mesh Messaging Apps Landscape

Application	Communication			Anonymity			Trust System			Revocable IDs	
	O2O	S2S	O2M	SRA	FA	PCA	DT	DTM	TT	SR	HR
Firechat [9]	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗
Bridgefy [11]	✓	✓	✓	✗	✗	✗	✓	✗	✗	✗	✗
Briar [10]	✓	✓	✗	✗	✗	✗	✓	✓	✗	✗	✗
1am [25]	✓	✓	✗	✗	✗	✗	✓	✗	✗	✗	✗
Moby [22]	✓	✗	✗	✓	✓	✗	✓	✗	✗	✗	✗
Perry et. al. [26]	✓	✓	✗	✓	✗	✗	✓	✓	✗	✗	✗
ASMesh [23]	✓	✗	✗	✓	✓	✓	✓	✗	✗	✗	✗
Rangzen [7]	✓	✗	✓	✓	✓	✗	✓	✗	✓	✗	✗

Existing apps lack desirable properties

The Mesh Messaging Apps Landscape

Application	Communication			Anonymity			Trust System			Revocable IDs	
	O2O	S2S	O2M	SRA	FA	PCA	DT	DTM	TT	SR	HR
Firechat [9]	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗
Bridgefy [11]	✓	✓	✓	✗	✗	✗	✓	✗	✗	✗	✗
Briar [10]	✓	✓	✗	✗	✗	✗	✓	✓	✗	✗	✗
1am [25]	✓	✓	✗	✗	✗	✗	✓	✗	✗	✗	✗
Moby [22]	✓	✗	✗	✓	✓	✗	✓	✗	✗	✗	✗
Perry et. al. [26]	✓	✓	✗	✓	✗	✗	✓	✓	✗	✗	✗
ASMesh [23]	✓	✗	✗	✓	✓	✓	✓	✗	✗	✗	✗
Rangzen [7]	✓	✗	✓	✓	✓	✗	✓	✗	✓	✗	✗
Anix	✓	✓	✓	✓	✓	✓	✓	✓✓	✓	✓	✓

Anix

- Anonymous blackout-resistant mesh messaging platform:
  - Based on selectively linkable one-time-use pseudonyms (PSUs)
  - Able to establish & manage trust relationships across the mesh
  - Able to prioritize microblogging-style messages vouched by trusted contacts via anonymous message endorsing

Anix's Operational Workflow

Alice: Sender

Long-term ID

Anix's Operational Workflow

Alice: Sender

PSUs:

- One-time-use
- Unlinkable
- Anonymous
  - Unless ID is known 

Anix's Operational Workflow

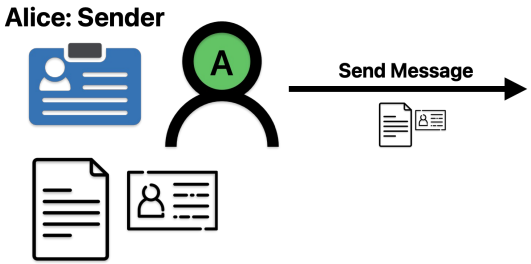

PSUs:

- One-time-use
- Unlinkable
- Anonymous
  - Unless ID is known

Anix's Operational Workflow

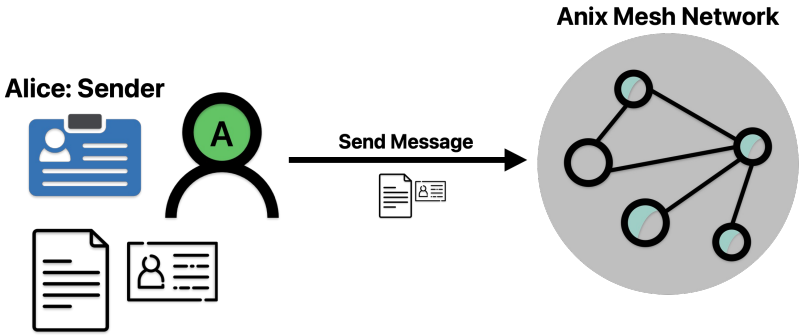

PSUs:

- One-time-use
- Unlinkable
- Anonymous
  - Unless ID is known

Anix's Operational Workflow

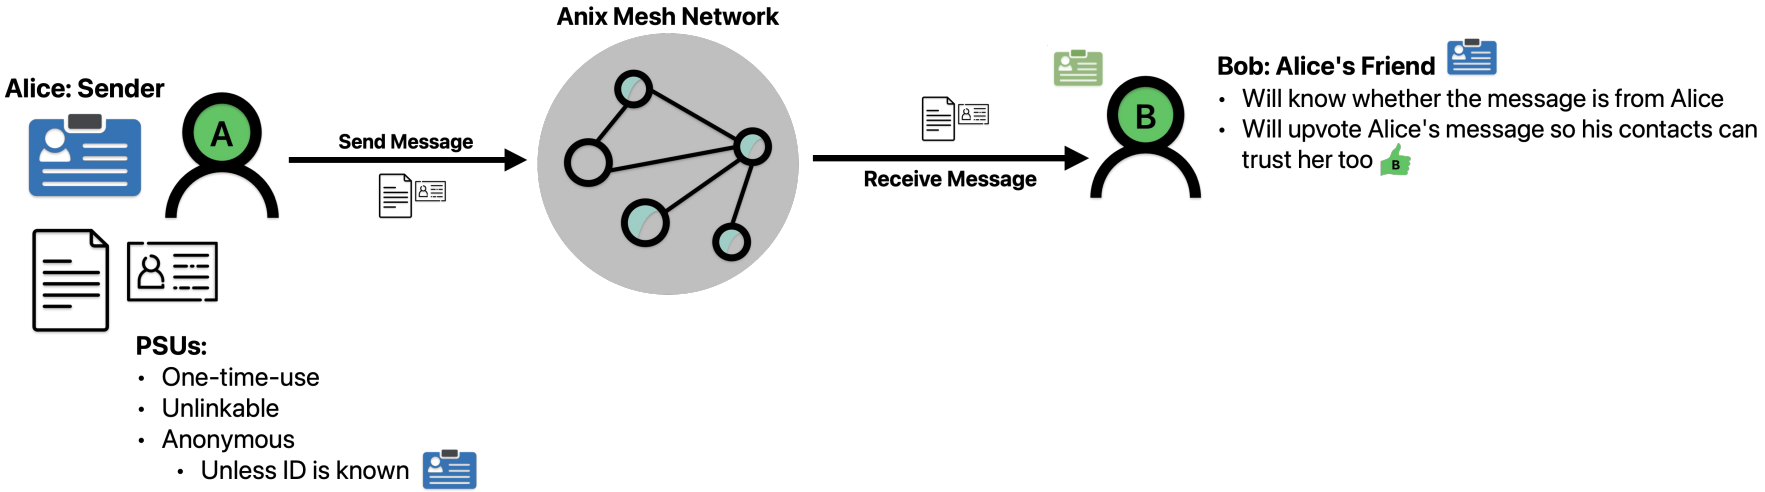

Anix's Operational Workflow

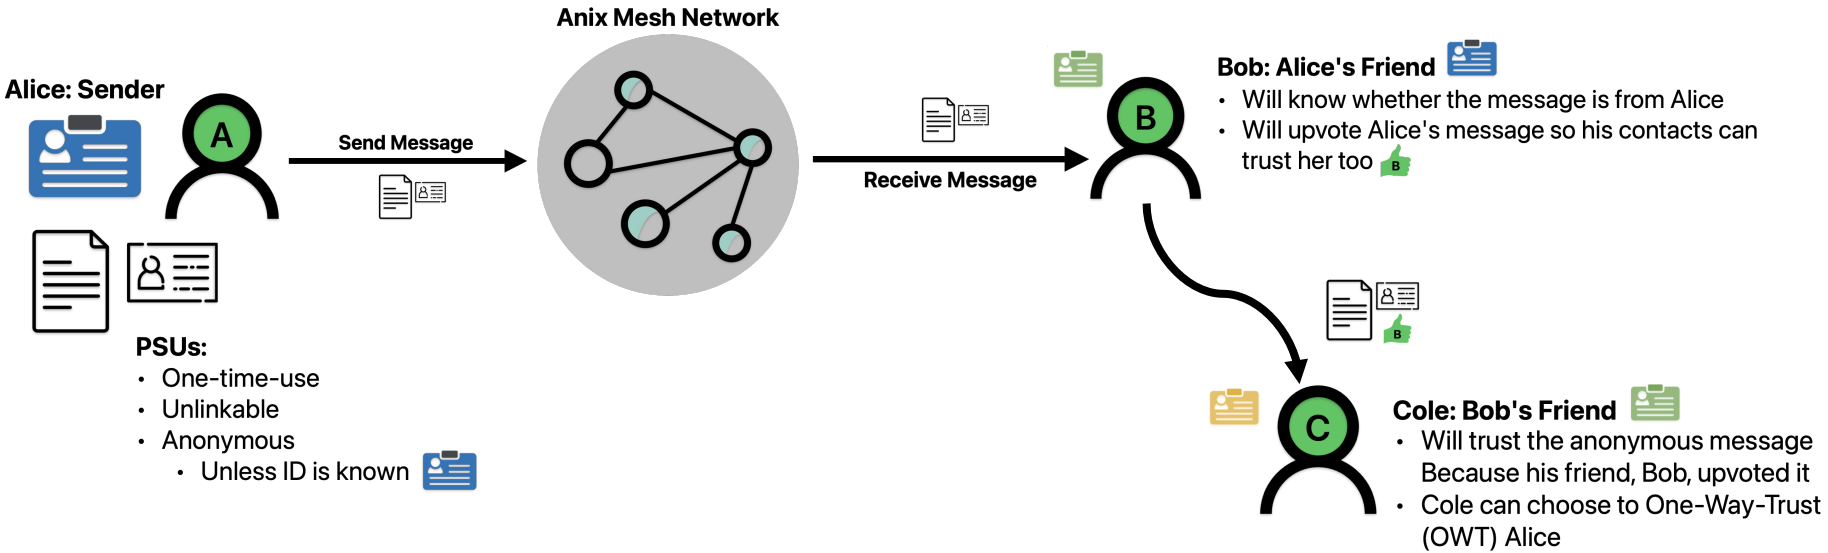

Anix's Operational Workflow

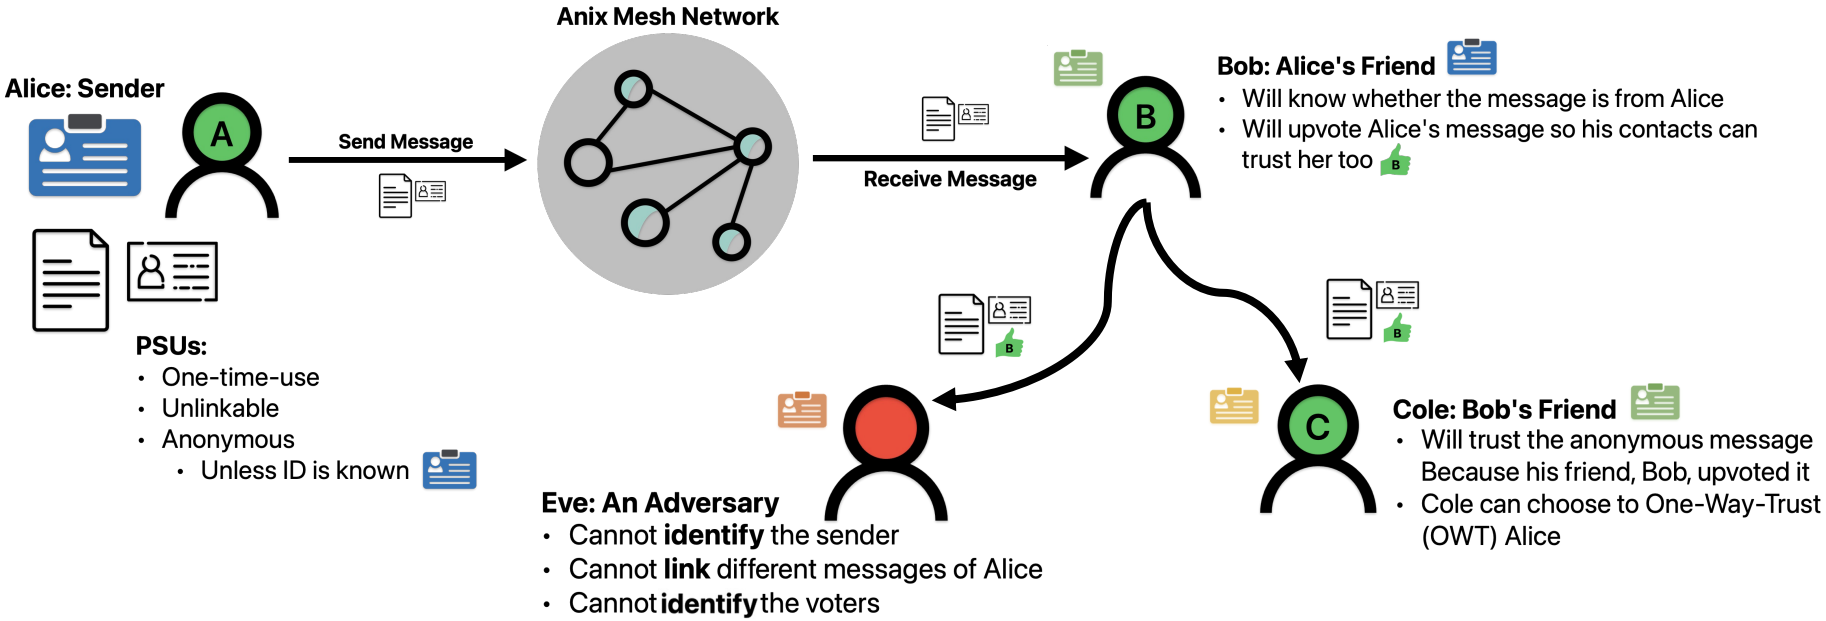

One-time-use Pseudonyms (PSUs)

- Each user holds two sets of key pairs:
  - Long term ID keys (kept secret)
  - One-time-use (OTU) keys

- Used to generate PSUs
- Allow selective linking of a user's messages/votes by trusted contacts

$$PSU = Pub_{OTU} \,||\, bSig_{Priv_{ID},\, bk}(Pub_{OTU})$$

, where $bSig$ is a public key-blinded signature scheme*

with blinding factor $bk = \text{Hash}(\text{Pub}_{\text{OTU}} || \text{Pub}_{\text{ID}})$

*"Key Blinding for Signature Schemes". *F. Denis, E. Eaton, T. Lepoint, and C. A. Wood.* **irtf-cfrg 2024**

Evaluation: Performance Micro-Benchmarks

- Implemented Anix on Android
- Avg. data exchange time: 11.58s
  - 100 messages * 10,000 votes (each)
- Avg. battery consumption: 1.5%/h

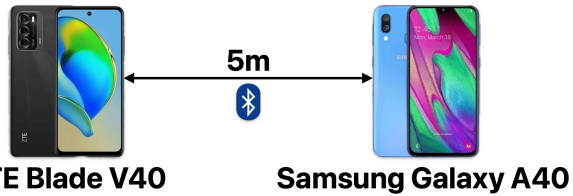

Computation time (in ms) for Anix operations

Op./Device	Gen. PSU	Create Msg.	Create Vote	Verify Sig.	BVer (Alg. 3)
Samsung A40	175.06 ± 1.05	46.30 ± 0.01	84.61 ± 1.14	61.33 ± 0.21	67.68 ± 0.21
ZTE Blade V40	64.95 ± 0.29	19.75 ± 0.01	38.76 ± 0.32	43.29 ± 0.28	47.30 ± 0.48

Evaluation: Simulation Testbed

- Simulated a scaled-down city environment with 600 users
- Blackout duration of 5 days (120 simulation steps)
- Most users are benign (98%), but a fraction are malicious (2%):
  - Drop benign messages
  - Attempt to gain the trust of benign users
  - Spread misinformation

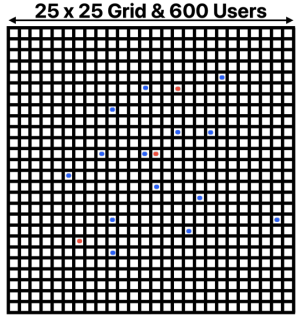

Evaluation: Simulation Testbed

- Simulated a scaled-down city environment with 600 users
- Blackout duration of 5 days (120 simulation steps)
- Most users are benign (98%), but a fraction are malicious (2%):
  - Drop benign messages
  - Attempt to gain the trust of benign users
  - Spread misinformation

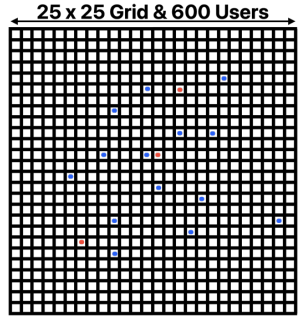

Parameter / Category	Description	Value
A and B	Dimensions of the simulation world ($A \times B$ grid)	25×25
m	Maximum distance that a user can move in a simulation step	2
N	Total number of users	600
β	Connectivity of the network given by the Watts-Strogatz model	0.5
K	Average social graph degree in the Watts-Strogatz model	15
T	Total steps of the simulation (1 step = 1 hour)	120
Adv	Fraction of adversarial nodes amongst all users	2% – 25%
S	Maximum device storage space allotted to the Anix app	3 GB
P_{inter}	Probability of a given user interacting with the Anix app at any step	0.15
C_m	Probability for a user to send out a message in a given step	0.05
OWT_{ud}	Required ratio of a message's known upvotes/downvotes to OWT the author	0.66
U_{ud}	Required ratio of a message's known upvotes/downvotes to upvote it	0.55
R	Ratio of an adversary's friends to benign user's friends	0.1 – 0.9
UV	Probability of a user who has no information about a message to vote on it	0.01 – 0.2
UM	Probability that a user upvotes a message containing misinformation	0.1 – 0.5
UN	Probability that a user upvotes a benign message	0.5 – 0.8
tp_m	Persistence time of a message on a user's device	24h

Check the paper for results under multiple settings!

Coverage and Resilience to Misinformation

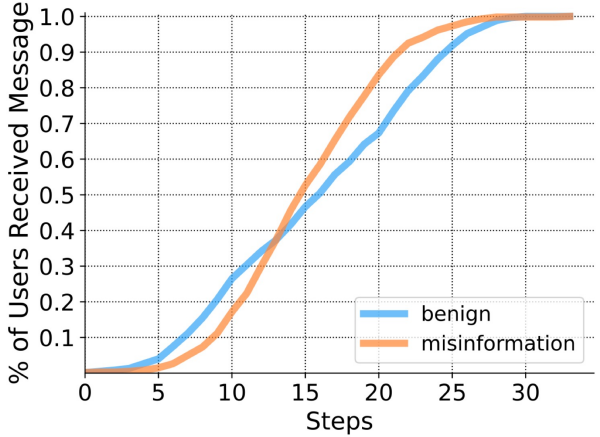

Benign messages take ~1 day to reach >90% of users

Messages up/downvoted by the majority of users

Scenario $(Adv = 0.02)$	**Misinformation** Upvoted	Downvoted
Very naive	204	1164
Naive	40	1301
Default	25	1320
Aware	15	1314
Very Aware	5	1297

Anix users can weed out misinformation

Intermission

- Internet shutdowns are becoming prevalent, and existing blackout-resistant mesh networking apps cannot sufficiently address users' needs
- **Anix** is an anonymous mesh-based microblogging platform
  - Enables trusted users to exchange data while remaining anonymous to untrusted users
  - Resilient to adversaries aiming to spread misinformation
- Future work:
  - Automate identity revocation (*dead-man's switch*?)
  - Optimize vote exchange (*dynamic group signatures*?)

Intermission

- Internet shutdowns are becoming prevalent, and existing blackout-resistant mesh networking apps cannot sufficiently address users' needs
- **Anix** is an anonymous mesh-based microblogging platform
  - Enables trusted users to exchange data while remaining anonymous to untrusted users
  - Resilient to adversaries aiming to spread misinformation
- Future work:
  - Automate identity revocation (*dead-man's switch*?)
  - Optimize vote exchange (*dynamic group signatures*?)

Now what?

Other usage scenarios?

- Prior works have focused on messaging and microblogging.
- A recent study* has identified other needs shared by users, namely on accessing:
  - "work-related resources"
  - "educational materials like Wikipedia"
  - Reference "news or articles"

Accessing knowledge sources remains unfeasible

* "Bridging barriers: A survey of challenges and priorities in the censorship circumvention landscape". *Diwen Xue et al.*

In a similar (online) vein…

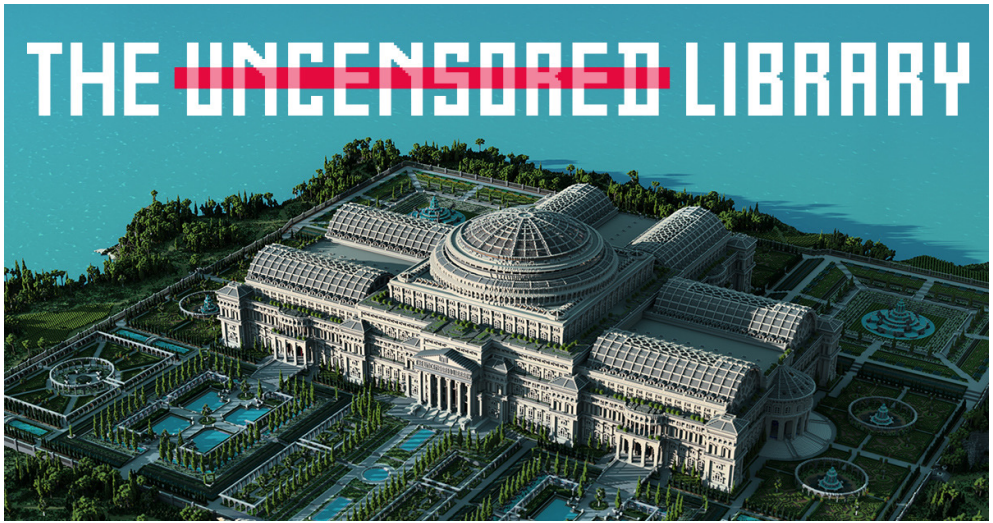

But even where almost all media is blocked or controlled, the world's most successful computer game is still accessible. Reporters Without Borders (RSF) uses this loophole to bypass internet censorship to bring back the truth – within Minecraft.

https://www.uncensoredlibrary.com/en

CTTF: Cache to the Future

- Provide access to cached web content via a distributed internet archive
  - Resilient against a wide range of attacks
  - Evaluated on real-world citywide smartphone GPS data
- Two-phased approach:
  - Pre-blackout: users rate webpages, which will determine which pages are cached
  - During blackout: users request pages and receive them from nearby devices

"Cache to the Future: A Distributed Webpage Archive for Internet Blackouts". *Ross Evans, Diogo Barradas.*
Under submission to IEEE S&P'26.

Pre-blackout Phase (I)

Start the cache

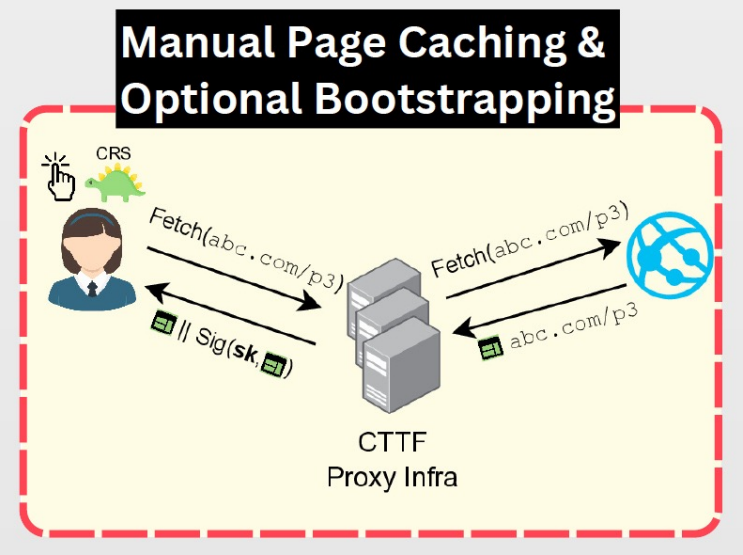

- Users cache pages through a trusted proxy which provides a digital signature.
- The signature prevents adversaries from tampering page contents during blackouts.
- Page fetching can resort to a CRS.

Pre-blackout Phase (II)

Exchange page ratings

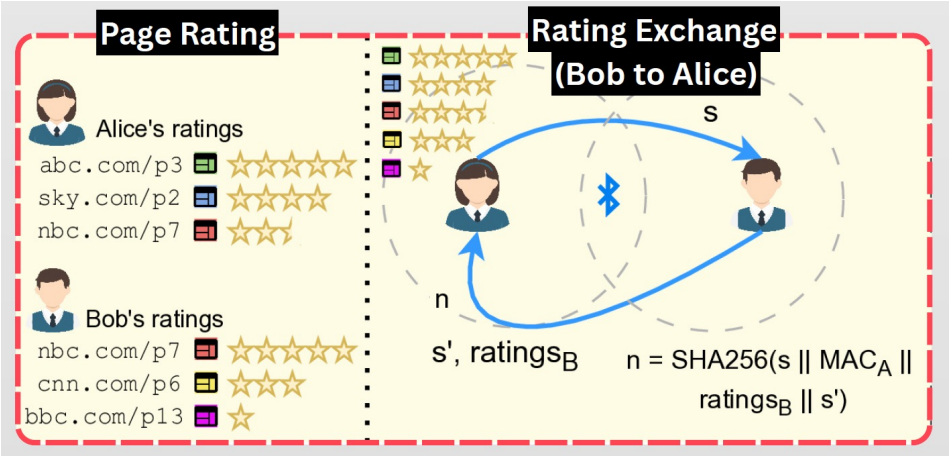

- Users rate pages based on their perceived utility during a blackout, affecting cache replication.
- Ratings are automatically shared among users when they come in close contact.
- A Proof-of-Work scheme ensures that adversaries cannot excessively manipulate local averages.

Pre-blackout Phase (III)

Automatic caching

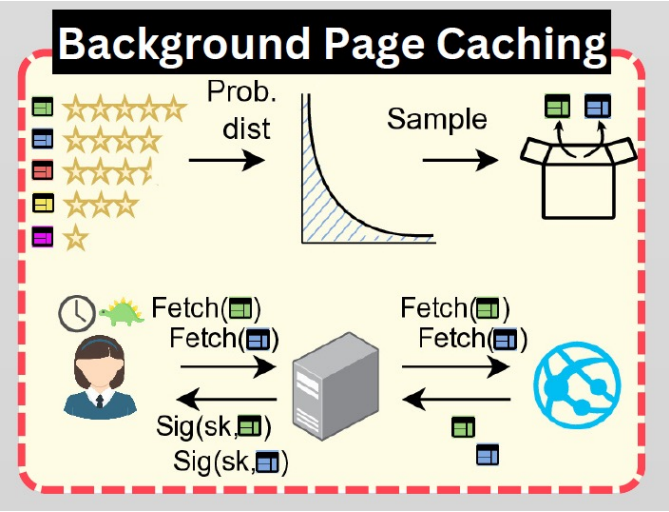

- Cache replication:
  - Occurs automatically
  - Ratings transformed into a Zipf
  - Sampling and storing according to leftover space

Blackout Phase

Request & receive pages

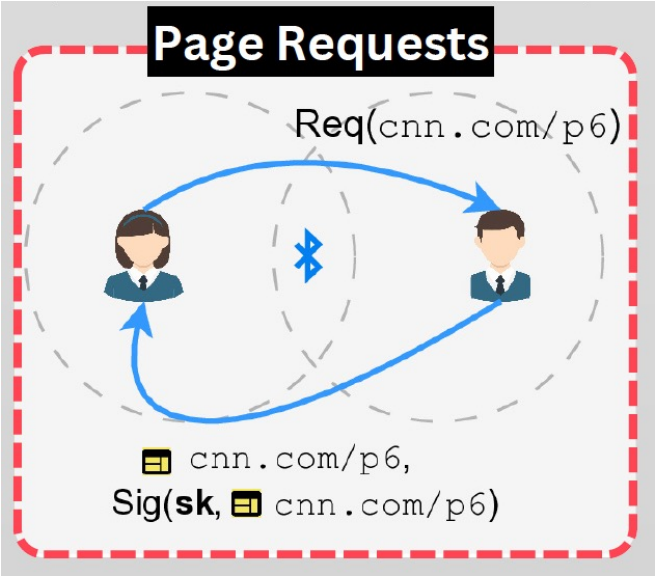

- Users request and receive pages via Bluetooth
- Authenticity is validated by checking the proxy's signature (or some *fallback mechanism*)
- Freshness-based request resolution:
  - Old vs new timestamp comparison for deciding whether to keep new copies

Evaluation Setup

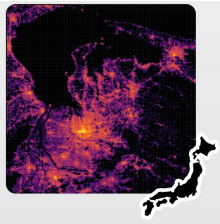

Dataset:

YJMob100K: real-world GPS data from 25,000 people in a Japanese city.

Cell area = 500m x 500m | **Positioning** data every 30min

Evaluation Setup

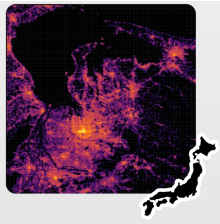

Dataset:

YJMob100K: real-world GPS data from 25,000 people in a Japanese city.

Cell area = 500m x 500m | **Positioning** data every 30min

User distribution:

Seeder users (24%): Actively cache and rate pages.
Leech users (74%): Just cache pages in the background.
Adversaries (2%): Rating manipulation, signal jamming, stalk.

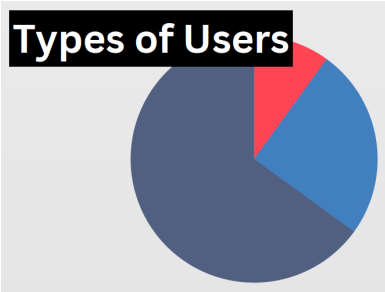

Evaluation Setup

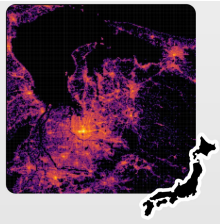

Dataset:

YJMob100K: real-world GPS data from 25,000 people in a Japanese city.

Cell area = 500m x 500m | **Positioning** data every 30min

User distribution:

Seeder users (24%): Actively cache and rate pages.
Leech users (74%): Just cache pages in the background.
Adversaries (2%): Rating manipulation, signal jamming, stalk.

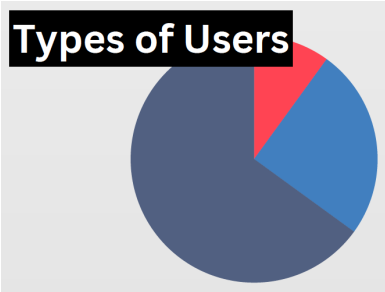

Baseline parameters:

- Seeders rate 500 pages prior to blackout;
- Each user requests a new page every 2h (in avg.)
- Users in the same cell interact with 100% prob, and can exchange 1000 ratings + 4 pages per 1min interaction

Parameter	Description	Value	Rationale
	Users' Distribution Parameters		
SEEDER_PERCENT	Percent of seeder users in the simulation	5%, 10%, **25%**, 50%, 75%, 90%	Conservative estimate [64], [65]
LEECHER_PERCENT	Percent of leecher users in the simulation	5%, 10%, 25%, 50%, **75%**, 90%	Conservative estimate [64], [65]
ADVERSARY_PERCENT	Percent of adversaries in the simulation	**0%**, 1%, 2%, 5%, 10%, 25%	Increasingly stringent – Stasi used 2% of citizens as informants [22], [89]
	Page Rating Assumptions Parameters		
PAGE_COUNT	Number of pages in the simulation universe	1 million	Top million sites capture over 95% of all page loads [66]
PAGES_STORED	Number of pages that can be cached on one device	1 200	3GB storage / 2.5MB median mobile page weight [67]
PAGES_RATED	Number of unique pages each proactive user has rated pre-blackout	**500**, 750, 1 000	Near average page visits per day [68]
RATING_CAP	Number of ratings any user can transmit to another	1 000	Mitigates adversarial manipulation
INDIVIDUAL_NOISE	Additive noise added to each seeders ratings (exponentially decays)	0.5	Accounts for differences in ratings
	Page Request Frequency Parameters		
PAGE_REQUEST_PROBABILITY	Probability an individual requests a page at any given timestep	0.25	Average one request per 2 hours
CONTACT_PROBABILITY	Probability two users will connect when in the same cell for a timestep	**100%**, 75%, 50%, 25%, 10%, 5%	Prior works [18], [20], [24] fix at 100% during evaluation
FORWARDING_LIMIT	How many pages can be exchanged during communication	4	One minute of connection time allows approximately 10MB of data to be exchanged (§6.5)
	Grid-Based Parameters		
GRID_SIZE	Dimensions of the simulated world in grid cells	**200 × 200**, 25 × 25	YJMOB100K [63], ASMesh [24], Anix [20]
TOTAL_USERS	Total number of users in the simulation	**25 000**, 600	YJMOB100K [63], ASMesh [24], Anix [20]
MOVEMENT_DISTANCE	Maximum number of grid cells a user can move in a simulation timestep	**N/A**, 2	YJMOB100K [63], ASMesh [24], Anix [20]
	Adversary Behaviour Parameters		
ADVERSARY_STRATEGY	Strategy of adversary for it's ratings	Decrease top 500, Increase bottom 500	Impede useful page caching and promote useless page caching
TOP_K_JAMMING_LOCATIONS	How many of the most popular cells will adversaries jam	**0**, 10, 100, 1 000, 10 000	Increasingly stringent
	Epidemic Routing Parameters		
SPACE_FOR_FORWARDING	Space dedicated to storing and forwarding non-personally-requested pages	300	750MB with 2.5MB median mobile page weight [67]
ADVERSARY_FORCE_MULTIPLIER	Number of useless page requests adversaries make and forward	**N/A**, 1, 8, 32, 128, 256, 512	Increasing spam
	Proof-of-Work Parameters		
STALKING	Whether adversaries engage in stalking for rating manipulation, necessitating PoW mitigation	**False**, True	Stalking and rapidly spoofing Bluetooth MAC address is an extreme threat model
POW_LIMIT	Avg. time to compute the PoW, which limits how many rating exchanges can happen in one timestep	**N/A**, 1 second, 30 seconds, 1 minute	Investigating appropriate parameters

tion Setup

GPS data from 25,000 people in a Japanese city.

00m | Positioning data every 30min

Types of Users

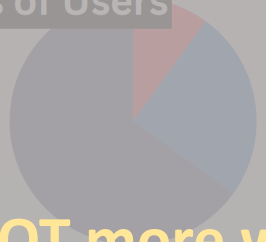

rate pages.
the background.
gnal jamming, stalk.

(in avg.)
% prob, and can
n interaction

Trust me, there's a LOT more where these came from…

Evaluation (Benign Scenario)

1 week cache warm-up, 1 week of requests, ~1 month steady-state

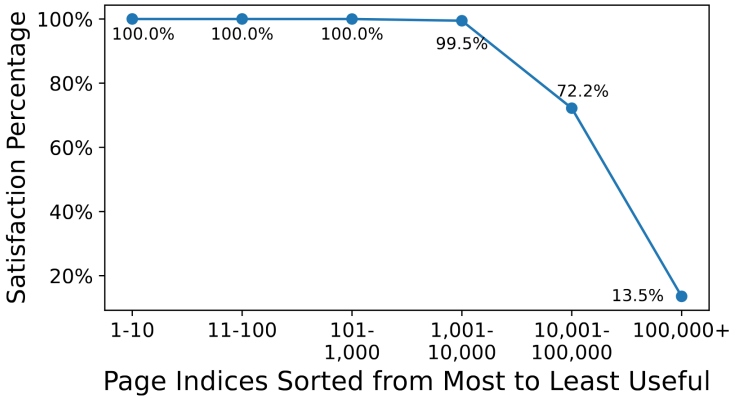

Users can fetch top 100K pages with high probability

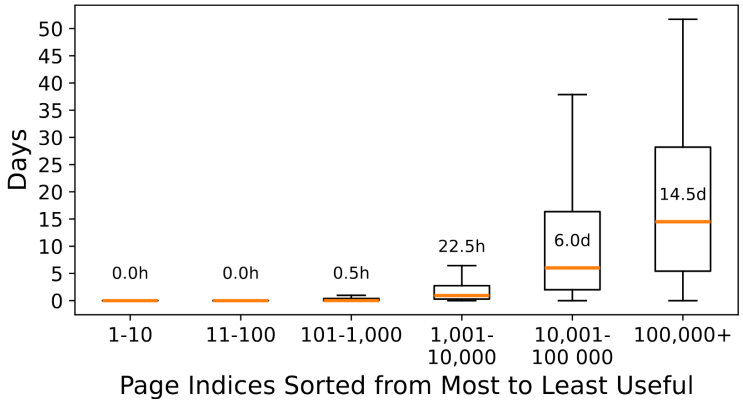

On avg., users can fetch top 10K pages within a day, 10K-100K within a week

Evaluation (Adversarial Scenario)

TLDR: Aggressive jamming and adversary participation

Evaluation (Adversarial Scenario)

TLDR: Aggressive jamming and adversary participation

YJMob100K area with cyan cells jammed, prioritizing those with most points of interest

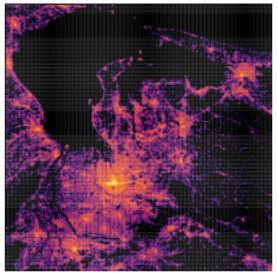

(a) Jamming 0 km^2

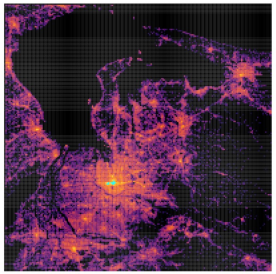

(b) Jamming 2.5 km^2

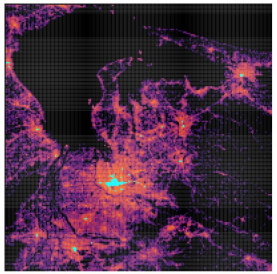

(c) Jamming 25 km^2

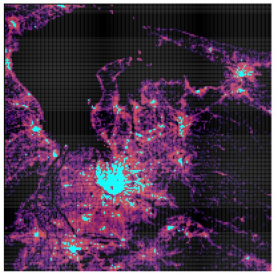

(d) Jamming 250 km^2

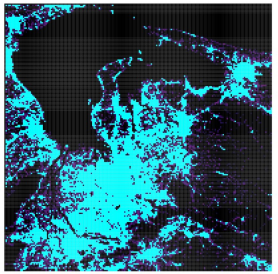

(e) Jamming 2 500 km^2

Evaluation (Adversarial Scenario)

TLDR: Aggressive jamming and adversary participation

YJMob100K area with cyan cells jammed, prioritizing those with most points of interest

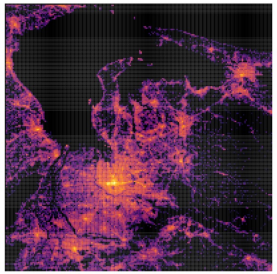

(a) Jamming 0 km^2

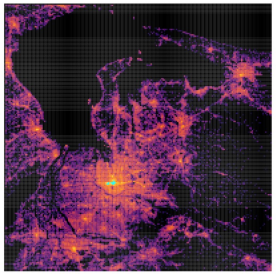

(b) Jamming 2.5 km^2

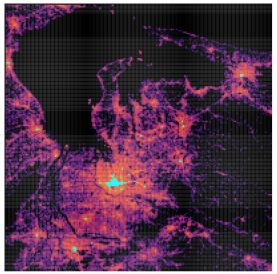

(c) Jamming 25 km^2

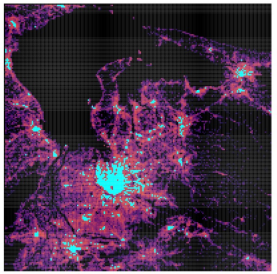

(d) Jamming 250 km^2

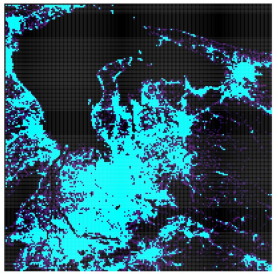

(e) Jamming 2 500 km^2

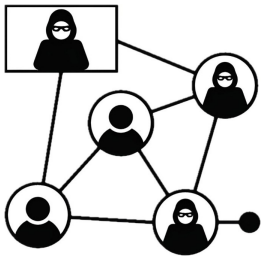

+ Up to 25% Sybils

Evaluation (Adversarial Scenario)

1 week cache warm-up, 1 day of requests, ~1 week steady-state

Req. satisfaction after a week

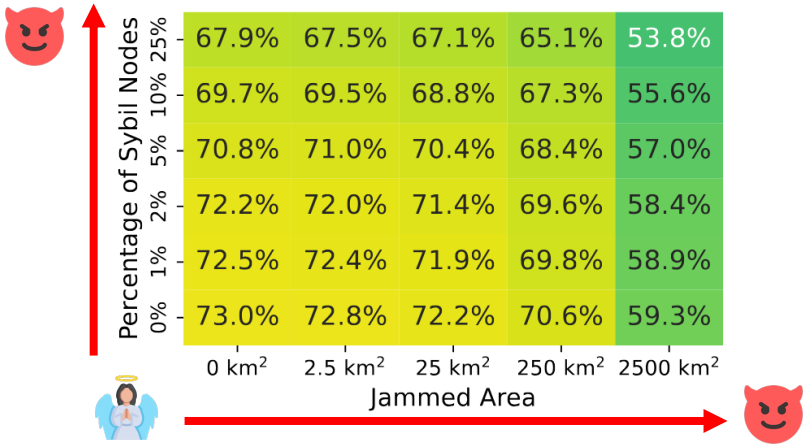

Users receive 70% of the pages they request

90th percentile – latency for top 10K pages

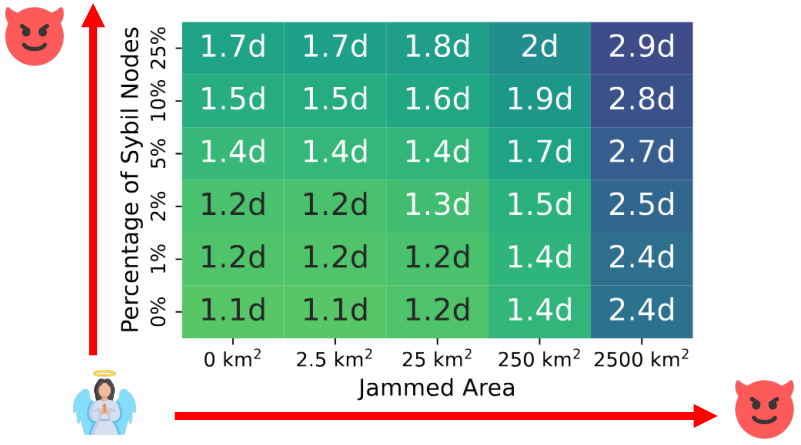

Users receive 90% of requested pages within 1-2 days

Intermission

- Blackout-resistant technologies are desirable for more than messaging
- CTTF provides access to cached web content via a distributed internet archive
  - Resilient against a wide range of attacks
  - Evaluated on real-world citywide smartphone GPS data
- Future work:
  - Resist user fingerprinting and information leakage:
    - Shield rating exchanges with *differential privacy (DP)*?
    - Shield requested pages with *private information retrieval (PIR)*?

Takeaways

- Blackout-resistant technologies lack anonymity and trust propagation primitives
  - Our work on **Anix** enables trust formation across an anonymous mesh of devices
- Blackout-resistant technologies are desirable for more than messaging
  - We developed **CTTF** as a distributed blackout-resistant Internet archive

Diogo Barradas
diogo.barradas@uwaterloo.ca

Thank you!