I Was Told There Would be Blockchain: Five Years of Real-World Cryptography at DARPA

Dr Josh Baron

Program Manager, Information Innovation Office

Real World Crypto

March 2023

Blockchain:
graphy at DARPA

ovation Office

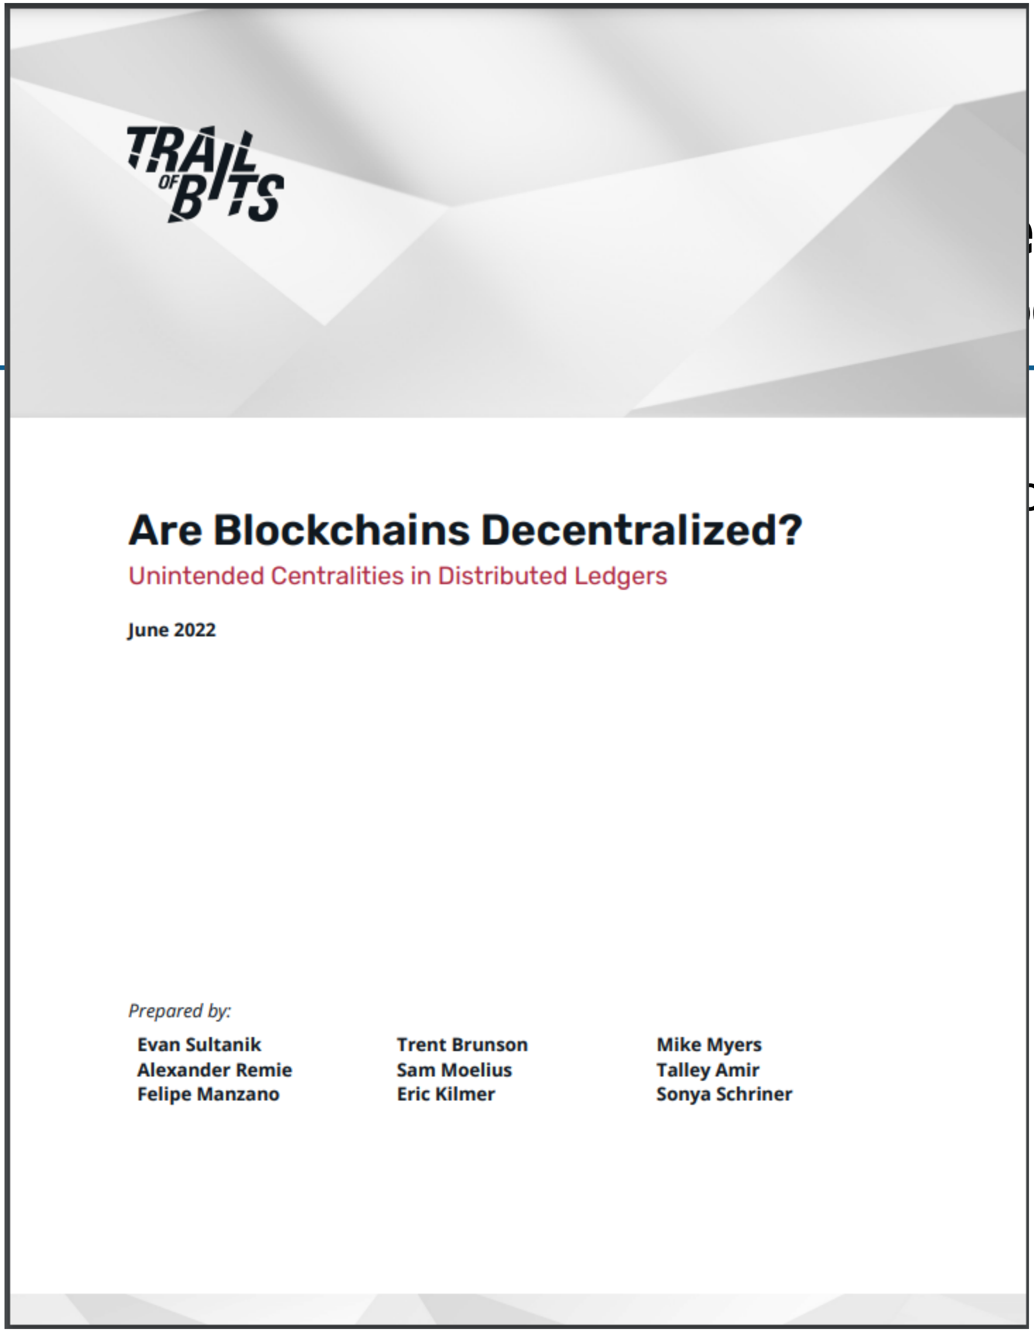

Are Blockchains Decentralized?

Unintended Centralities in Distributed Ledgers

June 2022

Prepared by:

Evan Sultanik
Alexander Remie
Felipe Manzano
Trent Brunson
Sam Moelius
Eric Kilmer
Mike Myers
Talley Amir
Sonya Schriner

https://blog.trailofbits.com/2022/06/21/are-blockchains-decentralized/

Cryptographic Programs at DARPA

- Proceed – Computation on encrypted data (Completed)
  - Fully Homomorphic Encryption (FHE), Secure Multiparty Computation (MPC)
- SAFER – Safe, resilient communications over the internet (Completed)
  - Pluggable Transports, Decoy Routing, Three-Party MPC
- Brandeis – Build privacy-aware systems (Completed)
  - MPC, Differential privacy, human factors
- SAFEWARE – Provably-secure software obfuscation (Completed)
  - Indistinguishability Obfuscation
- RACE – Secure, distributed messaging in contested network environments (Ongoing)
  - MPC, Obfuscated Communications
- SIEVE – Zero knowledge (ZK) proofs for DoD applications (Ongoing)
  - Translate DoD-relevant problems into nondeterministic polynomial time (NP) problems, ZK for large circuits
- Cooperative Secure Learning – Privacy-preserving machine learning (Completed)
  - FHE, MPC, Differential privacy
- DPRIVE- Hardware accelerator for FHE (Ongoing)
- MICE – AI-enabled censorship measurement (Completed)

Cryptographic Programs at DARPA

- Proceed – Computation on encrypted data (Completed)
  - Fully Homomorphic Encryption (FHE), Secure Multiparty Computation (MPC)
- SAFER – Safe, resilient communications over the internet (Completed)
  - Pluggable Transports, Decoy Routing, Three-Party MPC
- Brandeis – Build privacy-aware systems (Completed)
  - MPC, Differential privacy, human factors
- SAFEWARE – Provably-secure software obfuscation (Completed)
  - Indistinguishability Obfuscation
- RACE – Secure, distributed messaging in contested network environments (Ongoing)
  - MPC, Obfuscated Communications
- SIEVE – Zero knowledge (ZK) proofs for DoD applications (Ongoing)
  - Translate DoD-relevant problems into nondeterministic polynomial time (NP) problems, ZK for large circuits
- Cooperative Secure Learning – Privacy-preserving machine learning (Complete)
  - FHE, MPC, Differential privacy
- DPRIVE- Hardware accelerator for FHE (Ongoing)
- MICE – AI-enabled censorship measurement (Completed)

Brandeis

Coffeebreak- MPC on phones

- Brandeis tech, released in August 2022
- Multiparty computation on phones (up to 10 parties in seconds-minutes)
- Joint work with Stealth Software, TwoSix Technologies, Raytheon BBN

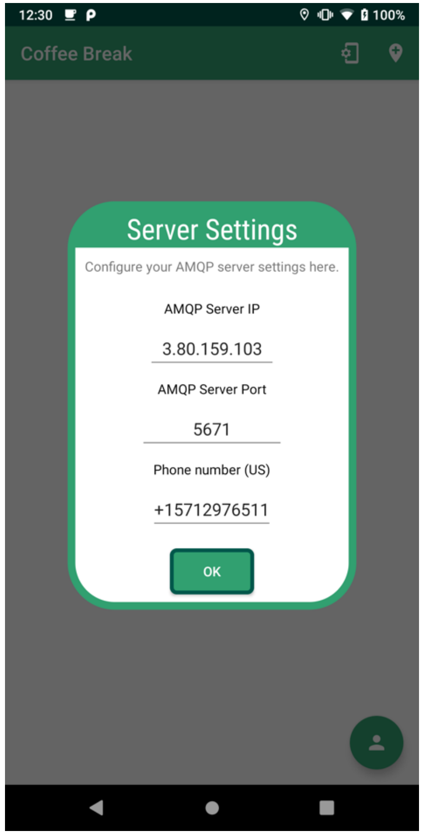

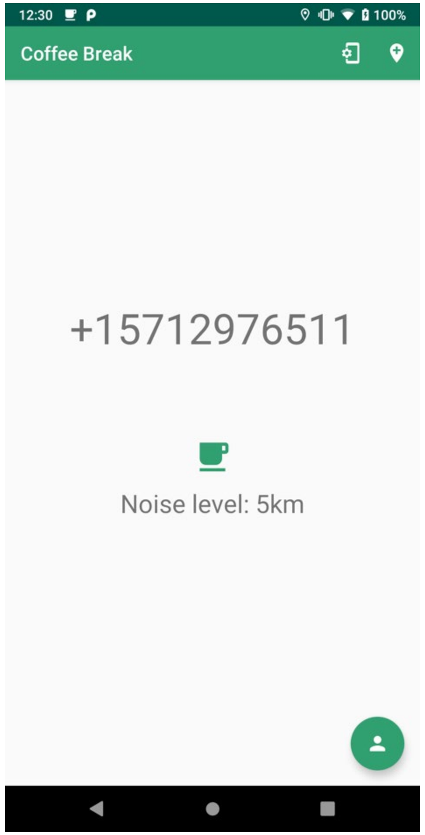

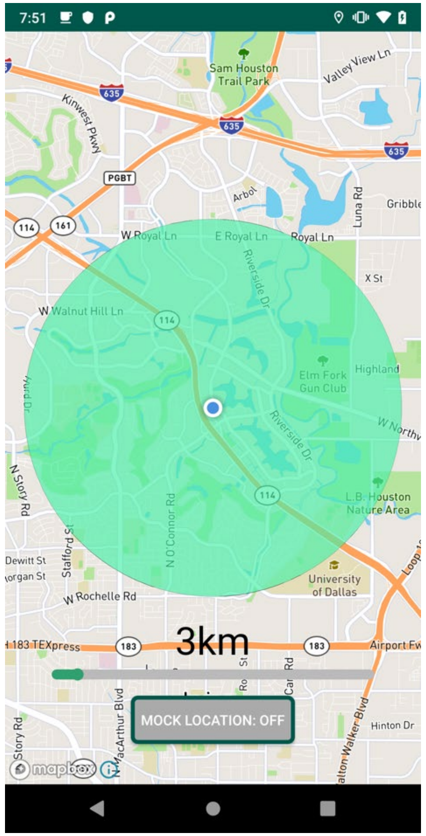

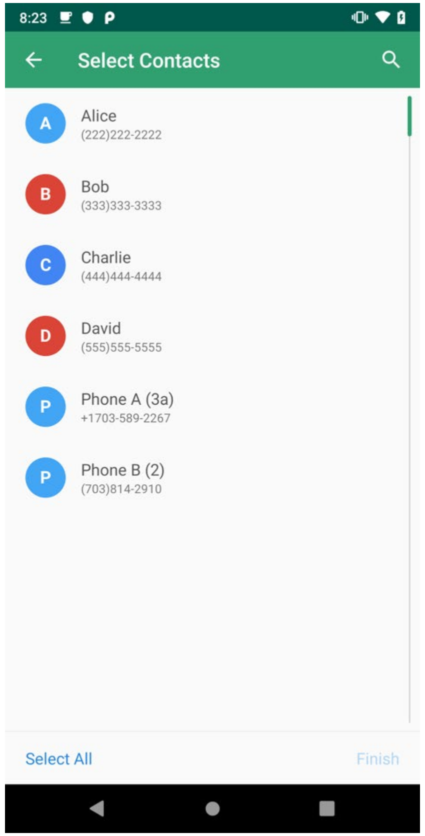

Source: Two Six labs demo

See https://github.com/twosixlabs/coffeebreak
https://github.com/stealthsoftwareinc/pulsar-mpc

Securing Information for Encrypted Verification and Evaluation (SIEVE)

SIEVE objective

Develop computer science theory and software to create mathematically verifiable public statements derived from hidden, sensitive information in order to publicly yet securely communicate about DoD capabilities

Enable verification while keeping secrets

SIEVE Technical Approach

Constructing Useful Zero Knowledge Statements

Building Efficient Zero Knowledge Proof Generation Compilers

Problem Statement

```c
#include "myMult.h"

void myMult(const double a[12], const double b[20], double c[15])
{
  int i0;
  int i1;
  int i2;
  for (i0 = 0; i0 < 3; i0++) {
    for (i1 = 0; i1 < 5; i1++) {
      c[i0 + 3 * i1] = 0.0;
      for (i2 = 0; i2 < 4; i2++) {
        c[i0 + 3 * i1] += a[i0 + 3 * i2] * b[i2 + (i1 << 2)];
      }
    }
  }
}
```

The statement is in its "native language" (e.g., plain English, code, math)

Encode

Intermediate Representation (IR)

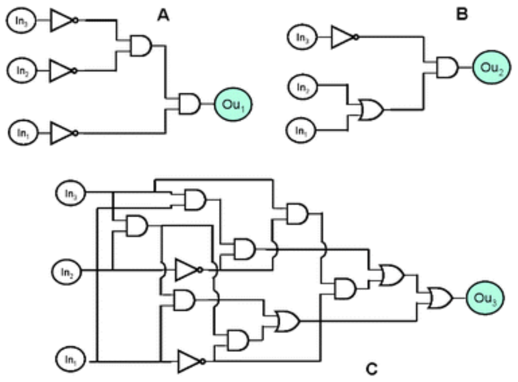

The statement is encoded in a lower-level language that will be used to construct the actual proof

In this case, the IR is a Boolean circuit made up of AND, OR, and NOT gates

Input

Zero Knowledge Proof

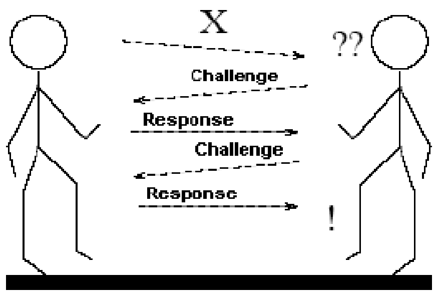

The proof is conducted between the prover and the verifier

Post-Quantum (PQ) Zero Knowledge

SIEVE Performers

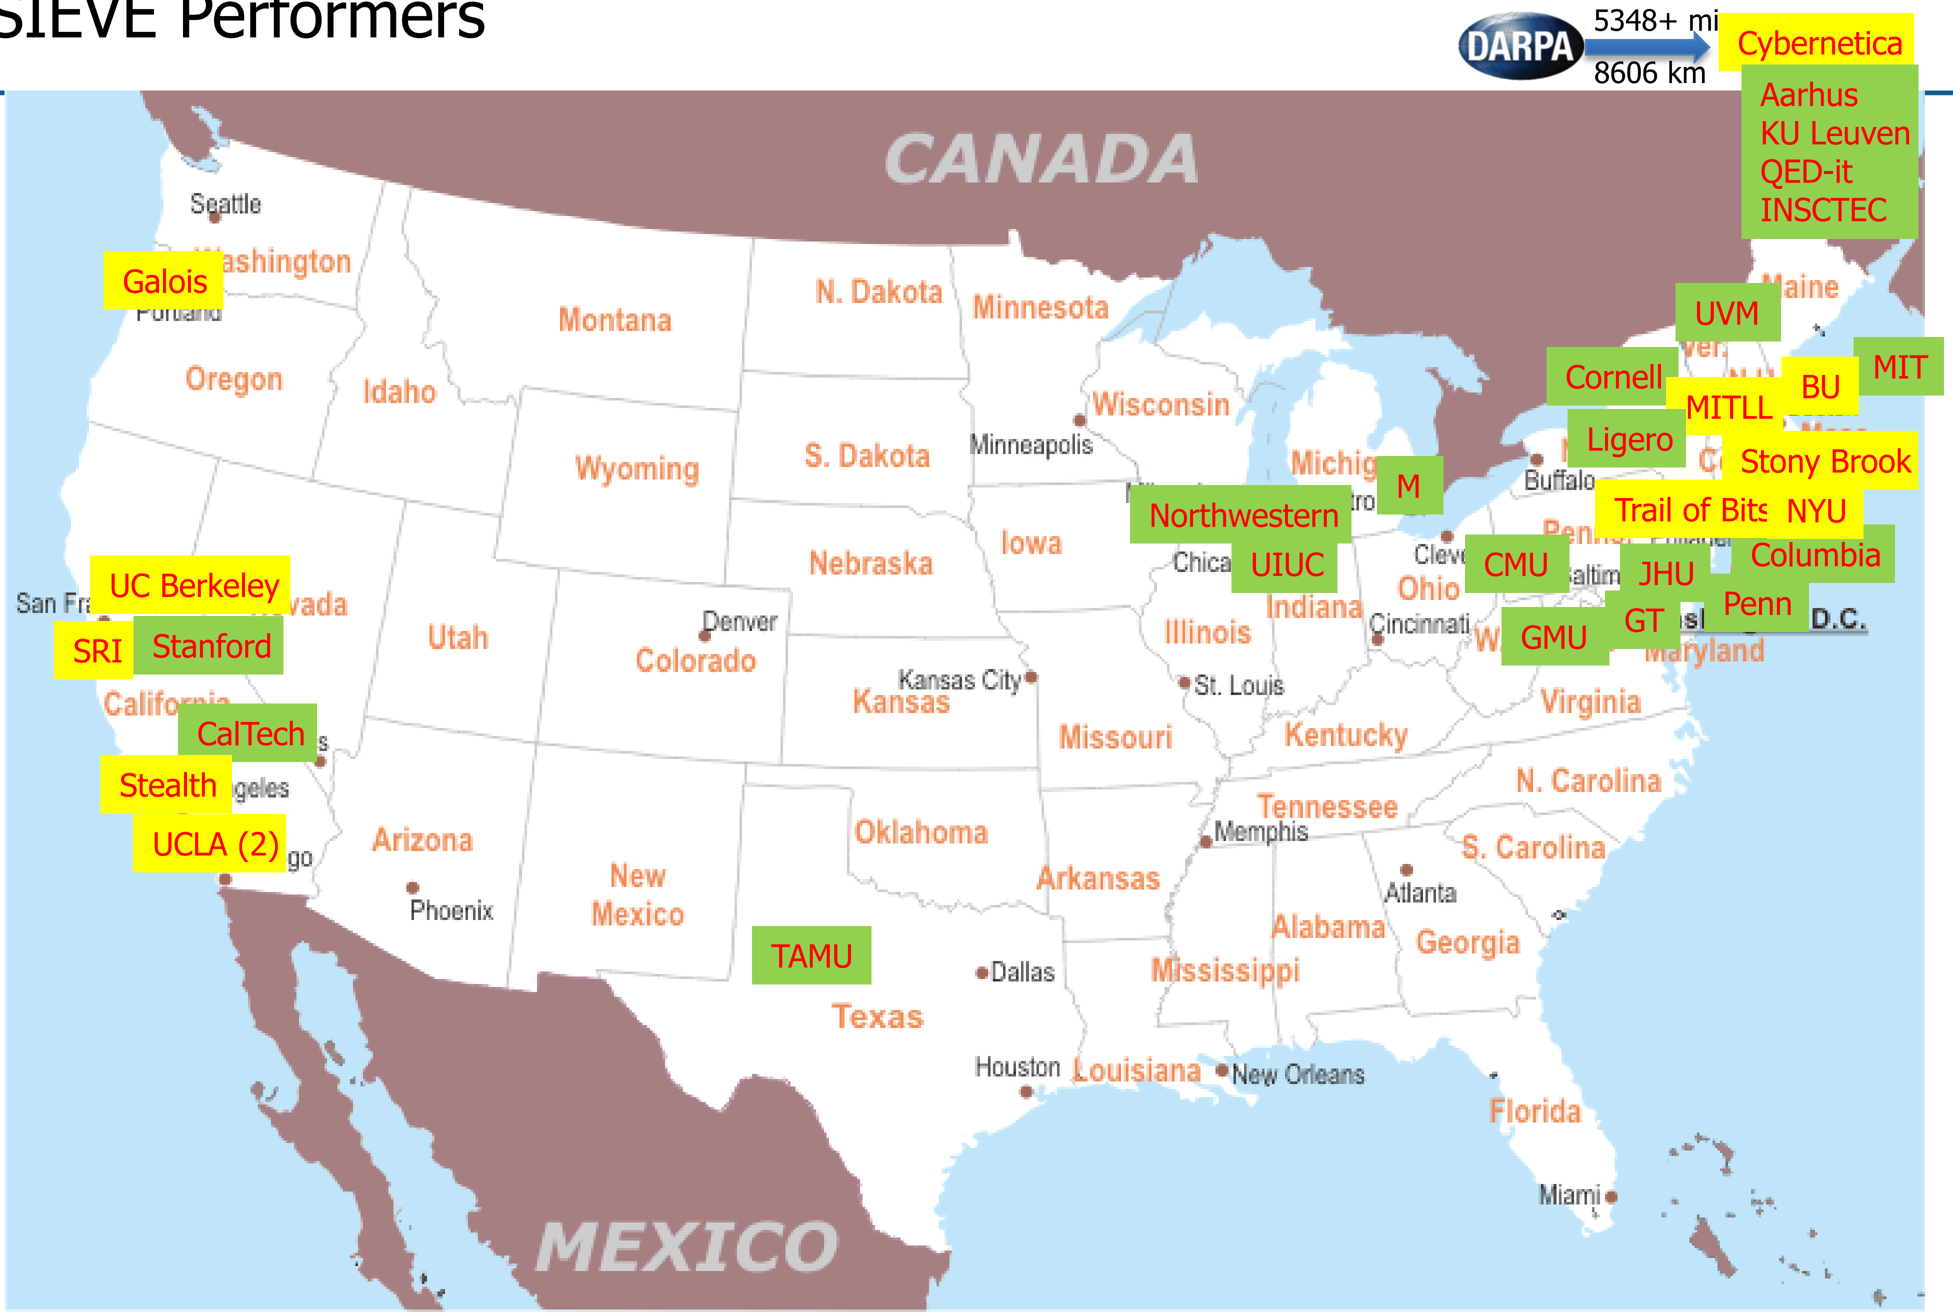

TA1 Performer Approaches

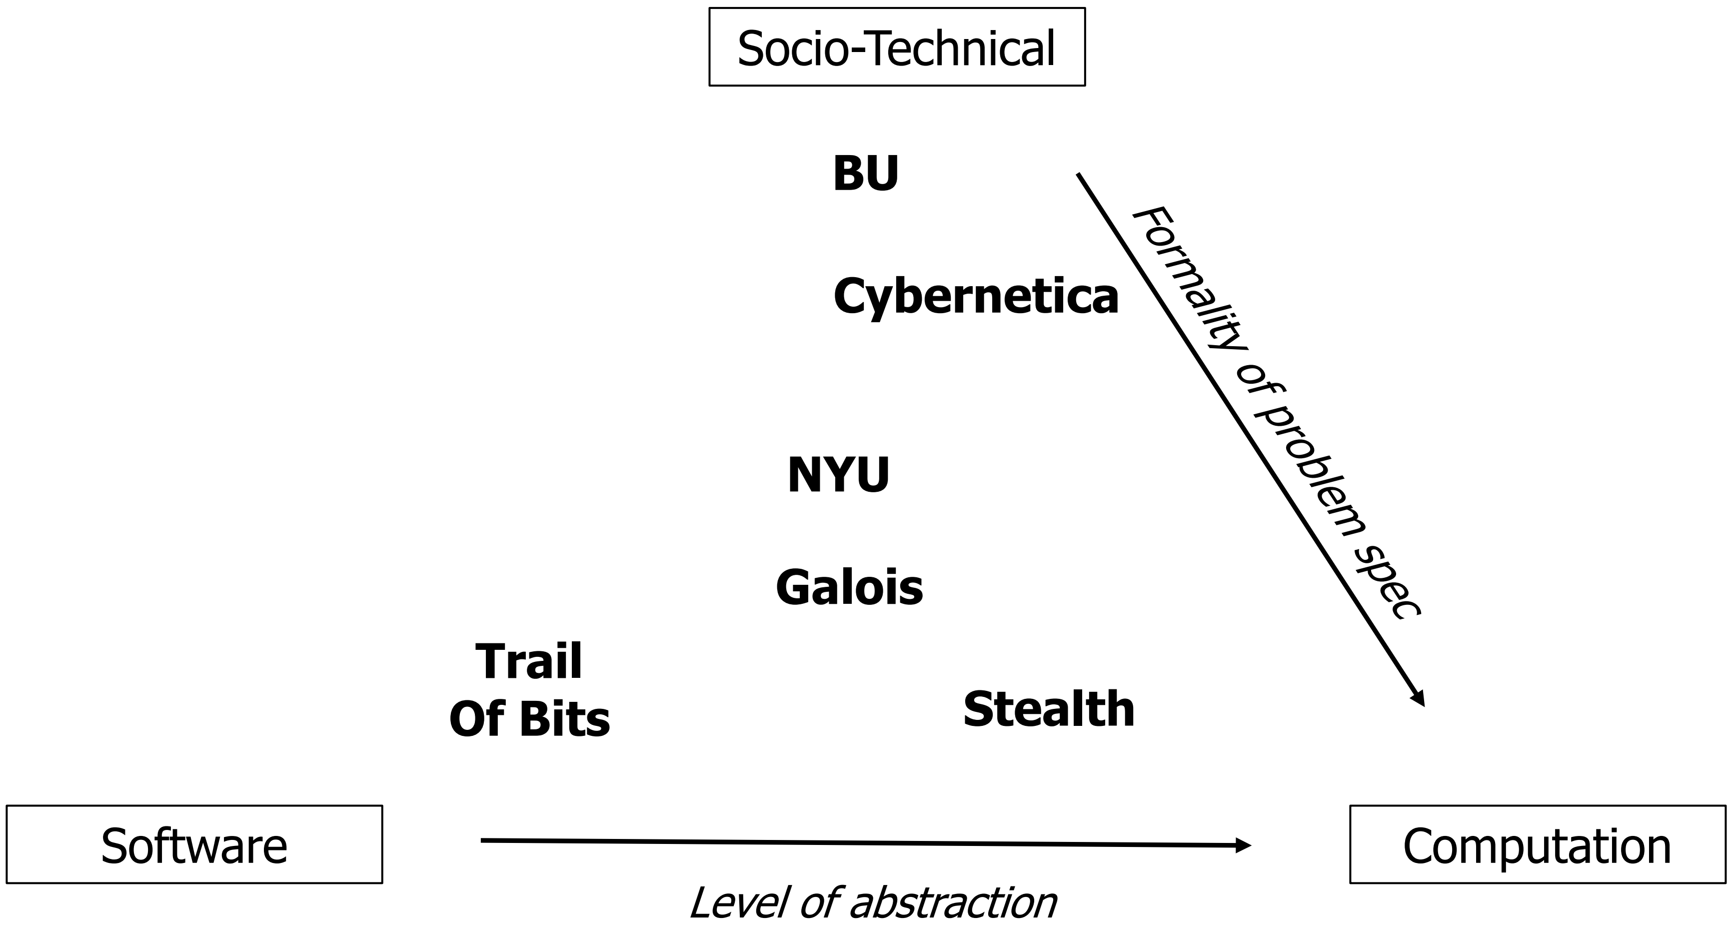

ZK Proofs of Exploitability

- **Idea:** We can embed a ZK proof of exploitability
- **First Approach:** Now: MSP430 Microcontroller. Coming: “x86” [Tiny86, 20k AND gates/cycle]
  - Green et at, “Efficient Proofs of Software Exploitability for Real-world Processors”, PoPETS 2023:

Table 1: Benchmarks for proofs of exploits (at 128 bits of security) for a representative subset of the Microcorruption exercises. The selected exercises cover the most important exploit categories, including buffer overflow, code injection, and bypassing memory protection. These exercises are ordered by the difficulty of the exercise, as estimated by the Microcorruption creators.

Exercise Name	Processor Cycles	Prover (sec)	Verifier (sec)	Size (mb)	Exploit Type
New Orleans	2392	22	7	295	Password embedded in binary
Hanoi	6199	25	18	322	Buffer overflow
Cusco	5178	21	15	269	Buffer overflow
Montevideo	6676	28	20	358	Code injection via `strcpy` bug
Johannesburg	6311	26	19	332	Stack cookie bypass
Santa Cruz	12835	754	39	680	Code injection via `strcpy` bug
Addis Ababa	5360	23	17	296	Format string vulnerability
Novosibirsk	19833	89	63	1100	Format string vulnerability
Vladivostok	50823	454	152	6048	ASLR bypass

- **Second Approach**: uncompiled C, C++, Rust code
  - Cuelar et al, “Cheesecloth: Zero-Knowledge Proofs of Real-World Vulnerabilities”, USENIX 2023

Program	Code size (K instrs)	Execution steps (K)	Mult gates (M)	Protocol time	Protocol memory
GRIT	3	5	26.7	3m 40s	845 MB
FFmpeg	24	79	672.7	1h 22m	19 GB
OpenSSL	340	1,300	17,049.5	36h 45m	460 GB

Table 1: Results for generating and running a ZK proof of software vulnerability for each case study.

Fast ZK cryptographic operations

Galois Team

	AND gate eval time / AES block eval
Start of SIEVE	450ns / 3000000ns
April 2022	7.5ns / 51200ns
Current	**4.5ns** / 30000ns
AES in software	~483ns (per block)

Evaluated Galois Mac'n'Cheese (CRYPTO 2021) on circuit provided by Trail of Bits

- ToB circuit:
  - One step of CPU
  - 21,592 AND gates, 2,666 XOR gates, 16,855 INV gates
- Performance: 3.0 ns per AND gate

⇒ 65.57 μs per CPU cycle

⇒ 15.25 kHz ZK Processor

Stealth Team (Ligero)

Performance on the browser (128-bit)

String length	Cons./state	Batch size	Prover speed (end to end)	Verifier Speed (end to end)
30	334820	1024	1.22 μs/g	14.1 ns/g
40	589676	1024	1.21 μs/g	12.9 ns/g
50	915684	1024	1.17 μs/g	11.8 ns/g

~1 Billion gates

And Non-interactive!

And Sublinear!!

And plausibly PQ secure!!

Performance on c6i.8x (32 vcpu,64 GB RAM)

String length	Cons. per state.	Batch size	Prover speed	Verifier Speed	Prover Memory	Verifier Memory	Proof Length
10	40804	16384	34.40 ns/g	2.39 ns/g	130MB	63MB	28MB
15	88282	16384	34.57 ns/g	2.28 ns/g	167MB	107MB	59MB
20	152012	16384	34.93 ns/g	2.41 ns/g	217MB	169MB	100MB

2.5 Billion gates

And Non-interactive!

See ZKProof 2022

ZK and Society

ACM CS/Law	**E. Balsa**, **H. Nissenbaum**, **S. Park**	Trust and Privacy: It's Complicated
	A. Bestavros, S. Dogan, P. Ohm, A. Sellars	Bridging the Computer Science-Law Divide
	D. Bitan, **R. Canetti**, **S. Goldwasser**, R. Wexler	Using Zero-Knowledge to Reconcile Law Enforcement Secrecy and Fair Trial Rights in Criminal Cases
	A. Cohen, S. Scheffler, **M. Varia**	Can the government compel decryption? Don't trust – verify
	S. Scheffler, E. Tromer, **M. Varia**	Formalizing Human Ingenuity: A Quantitative Framework for Copyright Law's Substantial Similarity
	J. Walsh, **M. Varia**, **A. Cohen**, A. Sellars, **A. Bestavros**	Multi-Regulation Computing: Examining the Legal and Policy Questions That Arise From Secure Multiparty Computation

See also: "Verification Dilemmas, Law, and the Promise of Zero-Knowledge Proofs" Kenneth Bamberger, Ran Canetti, Shafi Goldwasser, Rebecca Wexler, Evan Zimmerman, Berkeley Technology Law Journal, Vol. 37, No. 1, 2022

ZK proofs for electric vehicle subsidy terms (1/2)

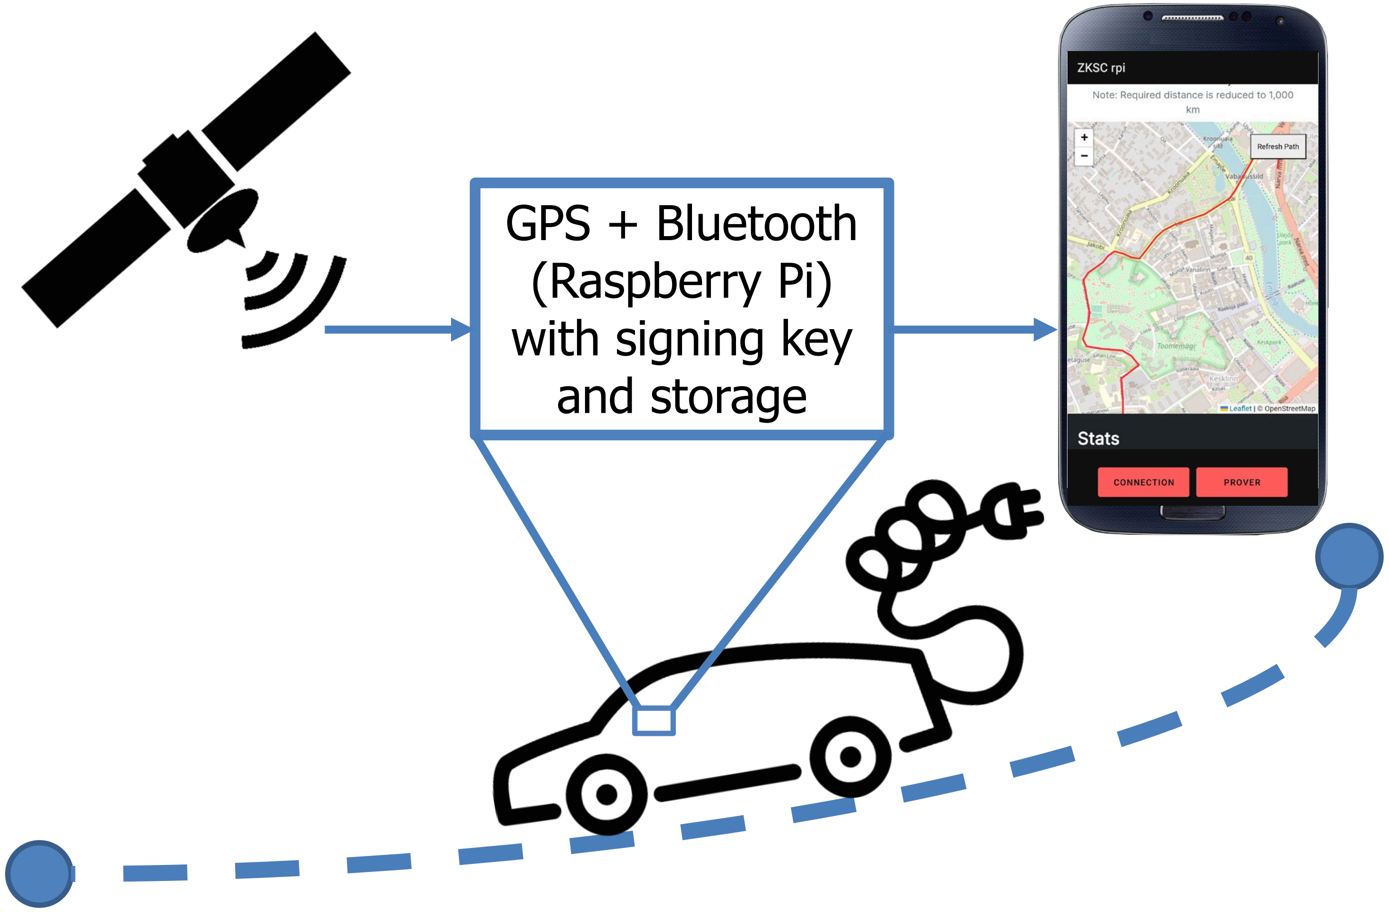

1. **Device** in vehicle stores GPS data
2. **Device** digitally signs GPS trail and sends to phone
3. **Phone** compiles a ZKP on mileage in a given country and the signature (terms of the EV subsidy!)
4. **Driver** preserves privacy of locations!

Proof constructed with Cybernetica's ZK-SecreC tools
https://arxiv.org/abs/2203.15448

ZK proofs for electric vehicle subsidy terms (2/2)

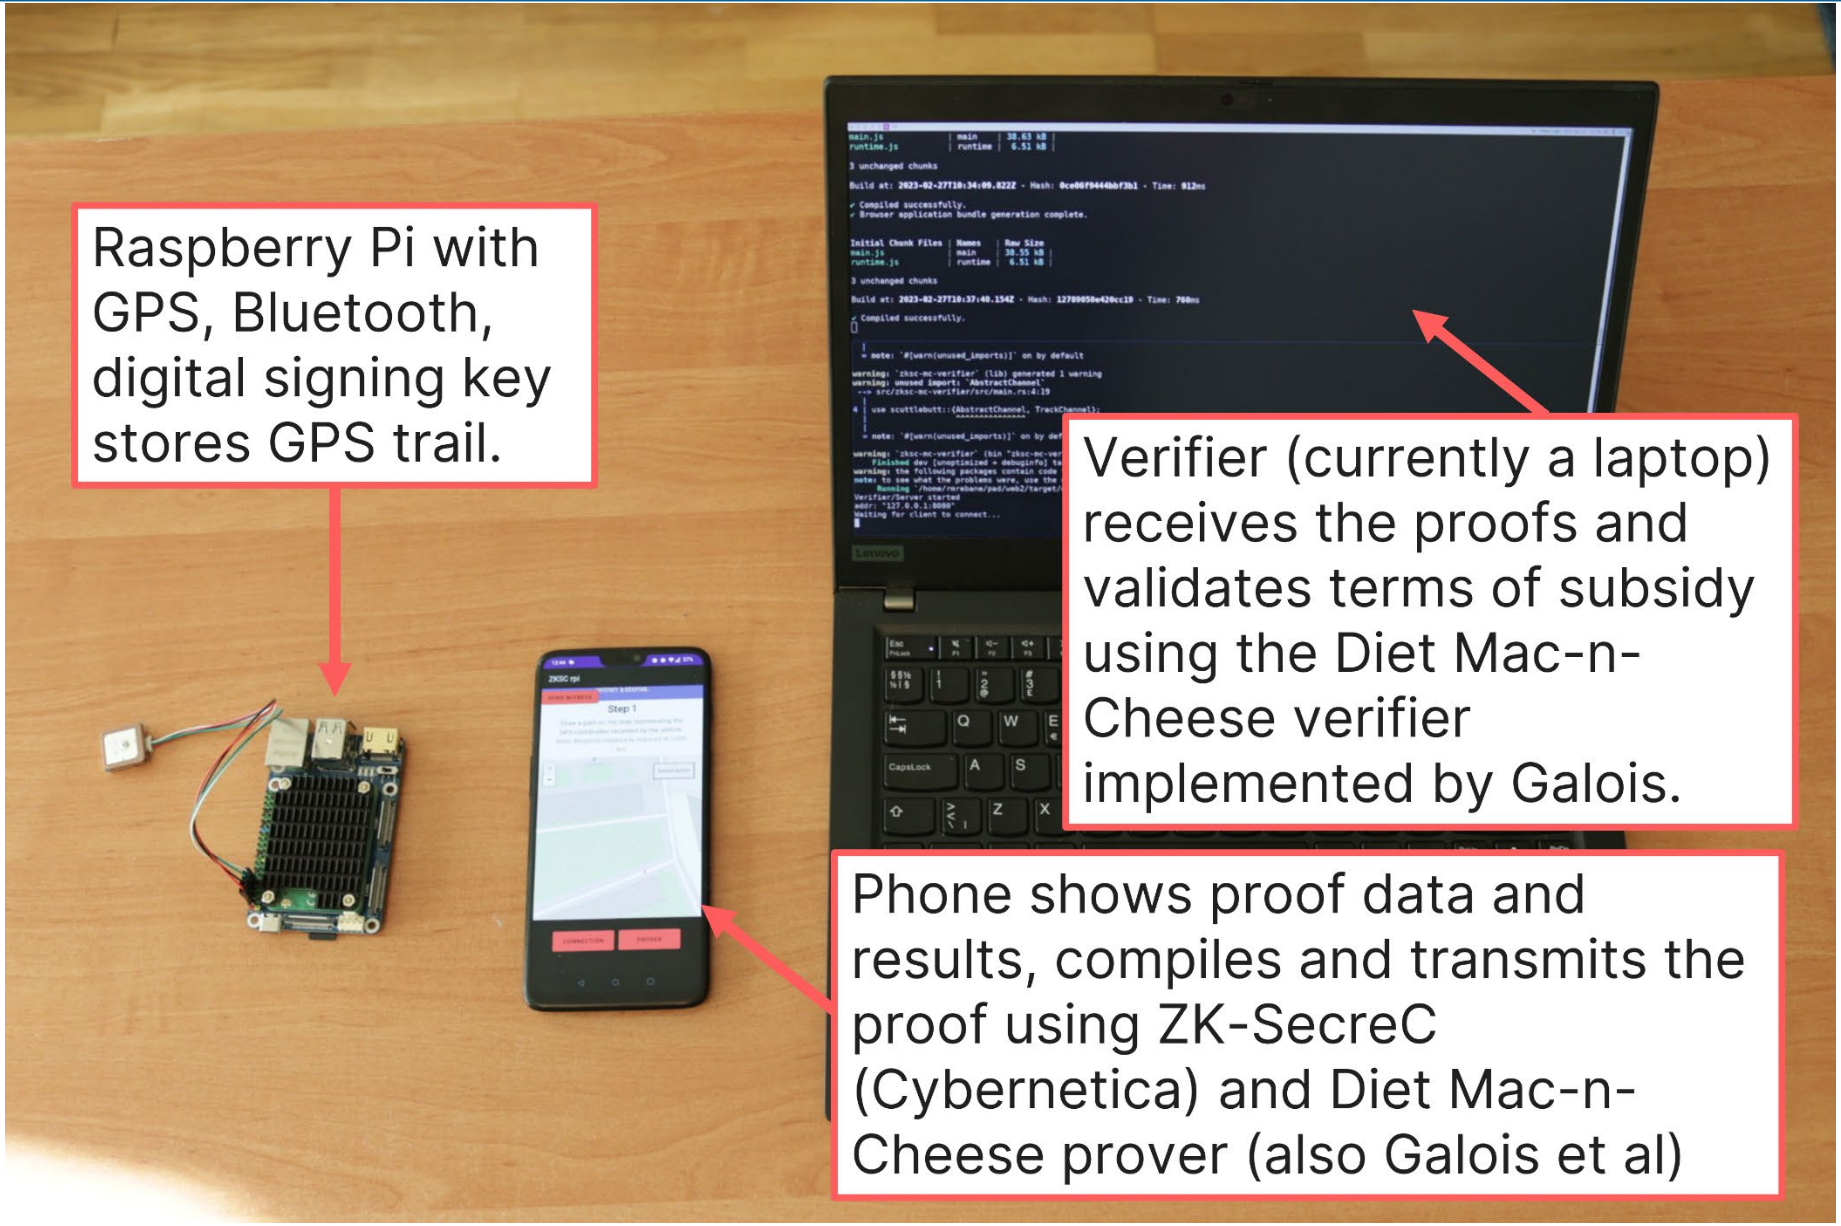

Resilient Anonymous Communciation for Everyone (RACE)

Resilient Anonymous Cryptography for Everyone (RACE)

Use cryptography and obfuscated communications to build an anonymous, attack-resilient mobile communication system that can reside completely within a country.

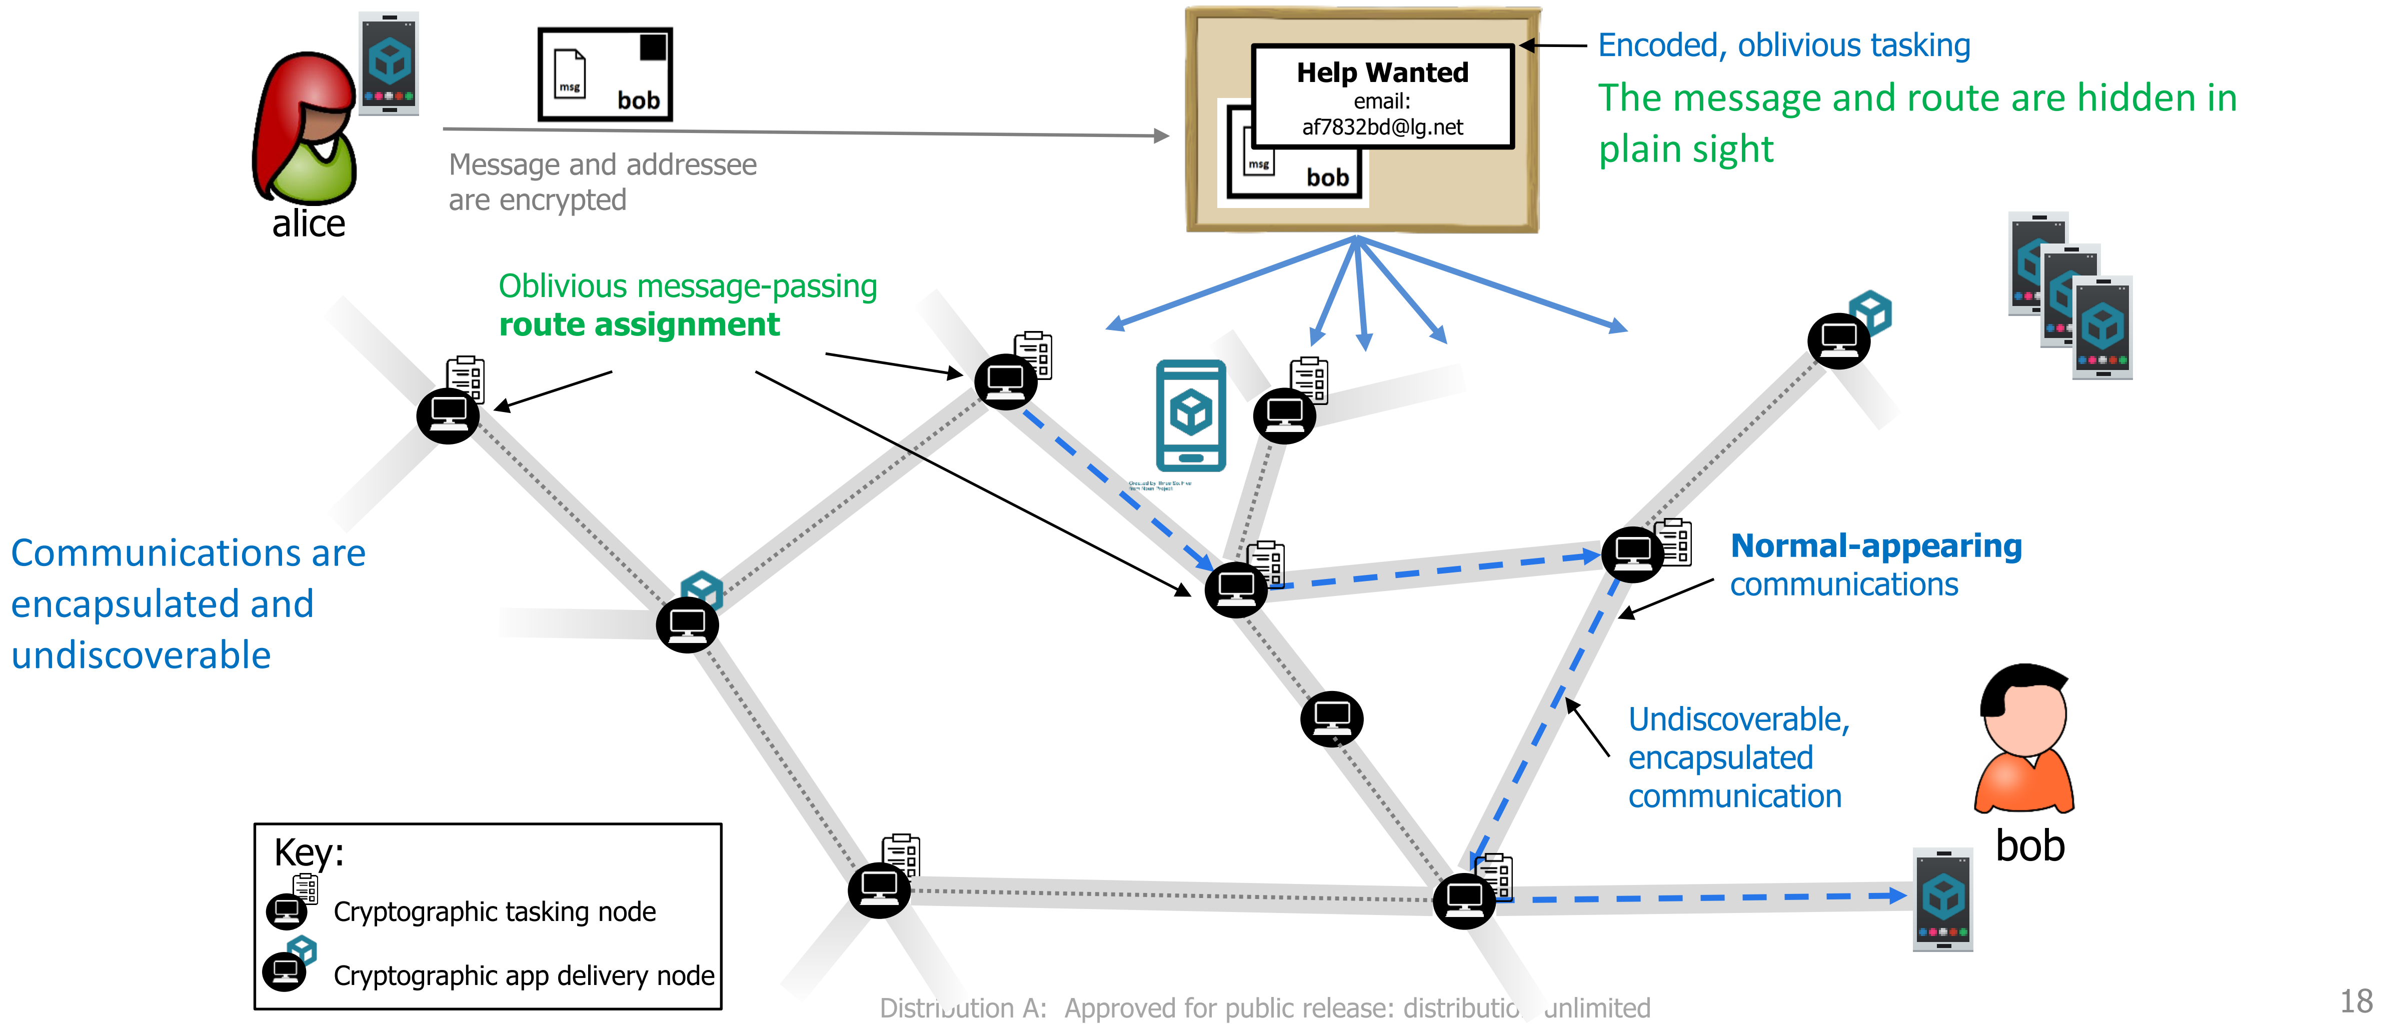

Technical Approach

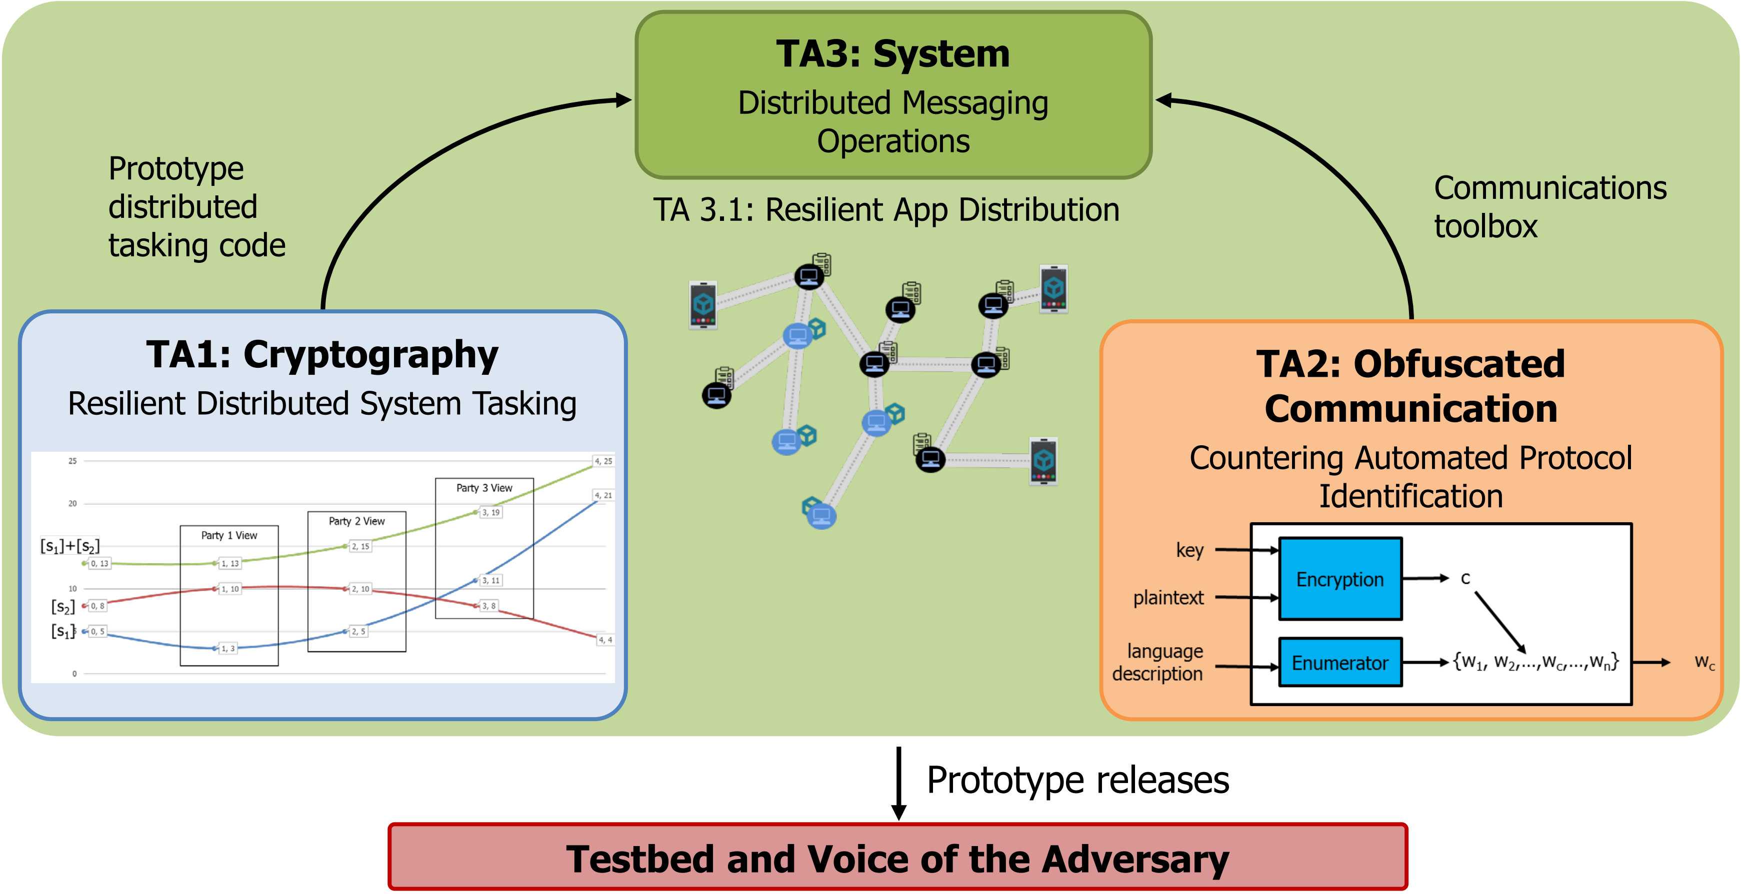

RACE transition partners: DoD,
State Department Democracy Human Rights and Labor (DRL), USAGM (Open Technology Fund)

RACE Accomplishments in Communication, 1/3

Mimicry

Embeds communications in ML generated protocols.

Georgetown RAVEN – A novel direct-service-use channel that mimics the statistics of real-world user emails to embed a channel into public email services.

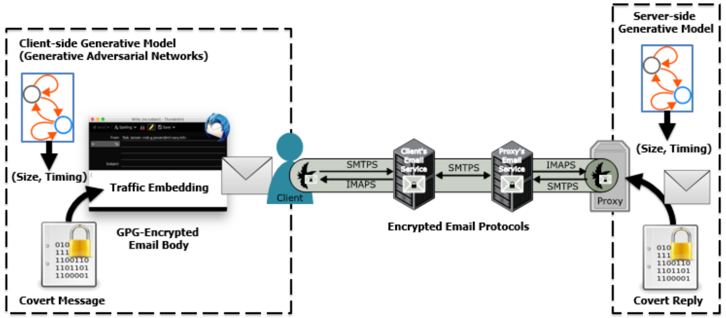

STR Semantic Steganography – Provides a novel framework to hide message in the content of the images and audio by embedding the hidden message by using it to generate synthetic images or audio.

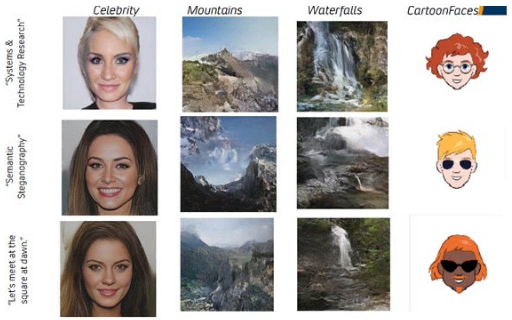

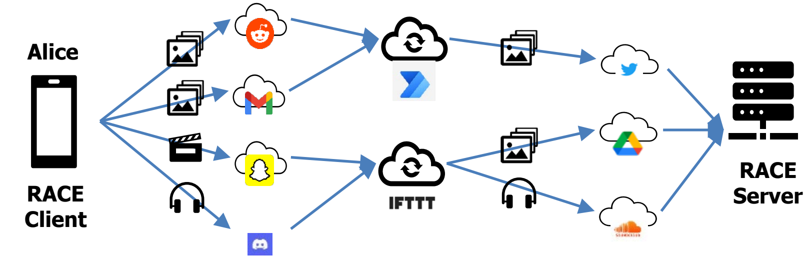

Galois Butkus – A novel communication channel that posts text into open web community forms using model-based Format-Transforming Encryption to encode hidden messages into the text of the posts or the structure of code or data files.

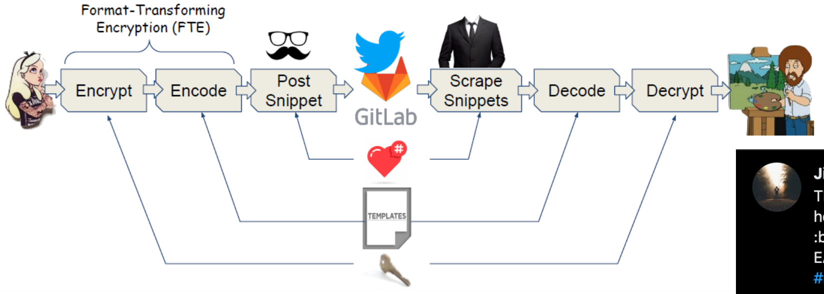

Source: RACE presentations

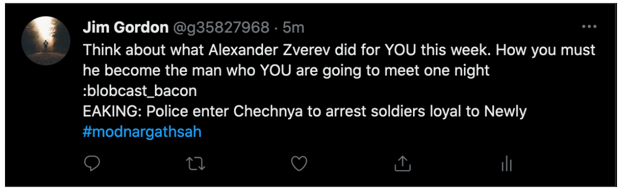

Source: Twitter

RACE Accomplishments in Communication, 2/3

Replacement / Hot Swap

Protocols that look the same on the wire as normal real-world connections

Galois Balboa – Embeds hidden data into protocols within existing secure channels such at web browsing or audio streaming while matching the statistics of a connection without the hidden channel. (USENIX 2021)

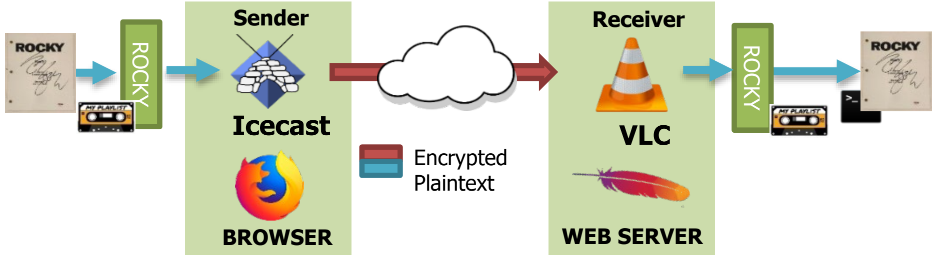

SRI MINECRAFT – A novel traffic substitution protocol that embeds a hidden channel into the actual Minecraft game play without modifying the game protocol itself. For any traffic produced by this channel, there exists a Minecraft game session that produces indistinguishable traffic.

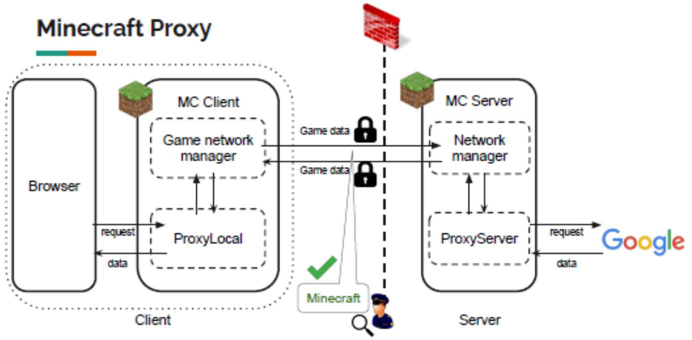

RACE Accomplishments in Communication, 3/3

Look-like-Nothing

Protocols when blocked cause so much collateral damage, they are prohibitive to censor

Georgetown OBFS – A look-like-nothing protocol that does not match a known protocol making it hard to develop a signature that catches it.

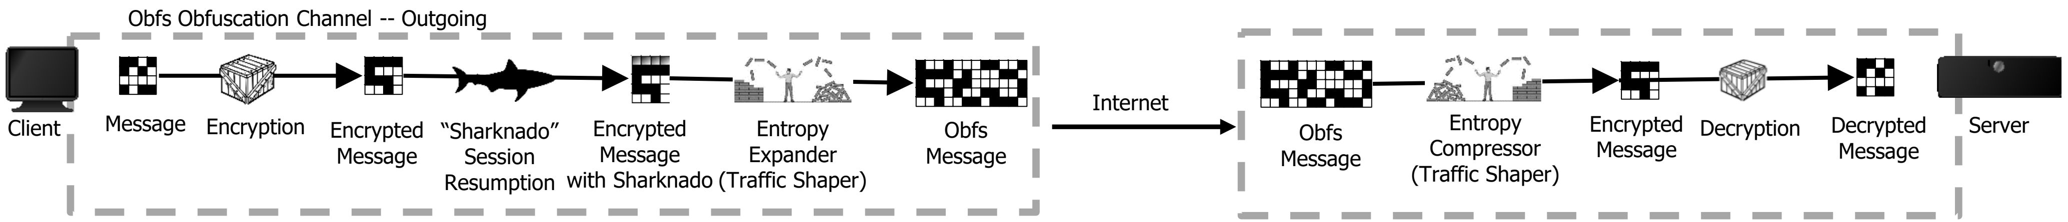

- TwoSix Technologies (TA3) is currently adapting low-rate channels over internet whiteboards developed under RACE for use as to replace domain fronting- initial application is for Tor.
  - Domain fronting is a technique in which a client conceals the true intended destination of an browser request from censors by "fronting" the request with a connection to a different domain. This is typically done with popular content delivery networks (CDN), who often require payment. Domain fronting organizations can also come under intense pressure from authoritarian governments.
  - Specifically, TwoSix is modifying Moat, which is an interactive tool used to get bridges from within a Tor Browser.

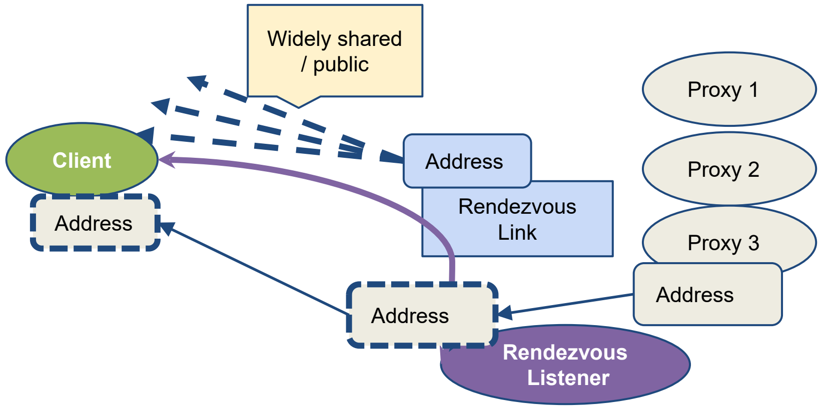

Kestrel- deployable, MPC-based messaging service

Kestrel Client App
- Lightweight: 10MB
- Easy to deploy: self-contained packages
- Behave like any chat app: WhatsApp, Messenger, Signal, ...

Kestrel Server App
- Lightweight: 10MB
- Heterogeneous: Platform-agnostic: mobile, IoT, raspberry pi, sensors
- Act as relays: Forward messages between them (and ultimately to Clients)
- Standalone: e.g. does not require Docker, external libraries, etc.

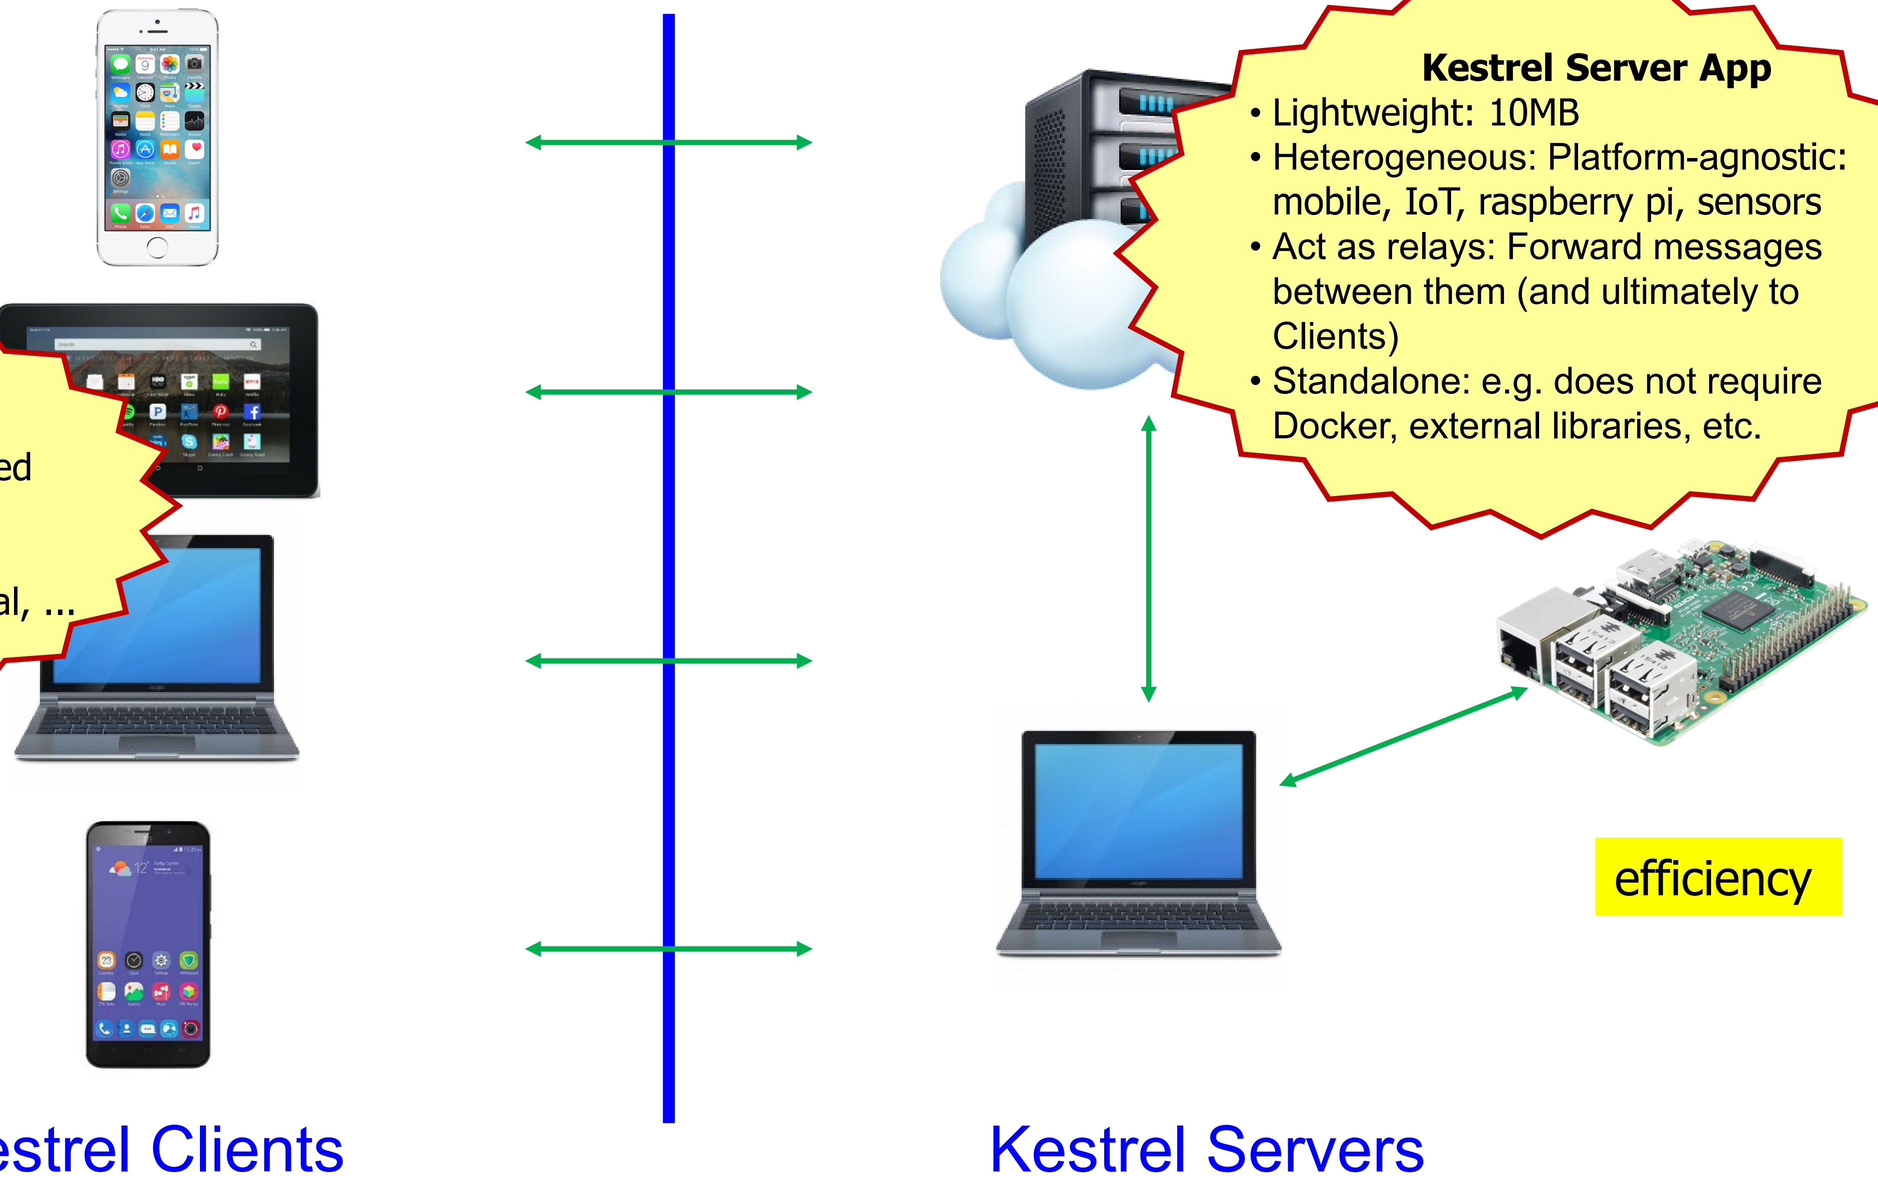

RACE Lessons Learned

- Distributed protocols are hard to engineer… especially with weird crypto and weird transports
- Channel security models really need to include user models
  - For more theory though, see Howes et al, “Security Foundations for Application-Based Covert Communication Channels”, IEEE Security & Privacy 2022
- Testing is HARD → moving to real-world test
- But, RACE really works!***
  - At end of Phase 2, successfully demonstrated a fully integrated system across all performers that meets end-to-end system phase 2 metrics (100 clients x 1000s, 61s and 118,557 messages/day).

Measuring the Information Control Environment (MICE)

Measuring the Information Control Environment (MICE)

- Idea: We have a number of open tools that measure censorship: The Open Observatory of Network Interference (OONI), Censored Planet, etc. Let's throw machine learning at the problem to make the analysis better and more scalable.
- **Censored Planet:** Used AI to clean up sensor data and enable more scalable analysis

  See also Wu et al, "TSPU: Russia's Decentralized Censorship System", IMC 2022
- **Psiphon:** Used AI to create real-time anomaly detection system for country-scale censorship detection
- **Thresher:** Build a predictive technique for what social media posts are likely to be censored in China
- **TwoSix Technologies/UMass:** Be on the lookout for Wu et al, "How the Great Firewall of China Detects and Blocks Fully Encrypted Traffic", USENIX 2023

What could be next?

1. Privacy-enhancing technology (PET) engineering for muggles

2. Outracing censors using AI

3. Clippy for Privacy
   - Browser version
   - Mobile phone version

www.darpa.mil