


Proactively Accountable Anonymous Messaging in Verdict

Henry Corrigan-Gibbs, David Isaac Wolinsky, and Bryan Ford

Department of Computer Science
Yale University

22nd USENIX Security Symposium
14 August 2013

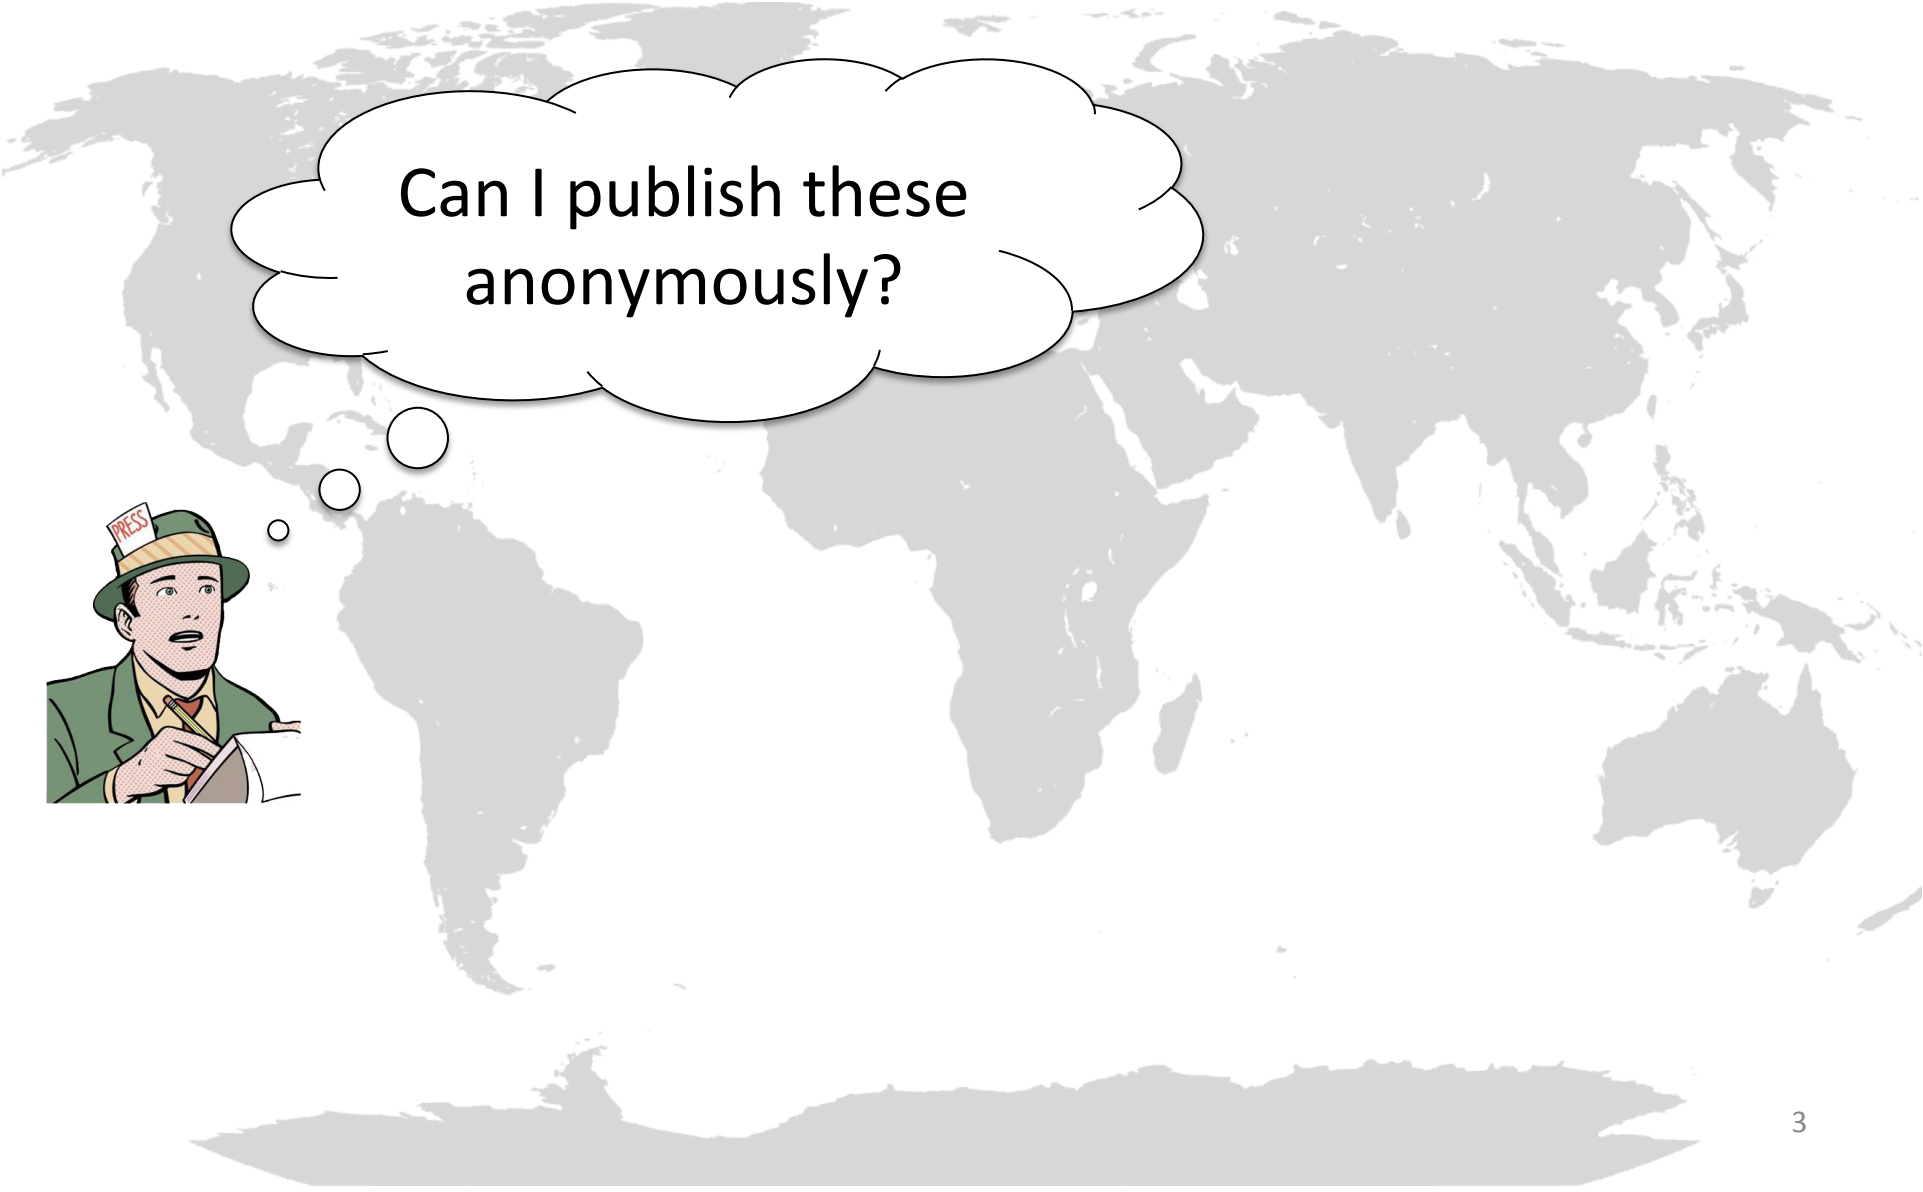


Acct #35139387
Acct #09836271



On the eve of an election in country X...
activist learns that the prime
minister is stashing stolen money
in a secret bank account.

**MUST PUBLISH this info
before the election**



Can I publish these
anonymously?




Can I publish these
anonymously?




Blog
Server



Possible Solution #1: Onion Routing



Possible Solution #1: Onion Routing




**Dining Cryptographers
networks (DC-nets)** are
resistant to traffic
analysis attacks



**Blog
Server**



Possible Solution #2: DC-nets




Possible Solution #2: DC-nets




DC-nets in 30 Seconds

Implement an
anonymous group
broadcast primitive




David Chaum
“Dining Cryptographers Problem”
[J. Cryptography ‘88]


DC-nets in 30 Seconds



DC-nets in 30 Seconds




DC-nets in 30 Seconds



DC-nets are resistant to traffic analysis attacks

Primarily use fast symmetric-key crypto operations (PRNG, XOR)


DC-nets in 30 Seconds



Dissent: DC-nets made practical

- Splits nodes into clients and servers
- Scales to 1000s of nodes
- Handles client churn
- Anonymity set size = set of honest nodes

Possible Solution #2: DC-nets



Blog
Server





Possible Solution #2: DC-nets



Blog
Server









Possible Solution #3: Dissent

- Dissent can handle this sort of misbehavior
 - After a disruption occurs, participants run a **shuffle/e-voting protocol**
 - The anonymous sender sends an **accusation** through the shuffle
 - All nodes use the accusation to trace (“**blame**”) the disruptor


Possible Solution #3: Dissent




Possible Solution #3: Dissent




Possible Solution #3: Dissent




Possible Solution #3: Dissent




Possible Solution #3: Dissent



Possible Solution #3: Dissent



Possible Solution #3: Dissent



Verdict: Motivation

- Can we get
 - the traffic-analysis-resistance of DC-nets and
 - the scalability of Dissentwith **lower blame cost**?
- Idea: Group members **prove** that their messages are sending are correctly formed.
 - Identify disruptors **before** they jam the anonymous communication channel

“Verifiable” DC-nets

- In 2004 Eurocrypt paper, Golle and Juels propose applying zero-knowledge proof (ZKP) techniques to DC-nets
- Participants **prove correctness** of messages
- Drawbacks of Golle-Juels work: computationally expensive, inefficient in communication cost, uses pairings, requires trusted setup, ...
- **Never implemented...**

Verdict: Contributions

1. First (to our knowledge) implementation and evaluation of verifiable DC-nets
2. Two new verifiable DC-nets constructions which give 5.6x speedup over Golle-Juels approach
3. Optimizations to make verifiable DC-nets fast
 - for long messages,
 - when there are no *active* disruptors, and
 - by optimistically using XOR-based DC-nets when possible (**138x speedup**)

Outline

- Background and Motivation
- **Verdict**
 - **Design Challenges**
 - Optimizations
- Evaluation
- Conclusion

Design Challenges

1. Resist traffic analysis attacks
2. Make sender's transmission indistinguishable
3. Prove that transmissions are well-formed

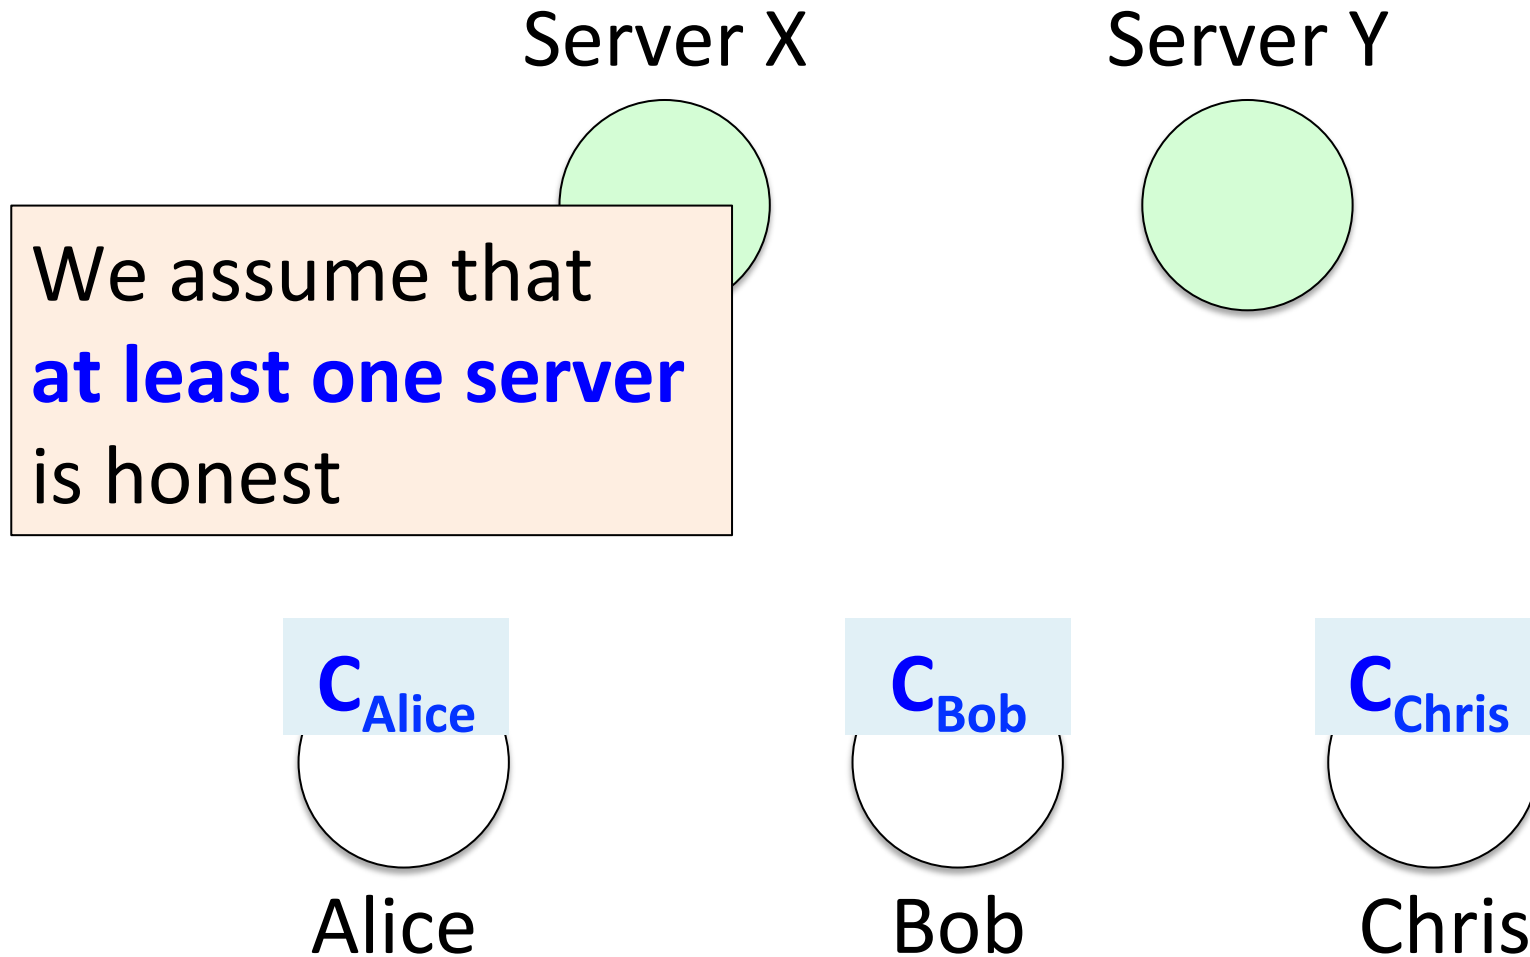
Design Challenges

- 1. Resist traffic analysis attacks**
2. Make sender's transmission indistinguishable
3. Prove that transmissions are well-formed


Challenge 1: Traffic Analysis Resistance

- Time is divided into **messaging rounds**
- One anonymous sender per messaging round
- Every client transmits the **same number of bits** in every messaging round
 - # of bits sent does not leak sender's identity
- Clients' ciphertexts are **cryptographically indistinguishable**
 - Content does not leak sender's identity


Challenge 1: Traffic Analysis Resistance



Challenge 1: Traffic Analysis Resistance



Challenge 1: Traffic Analysis Resistance



Design Challenges

1. Resist traffic analysis attacks
- 2. Make sender's transmission indistinguishable**
3. Prove that transmissions are well-formed


Challenge 2: Encoding Messages

- The transmitting client sends an encryption of arbitrary message: m
- Non-transmitting clients set $m = 1$
 - An encryption of the identity element
- Use an ElGamal-like scheme to encrypt


$$E(m, \sigma_1, \dots, \sigma_N) = mg^{\sigma_1 + \dots + \sigma_N}$$

... where the σ s are secrets shared between clients and servers.

Challenge 2: Encoding Messages




Clients and servers agree to k -bit shared secrets σ using DH exchange




$$C_{\text{Alice}} = 1 * g_t^{\sigma_{ax} + \sigma_{ay}}$$


g_t generates some group in which DDH is hard



$$C_{\text{Chris}} = m * g_t^{\sigma_{cx} + \sigma_{cy}}$$



$$C_X = g_t - \sigma_{ax} - \sigma_{bx} - \sigma_{cx}$$



Challenge 2: Encoding Messages


- In product of C s, every secret σ_{ij} is included as an exponent once with (+) sign and once with (-) sign:

$$C_{\text{Alice}} C_{\text{Bob}} C_{\text{Chris}} C_X C_Y = m$$

$$C_{\text{Alice}} = 1 * g_t^{\sigma_{ax} + \sigma_{ay}}$$

$$C_{\text{Chris}} = m * g_t^{\sigma_{cx} + \sigma_{cy}}$$

Without knowing the secrets σ , an attacker cannot tell whether Alice or Chris is the anonymous sender of m (by DDH assumption)



Design Challenges


1. Resist traffic analysis attacks
2. Make sender's transmission indistinguishable
- 3. Prove that transmissions are well-formed**

Challenge 3: Proving Correctness

- Clients attach non-interactive zero-knowledge proofs of knowledge to their ciphertexts
 - Use off-the-shelf ZKP techniques
Camenisch-Stadler [ETH Zurich TR-260, '97]
 - Servers check proofs before accepting client ciphertexts
- Servers prove validity of their ciphertexts too

Challenge 3: Proving Correctness

- Recall: one client transmits in each messaging round
- As in Dissent, we use a **key shuffle** to assign pseudonymous “owners” to messaging rounds
 - Each client submits a pseudonym public key to shuffle
 - Shuffle hides owner-to-pseudonym mapping



“A Verifiable Secret
Shuffle and its
Application to E-Voting”
– Neff [CCS '01]

Challenge 3: Proving Correctness

- When Bob submits a ciphertext in a messaging round owned by pseudonym g^d , Bob attaches a proof that:


PoK { C_{Bob} is a correct encryption of $m=1$
OR
I know the pseudonym secret d }

Challenge 3: Proving Correctness

- Clients and servers publish **commitments** to their shared secrets σ_{ij}

$$S_{ax} = \text{Commit}(\sigma_{ax}) = h^{\sigma_{ax}}$$

...using some generator h of group G for which no one knows $\log_{g_t}(h)$.



Challenge 3: Proving Correctness

- When Bob submits a ciphertext in a messaging round owned by pseudonym g^d , Bob attaches a proof that:

PoK { C_{Bob} is a correct encryption of $m=1$
OR
I know the pseudonym secret d }

Challenge 3: Proving Correctness

- When Bob submits a ciphertext in a messaging round owned by pseudonym g^d , Bob attaches a proof that:

$$\text{PoK} \left\{ \begin{array}{l} \sigma_{bx} \\ \sigma_{by} \\ d \end{array} : \left(\text{AND} \begin{array}{l} C_{\text{Bob}} = g_t^{\sigma_{bx} + \sigma_{by}} \\ S_{bx} S_{by} = h^{\sigma_{bx} + \sigma_{by}} \end{array} \right) \text{OR} \begin{array}{l} D = g^d \end{array} \right\}$$

Challenge 3: Proving Correctness

- When Bob submits a ciphertext in a messaging round owned by pseudonym g^d , Bob attaches a proof that:

$$\text{PoK} \left\{ \begin{array}{l} \sigma_{bx} \\ \sigma_{by} \\ d \end{array} : \left(\text{AND} \begin{array}{l} C_{\text{Bob}} = g_t^{\sigma_{bx} + \sigma_{by}} \\ S_{bx} S_{by} = h^{\sigma_{bx} + \sigma_{by}} \end{array} \right) \text{OR} \begin{array}{l} D = g^d \end{array} \right\}$$

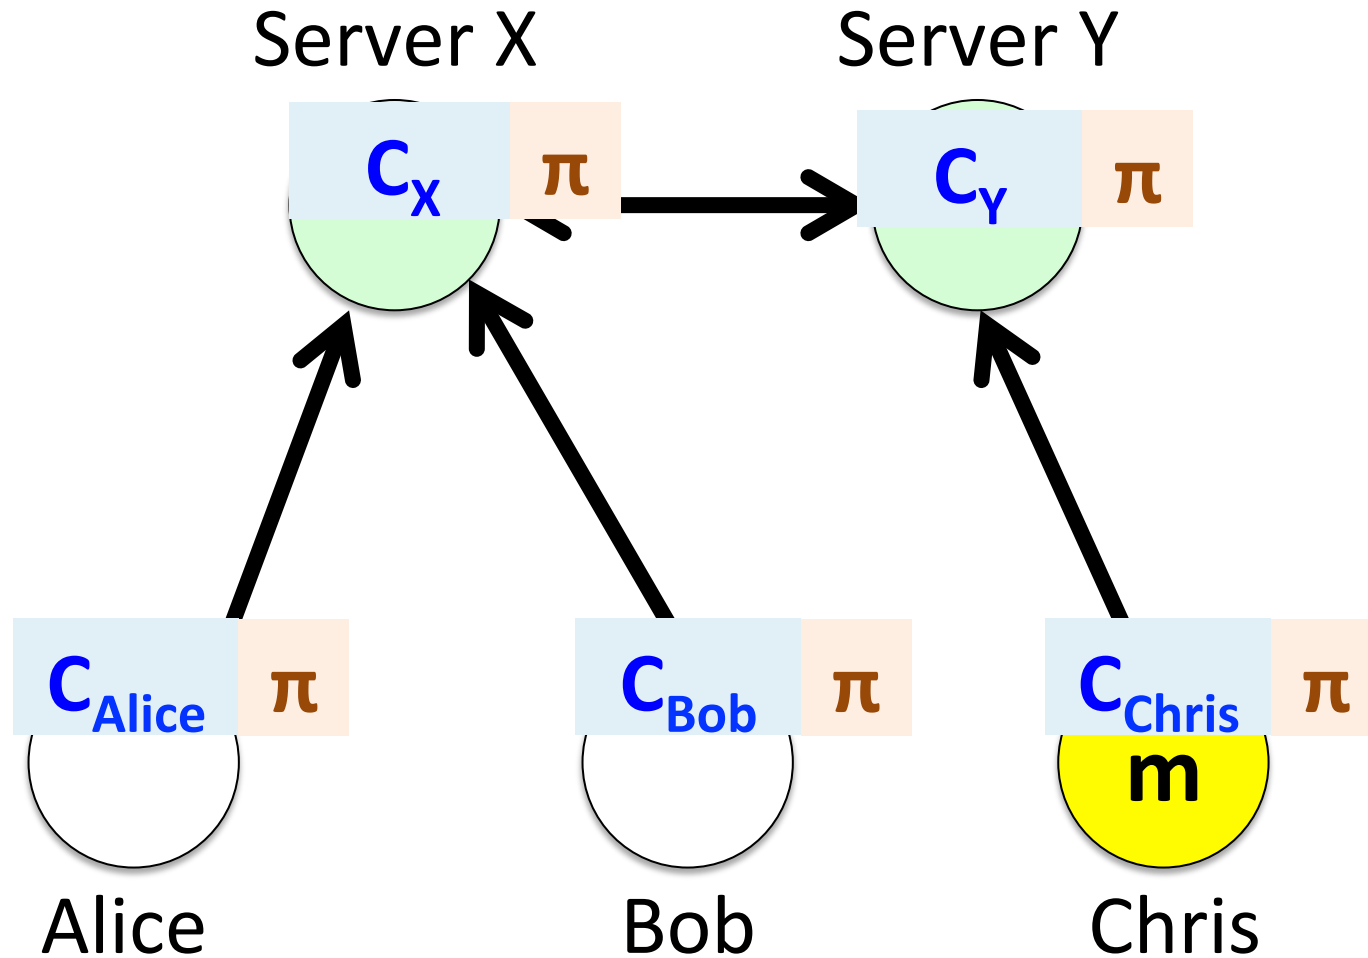
Challenge 3: Proving Correctness

- When Bob submits a ciphertext in a messaging round owned by pseudonym **g** a proof that:


Proof is just a bit-string
—can send it over the
network, etc

$$\text{PoK} \left\{ \begin{array}{l} \sigma_{bx} \\ \sigma_{by} \\ d \end{array} : \left(\begin{array}{l} \text{AND} \\ C_{\text{Bob}} = g_t^{\sigma_{bx} + \sigma_{by}} \\ S_{bx} S_{by} = h^{\sigma_{bx} + \sigma_{by}} \end{array} \right) \text{OR} \right. \\ \left. D = g^d \right\}$$

Challenge 3: Proving Correctness



Challenge 3: Proving Correctness



Outline

- Background and Motivation
- **Verdict**
 - Design Challenges
 - **Optimizations**
- Evaluation
- Conclusion

Optimizations

1. Long messages
2. “Lazy” proof verification
3. Hybrid Dissent+Verdict DC-net

Optimizations

1. **Long messages**
2. “Lazy” proof verification
3. Hybrid Dissent+Verdict DC-net

Long Messages

- We use a public hash function to pick these generators

Golle-Juels scheme required a pairing to achieve a similar effect

Messages

$$C_{A,1} = g_1^{\sigma_{ax} + \sigma_{ay}}$$

$$C_{A,2} = g_2^{\sigma_{ax} + \sigma_{ay}}$$

$$\dots = \dots$$

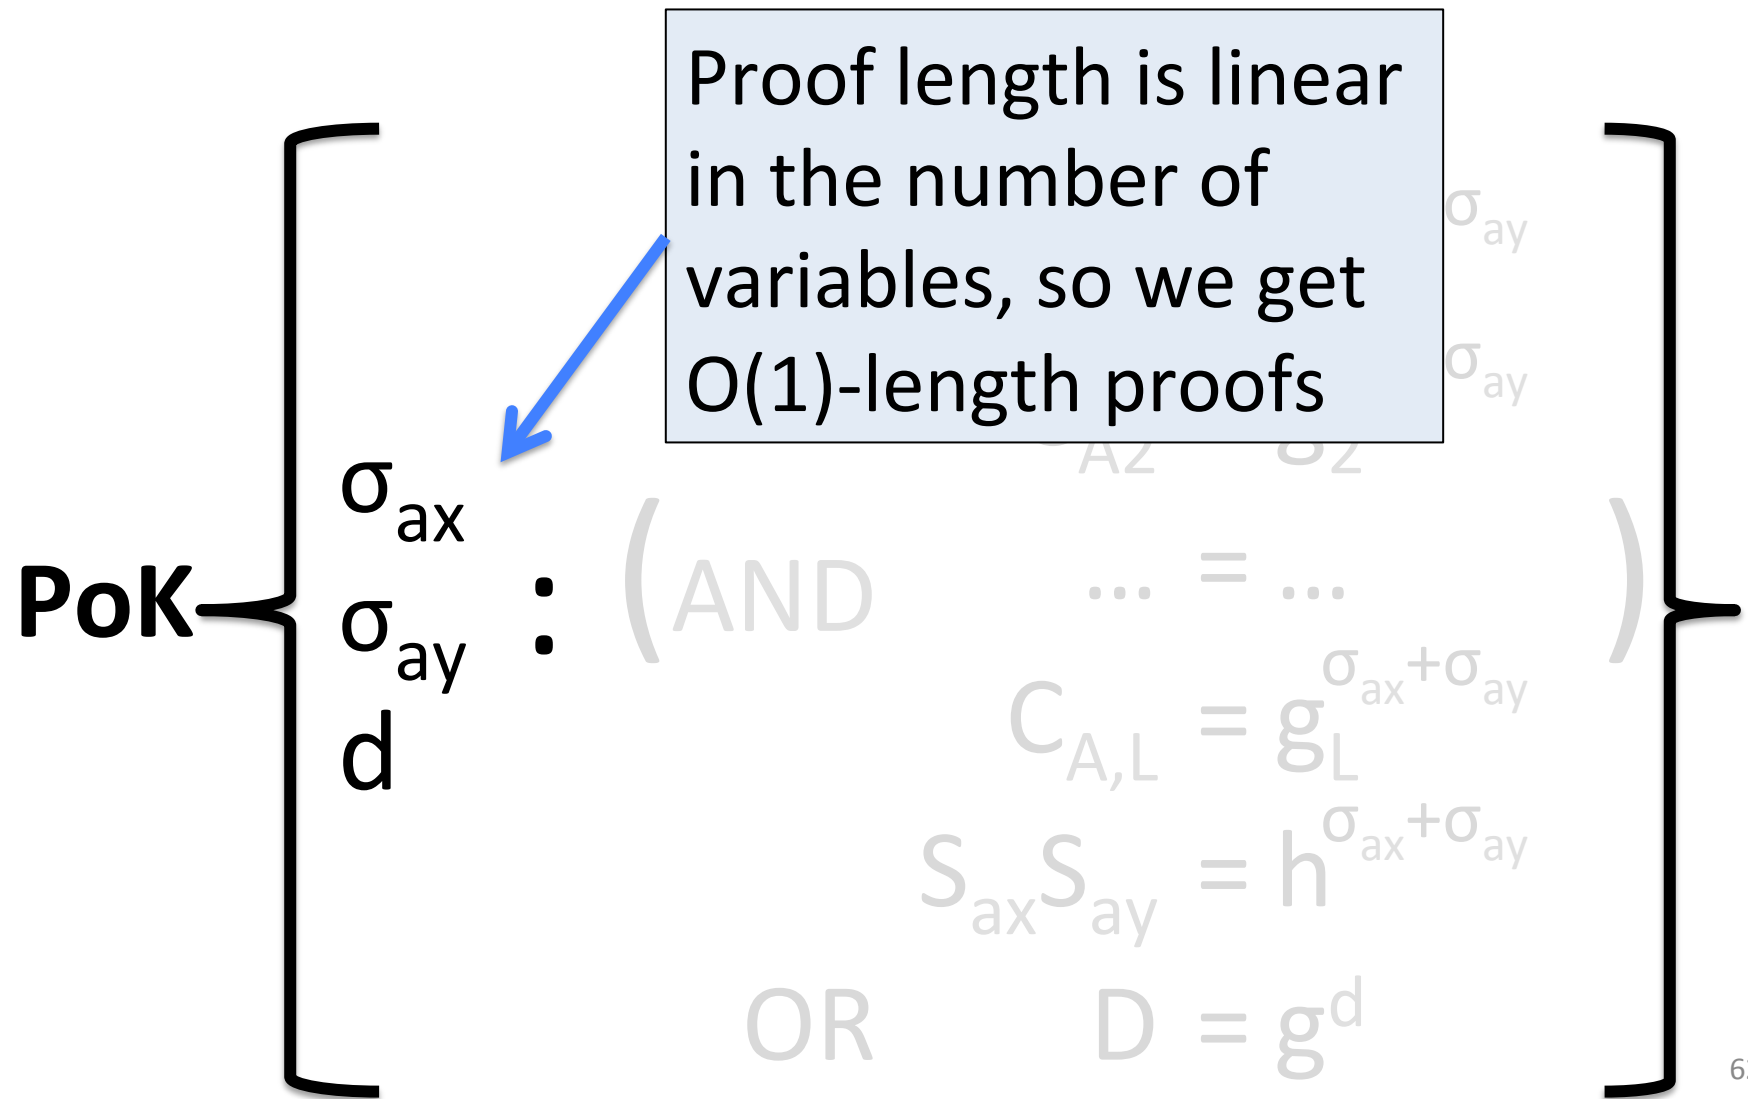
$$C_{A,L} = g_L^{\sigma_{ax} + \sigma_{ay}}$$

$$S_{ax} S_{ay} = h^{\sigma_{ax} + \sigma_{ay}}$$

OR

$$D = g^d$$

Optimization 1: Long Messages




Optimizations

1. Long messages
2. **“Lazy” proof verification**
3. Hybrid Dissent+Verdict DC-net

Optimization 2: “Lazy” Verification

- Checking proofs is expensive
- Servers defer checking proofs until **after a disruption occurs**



- Anonymous sender signs content with **pseudonym secret key**
- If sig check fails, servers know that disruption has occurred—then they check proofs

Optimizations

1. Long messages
2. “Lazy” proof verification
3. **Hybrid Dissent+Verdict DC-net**

Optimization 3: Hybrid DC-net


Verdict: heavy pub key crypto

Dissent/DC-nets: AES + XORs

- Recall: After a disruption in Dissent, the anonymous sender broadcasts an “accusation” using a **verifiable shuffle protocol**
 - Participants use the accusation to trace the disruptor
 - Over **99% of the “blame” process is spent in shuffle**
- **Idea**: Use Verdict to broadcast Dissent’s anonymous accusations → *hybrid* DC-net


Optimization 3: Hybrid DC-net

- Participants set up parallel Dissent and Verdict communication sessions




Optimization 3: Hybrid DC-net

- Participants set up parallel Dissent and Verdict communication sessions




Optimization 3: Hybrid DC-net

- Participants set up parallel Dissent and Verdict communication sessions



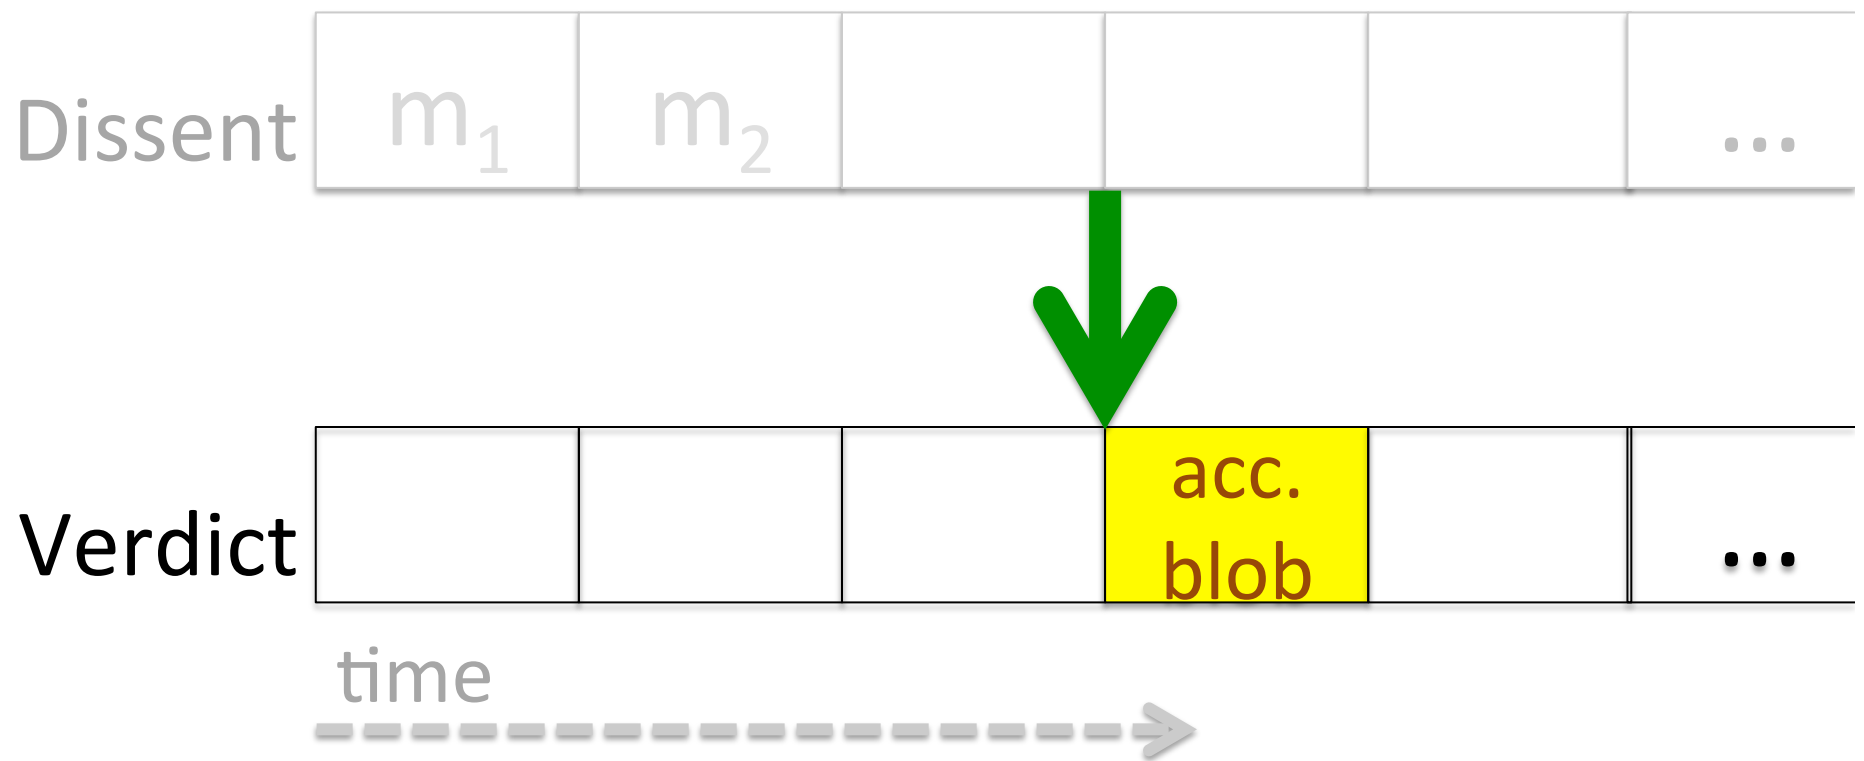
Optimization 3: Hybrid DC-net

- Participants set up parallel Dissent and Verdict communication sessions




Optimization 3: Hybrid DC-net

- Participants set up parallel Dissent and Verdict communication sessions




Optimization 3: Hybrid DC-net

- Participants set up parallel Dissent and Verdict communication sessions




Optimization 3: Hybrid DC-net

- Participants set up parallel Dissent and Verdict communication sessions



Optimization 3: Hybrid DC-net

- Participants set up parallel Dissent and Verdict communication sessions




Optimization 3: Hybrid DC-net

- Participants set up parallel Dissent and Verdict communication sessions

Normal case: Dissent XOR-based DC-net
Under disruption: Verdict (faster than shuffle)

Verdict



time




Outline

- Background and Motivation
- Verdict
 - Design Challenges
 - Optimizations
- **Evaluation**
- Conclusion

Implementation


- Implemented in C++ as an extension to Dissent
- Cryptographic primitives
 - OpenSSL, Crypto++, and Botan libraries
 - 256-bit NIST elliptic curve group
- Used the DeterLab testbed
 - Physical nodes: 8 servers, 128 clients
 - Ran many client processes per machine to simulate up to 1024 clients
- Source code: <https://github.com/DeDis/Dissent>

Encryption Throughput (CPU Cost)




Messaging Latency

Tweet-length messages



Blame Time



Outline

- Background and Motivation
- Verdict
 - Design Challenges
 - Optimizations
- Evaluation
- **Conclusion**

Details in the Paper

- Messaging protocol
 - Handling equivocation, dropped messages, etc.
- Proof constructions
 - The paper describes three variants
 - Implementation details
- Handling server failure
- Handling client churn

Conclusion

First practical verifiable DC-nets scheme

- Introduces two new verifiable DC-nets constructions
- Reduces the cost of finding DC-net disruptors by **two orders of magnitude**
- By reducing the cost of disruption, Verdict brings strong **traffic-analysis-resistant anonymity** closer to practicality

Acknowledgements

Thanks to:

- the anonymous reviewers,
- our shepherd, Micah Sherr,
- the DeterLab staff,
- Aaron Johnson, Ewa Syta, Michael J. Fischer, Michael Z. Lee, Michael “Fitz” Nowlan, Ramki Gummadi, and
- all of you for listening.

<https://dedis.cs.yale.edu/2010/anon/>

Shameless plug: **The Dissent project is hiring!**

