Privacy Proofs for Anonymous Communication Networks

ProTeCS @ EuroCrypt 2025

Christoph Coijanovic, D. Schadt, T. Strufe | 03.05.25

KASTEL

Anonymous Communication

Privacy Proofs

What make proofs for anonymous communication networks interesting?

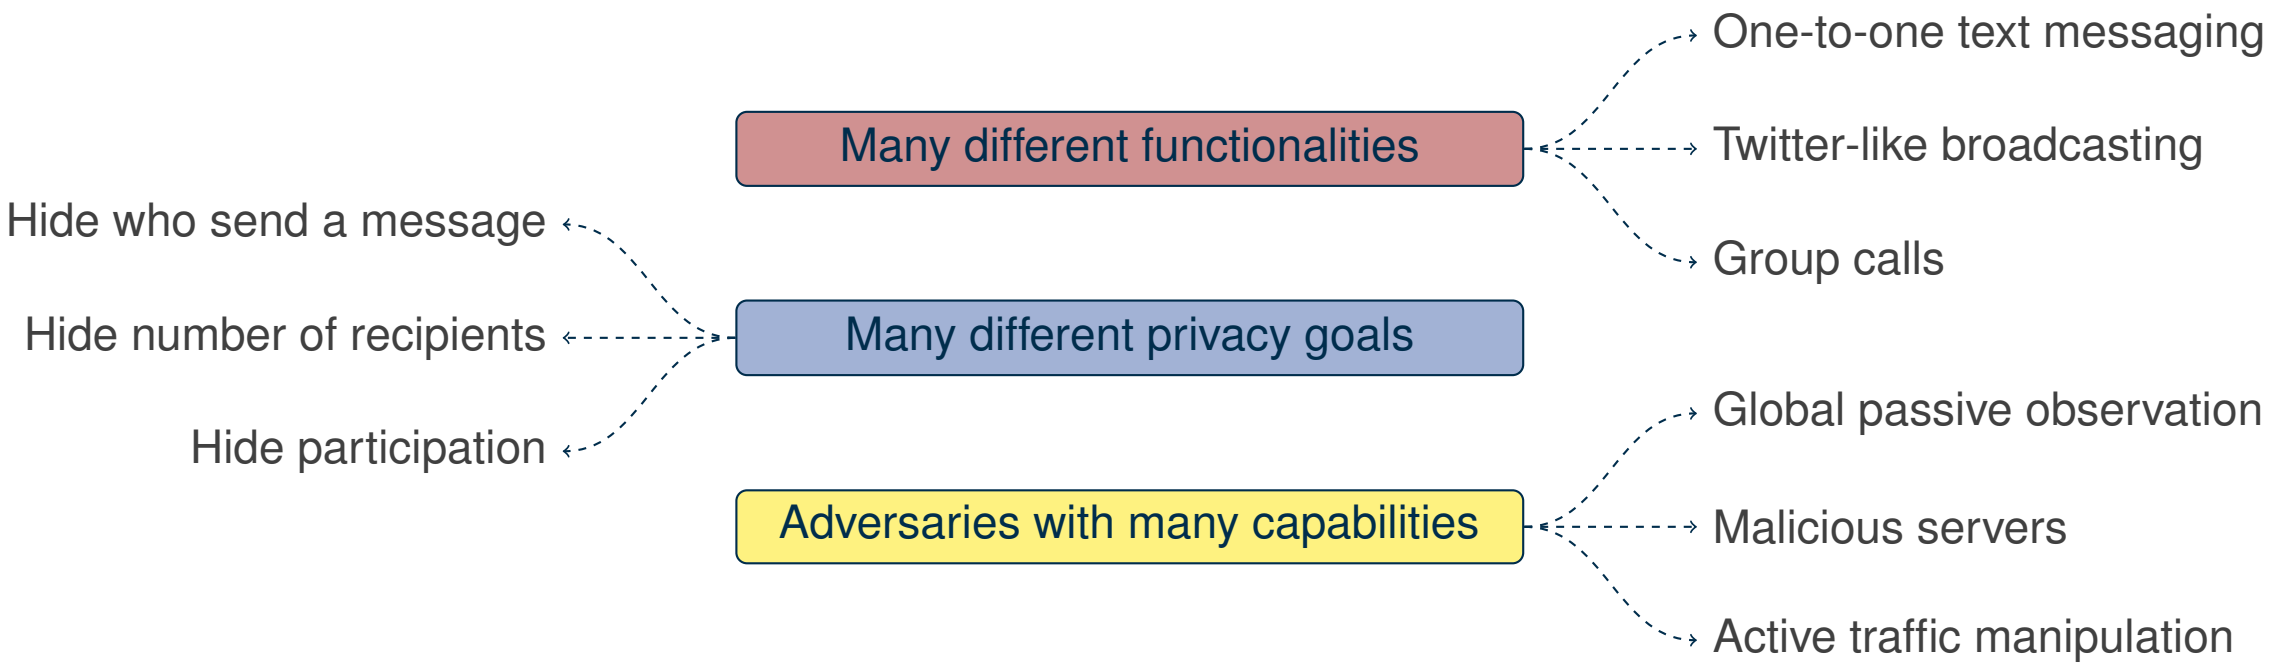

Roadmap

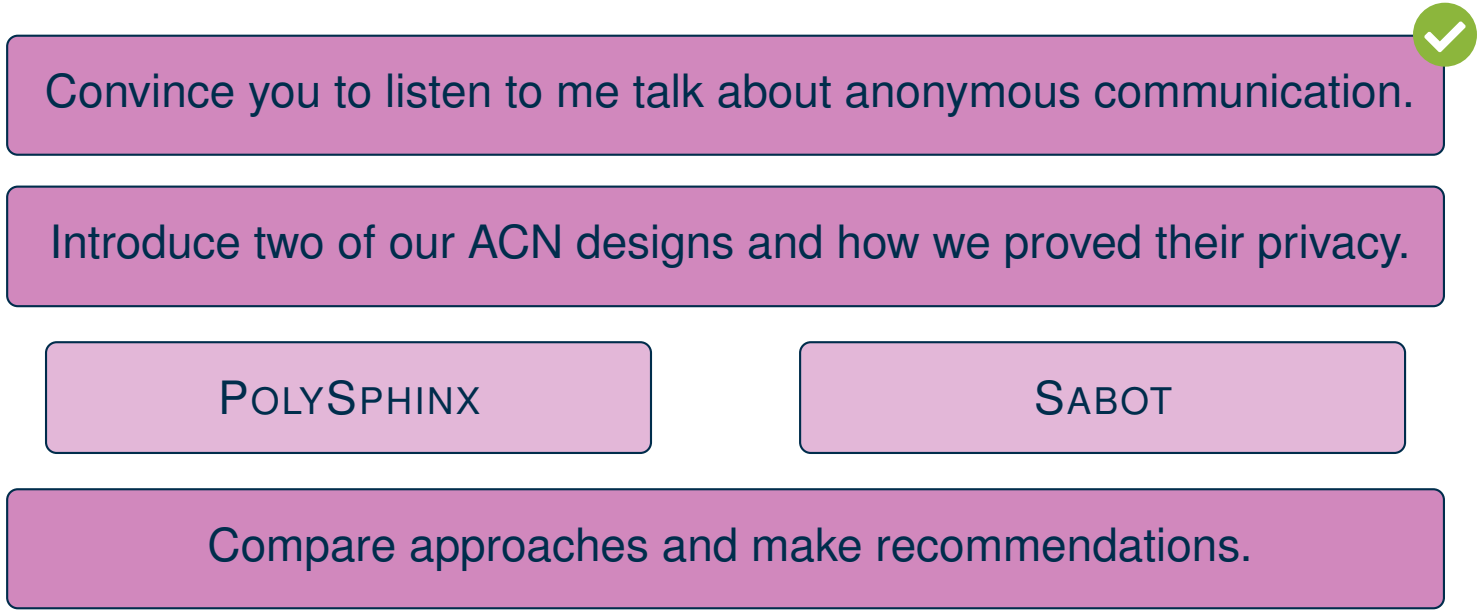
Convince you to listen to me talk about anonymous communication.
Introduce two of our ACN designs and how we proved their privacy.
PolySphinx
Sabot
Compare approaches and make recommendations.

PolySphinx

D. Schadt, C. Coijanovic, C. Weis and T. Strufe, “PolySphinx: Extending the Sphinx Mix Format With Better Multicast Support”, 2024 IEEE Symposium on Security and Privacy (SP)

PolySphinx

Functionality

Alice wants to send a message to both Bob and Carol.

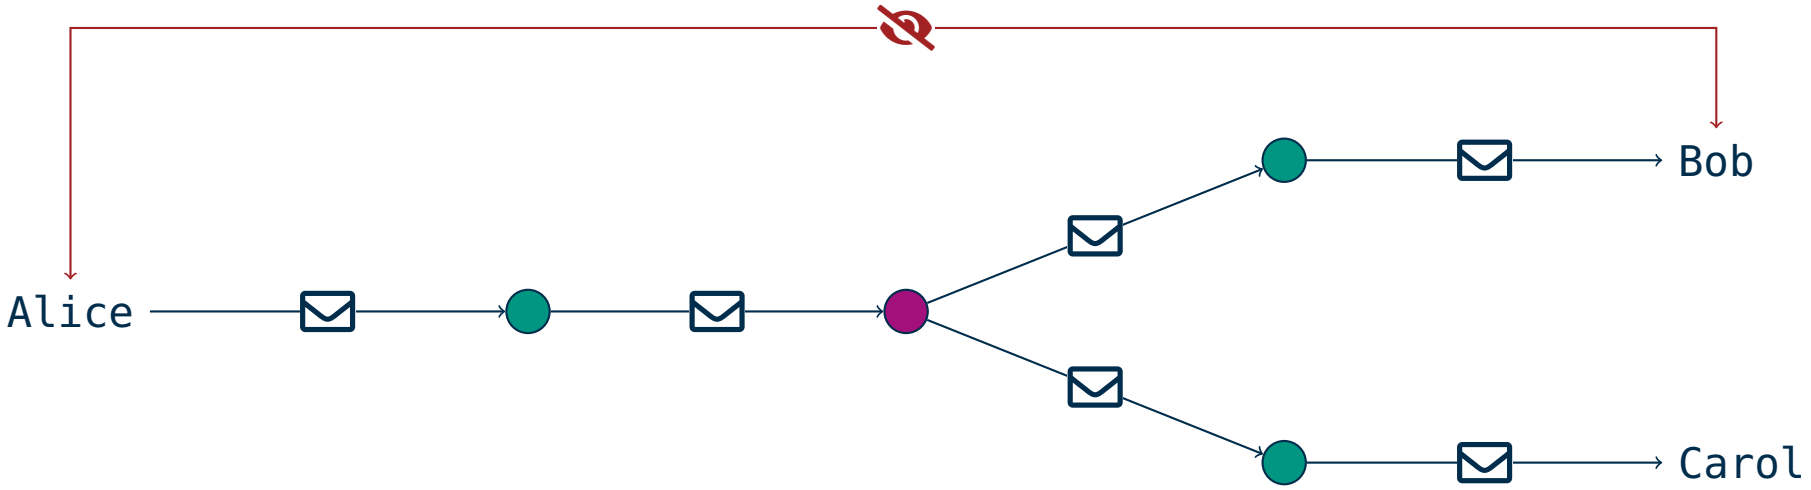

Privacy Goal: *Single Sender Anonymity*

Given at least one honest node between Alice and Bob, the adversary cannot tell that she sent to him.

PolySphinx

Observation

Points of attack to link Alice to Bob is quite limited!

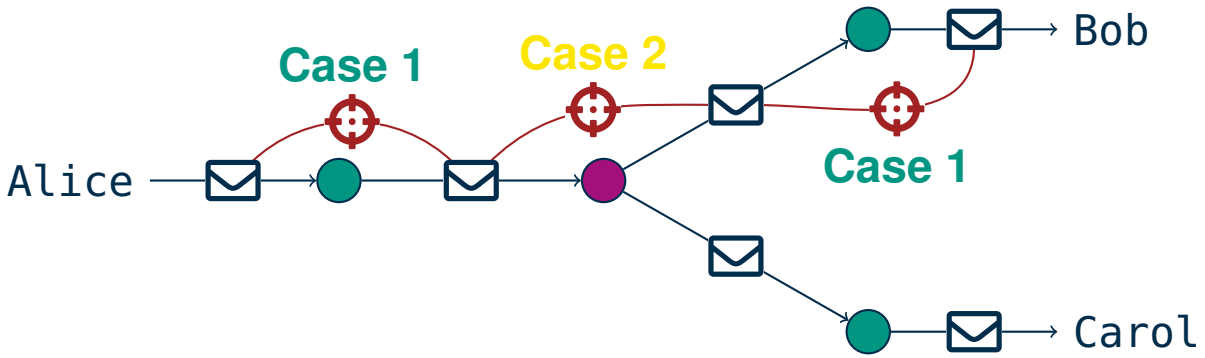

- **Case 1:** (Layer Unlinkability[a]) Prevents adversary from linking incoming packet to outgoing at „normal“ node.
- **Case 2:** Special case of Layer Unlinkability at replication node.

[a] Kuhn et al. „Breaking and (Partially) Fixing Provably Secure Onion Routing“, IEEE SP 2019

PolySphinx

Case 1: (Layer Unlinkability) Prevents adversary from linking incoming packet to outgoing at „normal“ node.

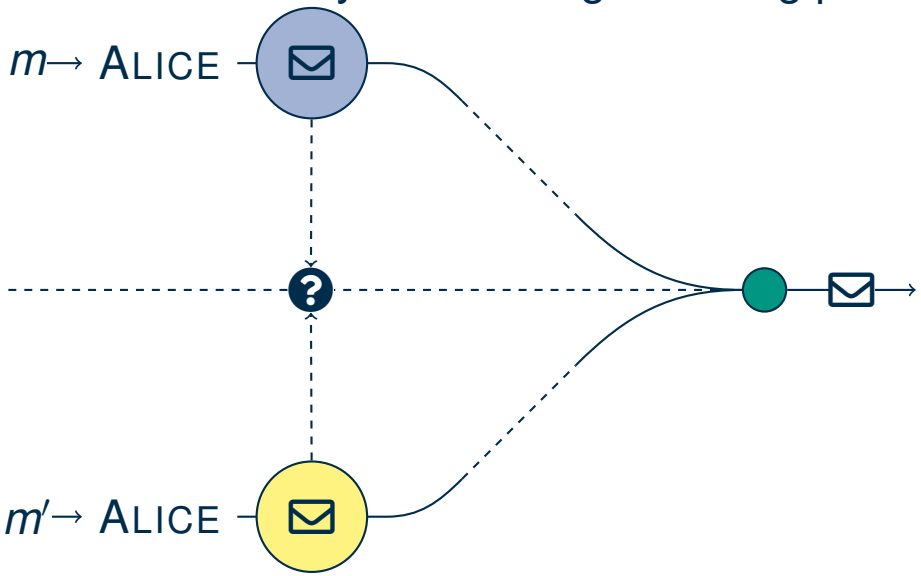

PolySphinx

Case 1: (Layer Unlinkability) Prevents adversary from linking incoming packet to outgoing at „normal“ node.

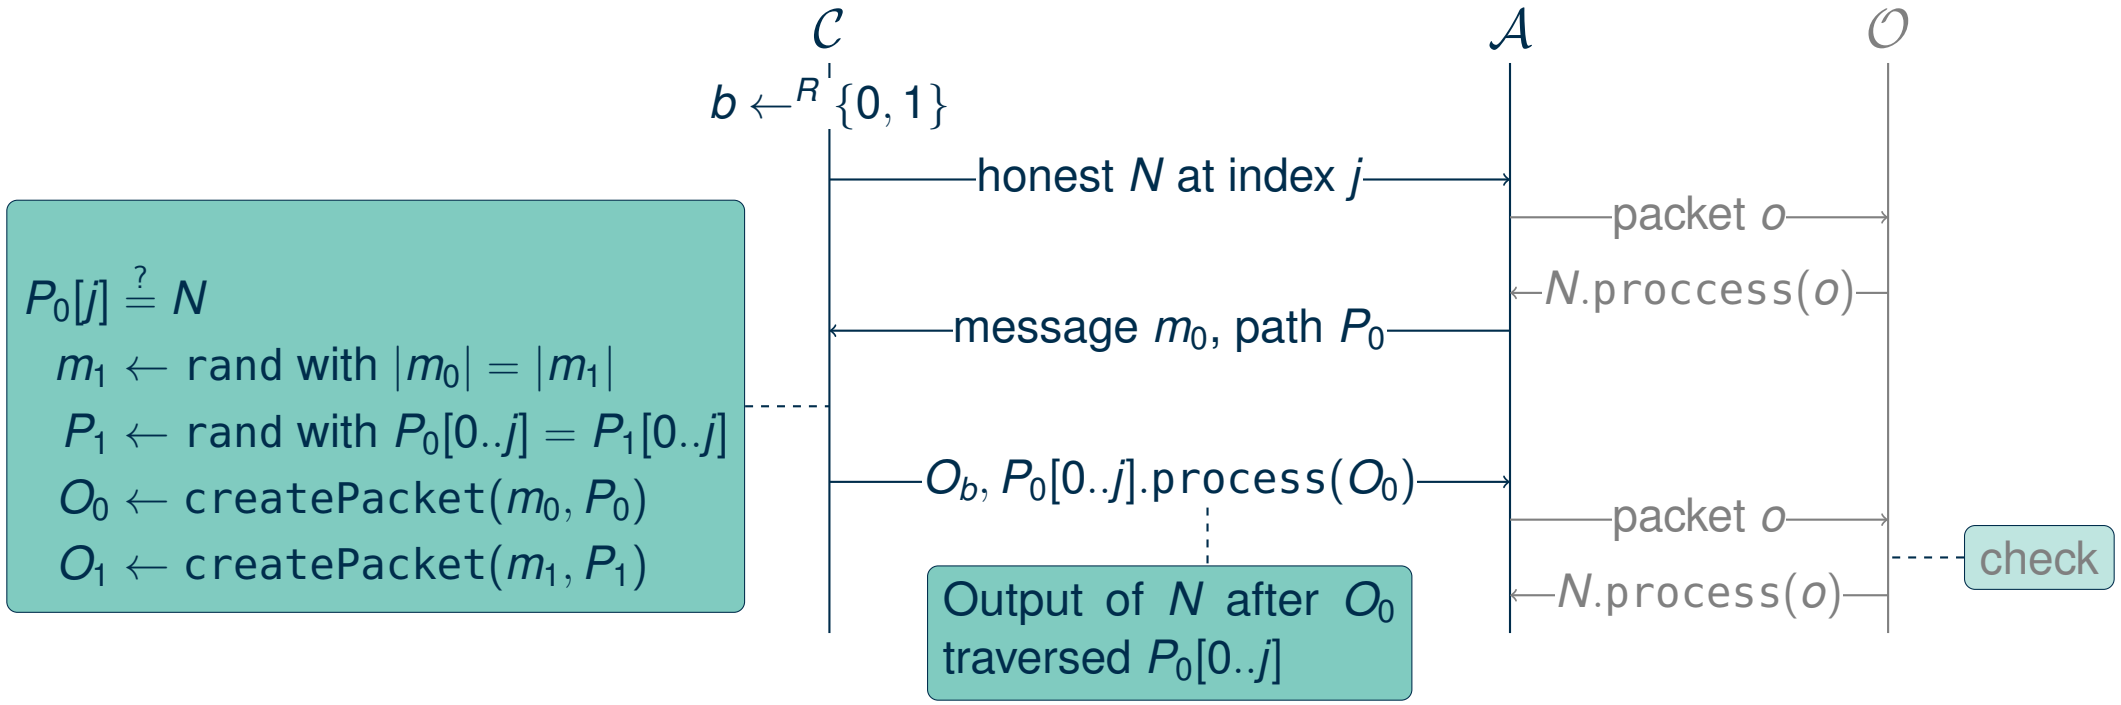

PolySphinx

Case 2: Adversary can link incoming packet to outgoing at replication node.

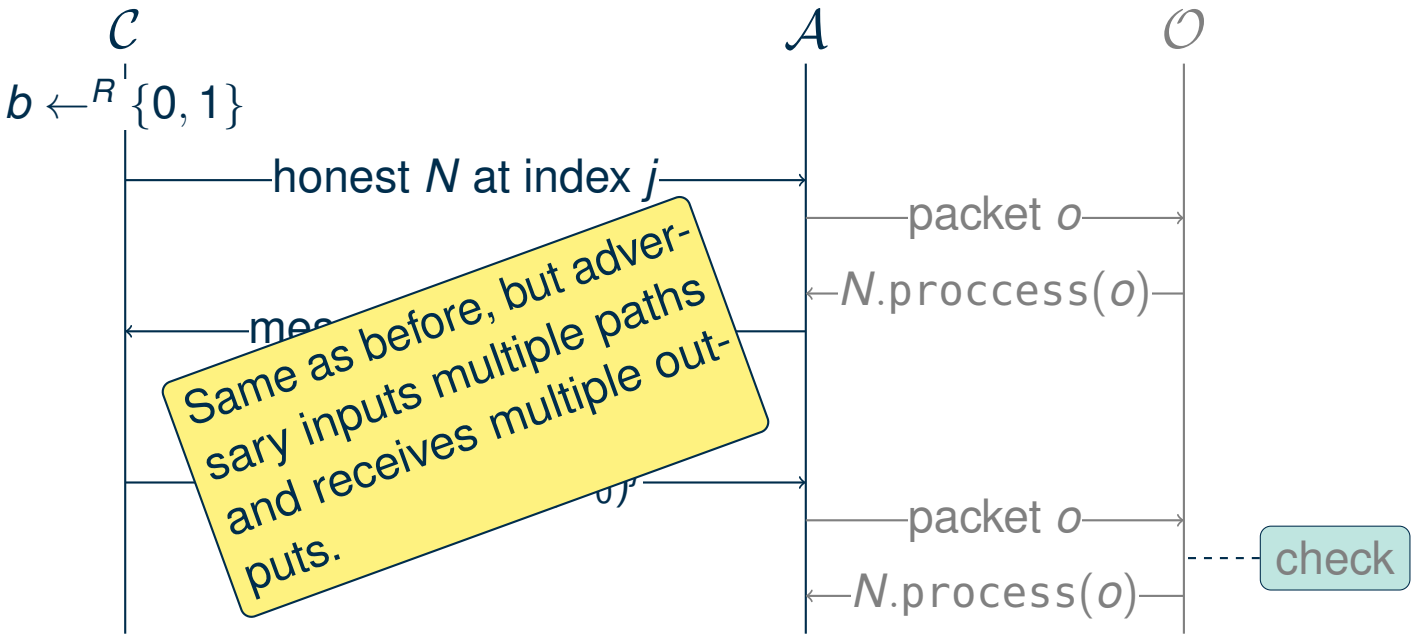

PolySphinx

There exists no PPT $\mathcal{A}$ who can win **Case 1**-game with non-negligible advantage over random guessing.

There exists no PPT $\mathcal{A}$ who can win **Case 2**-game with non-negligible advantage over random guessing.

Privacy Goal: *Single Sender Anonymity*

Given at least one honest node between Alice and Bob, the adversary cannot tell that she sent to him.

PolySphinx

There exists no PPT $\mathcal{A}$ who can win **Case 1**-game with non-negligible advantage over random guessing.

Proof strategy: Hybrid games where package creation is progressively replaced with randomness.

- H_0: Original game
- H_1–H_2: Key material does no longer depend on adversary's input
- H_3: Payload does no longer depend on adversary's input
- H_4: Header does not depend on adversary's input

Sabot

C. Coijanovic, L. Hetz, K. Paterson and T. Strufe, “Sabot: Efficient and Strongly Anonymous Bootstrapping of Communication Channels”, 2025, in submission.

Contact me for preprint!

Sabot – Abstract

Functionality

Alice wants to *start* communicating with Bob in some anonymous communication network.

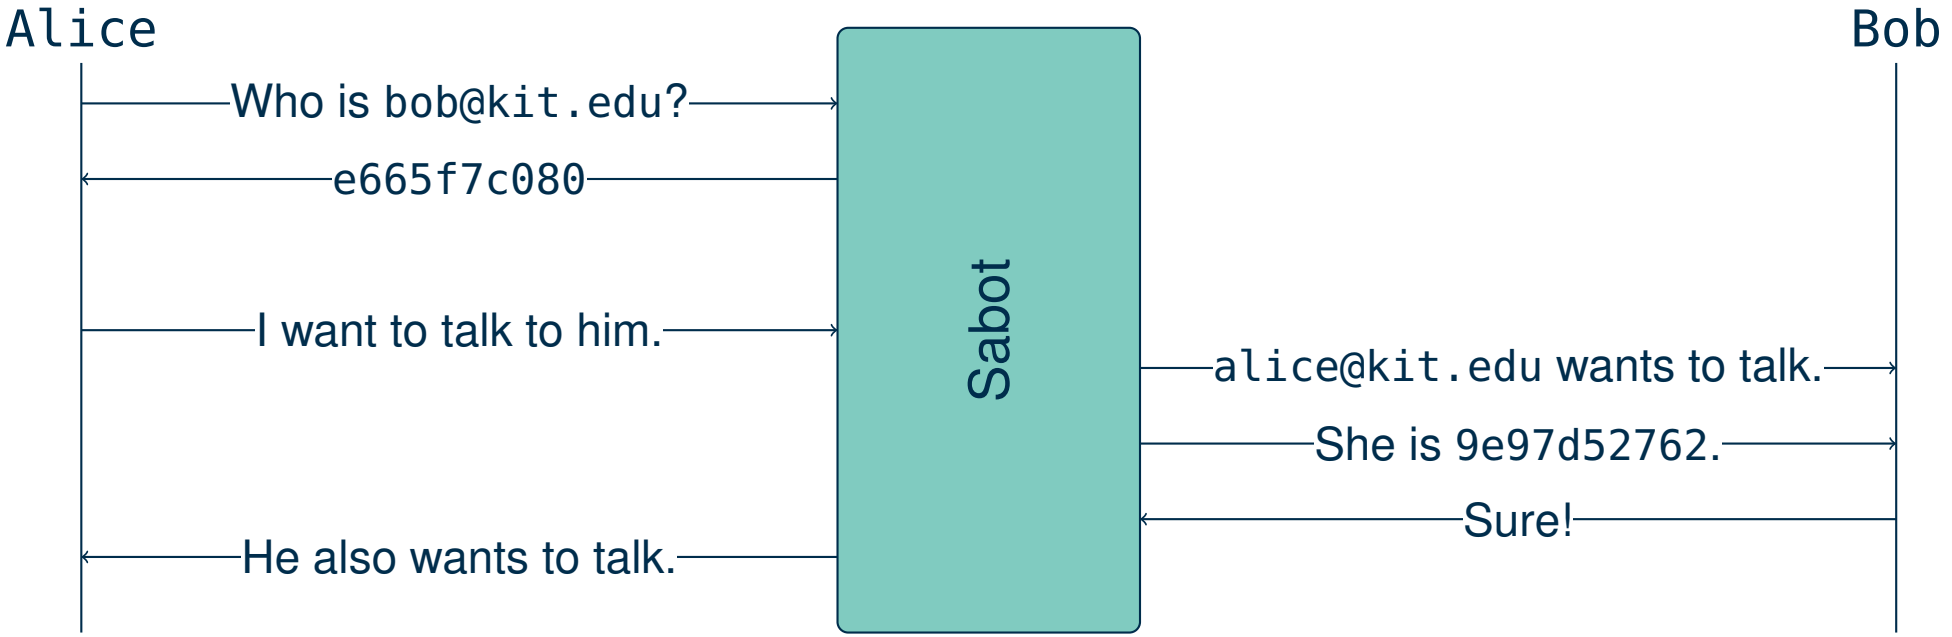

Sabot – Abstract

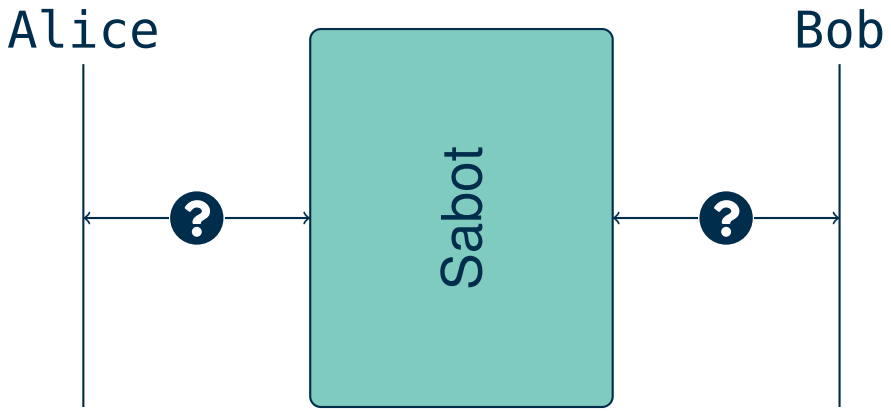

- Unlinking Alice from Bob (like in PolySphinx) is not enough
- How many (if any) people does Alice want to contact?
- Does anyone else want to contact Bob?

Informal Privacy Goal: *Communication Unobservability*

The adversary should not be able to gain information about the communication patterns of honest clients.

Sabot – Abstract

Observation

There are many more and vaguer points of attack!

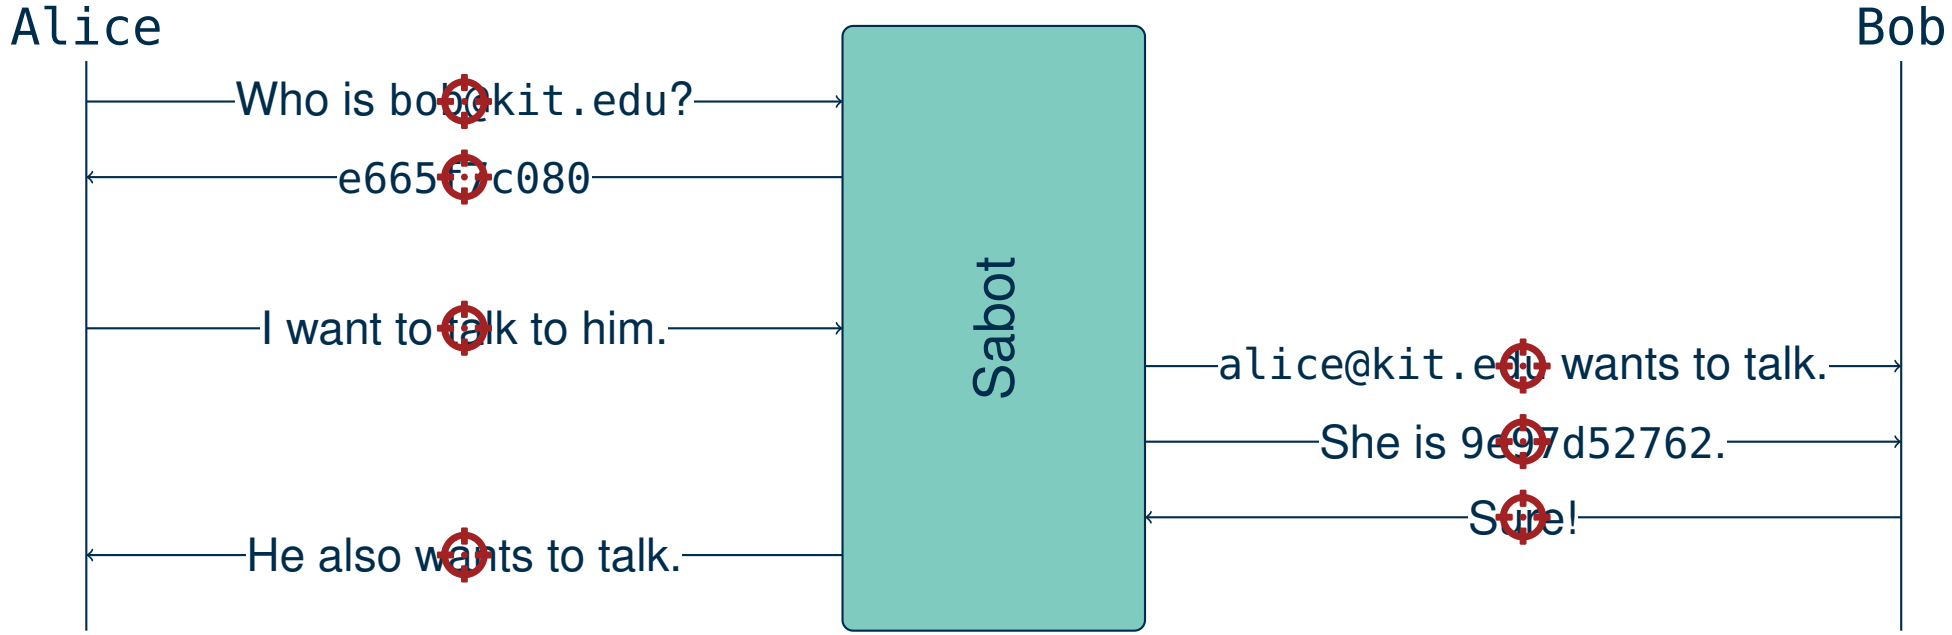

Sabot

Privacy Goal: *Communication Unobservability*[1]

The adversary should not be able to gain information about the communication patterns of honest clients.

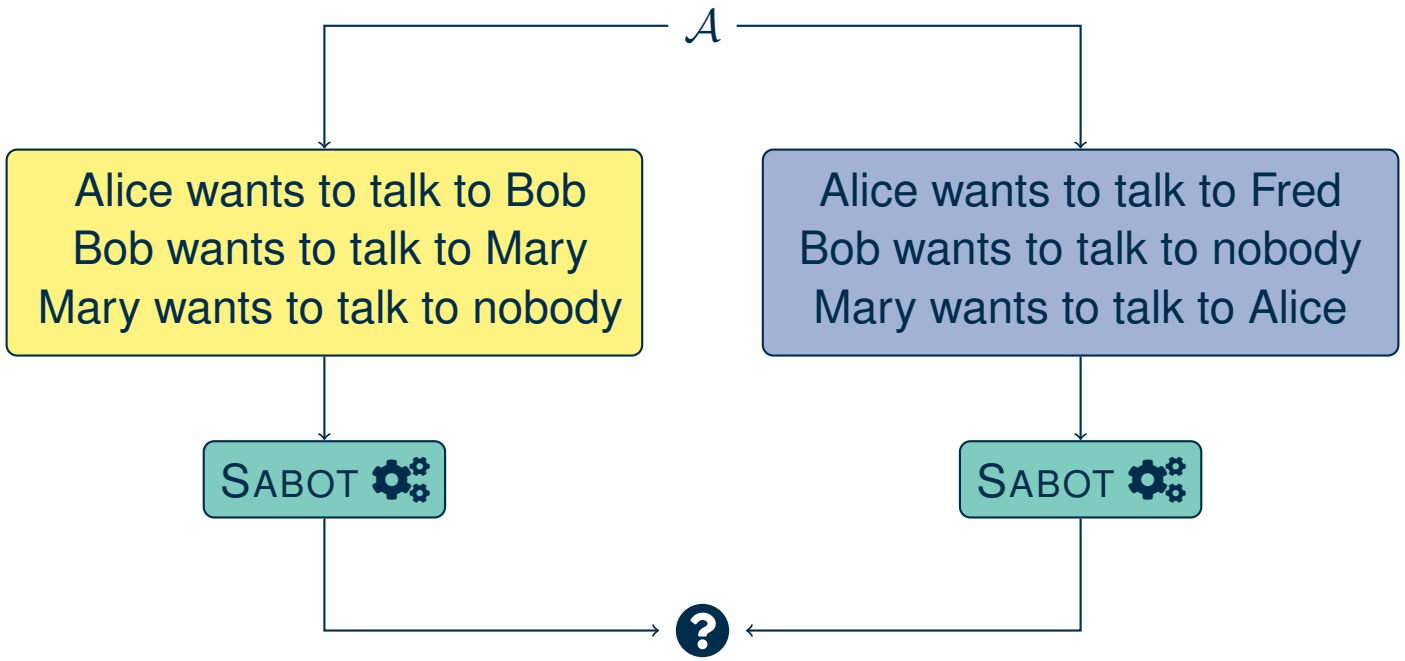

[1] Kuhn et al. „On Privacy Notions in Anonymous Communication“, PoPETS 2019

Sabot

Privacy Goal: *Communication Unobservability*[2]

The adversary should not be able to gain information about the communication patterns of honest clients.

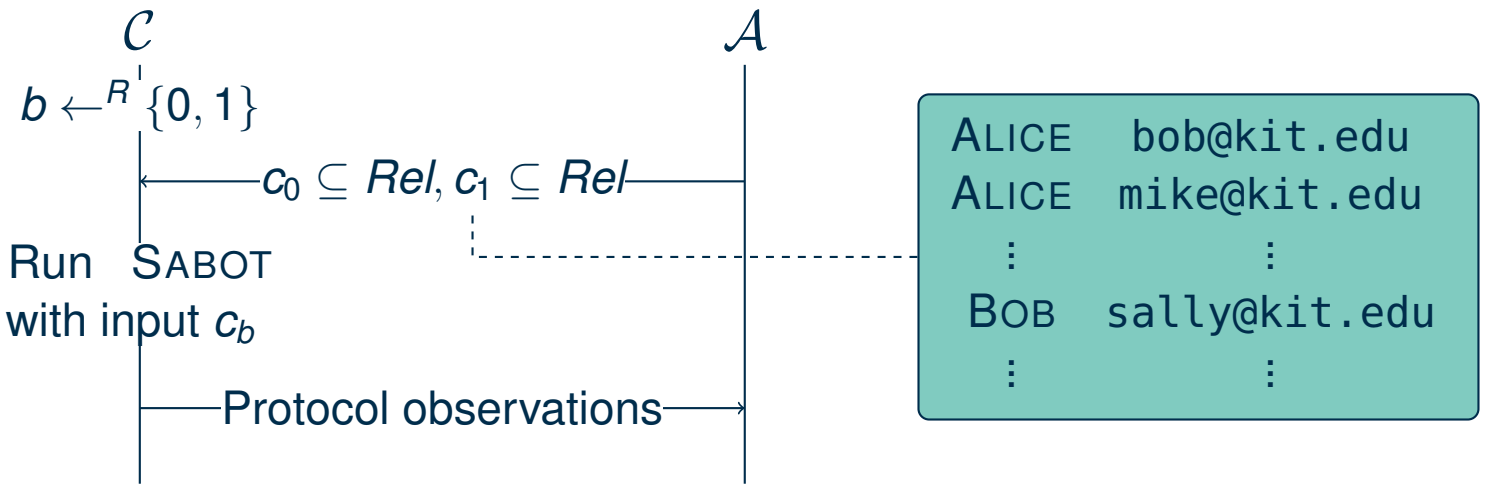

[2]Kuhn et al. „On Privacy Notions in Anonymous Communication“, PoPETS 2019

Sabot

There exists no PPT $\mathcal{A}$ who can win Communication Unobservability game with non-negligible advantage over random guessing.

Proof Strategy: Hybrid games where protocol steps are progressively replaced by random behavior.

- H_0: Game with „normal“ Sabot
- H_1: Like H_0, but senders retrieve random information
- H_2: Like H_1, but senders notify nobody
- H_3: Like H_2, but receivers retrieve random information
- H_4: Like H_3, but receivers notify nobody

Comparison

Comparison

Complexity shifts between high-level and low-level approaches.

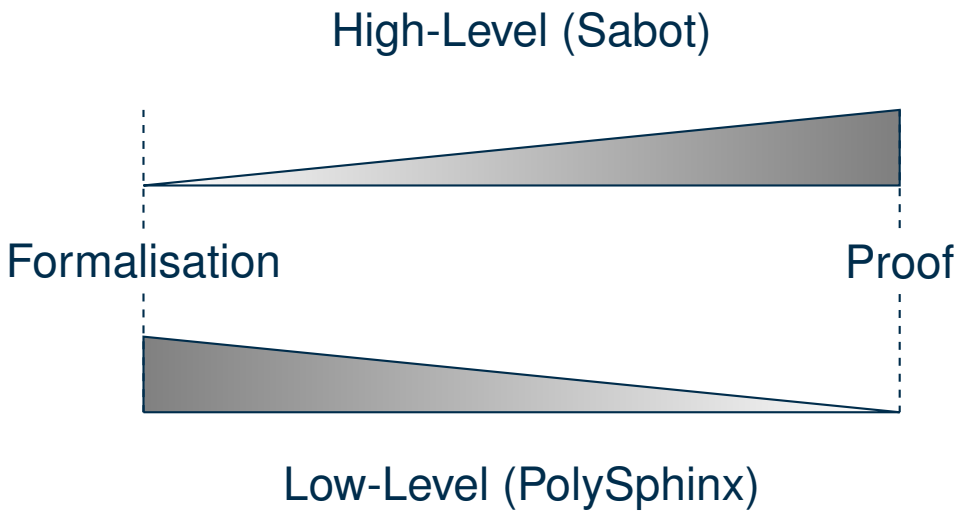

High complexity makes it difficult to

- ensure full *coverage*
- understand & verify correctness

Comparison

Comparison

High-level approach makes it easier to compare a wide range of protocols.

- Low-level properties are protocol-specific (or close to)
- If properties change, protocols are hard to compare

- High-level notions can be used for any* unicast ACN
- If x achieves notion and y achieves, both provide (at least) the same privacy protection

Comparison

Comparison

Low-level approach better covers active adversaries

- Low-level approach has active adversary „build in“ through access to key material of malicious nodes and oracle
- Current high-level approach has the ability to express active attacks, but it’s not well formalized
- Including active attacks makes already complex proofs even more complex

Recommendations

If there is an established way to formalized privacy/security for your functionality, *use it!*

If not, try to stick to established patterns (e.g., indistinguishability games).

Try to match the abstraction level of your formalization to that of the functionality.

Conclusion

Anonymous communication networks are cool!

There is no *one size fits all* solution for privacy formalization.

Provable privacy works very similarly to provable security.