How much privacy is enough?

Evaluating tradeoffs in privacy and scalability

Ian Miers
@secparam
Cornell Tech/ Zcash Foundation

WikiLeaks

@wikileaks

Follow

WikiLeaks now accepts anonymous Bitcoin donations on 1HB5XMLmzFVj8ALj6mfBsbifRoD4miY36v

8:12 PM - 14 Jun 2011

ELLIPTIC

CHAINALYSIS

An Analysis of Anonymity in the Bitcoin System

Fergal Reid and Martin Harrigan

Abstract

Anonymity in Bitcoin, a peer-to-peer electronic currency system, is a complicated issue. Within the system, users are identified by public-keys only. An attacker wishing to de-anonymize its users will attempt to construct the one-to-many mapping between users and public-keys and associate information external to the system with the users. Bitcoin tries to prevent this attack by storing the mapping of a user to his or her public-keys on that user's node

When the cookie meets the blockchain: Privacy risks of web payments via cryptocurrencies

Steven Goldfeder*, Harry Kalodner*, Dillon Reisman†, Arvind Narayanan*
Princeton University
*{stevenag, kalodner, arvindn}@cs.princeton.edu
†dillon@lonlon.io

***Abstract*—We show how third-party web trackers can deanonymize users of cryptocurrencies. We present two distinct but complementary attacks. On most shopping websites, third party trackers receive information about user purchases for purposes of advertising and analytics. We show that, if the user pays using a cryptocurrency, trackers typically possess enough information about the purchase to uniquely identify the transaction on the blockchain, link it to the user's cookie, and further to the user's real identity. Our second attack shows that if the tracker is able to link two purchases of the same user to the blockchain in this manner, it can identify the user's entire cluster of addresses and transactions on the blockchain, even if the user employs blockchain anonymity techniques such as CoinJoin. The attacks are passive and hence can be retroactively applied to past purchases. We discuss several mitigations, but none are perfect.**

I. Introduction

Eight years after Bitcoin's introduction, the ability to pay online using cryptocurrencies is common: prominent merchants such as Microsoft, Newegg, and Overstock support it. Cryptocurrency users tend to value financial privacy, and it is a major reason for choosing to pay with Bitcoin [1]. Yet, websites including shopping sites are known to be rife with

from web pages even if it is not leaked to them by default. We show that this is possible on the vast majority of merchant sites.

Of course, Bitcoin does not guarantee unlinkability of transactions. But while linking of a user's Bitcoin addresses *with each other* is well known [3]–[6], our attack shows how to link addresses to external information, including identity.

The main defense against linkage attacks is mixing [7], [8]. The best known mixing technique is CoinJoin, in which users send coins to each other in a way that hides the link between their old and new coins. Our second main contribution is showing the effectiveness of the *cluster intersection attack*, a previously known attack against mixing. Specifically, we show that a small amount of additional information, namely that two (or more) transactions were made by the same entity, is sufficient to undo the effect of mixing (see Figure 1). While such auxiliary information is available to many potential entities — merchants, other counterparties such as websites that accept donations, intermediaries such as payment processors, and potentially network eavesdroppers — web

A Fistful of Bitcoins: Characterizing Payments Among Men with No Names

Sarah Meiklejohn Marjori Pomarole Grant Jordan
Kirill Levchenko Damon McCoy† Geoffrey M. Voelker Stefan Savage
University of California, San Diego George Mason University†

ABSTRACT

Bitcoin is a purely online virtual currency, unbacked by either physical commodities or sovereign obligation; instead, it relies on a combination of cryptographic protection and a peer-to-peer protocol for witnessing settlements. Consequently, Bitcoin has the unintuitive property that while the ownership of money is implicitly anonymous, its flow is globally visible. In this paper we explore this unique characteristic further, using heuristic clustering to group Bitcoin wallets based on evidence of shared authority, and then using re-identification attacks (i.e., empirical purchasing of goods and services) to classify the operators of those clusters. From this analysis, we characterize longitudinal changes in the Bitcoin market, the stresses these changes are placing on the system, and the challenges for those seeking to use Bitcoin for criminal or fraudulent purposes at scale.

Categories and Subject Descriptors

K.4.4 [**Electronic Commerce**]: Payment schemes

Keywords

Bitcoin; Measurement; Anonymity

By far the most intriguing exception to this rule is Bitcoin. Fin
deployed in 2009, Bitcoin is an independent online monetary sy
tem that combines some of the features of cash and existing onlin
payment methods. Like cash, Bitcoin transactions do not explicitl
identify the payer or the payee: a transaction is a cryptographically
signed transfer of funds from one public key to another. Moreove
like cash, Bitcoin transactions are irreversible (in particular, there i
no *chargeback* risk as with credit cards). However, unlike cash, Bi
coin requires third party mediation: a global peer-to-peer networ
of participants validates and certifies all transactions; such decer
tralized accounting requires each network participant to maintai
the entire transaction history of the system, currently amounting t
over 3GB of compressed data. Bitcoin identities are thus *pseudo
anonymous*: while not explicitly tied to real-world individuals c
organizations, all transactions are completely transparent.[2]

This unusual combination of features has given rise to consider
able confusion about the nature and consequences of the anonymit
that Bitcoin provides. In particular, there is concern that the comb
nation of scalable, irrevocable, anonymous payments would prov
highly attractive for criminals engaged in fraud or money launde
ing. In a widely leaked 2012 Intelligence Assessment, FBI ar
alysts make just this case and conclude that a key "advantage
of Bitcoin for criminals is that "law enforcement faces difficu
ties detecting suspicious activity, identifying users and obtainin

Privacy enhancements: so many choices?

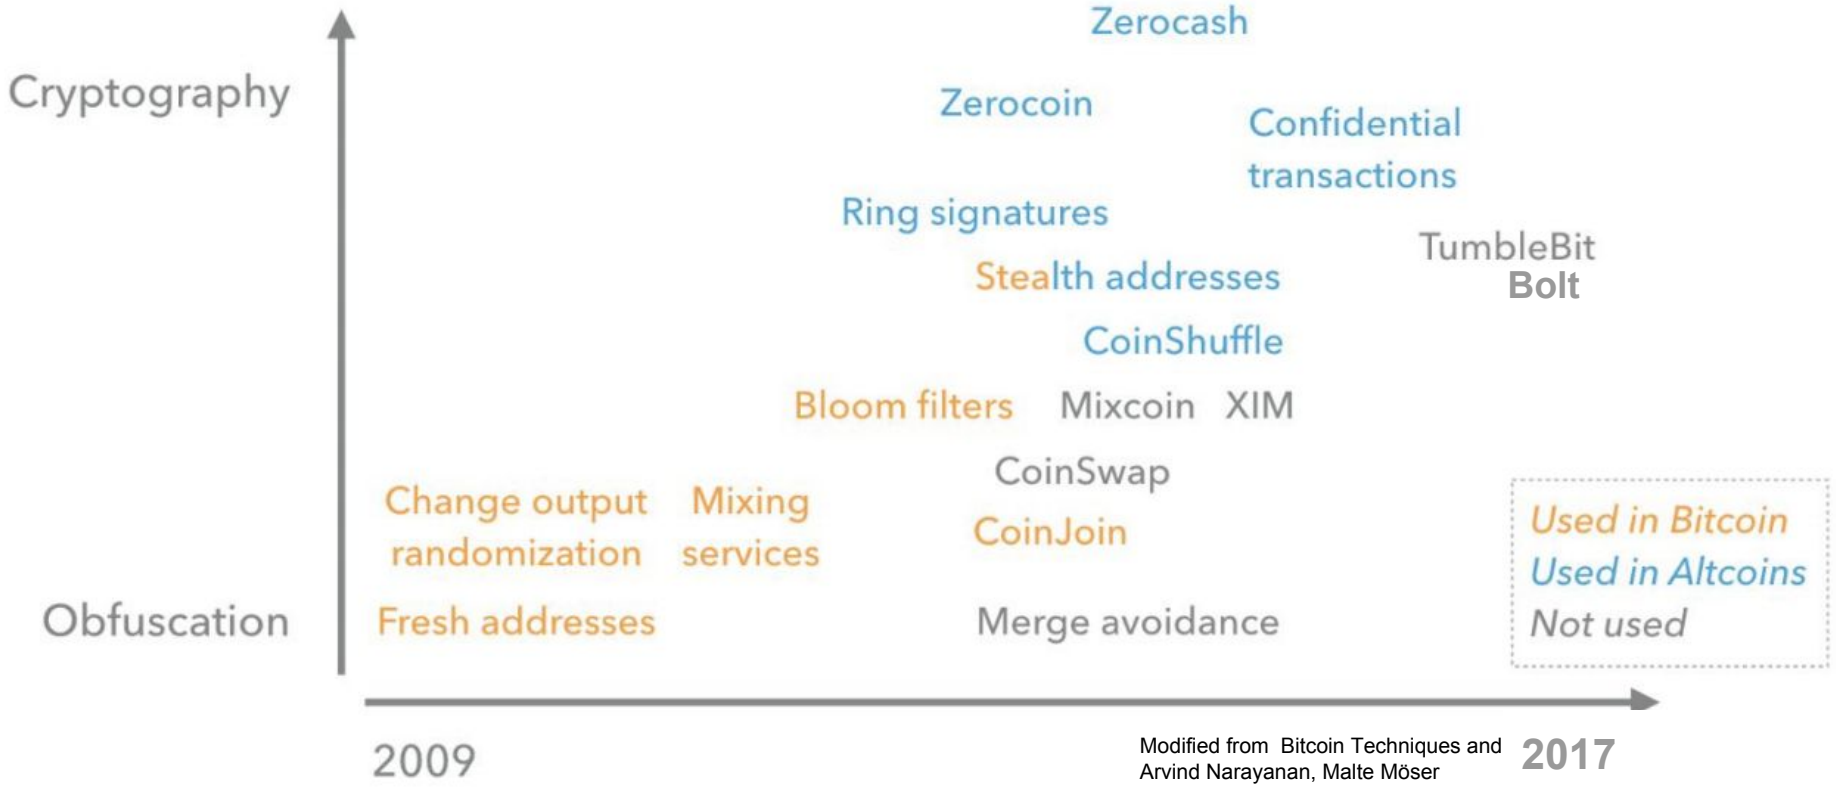

Modified from Bitcoin Techniques and Arvind Narayanan, Malte Möser

Where does a given technique fall?

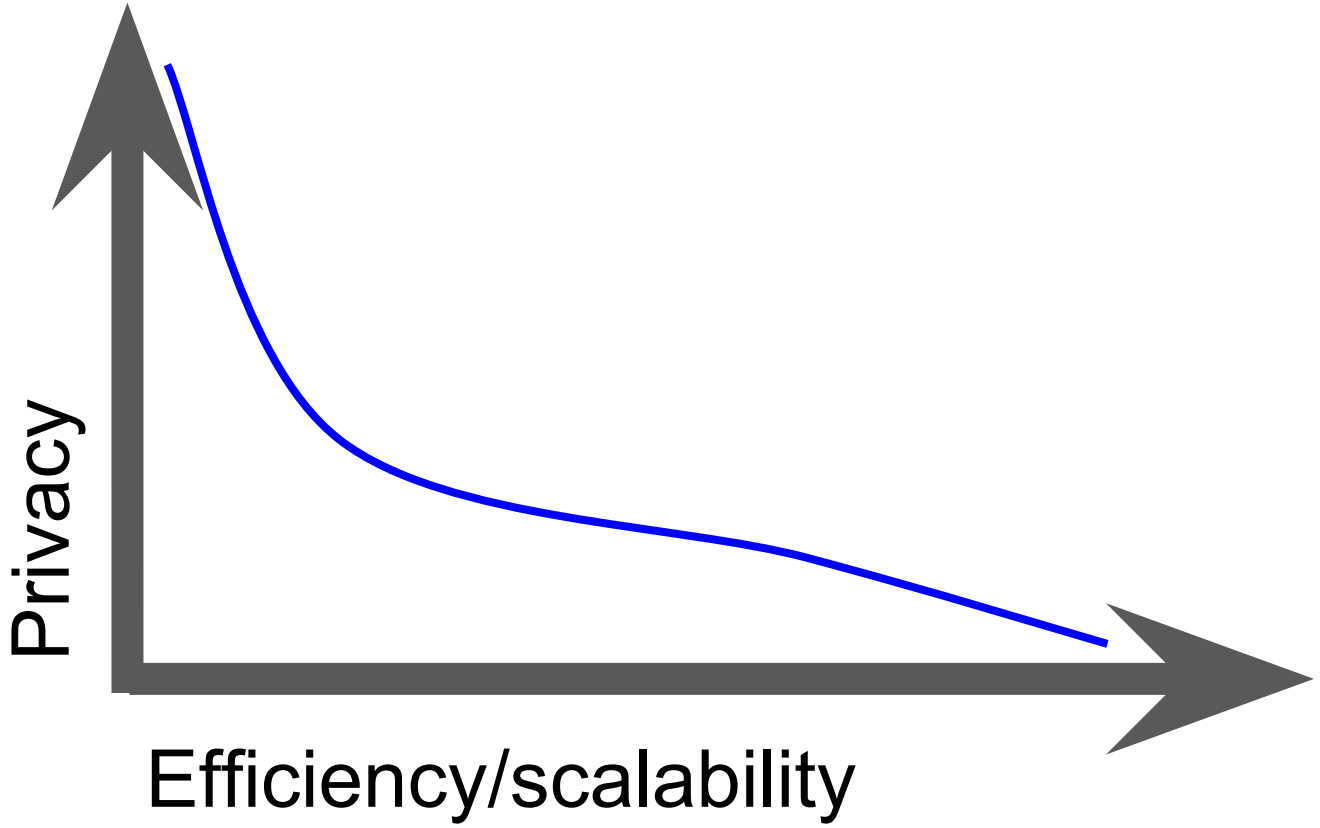

Evaluating privacy?

- Akin to evaluating privacy issues on the Internet in 1992
- Cannot measure with empirical attacks
  - Almost all transactions are speculative
  - Limited usage in daily lives
  - Researchers have data, cost, and ethics limitations
- Need to use thought experiments
- To do so, we must understand realistic threats

Some Real World Privacy Threats

This is your threat model:

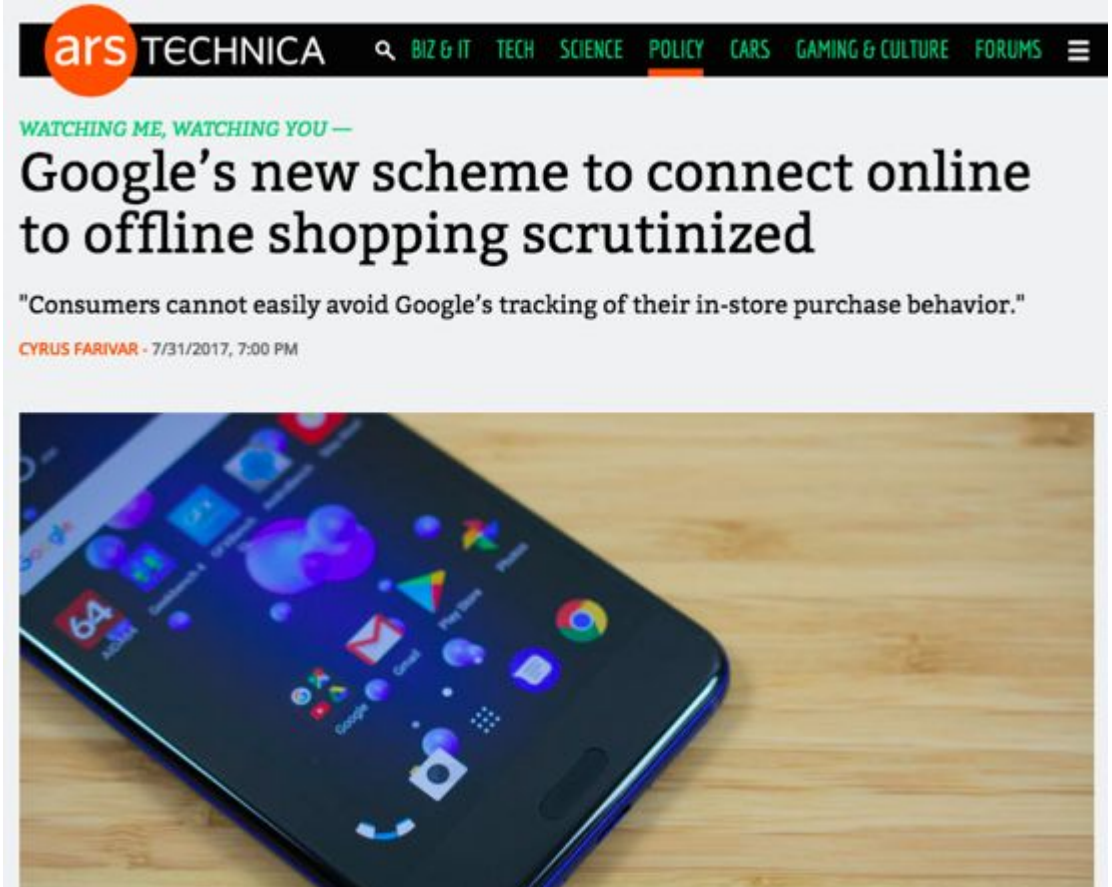

ars TECHNICA BIZ & IT TECH SCIENCE POLICY CARS GAMING & CULTURE FORUMS

WATCHING ME, WATCHING YOU —

Google's new scheme to connect online to offline shopping scrutinized

"Consumers cannot easily avoid Google's tracking of their in-store purchase behavior."

CYRUS FARIVAR - 7/31/2017, 7:00 PM

Google and Mastercard are secretly tracking your offline purchases

by BRYAN CLARK — 4 weeks ago in GOOGLE

How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did

Kashmir Hill, FORBES STAFF
Welcome to The Not-So Private Parts where technology & privacy collide **FULL BIO**

Every time you go shopping, you share intimate details about your consumption patterns with retailers. And many of those retailers are studying those details to figure out what you like, what you need, and which coupons are most likely to make you happy. Target TGT -0.66%, for example, has figured out how to data-mine its way into your womb, to figure out whether you have a baby on the way long before you need to start buying diapers.

Target has got you in its aim

And this

Galore

A GUIDE TO STALKING YOUR EX ON VENMO

Pop

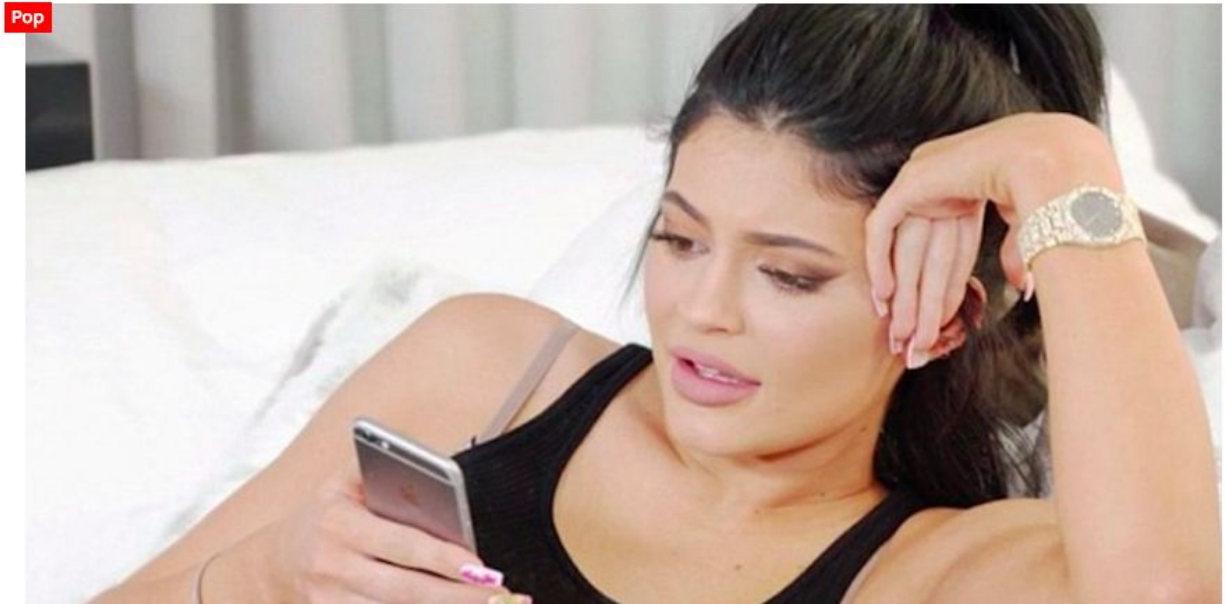

Fungibility:

- Freshly mined coins sell for a premium
- Exchanges blocking customers based on transaction history
- Exchanges are not mere third party observers:
  - Know more than just the transaction graph
  - Make transactions on user’s behalf
- Akin to being private on the internet while using Google/Gmail/Maps/Android

What are the defenses?

In a world of AI/ML and targeted ads, ***plausible deniability*** is not a **plausible defense**.

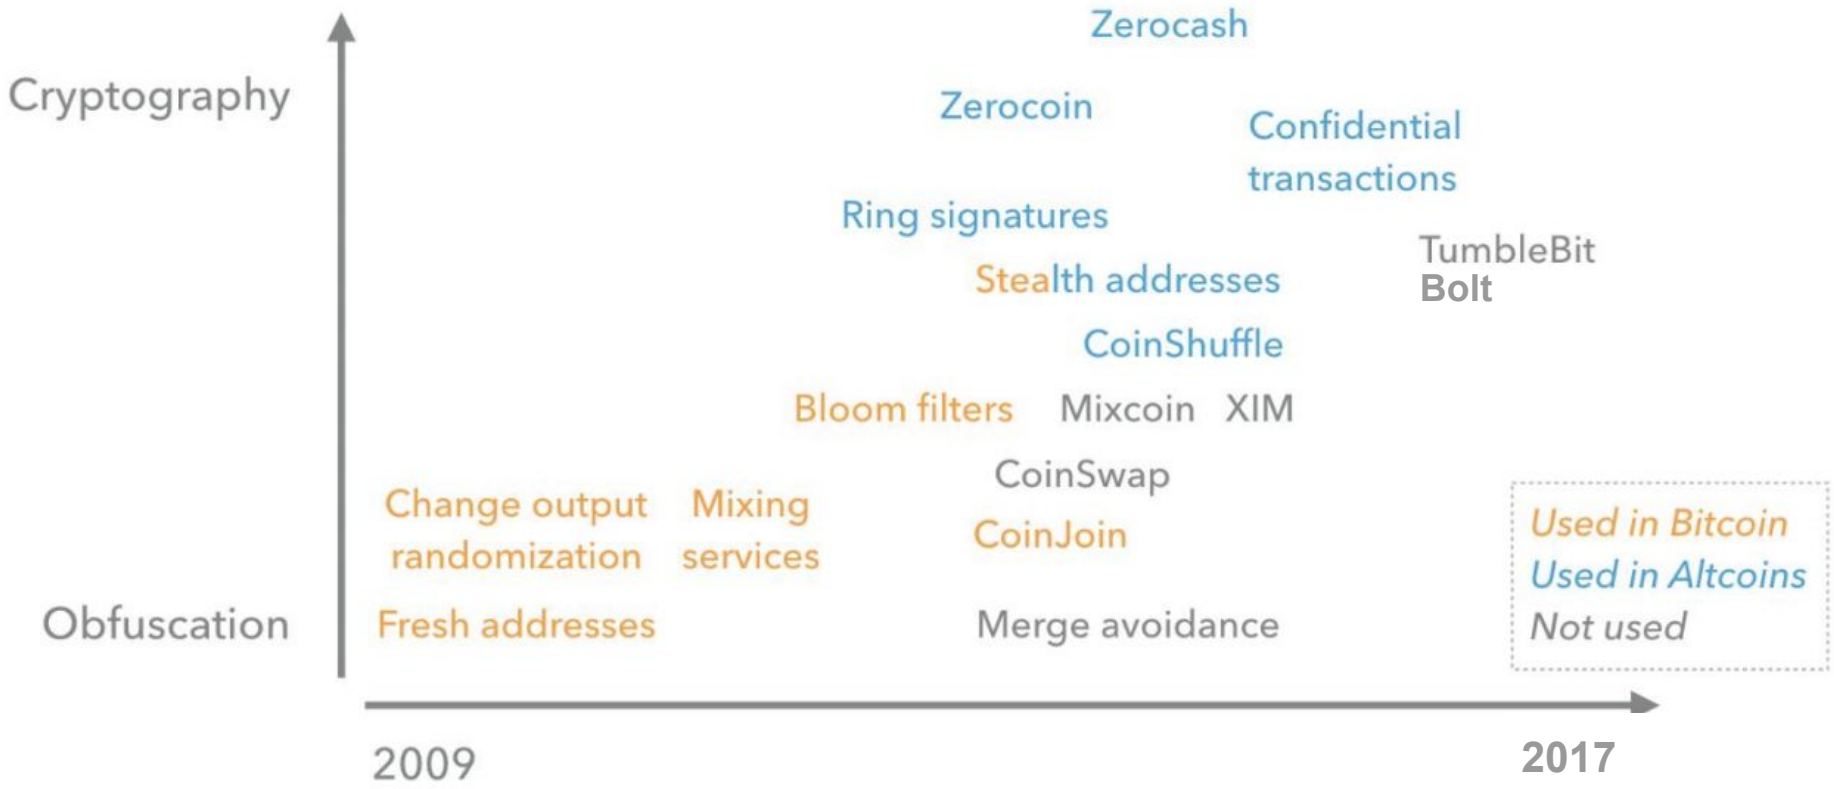

Modified from Bitcoin Techniques and Politics
Arvind Narayanan, Malte Möser

Blockchain privacy is not intuitive.

- Only threat is NOT a third party passive observer
- Must consider ***active*** attackers who:
  - Receives payments from targeted users
  - Sends payments to targeted users
  - Interact with third parties
- Consider obvious attacks:
  - Merchants who try and track customers
  - Users who try and identify a recipients real identity
  - Exchanges who ban customers for certain transaction types

Modified from Bitcoin Techniques and Politics
Arvind Narayanan, Malte Möser

Privacy approaches

- Bitcoin (vanilla): explicitly identify origin of payment
- Decoy transaction based systems:
  - Pick e.g. 5 decoy source transactions to hide real origin
    - Coinjoin, Mimblewimble, etc. (decoys sampled from current transactions)
    - Cryptonote/RingCT(e.g Monero, etc) (decoys sampled from all of history)
- Zerocoin and Zerocash like (e.g. Zcash, etc)
  - Private transactions have no identified origin.

Payments in Bitcoin:

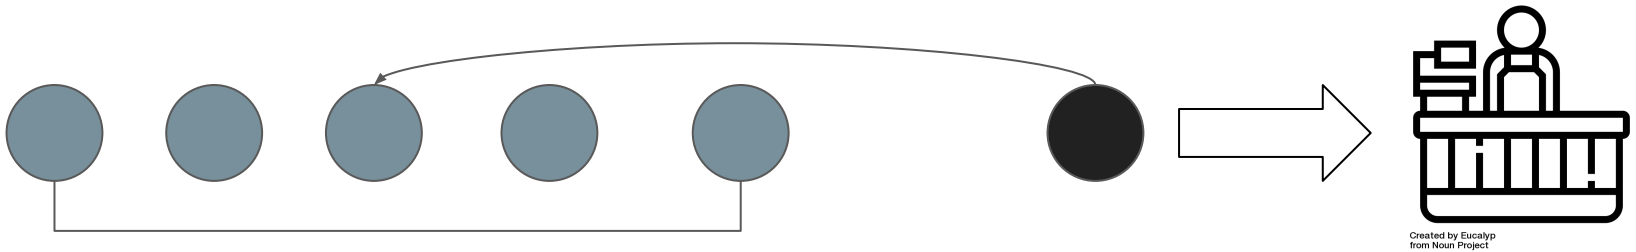

Payments in Decoy Systems (coinjoin/monero/etc)

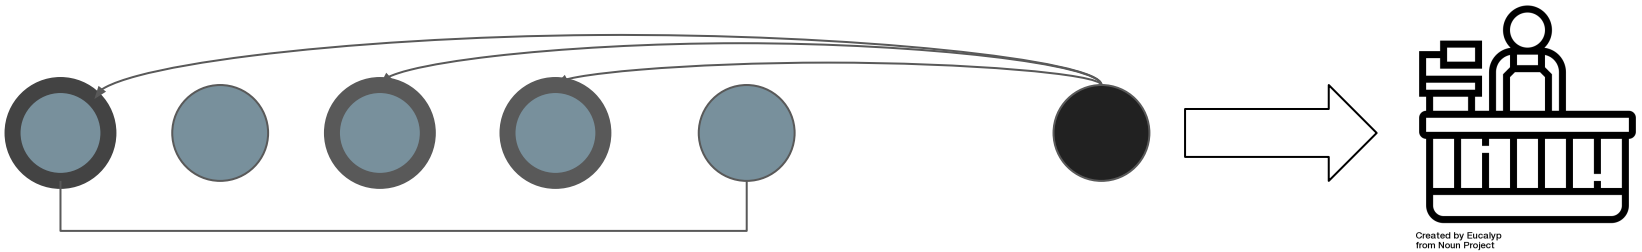

Blockchain

- Coinjoin, etc : decoy set sampled from current transactions
- Cryptonote/RingCT(e.g Monero): decoy sampled from all of history)

Decoy transactions

Payments in Zerocash

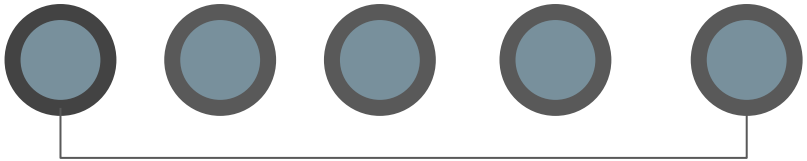

Blockchain

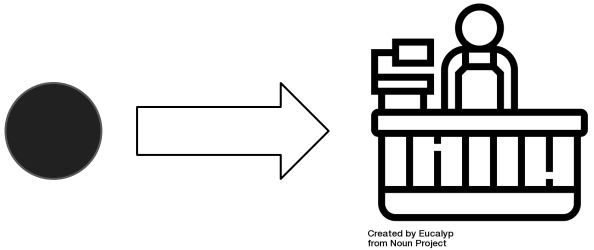

Bitcoin

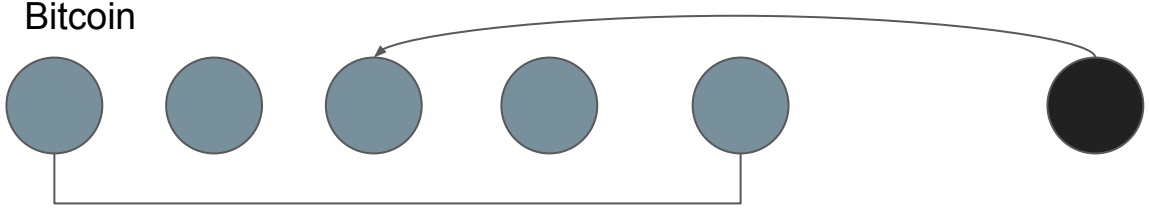

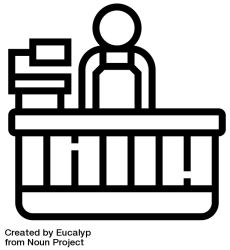

Created by Eucalyp
from Noun Project

Coinjoin/RingCT/etc

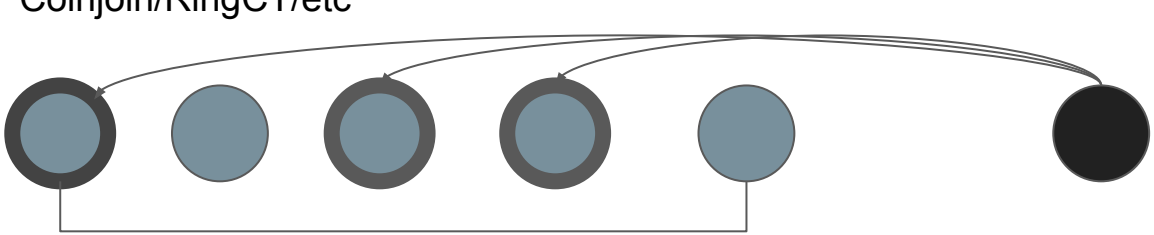

Created by Eucalyp
from Noun Project

Zerocash

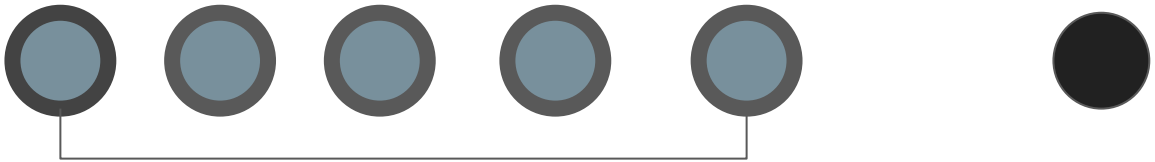

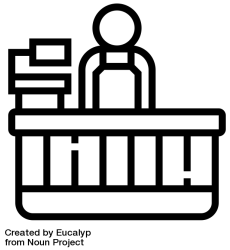

Created by Eucalyp
from Noun Project

Are decoy systems private?

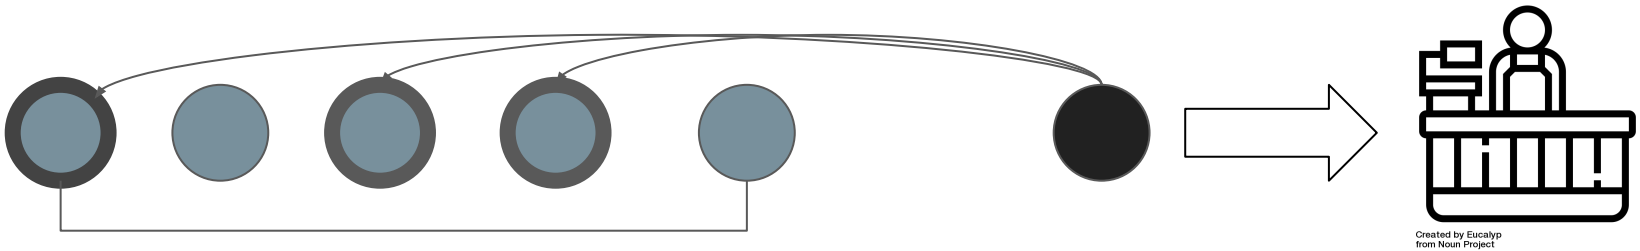

Taint tree: possible ancestor payments

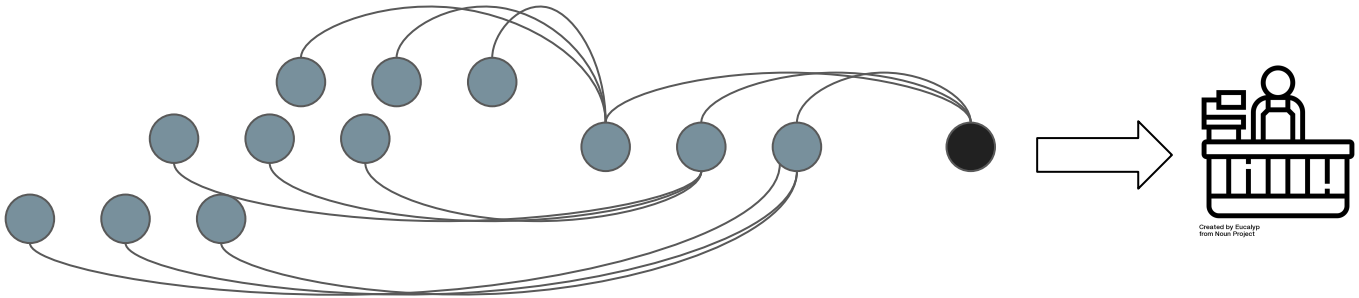

Taint free: following your money

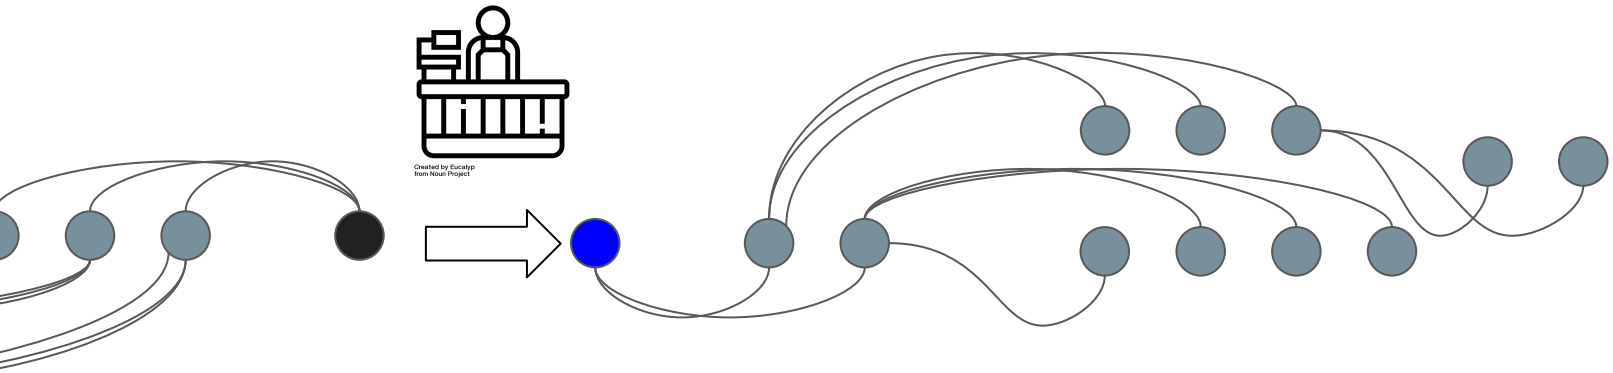

Attacks

Tracking customers

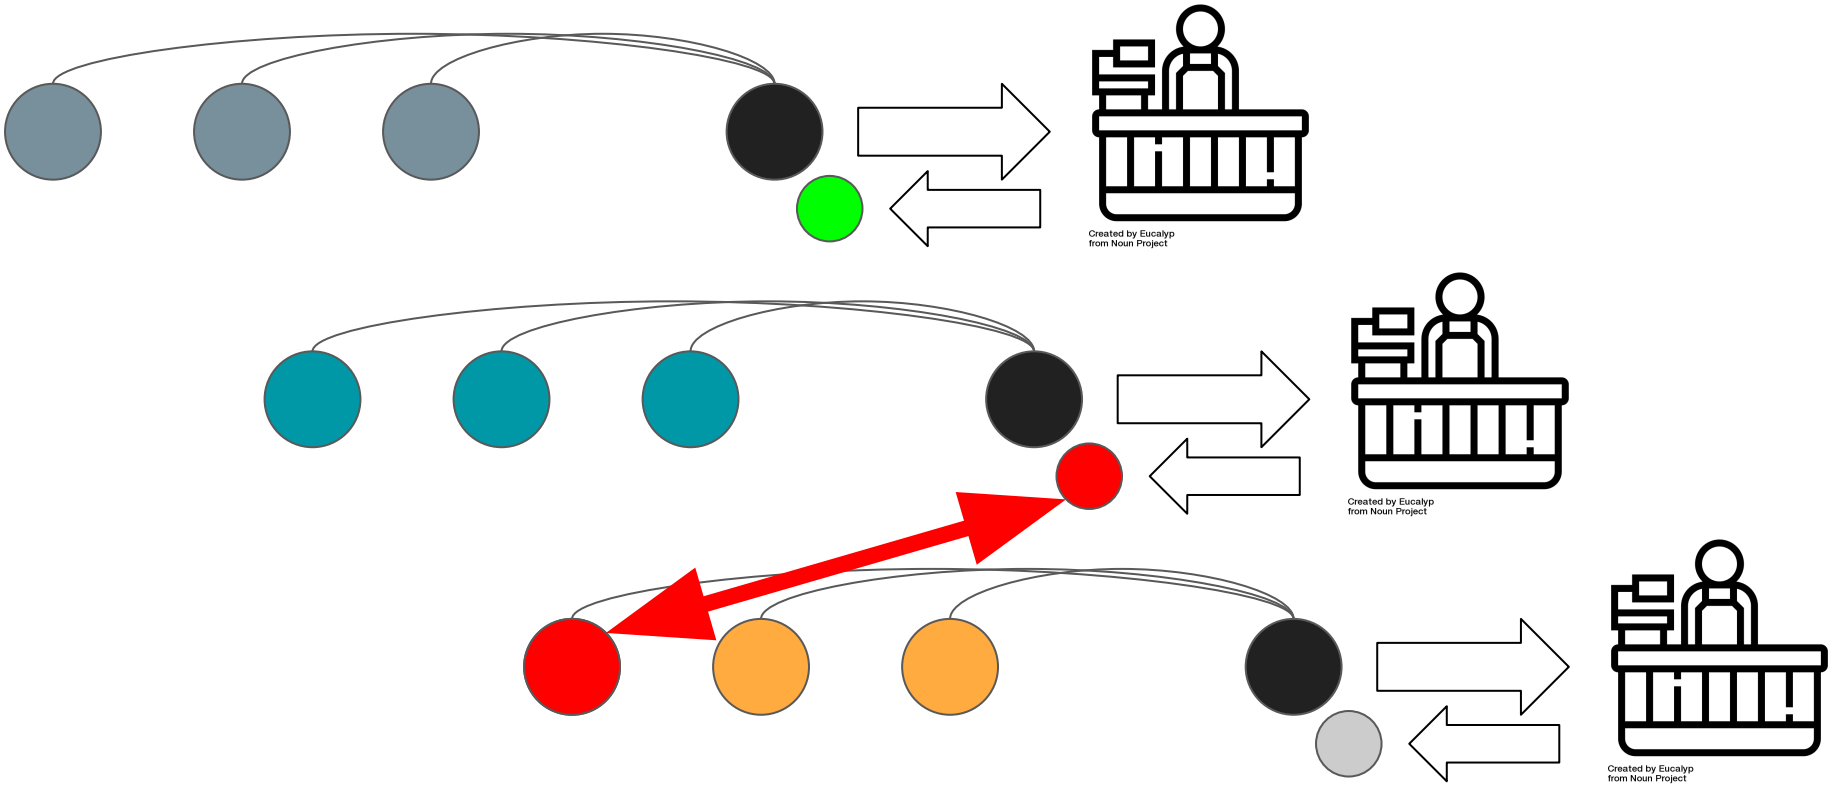

Tracking customers

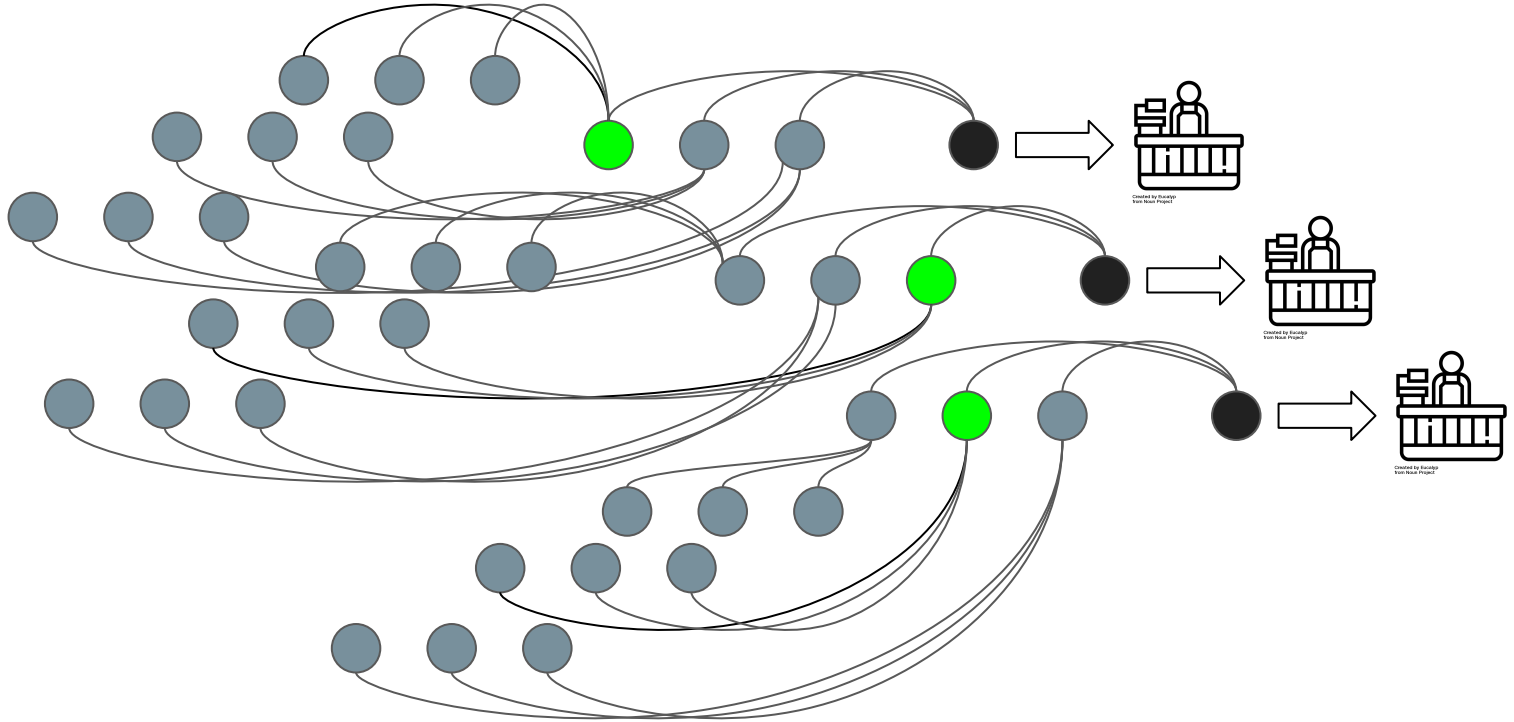

Tracking customers

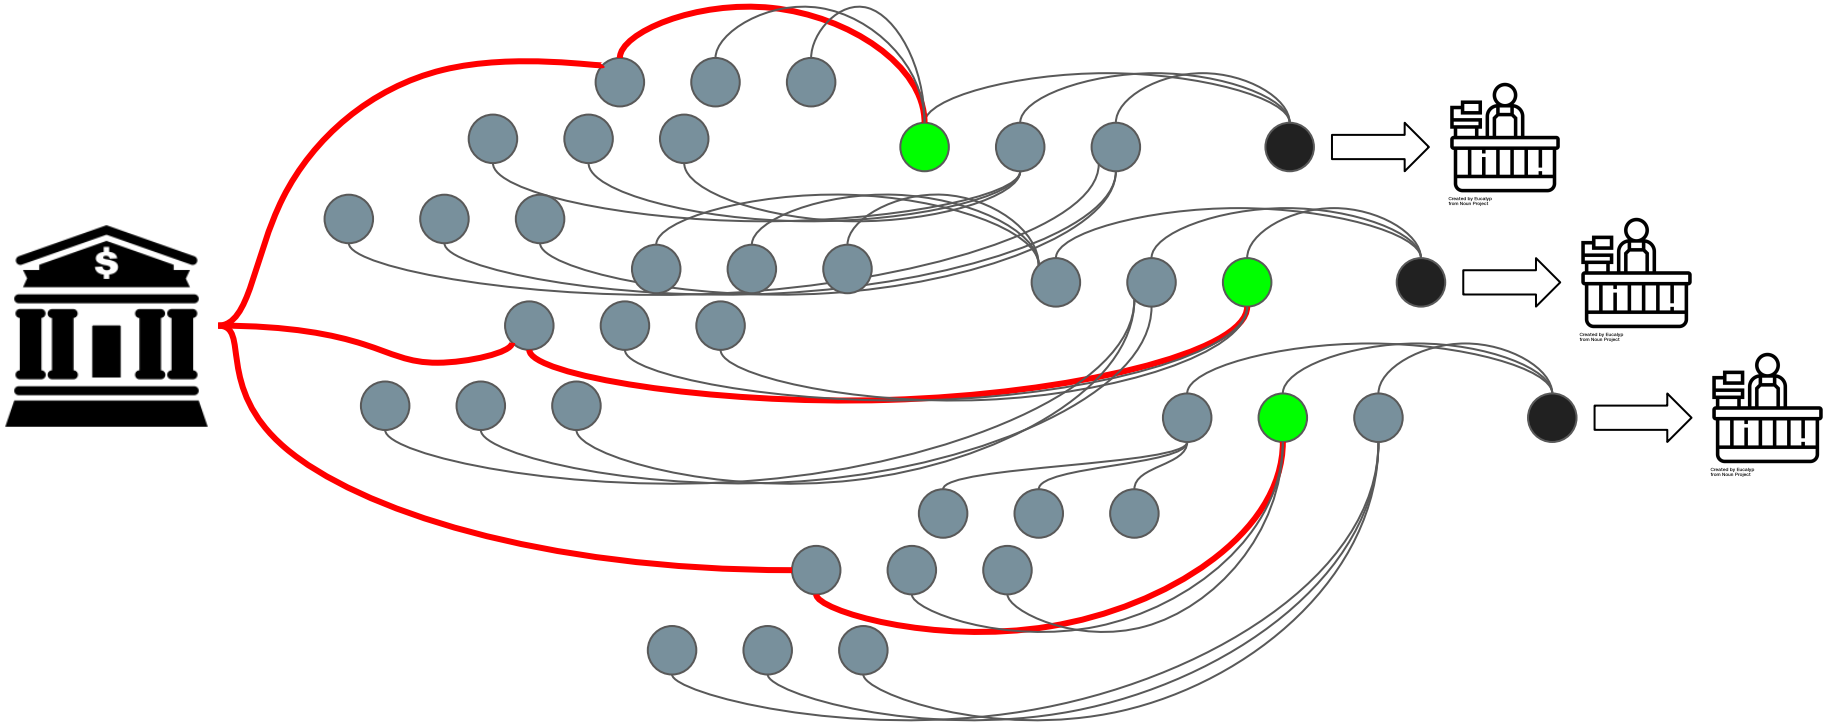

Identifying anonymous merchants

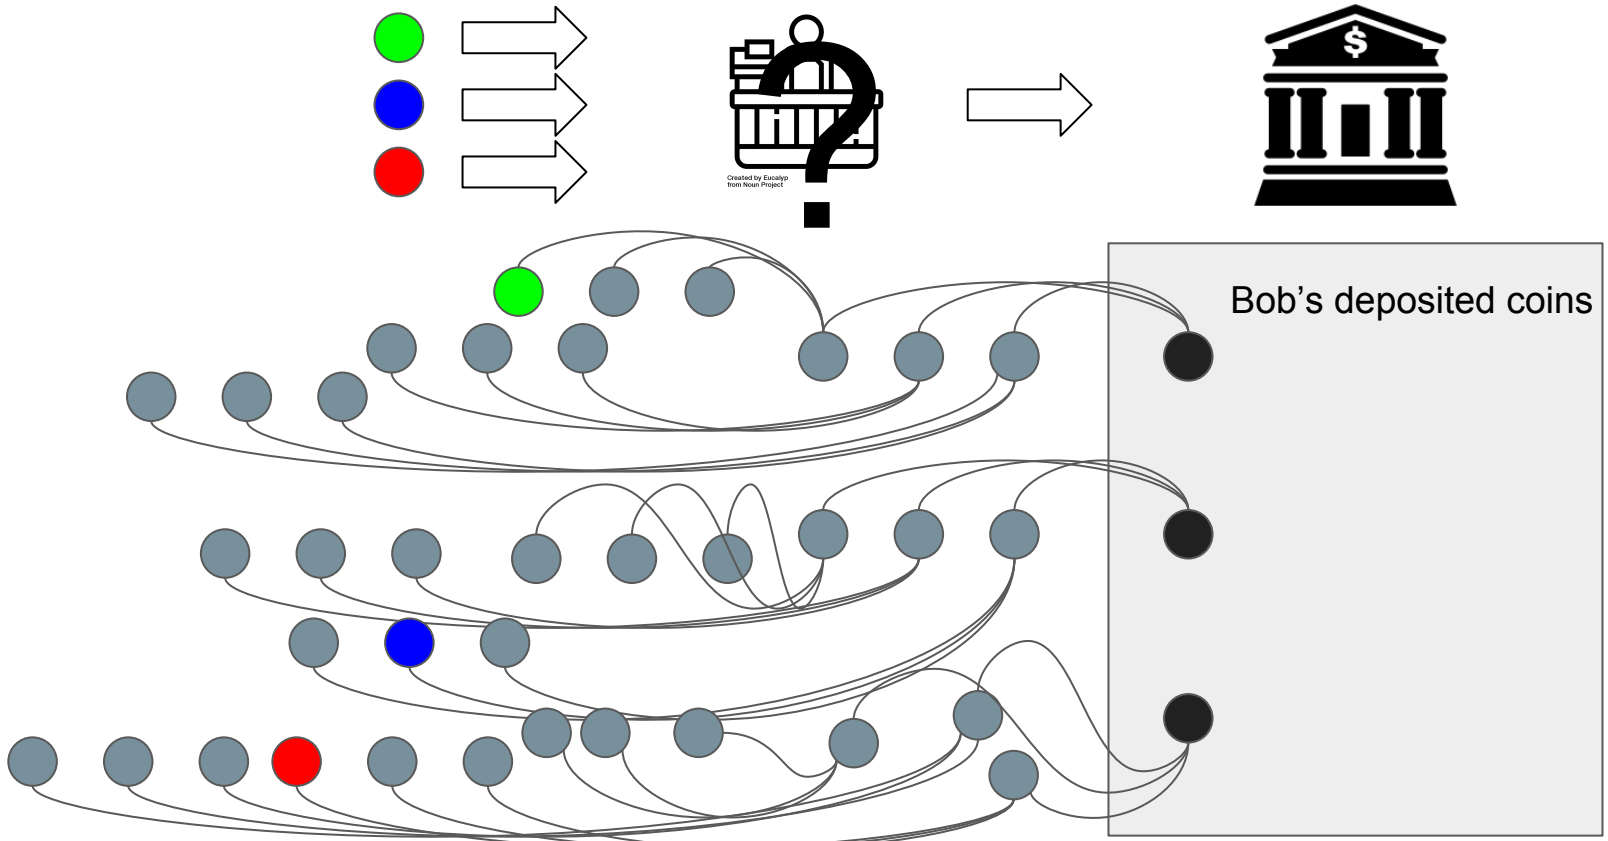

Repeated interactions with a malicious sender/recipient are dangerous

Taint tree: following your money

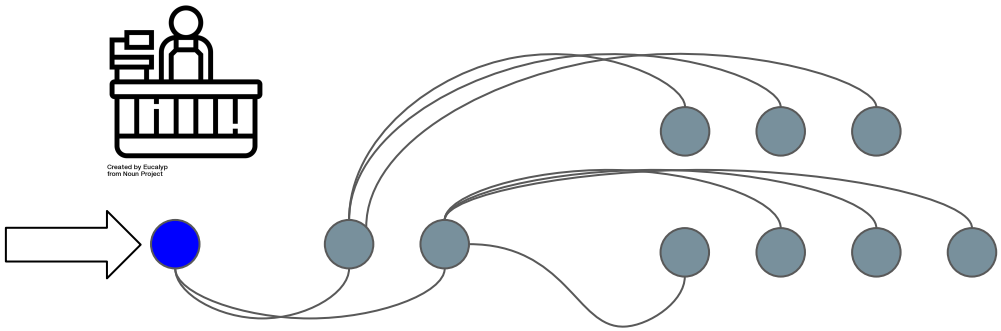

Dust attack: confirming where money is spent

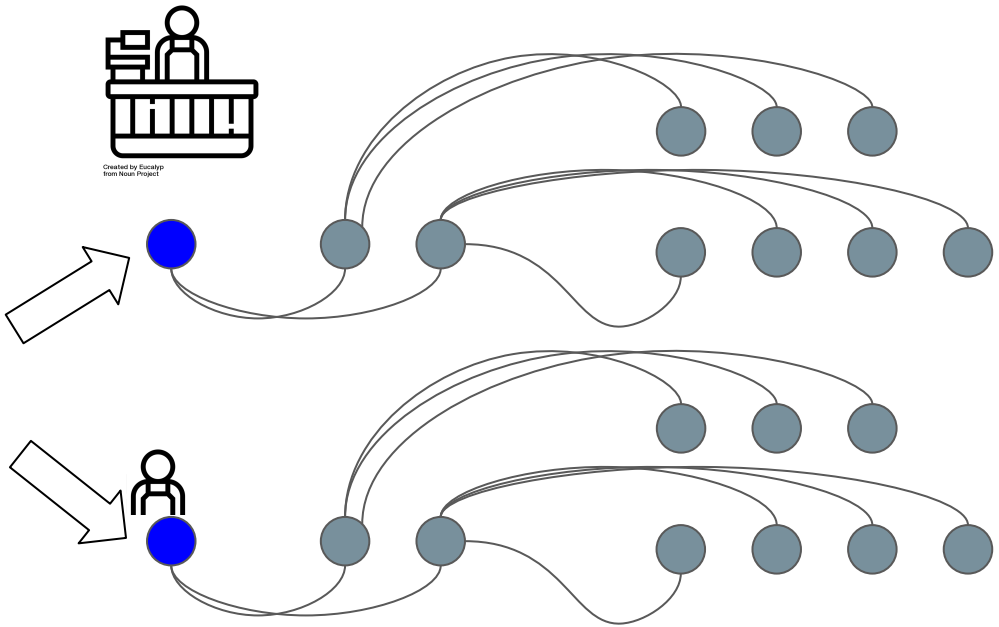

Dust attack: confirming where money is spent

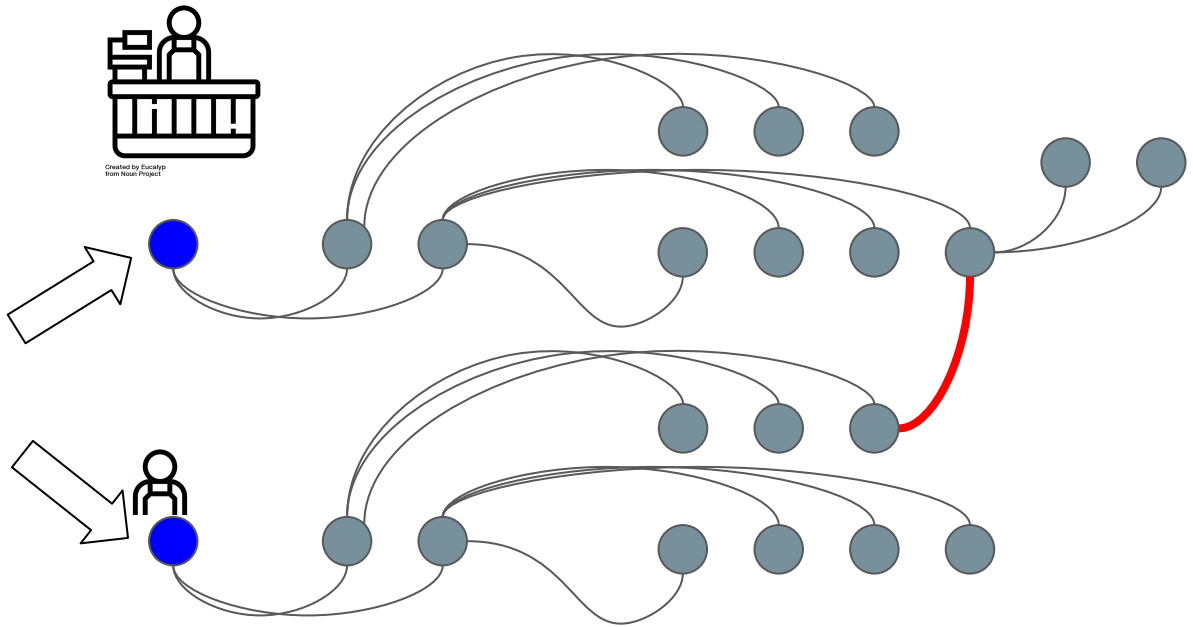

Dust attack: confirming where money is spent

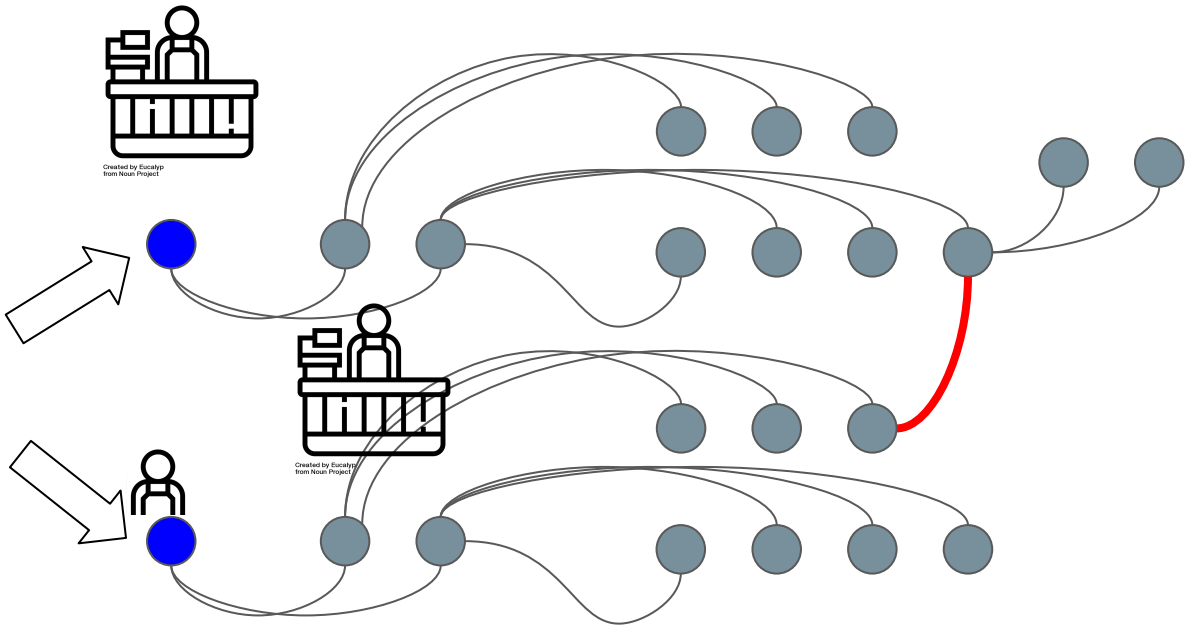

Dust attack: confirming where your money is spent

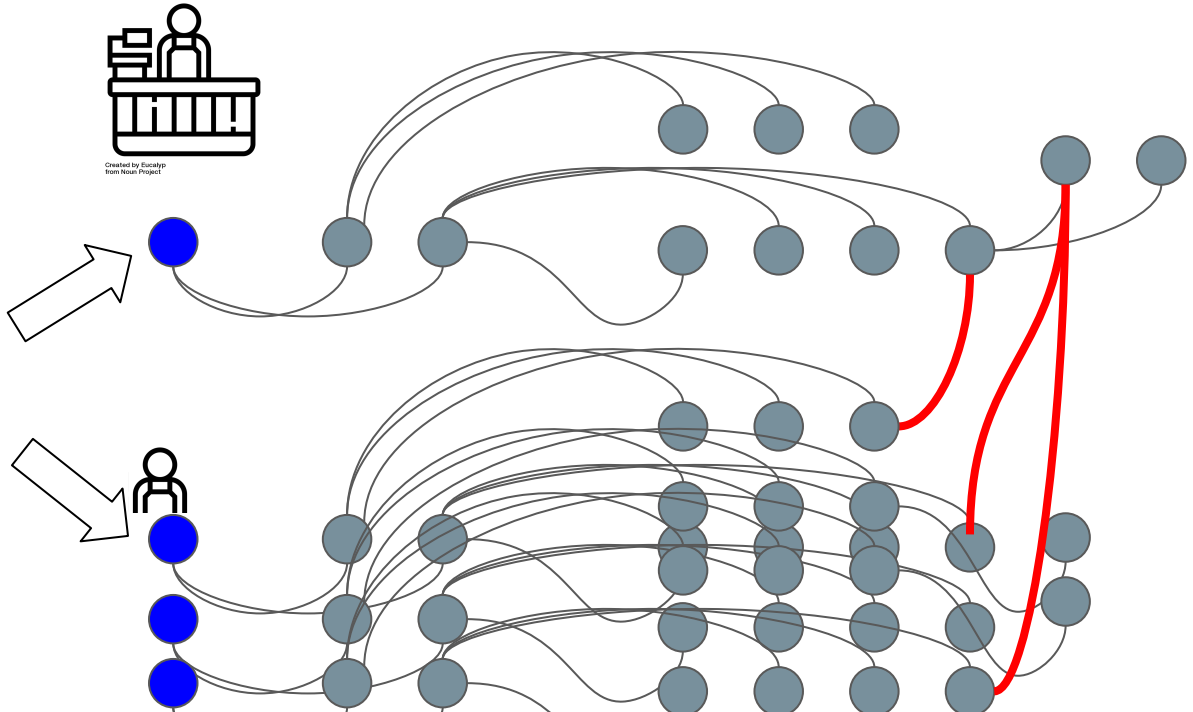

Dust attack: confirming where your money is spent

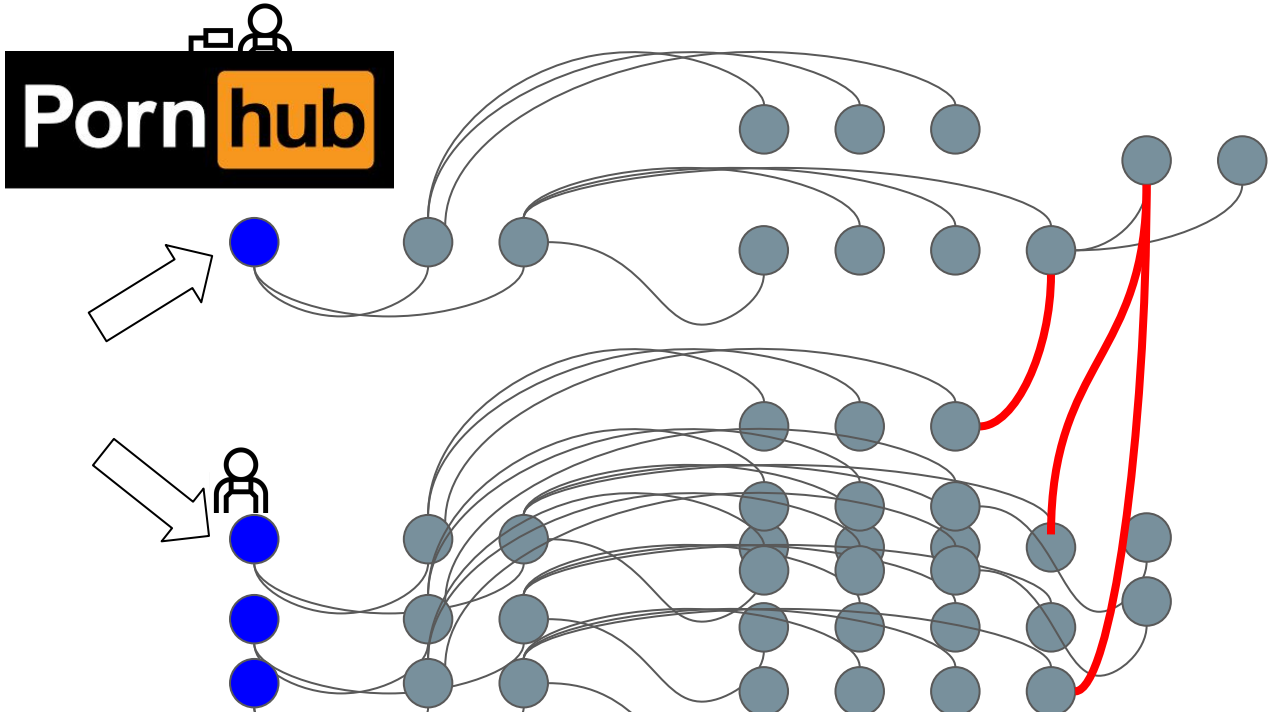

Limitations of decoy approaches:

- Customers can be tracked
  - Use of change transactions
  - Common origins in taint tree
- Anonymous Merchants can be identified
- Third parties can see where your money goes

Privacy approaches: perception

- Bitcoin (vanilla): explicitly identify origin of payment
- Decoy transaction based systems:
  - Pick e.g. 5 decoy source transactions to hide real origin
    - Coinjoin, Mimblewimble, etc. (decoys sampled from current transactions)
    - Cryptonote/RingCT(e.g Monero) (decoys sampled from all of history)
- Zerocoin and Zerocash like (e.g. Zcash, etc)
  - Private transactions have no identified origin.

NOT PRIVATE

PRIVATE

Privacy approaches: reality

NOT PRIVATE

- Bitcoin (vanilla): explicitly identify origin of payment
- Decoy transaction based systems:
  - Pick e.g. 5 decoy source transactions to hide real origin
    - Coinjoin, Mimblewimble, etc. (decoys sampled from current transaction
    - Cryptonote/RingCT(e.g Monero) (decoys sampled from all of history)
- Zerocoin and Zerocash like (e.g. Zcash, etc)
  - Private transactions have no identified origin.

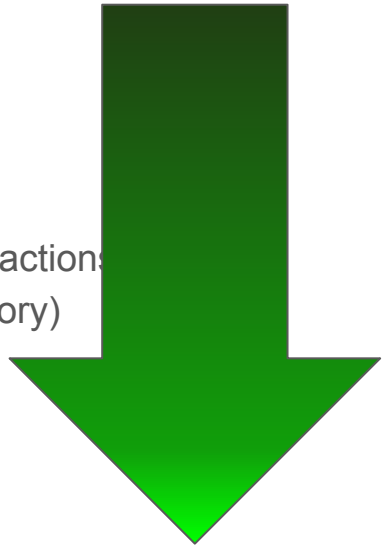

PRIVATE

If you do use decoy schemes

- Decoy systems might work if
  - If your decoy set is very large (i.e. 5,000,000 instead of 5)
  - **Decoy sets substantially overlap across all recent transactions**
  - Decoys are sampled really carefully
- But:
  - We need much more rigorous analysis
  - A careful understanding of when things fail
  - Acknowledge limitations

Scalable decoy schemes

- Cannot have O(decoy set size) sized transactions
- Need logarithmic scaling for size/ transaction generation
- Use a Zerocash style system:
  - Transactions outputs are commitments to (value, recipient address)
  - Merkle Tree over some fraction of UTXO set
  - Zk-proof that origin exists in the UTXO merkle tree.
- Pick a zk-proof technology you like (zkSNARKs, bullet proofs, STARKs, MPC in head,etc)
- Pick a merkle tree depth **d** that the zk circuit is efficient
- Your decoy set is now 2^**d**
- **Somehow sample decoys securely**

Strongly private protocols are getting faster

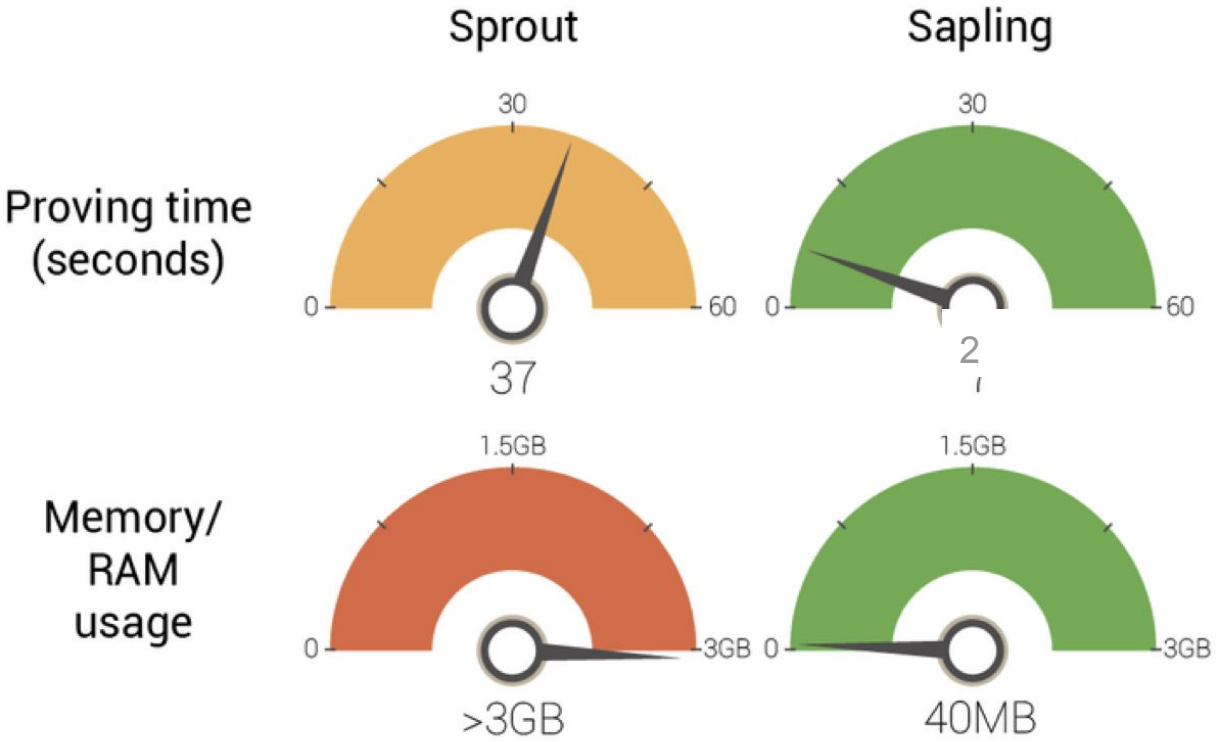

Think critically about scalability vs privacy

- Cryptocurrencies need some privacy solution:
  - Maybe on chain
  - Maybe in layer two
- By all means prioritize scaling over privacy, but understand the limitations of what you have:
  - Your threat model isn’t just passive observers
  - Adding some privacy doesn’t make a protocol private
  - Attacks only get better
- Privacy problems don’t magically go away with small tweaks

Questions?

Ian Miers

@secparam

Cornell Tech/ Zcash foundation