CS 5594: BLOCKCHAIN TECHNOLOGIES

Spring 2024

THANG HOANG, PhD

VIRTUAL MINING

Overview

Proof of Work (Recap)

Proof of Useful Work

Proof of Stake

Proof of Authority

Proof of Elapsed Time

Proof of Burn

…

Proof of Work (Recap)

First consensus algorithm in blockchain

Miners compete against each other to update the blockchain

- Solving hash puzzle
- Purely depending on raw computational power

PoW Issues

Reward and voting power proportional to mining power

Miners with better equipment get more reward

<u>Power usage and resource wasteful!!!</u>

Need high computation power

Incur high energy consumption

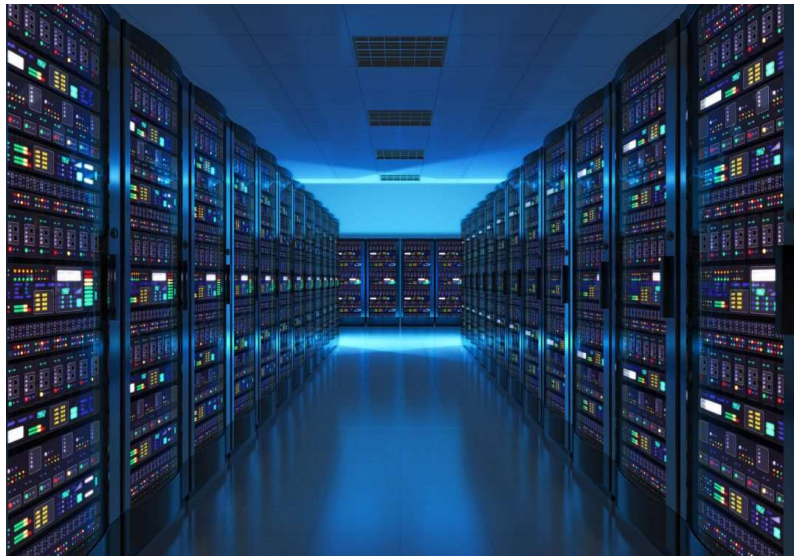

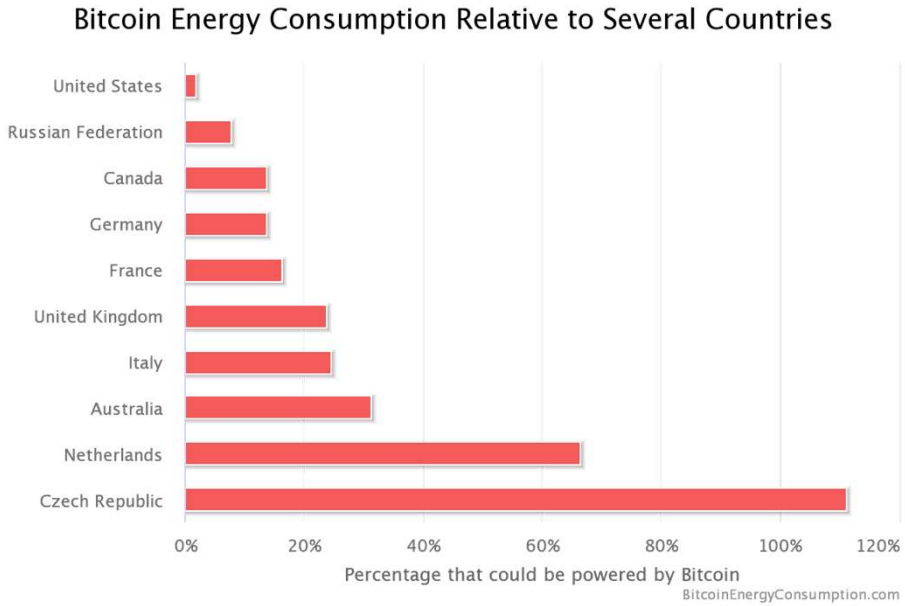

Bitcoin devours more electricity than many other countries!

Image credited to https://digiconomist.net/bitcoin-energy-consumption/

Proof of Useful Work

Spending computing resource to solve hash puzzle is wasteful

Can we recycle this to do something more meaningful?

- Protein folding? Find a low energy configuration
- Search for aliens? Find anomalies in radio signals
- Break crypto?

Challenges

- Randomly chosen instances must be hard!

Project	Founded	Goal	Impact
Great Internet Mersenne Prime Search	1996	Finding large Mersenne primes	Found new largest prime number twelve straight times, including 2^57885161 -1
Distributed.net	1997	Crypto brute-force demo	First successful public brute force of a 64-bit crypto key
SETI@home	1999	Identifying signs of extraterrestrial life	5 mil. Participants
Folding@home	2000	Atomic-level simulation of protein folding	118 scientific papers

Primecoin

Find sequence of large prime numbers

Cunningham chain:

$p_1, p_2, \ldots p_n$ where $p_i = 2^i a + 1$

Each p_i is a large (probable) prime

p_1 is divisible by H(prev || mrkl_root || nonce)

(2, 5, 11, 23, 47) is a Cunningham chain of length 5

Many large known Cunningham chains actually came from Primecoin miners

106680560818292299253267832484567360951928953599522783616513856655224358588041233 92×61# − 1 (2014, block #368051)

Hard problem?

Usefulness?

Permacoin

Proof of Storage / Proof of Retrievability

Replicated storage system

Each user stores a random subset of an extremely large file **F**

1. Build a Merkle tree, where each leaf is a segment of **F**

2. Generate a public signing key PK determining a random subset of file segments

3. Each mining attempt:

a) Select a random nonce
b) h1 := H(prev || mrkl_root || PK || nonce)

c) h1 identifies *k* segments from subset

d) h2 := H(prev || mrkl_root || PK || nonce || F_x)
e) Win if h2 < TARGET

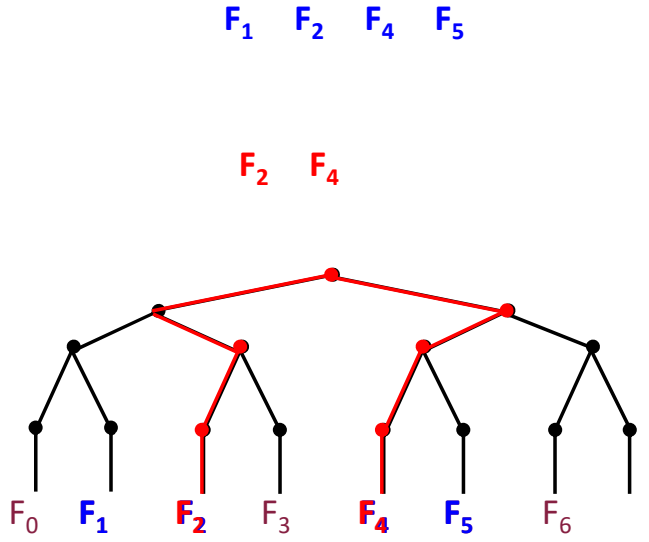

PoW Issues

Accessibility: Reward and voting power proportional to mining power

- High barrier to entry to becoming a miner
- Miners with better equipment get more reward

Scalability: Long waiting time for blocks to be verified

- Get worse if number of TXs increases

PoW Issues

Centralization: Mining pools make blockchain somewhat more centralized

Miners create a mining pool to combine computing powers and share profits

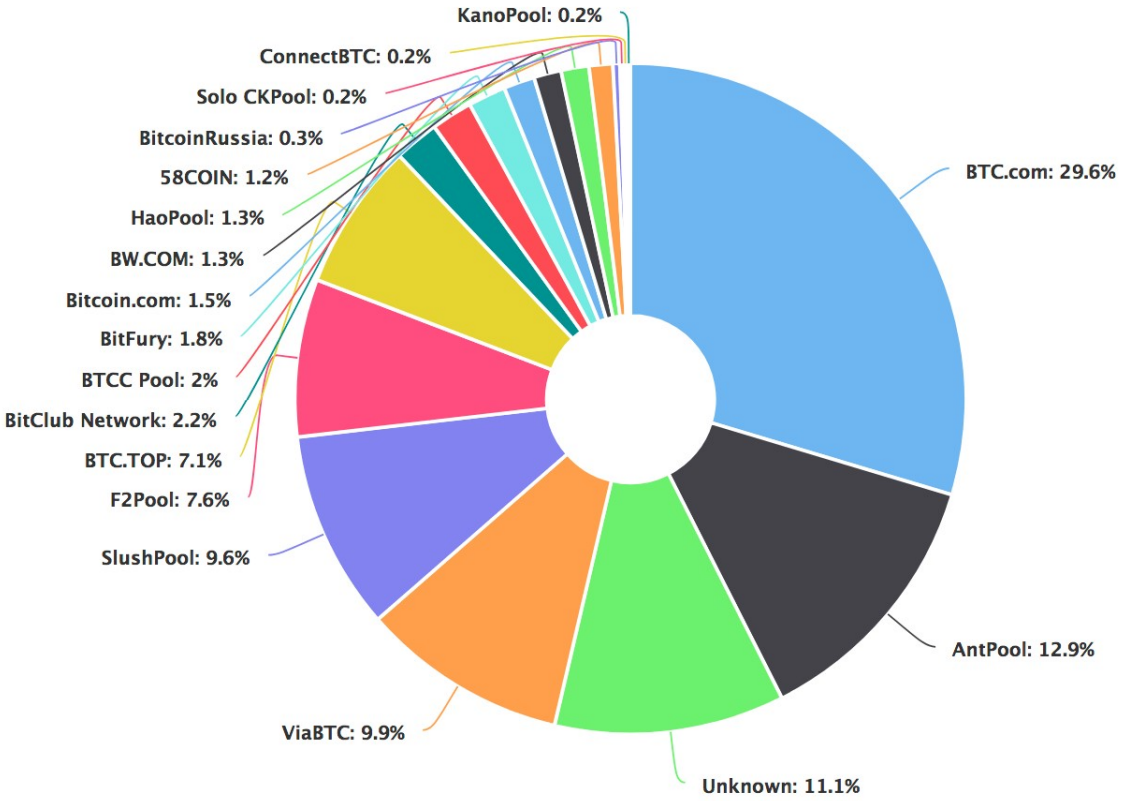

Image credited to https://medium.com/otrovio/bitcoin-mining-pools-52e2926dd50e

51% attack can be realistic if three biggest mining pools combine!

PoW Mining

What is the real implication of spending money on power and equipment?

Can we remove that step?

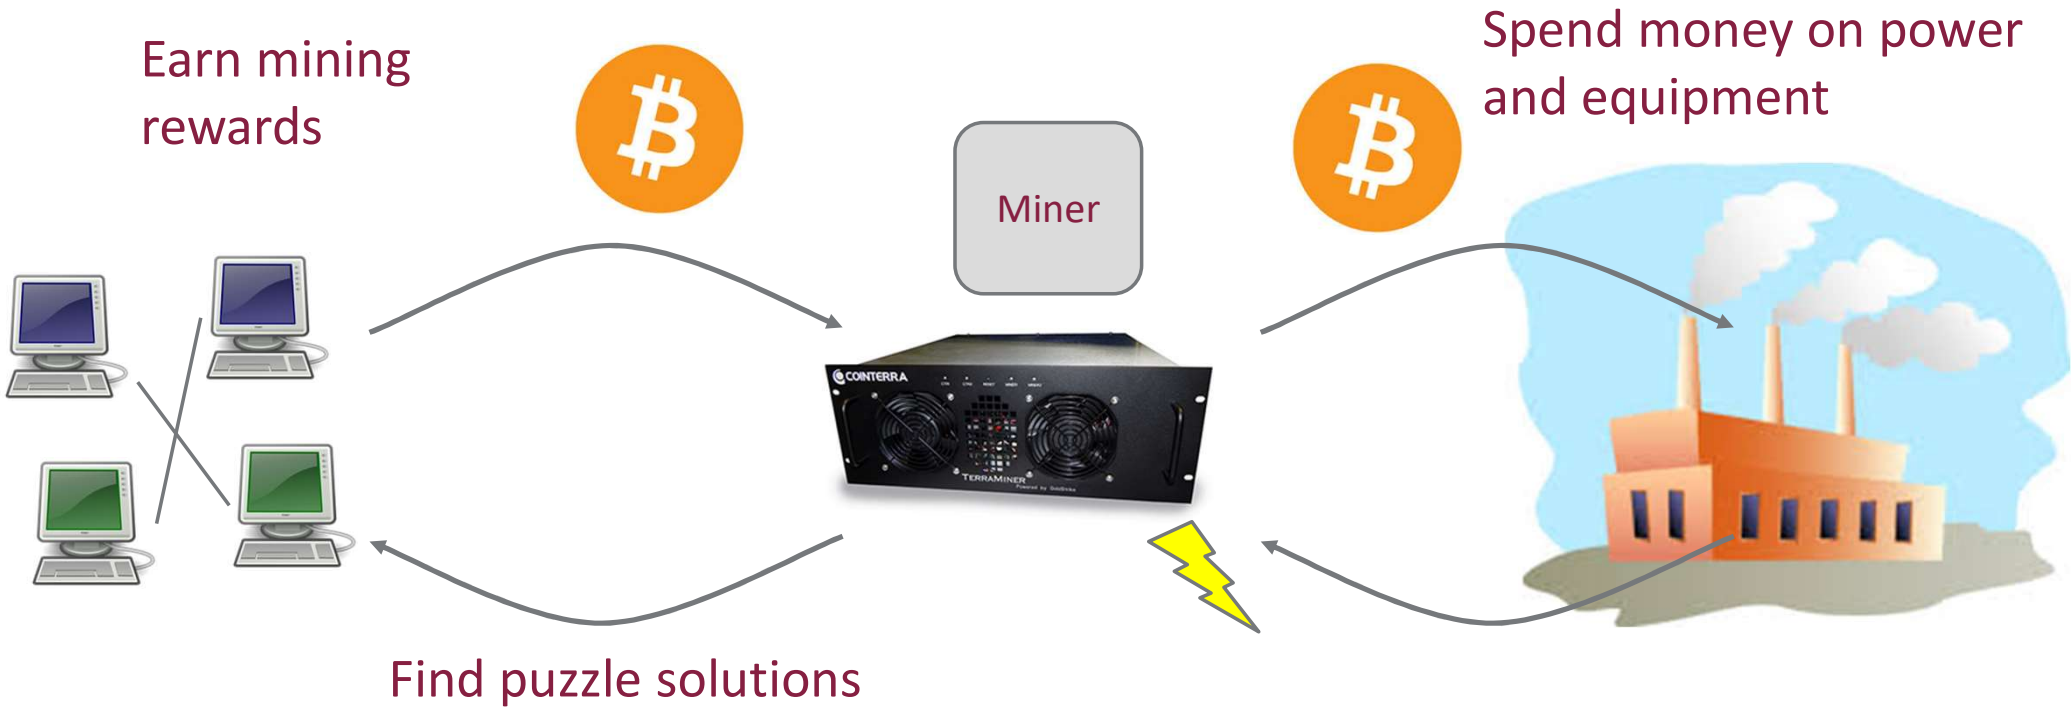

Virtual Mining

Virtual Mining

Allocate *mining power* directly to all currency holders in proportion to the resource (X) they hold

“Proof-of-X”, where X = {Stake, Deposit, Activity,…}

Winners chosen at random by lottery

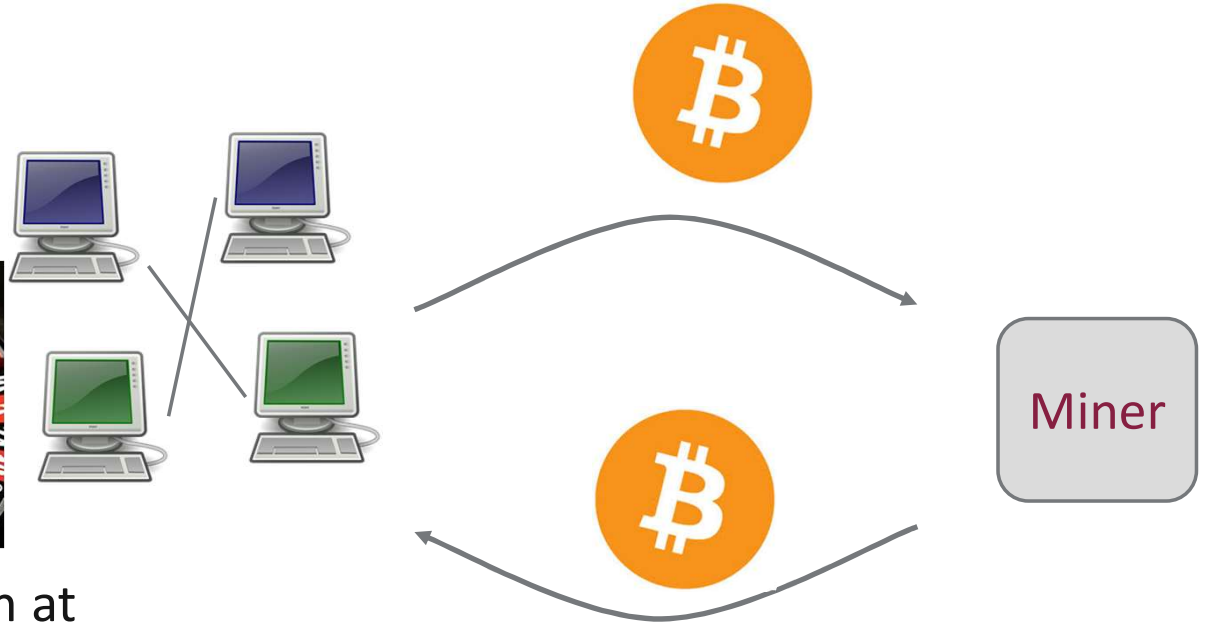

“Mine” by sending money to a special address

Virtual Mining Benefits

Lower overall costs

- No harm to the environment
- Savings distributed to all coin holders

Stakeholder incentives

No ASIC advantage

51% attack is even harder

Proof of Stake (PoS)

First proposed by a user named *QuantumMechanic* in 2011

Goal: Making blockchain more sustainable

Instead of competing computing power, validators (miners) are chosen based on their stake

- Attributing mining power to the proportion of coins (stake) held by the stakeholders

The chance of **minting/forging** a block rely on how much of a stake (coins) the validator owns

- If Alice own 1% of coins, she can mint/forge 1% of all of transactions in the network

PoS vs. PoW

- Vote ~ computing power
- Miner / mining
- New coin created
- Competition between miners
- Slow (~10 mins)

- Vote ~ stakes / coins
- Validator / Forging or minting
- No new coin created
- Deterministic validator selection
- Fast (< 1min)

PoS vs. PoW

Better energy efficiency

- No need to use lots of energy for mining blocks

Availability: Lower barriers to entry, reduced hardware requirements

- No need of elite hardware to stand a chance of mining blocks
- Anyone who holds the base coin(s) can become the miners

Stronger immunity against centralization (debatable!?)

- PoS (in theory) should lead to more nodes in the network

How PoS Works?

Anyone can become a validator by depositing a certain number of coins into the network (security deposit)

Everyone has a certain chance to be selected as validator for next round

Proportional to account balance

```
SHA256(prevhash + address + timestamp) <= 2^256 * balance / diff
```

Undesirable centralization

Permanent advantage for richest guys

Alternative Validator Selection

Alternative 1: Coin Age

Balance multiplied by the number of days the coins have been held

Once a stake of coins is used, starts over with *zero* coin age

```
SHA256(prevhash + address + timestamp) <= 2^256 * age * balance/ diff
```

Alternative 2: Proof of Deposit

Mirror with Coin Age

Rewards to those who are willing to keep coins unspent for long time into the future

How PoS Works?

Selected validator mint/forge a block by selecting transactions from transaction pool

Selected validator verified transactions in the block and collect transaction fee

- No block reward as in PoW
- May lose part of stake if verify fraud transactions
  - If stake > TX fee, validator likely does the job honestly

Verified block appended to blockchain

Proof of Stake (PoS)

If stopped being a validator

- Stake and transaction fees will be released after a certain period of time

PoS needs bootstrapping

- PoS only works if some nodes established stakes in the network
- Bootstrapping to initiate consensus committee
  - Sale initial coins
  - Start with PoW then transition to PoS

51% Attacks

51% attack: If one can buy majority in the network, they can compromise it

51% * (Bitcoin market cap = 600 billion U.S dollars)

Seem impractical

Need 300 billion dollars to compromise the network

Nothing at Stake Attack

Also called **Stake Grinding Attack**

- Nothing to lose from behaving badly
- Continue to participate on longest chain while simultaneously attempting to fork the chain
- Validator validates on every branch of the chain to optimize reward regardless of the outcome of the fork

Countermeasure

- Wrong voting penalty

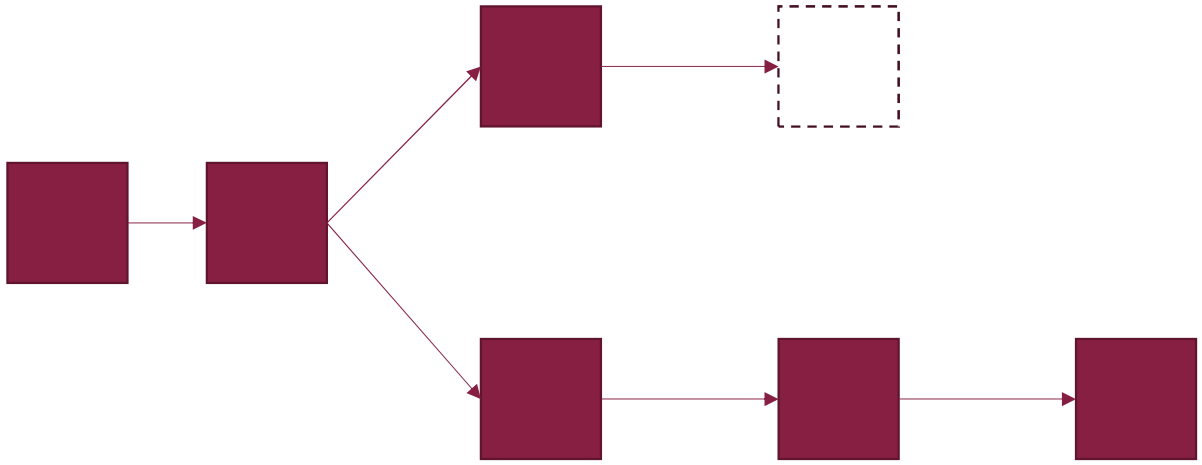

Long Range Attack

Also called **History Revision Attack**

Adversary forks the chain at the (past) point where it had large stake in the network and starts forging new blocks

Overtakes the main chain after some time

Longest chain rule not enough to counter the attack

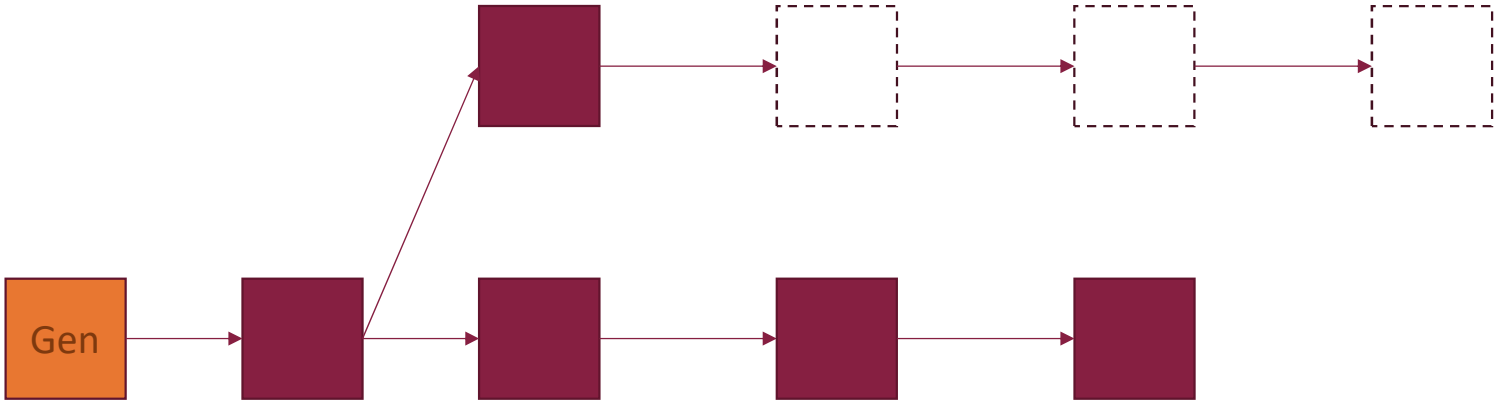

https://eprint.iacr.org/2018/248.pdf

Long Range Attack

Assume there are three validators: Alice Bob and Malory each having 1/3 stakes

Simplest attack

Malory goes back to genesis block, forks the chain and mint block on its branch

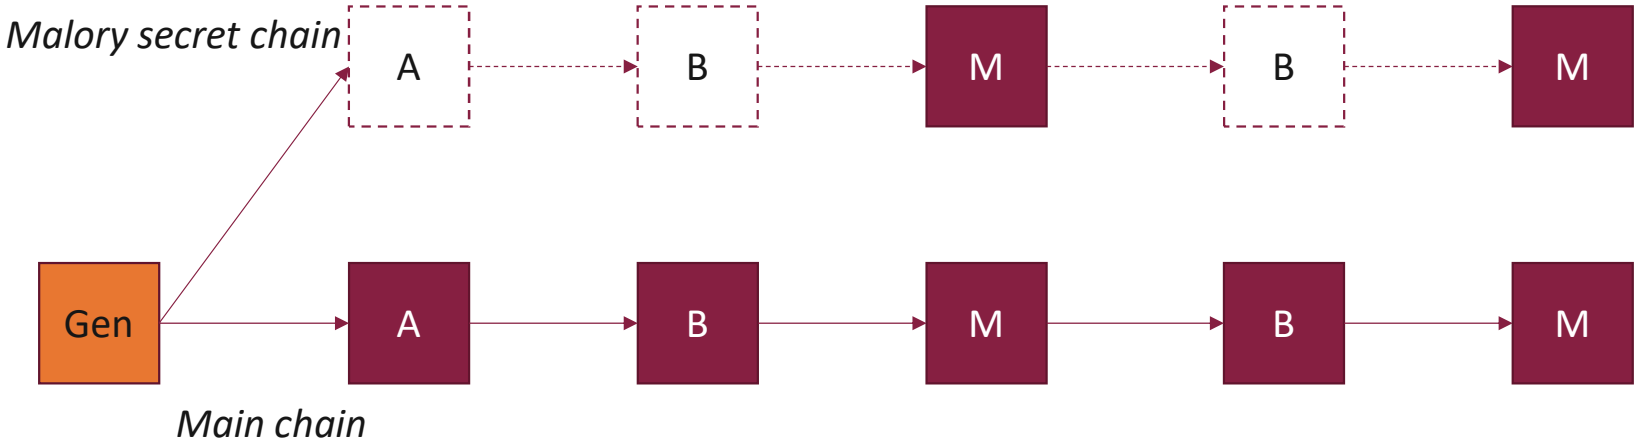

Long Range Attack

Assume there are three validators: Alice Bob and Malory each having 1/3 stakes

Simplest attack

Malory goes back to genesis block, forks the chain and mint blocks on its branch

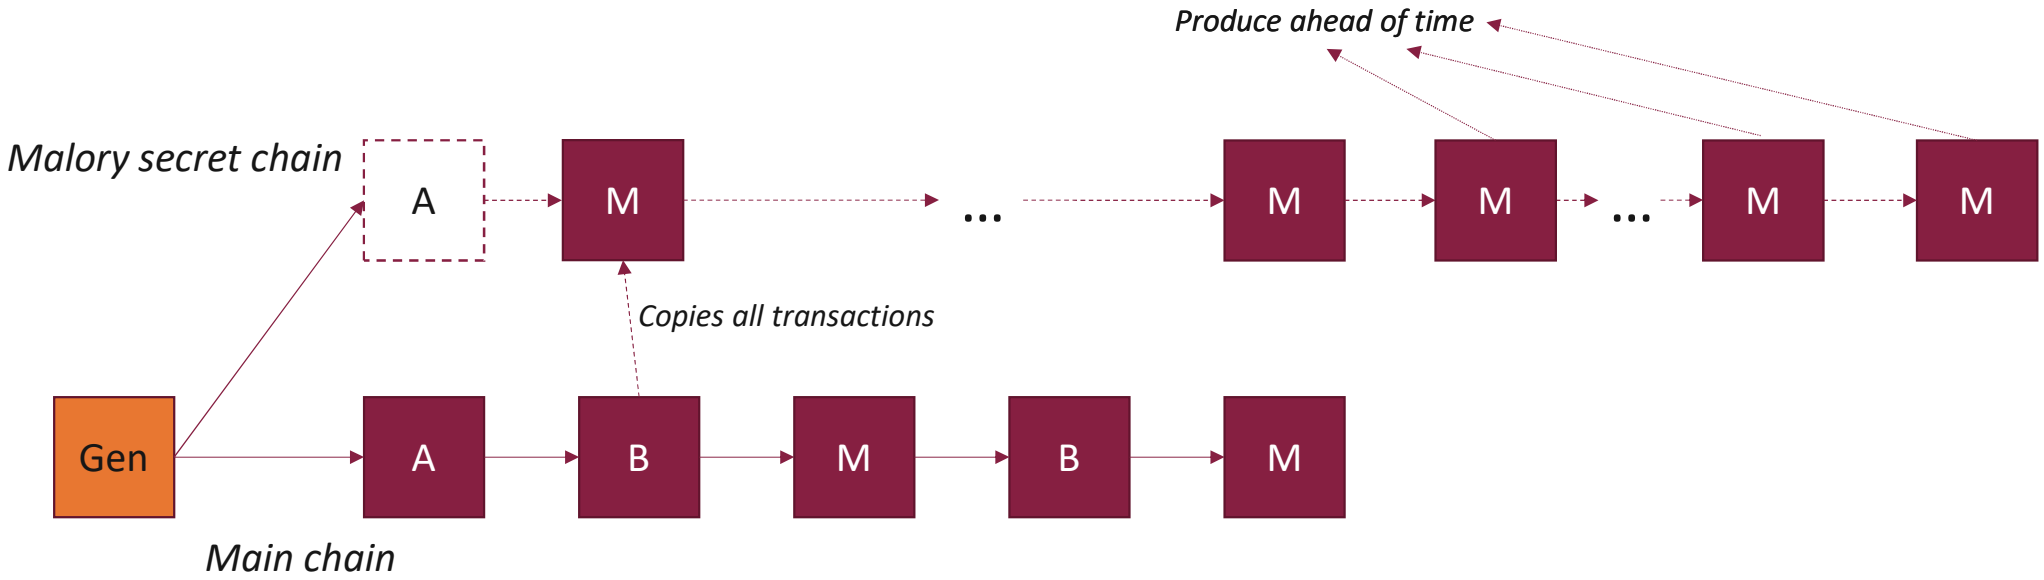

Countermeasure

Timestamping every block to reject chains with timestamp far ahead of time

Long Range Attack

Assume there are three validators: Alice Bob and Malory each having 1/3 stakes

Posterior Corruption

Bob retired and Malory corrupted Bob's private key

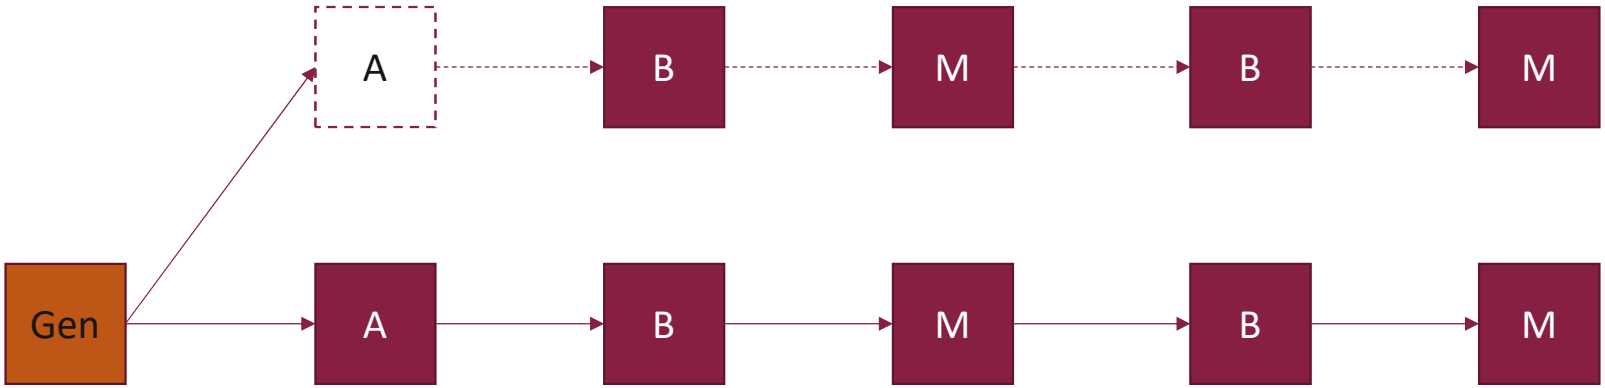

Countermeasure

Key Evolving Signatures, Moving Checkpoints

Long Range Attack

Assume there are three validators: Alice Bob and Malory each having 1/3 stakes

Stake Bleeding

Malory stalls the main chain and works on its branch

Lose stake in main chain, but start to increase stake in forked branch

Malory copies transactions off the main chain

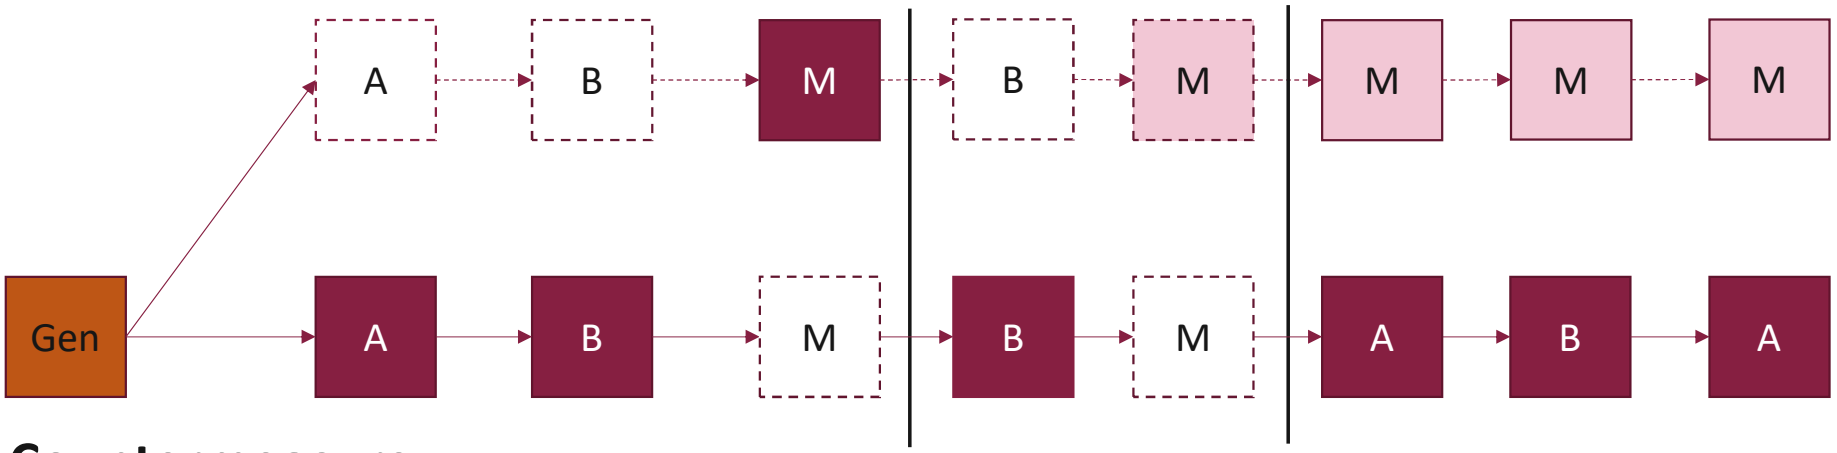

Countermeasure

Moving Checkpoints, Plenitude Rule, Context Awareness Transactions

Long Range Attacks

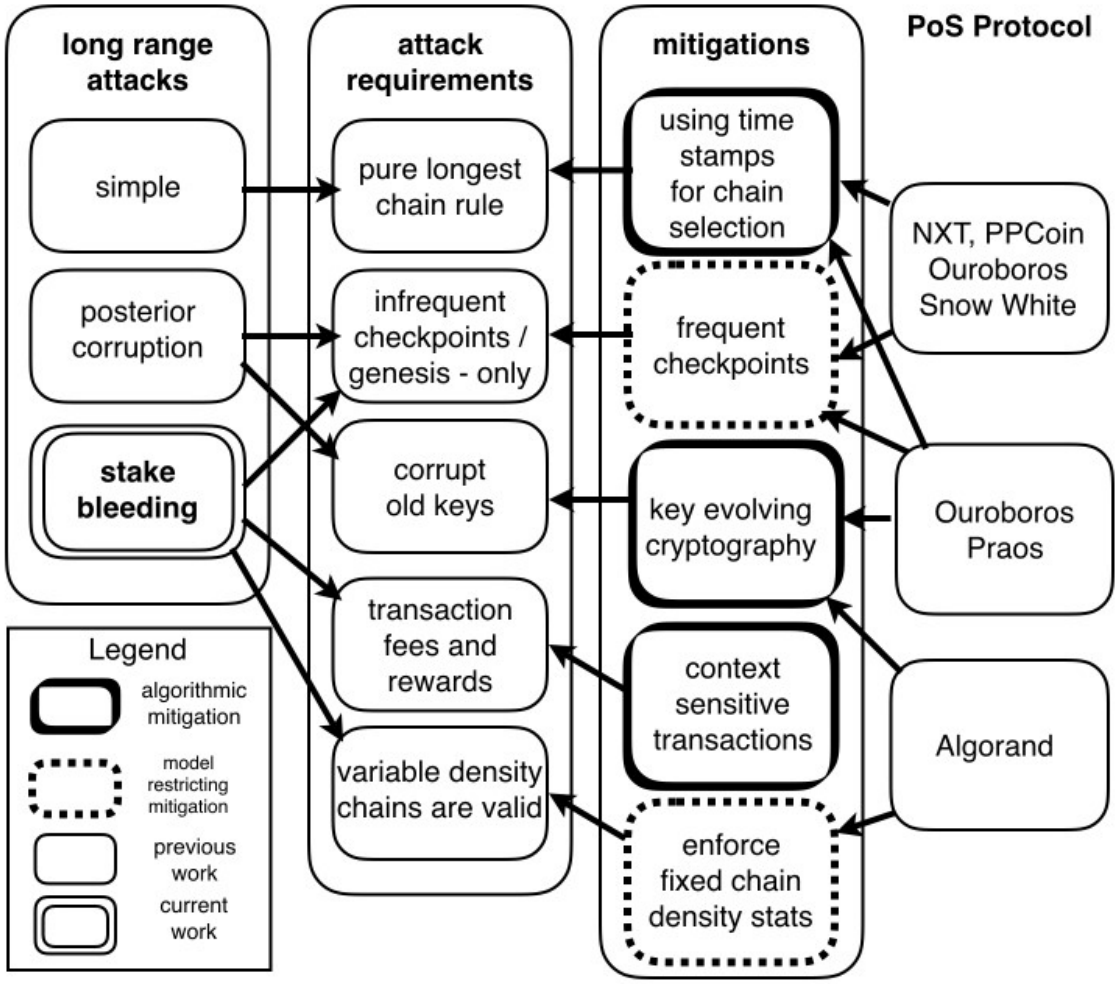

Image credited to https://eprint.iacr.org/2018/248.pdf

Delegated PoS

Evolution of PoS

Nodes elect witnesses (delegates) to validate the next block

Can withdraw their vote in case of improper witness's behavior

Rely on a group of delegates to validate blocks on behalf of all nodes in the network

Scalable, energy efficient, lower transaction fee

Somewhat centralized

Use-case: EOS, Bitshares, Steemit

Proof of Authority

Identity (instead of coins) as stake

- Somewhat centralized
- More suitable for private blockchain

Stakes social capital rather than financial capital

Nodes stake their reputation

Proof of Elapsed Time (PoET)

Nodes need to be identified and authorized to participate in the network (permissioned/private blockchain)

Trusted Execution Environment (TEE) for fair lottery

- Each node generates a (pseudo) random number for how long it must wait
- PRNG with secure hardware (e.g., Intel SGX)

Use-cases: Hyperledger Sawtooth

Proof of Burn

Bootstrap one cryptocurrency from another

Burn coins by sending them to a verifiably unspendable address (eater address)

<u>Use Case:</u>

- Counterparty
  - User transfers Bitcoin to an eater address and receive tokens in return
- Slimecoin

Open Questions with Virtual Mining

Virtual mining remains somewhat controversial

Is burning real resource really needed for security ?

- If yes, then *waste* is the cost to provide security for the system
- If not, PoW may not be necessary
- Both cases not proven though

Conclusion

Many possible design goals

- Prevent ASIC miners from dominating
- Prevent large pools from dominating
- Intrinsic usefulness
  - Eliminate the need of mining hardware at all

So far none of the alternatives demonstrated theoretical soundness and practical adaptation

Best tradeoff unclear