TRUTH JUSTICE ALGOS

Cryptocurrencies:
Transaction fees, Pools and PoS

Teachers: Ariel Procaccia and Alex Psomas (this time)

TODAY'S MENU

- Miner's dilemma
- Instability without the transaction fees
- Proof of Stake

MINING POOLS

The miner's dilemma [Eyal 15]

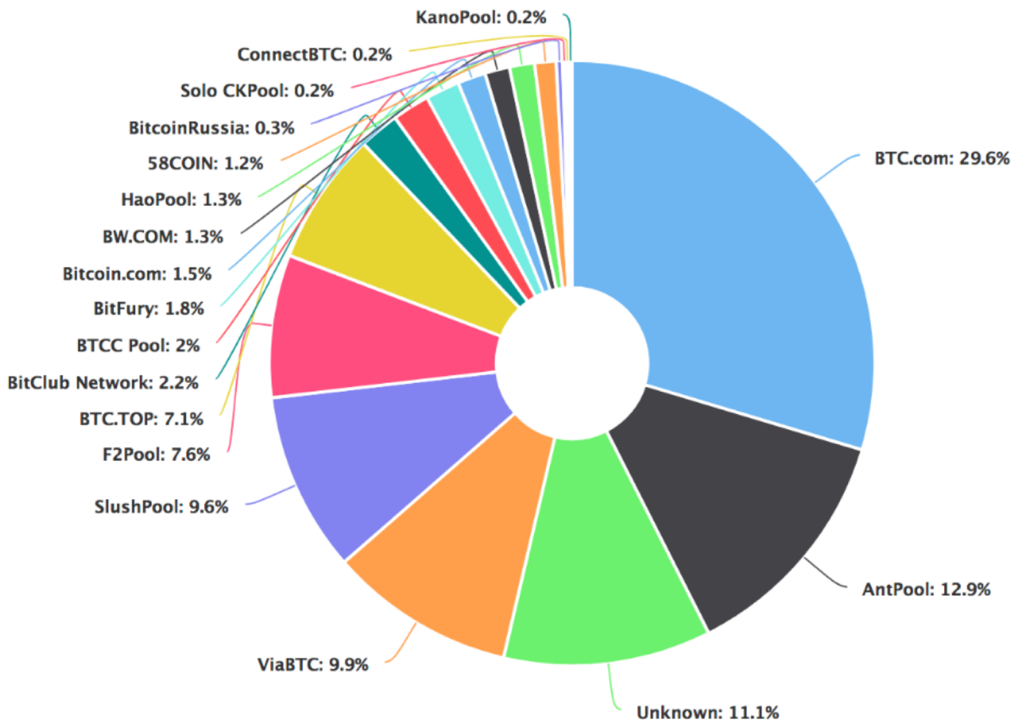

June 2018

MINING POOLS

- How pools work:
  1. Manager giver her ID, ManagerID, to all participants
  2. Participants try to find valid block with minerID=ManagerID
  3. Participants send "partial proofs" to manager
  4. Manager maintains "shares" in order to compute the contribution of each participant
  5. Manager distributes rewards (at say the end of the week) according to f($shares$)
     - Designing a good reward function is tricky [SBBR16] but don't worry about it for now

MINING POOLS

- Attack:
  - Send all “partial proofs”, but throw away actual blocks
- Sanity check: this does not actually “steal” blocks, because they are made with minerID = ManagerID
- This definitely hurts the pool
- It also hurts the attacker’s (pool) rewards
- Main result: if the attacker attacks and mines in the right proportions, then this is overall profitable

MINING POOLS

- Two pools, each has 50% of the total power

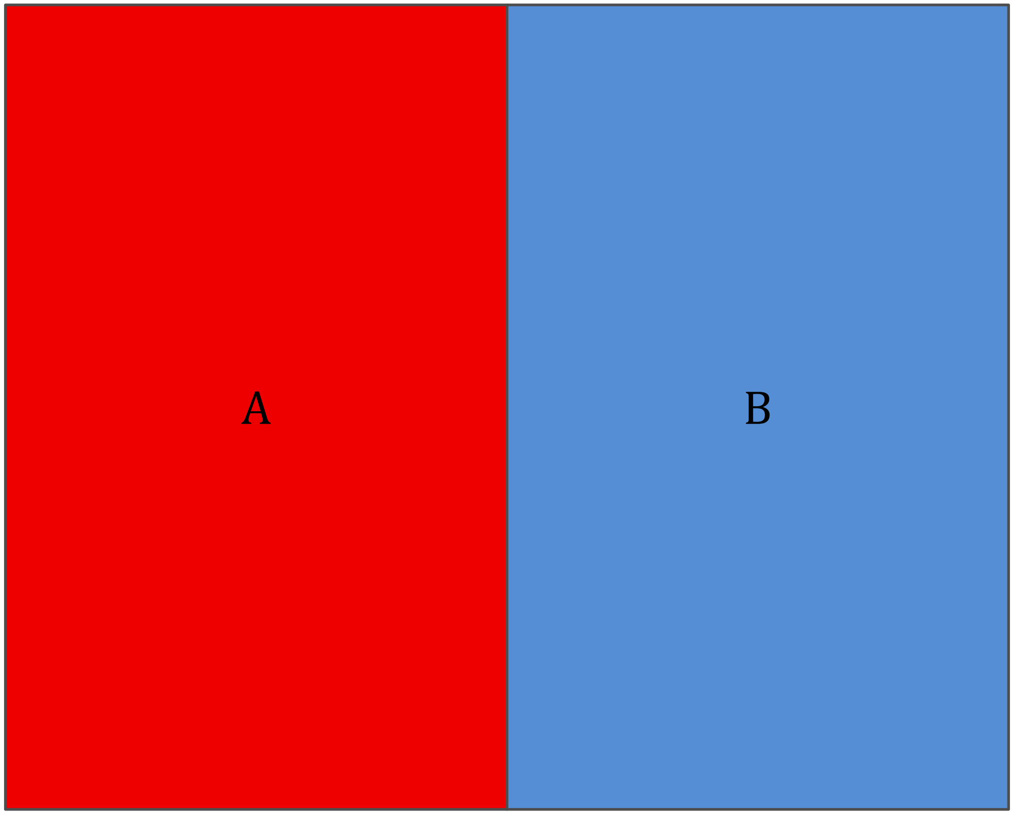

MINING POOLS

- Pool A attacks pool B with half of its mining power
- Pool A makes $\frac{1}{4}/\frac{3}{4} = 1/3$ of the total valid blocks
  - While pool B makes 2/3 of the total valid blocks
- Pool A also gets 1/3 of pool B's rewards
  - $\frac{1}{4}/(\frac{1}{4} + \frac{1}{2})$
- Therefore, pool A makes $\frac{1}{3} + \frac{1}{3} \cdot \frac{2}{3} = \frac{5}{9}$ of the total reward
  - As a bonus, participants get more "bang-per-buck" by joining A

MINING POOLS

- [Eyal 15] shows that this attack is profitable no matter how many pools and no matter the size!
  - "No attacking" is not an equilibrium
- In his model, the game between two pools reduces to a Prisoner's dilemma type of game, where "attack" is always a dominant strategy
  - Perhaps good news: the game is not played once, so cooperation could be a stable state

TRANSACTION FEES

TRANSACTION FEES

- Currently in Bitcoin, most of the mining rewards come from the block reward
  - Transaction fees are so small that it is reasonable for them to be 0 in an analysis of incentives in Bitcoin
- Plan: half the block reward every four years. Eventually all of the rewards will come from transaction fees
- Belief: "It doesn't matter if you make 12.5 bitcoins via block rewards or 12.5 bitcoins in expectation via transaction fees"
- Punchline: it does

SETUP

- Every miner has mining power x(m) with $\sum_m x(m) = 1$
- At all times miner m is aware of the whole tree $G(m)$
- Total of t transaction fees arrive in the interval $[0, t]$ for all t

GAME

- At Poisson clock with rate 1, miner m selected to mine block proportionally to $x(m)$
  - m creates a node B, points to any node in $G(m)$
  - Includes fees $F(B)$ subject to $\sum_{B' \in predecessor(B)} F(B') \leq t$
- Each time step every miner may broadcast any nodes in $G(m)$
- Game stops at time T

PROTOCOLS

- Longest Chain
  - Whenever selected to build a block point to the furthest node in $G(m)$
    - Break ties in favor of what you heard first
  - Include maximum possible transaction fees
  - Broadcast everything
- Petty Longest Chain
  - Whenever selected to build a block point to the furthest node in $G(m)$
    - Break ties in favor of most available fees
  - Include maximum possible transaction fees
  - Broadcast everything

PROTOCOLS

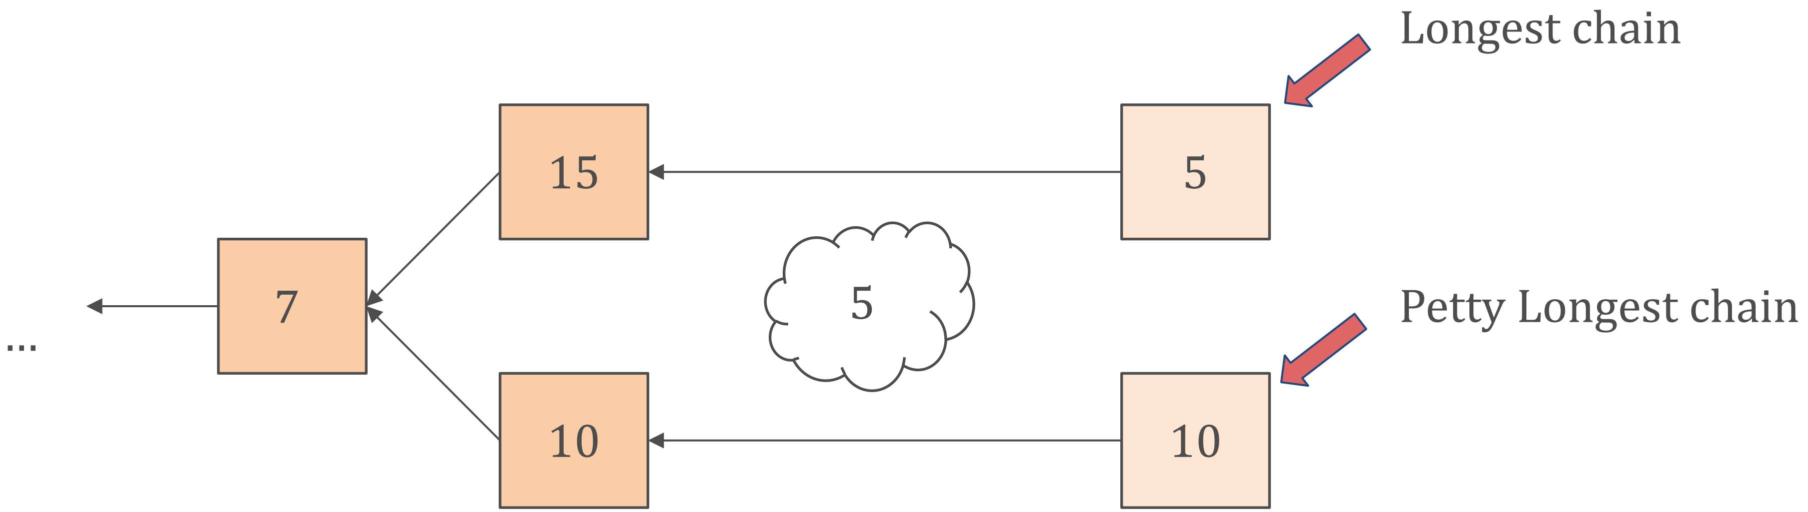

- What if you know others are using Petty Longest Chain?
- Extending Longest Chain gives 5

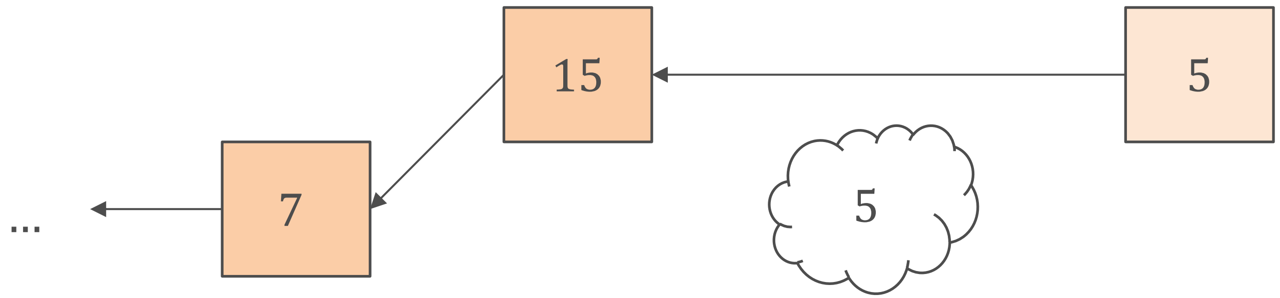

- What if you know others are using Petty Longest Chain?
- Extending Longest Chain gives 5
- Instead, build a new block and leave out some transactions!
    - (This made no sense when we had just block rewards)

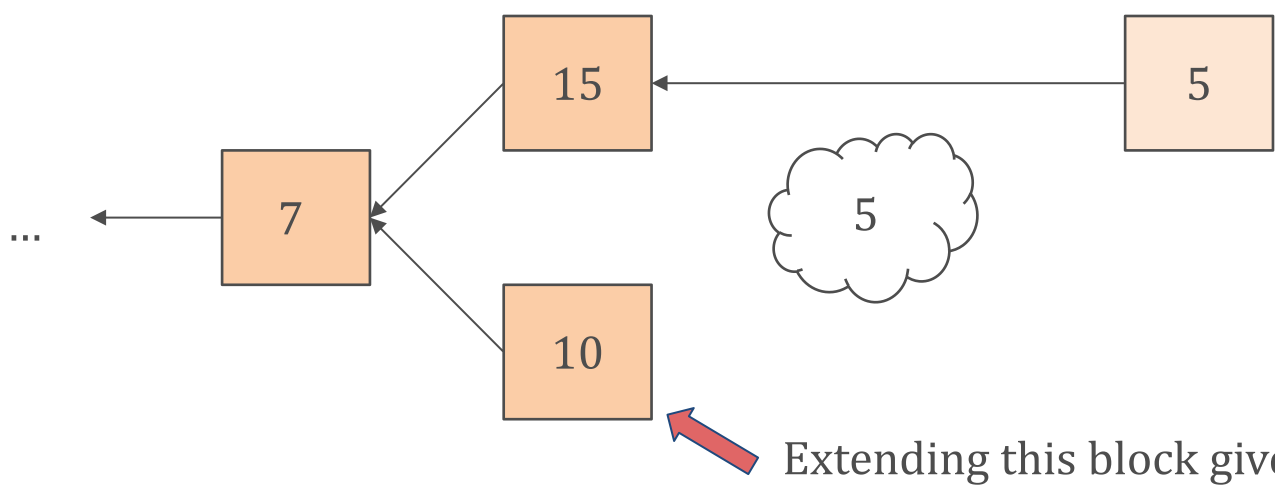

Extending this block gives 10, so everyone who uses Petty Longest Chain will extend this!

Theorem (informal): Undercutting (actively forking) is an equilibrium. Furthermore, there is a backlog of transactions (of size $\Theta(\sqrt{t})$)

Theorem (informal): Even if 2/3 of the miners play "honestly", it's still profitable to undercut

MORE PROBLEMS

- Currently, no incentives to broadcast transactions only you know about [BDOZ 12]
  - Incentives similar to MIT's DARPA red ballon challenge solution
  - [BDOZ12] give an incentive compatible and "Sybil-proof" reward scheme

TAKE AWAY

- Switching to all rewards coming from transactions creates new kind of incentive issues

PROOF OF STAKE

PROOF OF STAKE

- Proof of Work:
  - Random miner selected with probability proportional to their computational power
  - “One CPU, one vote”
- Proof of Stake:
  - Random miner selected with probability proportional to wealth rather than computational power
  - “One coin, one vote”

TODAY

1. A model for PoS cryptocurrencies
2. A set of properties such that every protocol in the model satisfies at least one property
3. An attack for each property

MODEL

Proof of Stake protocol blueprint

1. Protocol specifies an existing block
2. Protocol uses some method to pick a coin
3. Owner of the coin gets to add a new valid block of transactions on top of the existing block
4. Repeat

MODEL

Protocol

A Proof-of-Stake protocol is defined by two functions

1. A *validating function* V which takes as input a block and outputs 0 or 1
2. A *mining function* M which takes as input a block B, a coin c and a timestamp t, and outputs a valid block that extends B (if one exists)

- V should be efficiently computable by everyone
- M should be efficiently computable by the owner of c

MODEL

$$M(A, c, t) = B$$

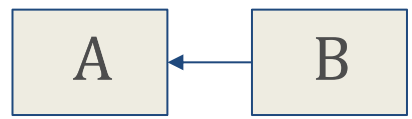

$$V(B) = 1$$

MODEL

Assumptions

1. Chain Dependence: Validity of block B at time t depends only on t and the predecessors of B
2. Monotonicity: If B is valid at time t then it is valid at all future times $t' > t$

- Without them an attacker can withhold messages to convince a victim invalid blocks are in fact valid (Eclipse attack).

MODEL

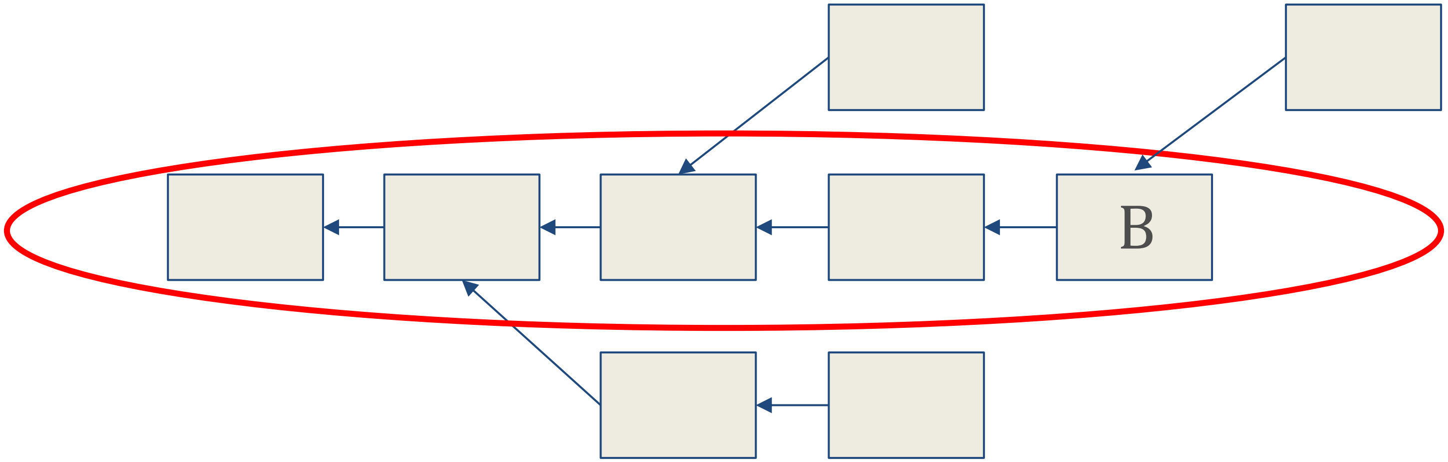

MODEL

Longest-Chain Protocol

A Longest-Chain protocol has a scoring functions S which takes as input a block and outputs a monotone increasing score:

If A is the predecessor of B then $S(A) < S(B)$

Miners are supposed to mine on top of A maximizing $S(A)$

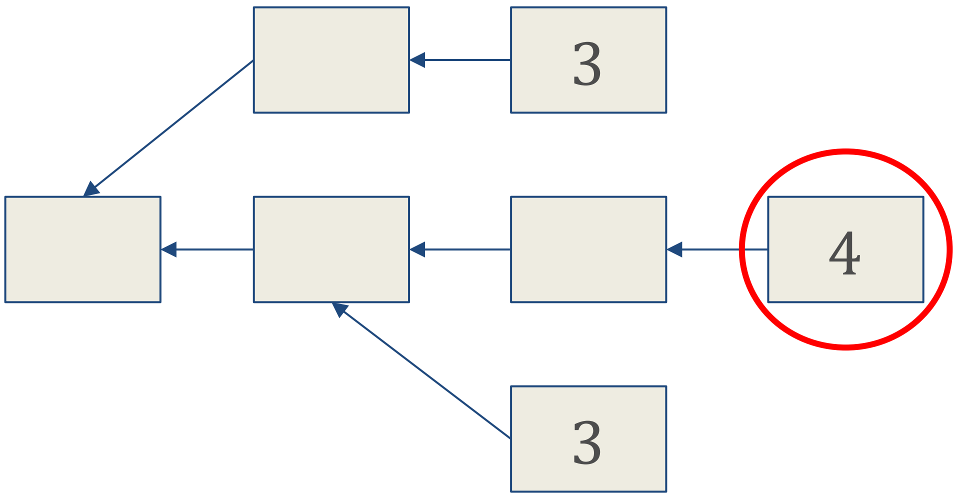

3
4
3

PROPERTIES

D-Locally Predictable

For a coin c, $Owner(c)$ can efficiently predict D blocks in advance if she is eligible to use c to mine a block

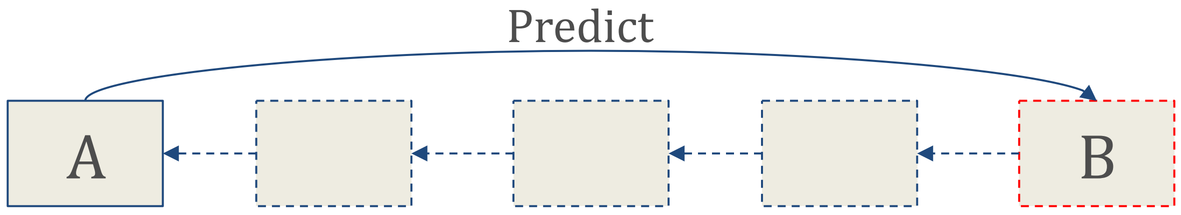

PROPERTIES

Observation

Every Proof-of-Stake protocol is 1-locally predictable

Proof. Just use the mining function M to efficiently predict whether you can mine the next block.

PROPERTIES

D-Globally Predictable

For a coin c,every protocol participant can efficiently predict D blocks in advance if $Owner(c)$ is eligible to use c to mine a block

Example: Let T be a threshold and H a hash function.

$$V(B) = 1 \Leftrightarrow H(c(B), t(B)) < T$$

PROPERTIES

D-Recent

The negation of D-locally predictable. $Owner(c)$ **cannot** efficiently predict D blocks in advance if she is eligible to use c to mine a block

Therefore, eligibility to mine a block depends on “recent history”

ATTACKS

Predictable Selfish Mining

Withhold a newly mined block B and secretly try to mine on top of it.

If you mine another block B', then you have the longest chain, even if other miners mine a block on $Pred(B)$

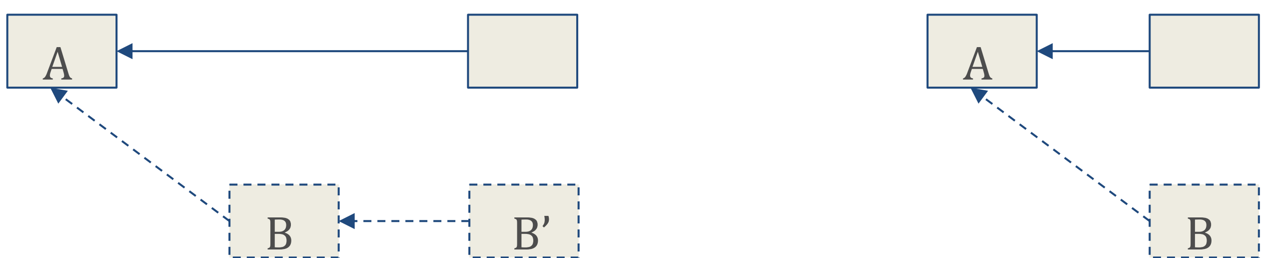

Attack Succeeds

Attack Fails

ATTACKS

Predictable Selfish Mining

- With global predictability there is no risk!
- Can predict precisely when you are able to mine k blocks faster than the rest of the miners

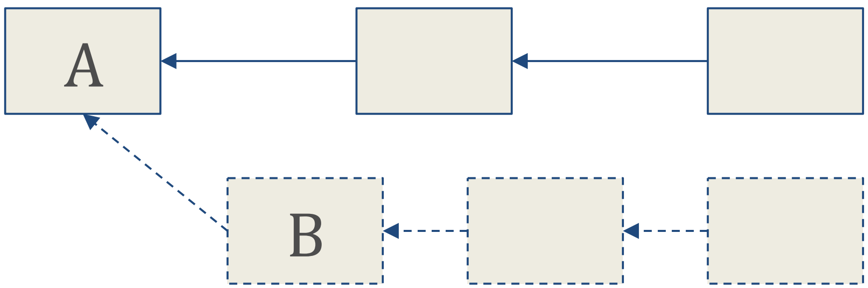

Launch Attack

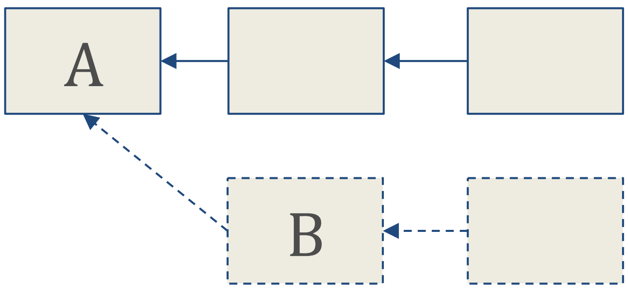

Abort Attack

ATTACKS

Predictable Selfish Mining

- Predict precisely how fast you will mine k blocks and then compare to the average rate
  - Even with 1-Local Predictability there is reduced risk

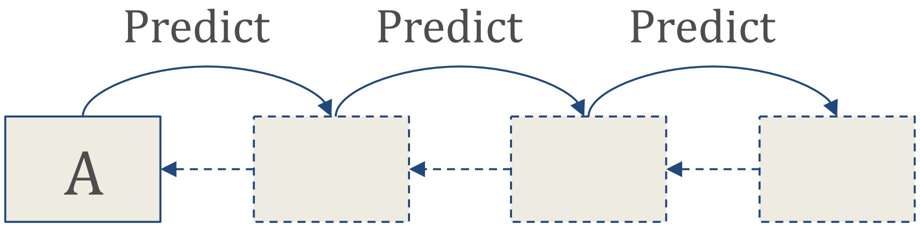

ATTACKS

Predictable Double Spending

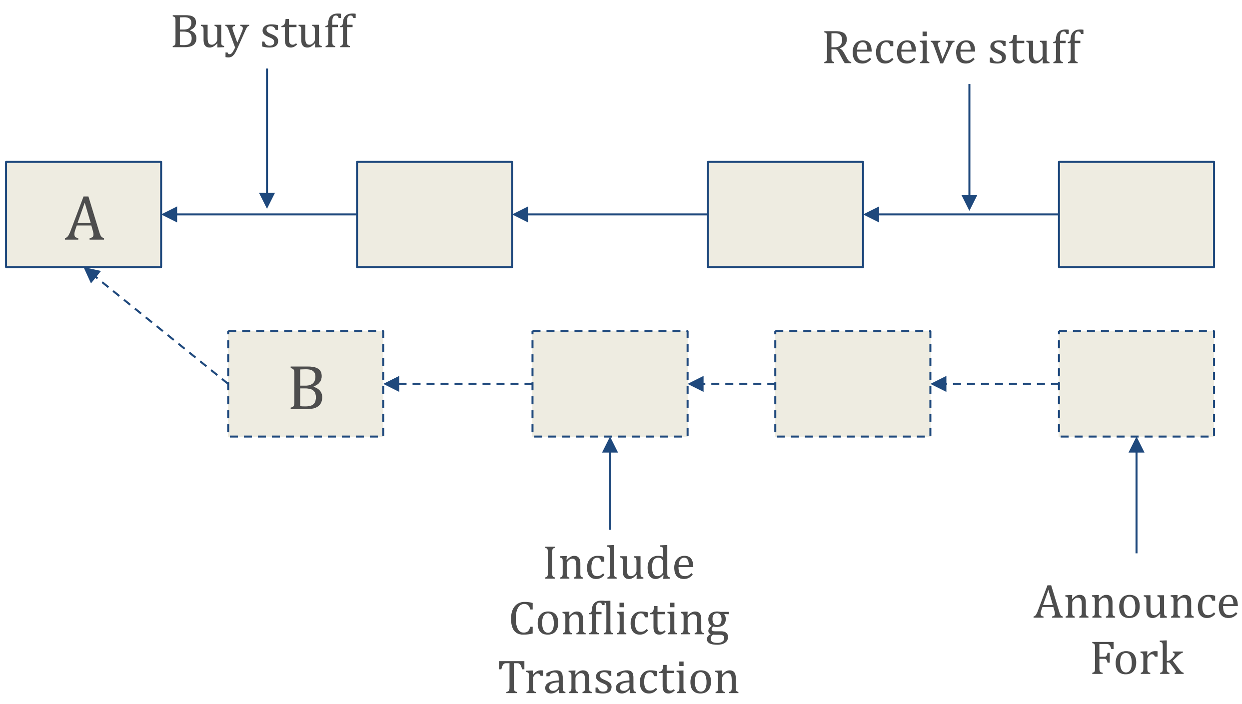

ATTACKS

Undetectable Nothing-at-Stake

- For D-Recent protocols, blocks A and B at the two ends of a length D fork are “independent”
    - A coin could “win” in A and “lose” in B
- Attempting to mine on both sides of the fork doubles your chances of successfully mining

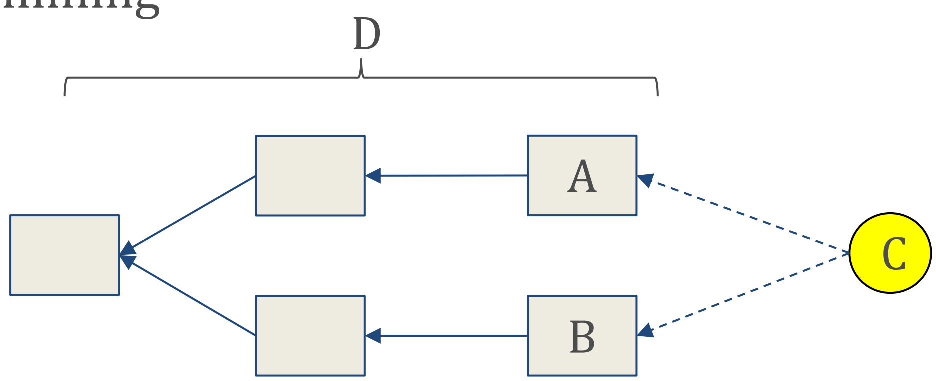

TAKE AWAY

- There are incentive-driven security issues for Proof-of-Stake protocols not present in Proof-of-Work
- There is a tradeoff between predictability and recency
- These attacks might not be devastating, but they are unavoidable for every protocol in our model

- Vitalik's response: https://ethresear.ch/t/formal-barriers-to-longest-chain-proof-of-stake-protocols/3509/2

CRYPTOCURRENCIES

- Selfish mining
- Incentive issues with mining pools
- Incentive issues with transaction fee rewards
- Incentive issues with Proof of Stake

REFERENCES

- The Miner's Dilemma, Ittay Eyal
- On the Instability of Bitcoin without the Block Reward, Miles Carlsten, Harry Kalodner, S. Matthew Weinberg, Arvind Narayanan.
- Formal Barriers to Proof-of-Stake Protocols, Jonah Brown-Cohen, Arvind Narayanan, Christos Alexandros Psomas, S.Matthew Weinberg
- Incentive compatibility of bitcoin mining pool reward functions, Okke Schrijvers, Joseph Bonneau, Dan Boneh, and Tim Roughgarden
- On bitcoin and red balloons. Moshe Babaioff, Shahar Dobzinski, Sigal Oren, and Aviv Zohar.