Bitcoin:
Concepts, Practice, and Research Directions

Part II
Security

Ittay Eyal, Emin Gün Sirer

Computer Science, Cornell University

DISC Bitcoin Tutorial, October 2014

Part 2 – Security

- Unbelievable security of core system
- The mining industry
- Classical attacks
- Centralization
- Misaligned incentives:
  - Transactions
  - Mining
- Reducing pool sizes
- User-side security

Core System Availability

Almost always on.

Despite no shortage of attack motivation.

The March 2013 Fork

- Miner with version 0.8.0 generated a large block.
- Old versions rejected it.

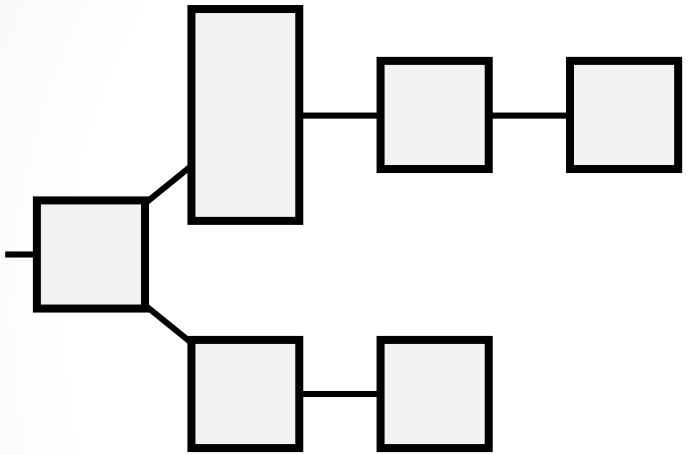

version 0.8.0 miners

pre-0.8.0 miners

Solution:

1. Major miners downgraded to pre-0.8.0.
2. Upgrade to 0.8.1 prevented large blocks.
3. 5 months later: Upgrade done right.

The Mining Industry

Mining

Difficulty rise:

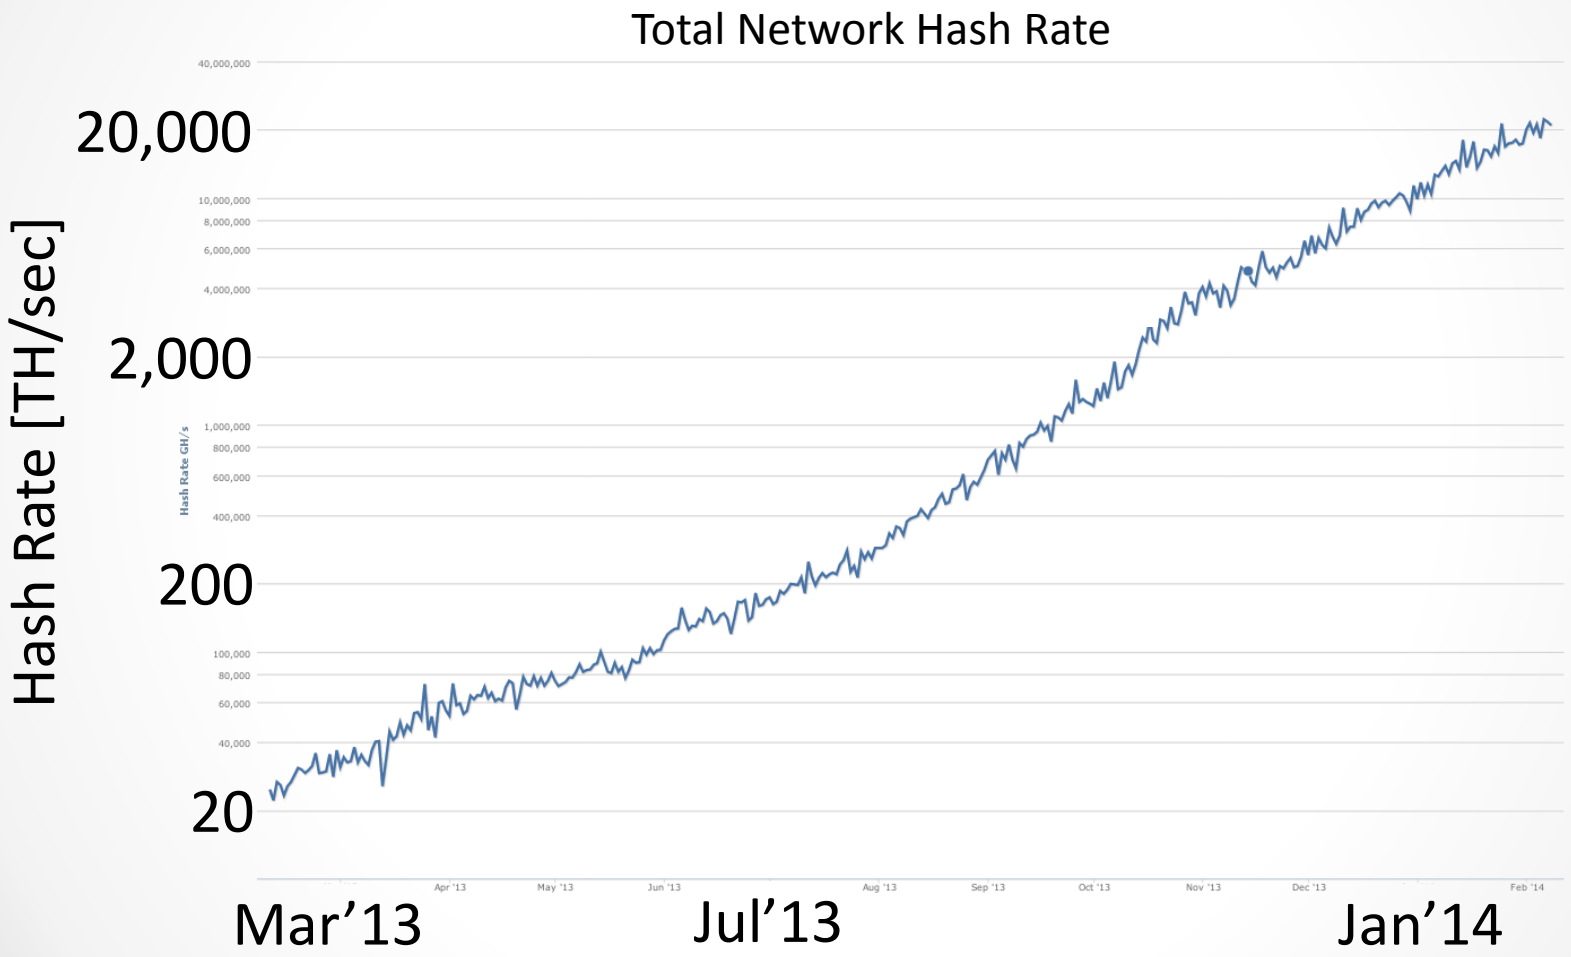

[Blockchain.info]

Mining Industry

Mining Industry

Mining Industry

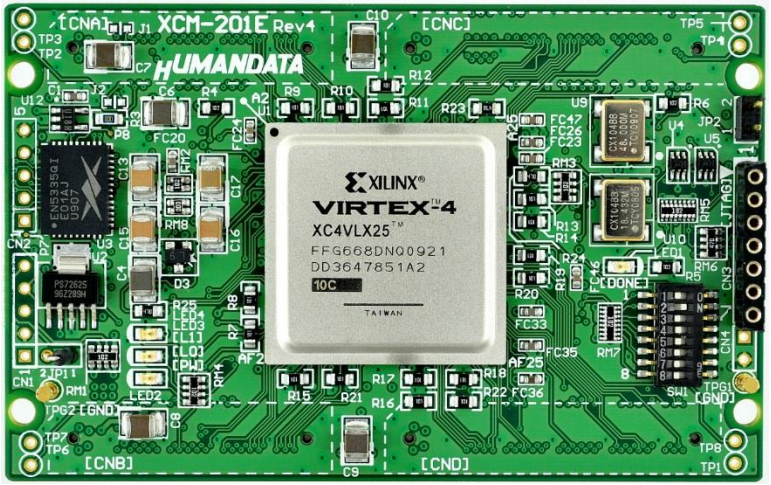

Mining Industry

Mining Industry

- Avalon
- ASIC Miner
- BitMine
- Butterfly Labs
- CoinTerra
- GAW Miners
- HashFast
- KnC Miner
- Spondoolies

Mining Industry

- Avalon
- ASIC Miner
- BitMine
- Butterfly Labs
- CoinTerra
- GAW Min ers
- HashFast
- KnC Miner
- Spondoolies

Mining Industry

- Avalon
- ASIC Miner
- BitMine
- ~~Butterfly Labs~~
- CoinTerra
- GAW Min ers
- HashFast
- KnC Miner
- Spondoolies

Mining Industry

Mining Industry

This is what makes Bitcoin secure.

Classical Attacks

Double Spending

Eve buys coffee from Bob but keeps her money:

1. 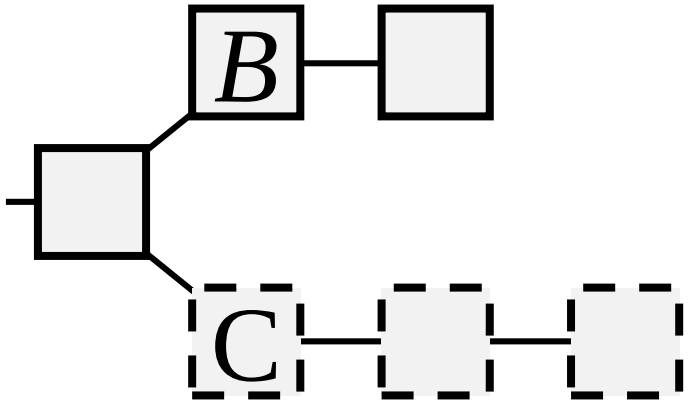

2. Bob provides product to Eve.

3. 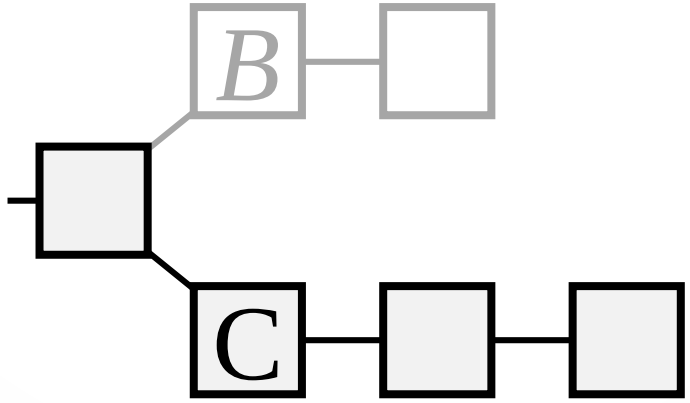

Similar, but more feasible: The Finney attack

Majority Attacker, aka 51%

Attacker produces the longest chain.

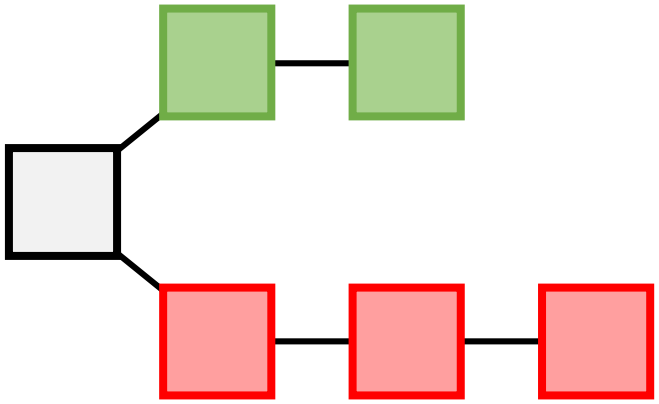

Attacker cannot steal.

Attacker can:

- Require excessive transaction fees,
- take ransom from a single user, or
- prevent all transactions (DoS).

Centralization

Centralization

One entity gains control of the blockchain:

- Single majority miner
- Consortium of pools

Breaks Bitcoin's essential premise.

Centralization

Pool GHash.IO (w/ CEX.IO) surpassed 50%.

Community raged.

DoS attacks on pool.

GHash promptly reduced its rate.

Centralization

Pool GHash.IO (w/ CEX.IO) surpassed 50%.
Community raged.
DoS attacks on pool.
GHash promptly reduced its rate.

(Almost) no good reason for such large pools.

- Nice interface.
- Good uptime.

Misaligned Incentives: Transaction Propagation

Transaction Propagation [1]

Nodes should propagate transactions.
But why would they?

Actual incentive: **don't propagate**.

[1] Babaioff, Dobzinski, Oren, and Zohar, On Bitcoin and Red Balloons. EC'12

Transaction Propagation [1]

DARPA Network Challenge ’09:
Find 10 red balloons in US.

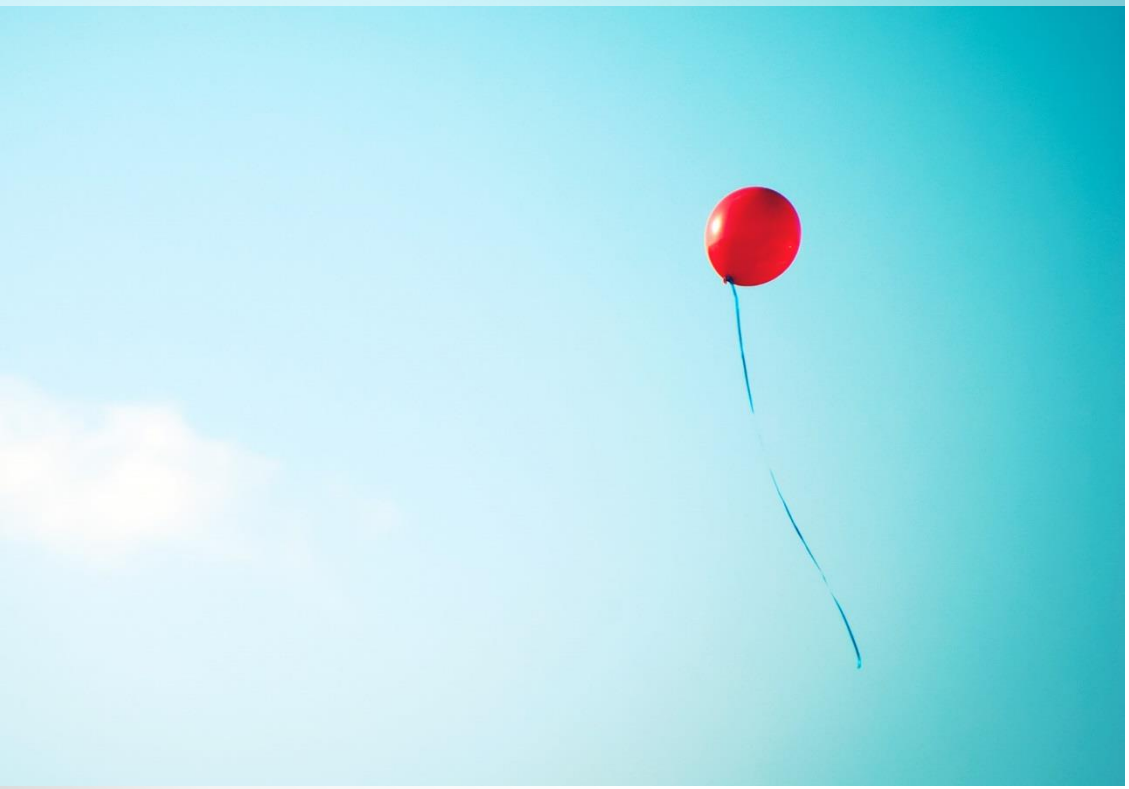

Winner: MIT Group

Technique:

$2000 to finder

$1000 to recruiter

$500 to 2nd recruiter

…

[1] Babaioff, Dobzinski, Oren, and Zohar, On Bitcoin and Red Balloons. EC’12

Transaction Propagation [1]

DARPA Network Challenge ’09:
Find 10 red balloons in US.

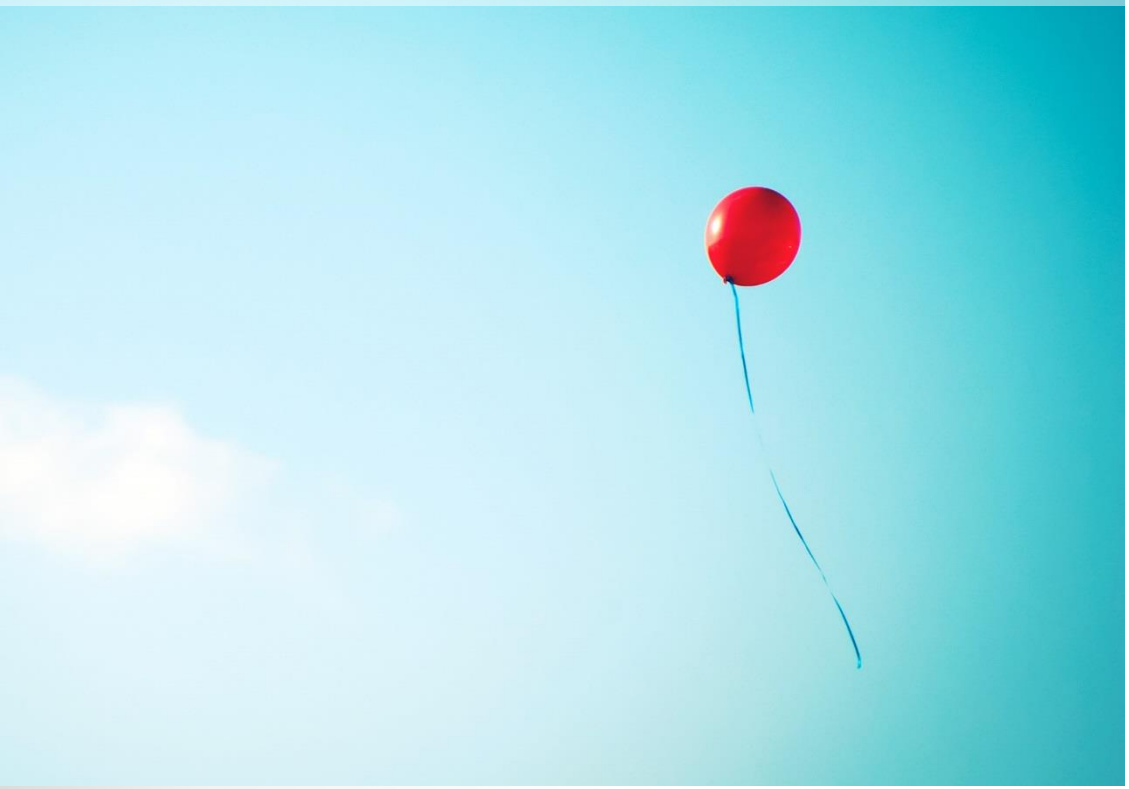

Winner: MIT Group

Technique:

$2000 to finder

$1000 to recruiter

$500 to 2nd recruiter

…

Applicable to Bitcoin?

[1] Babaioff, Dobzinski, Oren, and Zohar, On Bitcoin and Red Balloons. EC’12

Transaction Propagation [1]

Red balloons technique not applicable to Bitcoin.

- Why recruit your own competition?
  Unlike balloons case where you recruit far away.

- Can masquerade as your own recruits.
  Unlike balloons case where you physically show up.

[1] Babaioff, Dobzinski, Oren, and Zohar, On Bitcoin and Red Balloons. EC'12

Transaction Propagation [1]

Solution sketch:

Set integers H and β according to topology.
Then, for a chain of length l:

If $l > H$

- no reward.

Otherwise,

- miner gets $1 + (H - l + 1)\beta$,
- others get 1.

[1] Babaioff, Dobzinski, Oren, and Zohar, On Bitcoin and Red Balloons. EC'12

Misaligned Incentives: Selfish Mining

Common Wisdom

Nakamoto's Bitcoin mining protocol is incentive compatible (assuming an honest majority)

1. Best strategy: being honest

2. Revenue proportional to compute power

Selfish Mining [1]

Goal: Get more than fair share.

How: Maintain secret blocks, publish judiciously.

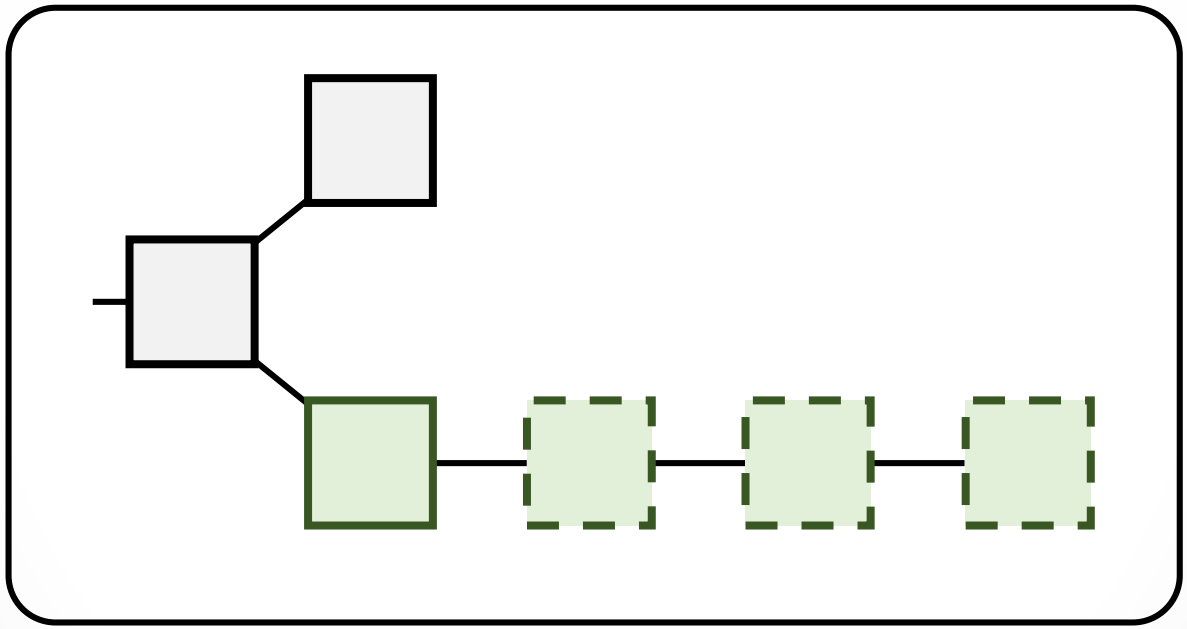

Intuition: Risk some work, others waste a lot.

[1] Eyal and Sirer: *Majority is not Enough: Bitcoin Mining is Vulnerable*, FC'14

Selfish Mining Algorithm

(a) Any state but two branches of length 1.
Pool finds a block.
Keep it secret. No revenue.

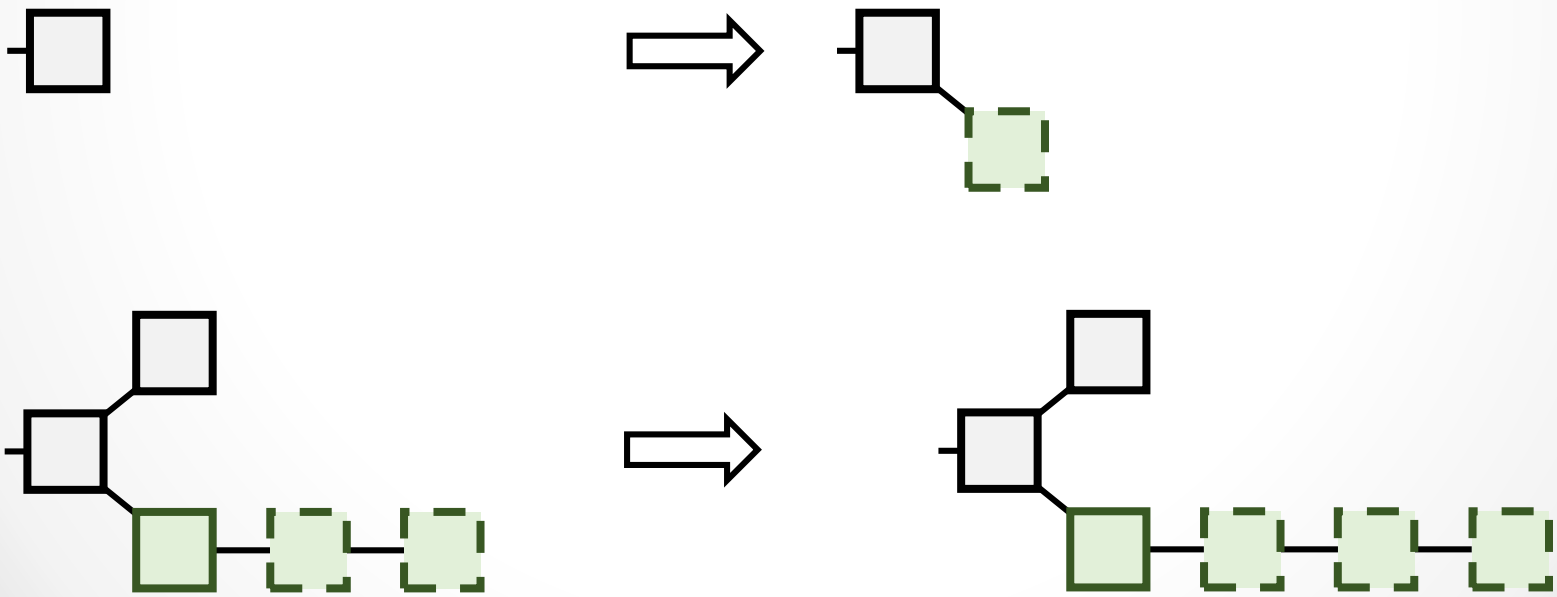

Selfish Mining Algorithm

(h) Lead more than 2.

Others find a block.

Publish one block. Selfish gets 1.

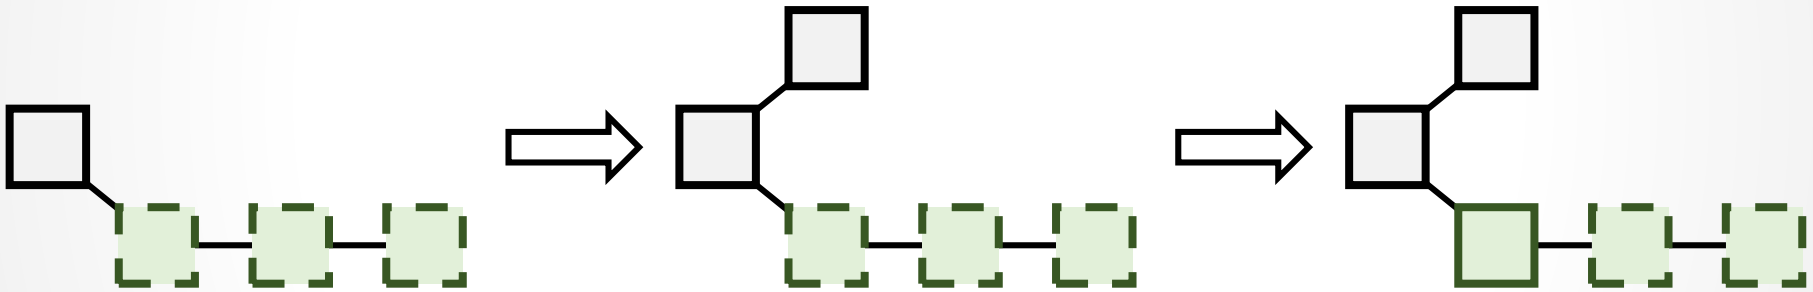

Selfish Mining Algorithm

(g) Lead of 2.

Others find a block.

Publish secret chain. Selfish gets 2.

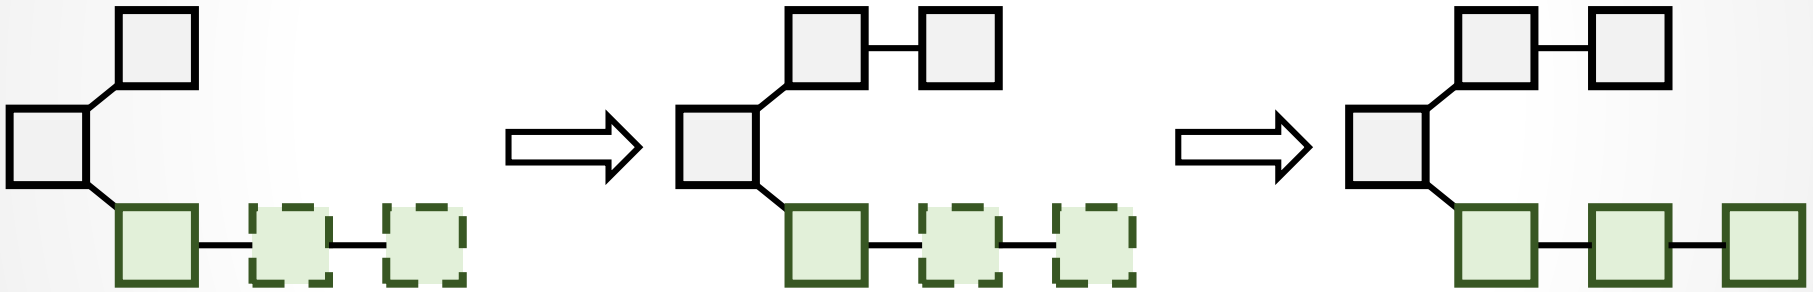

Selfish Mining Algorithm

(f) Lead of 1.

Others find a block.

Publish secret block. No revenue.

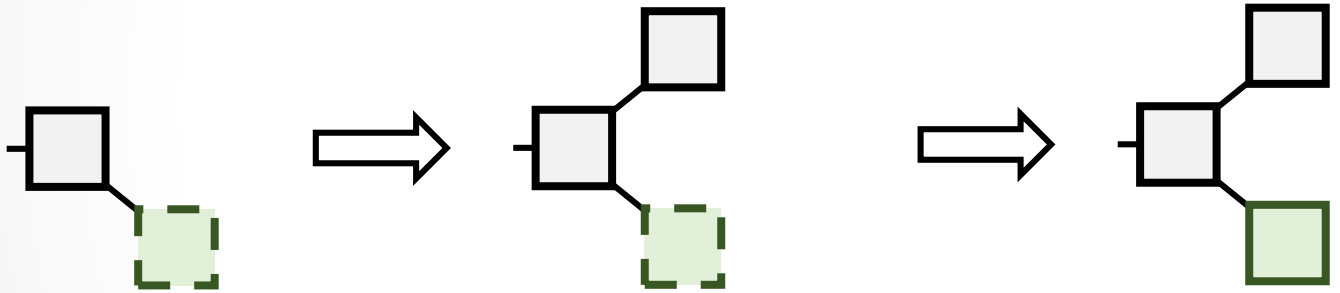

γ: Ratio of others that follow pool

Selfish Mining Algorithm

(b) Two branches of length 1.

Pool finds a block.

Publish branch. Selfish gets 2.

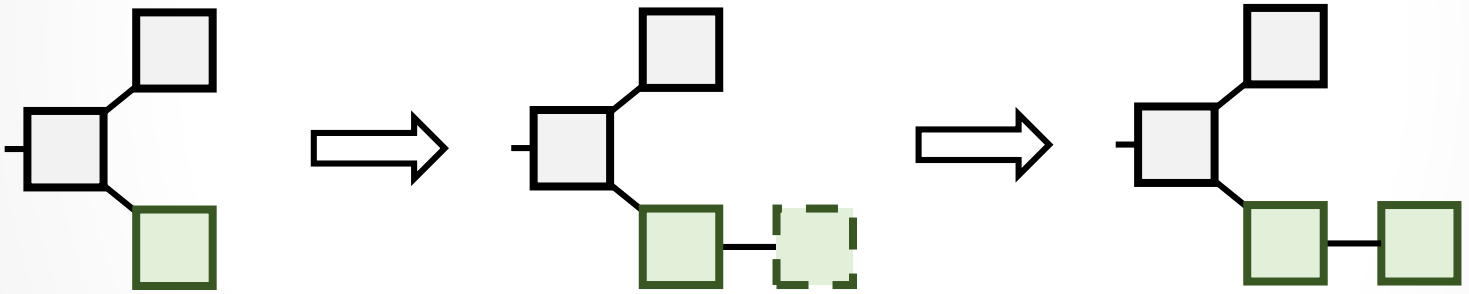

Selfish Mining Algorithm

(c) Two branches of length 1.

Others find a block after pool head.

Revenue: Each get 1.

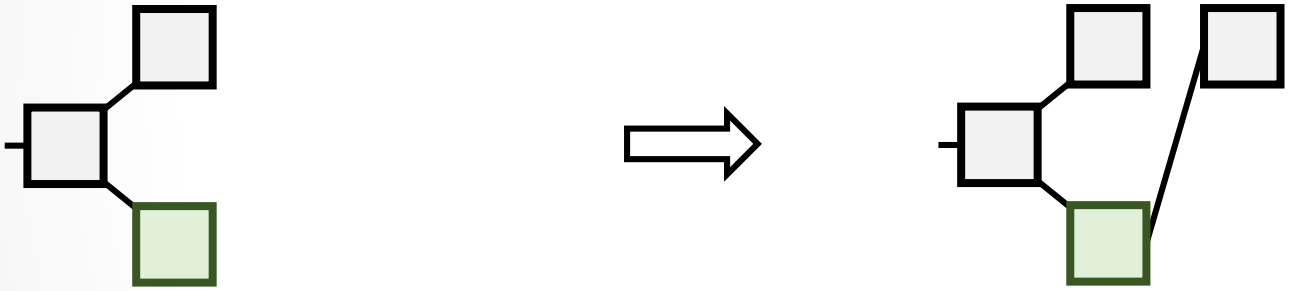

Selfish Mining Algorithm

(d) Two branches of length 1.

Others find a block after others' head.

Revenue: Others get 2.

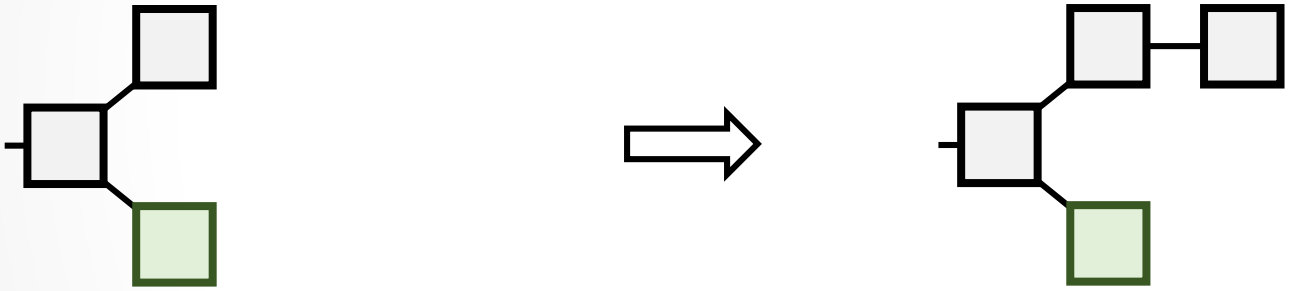

Selfish Mining Algorithm

(e) No private branch.

Others find a block.

Revenue: Others get 1.

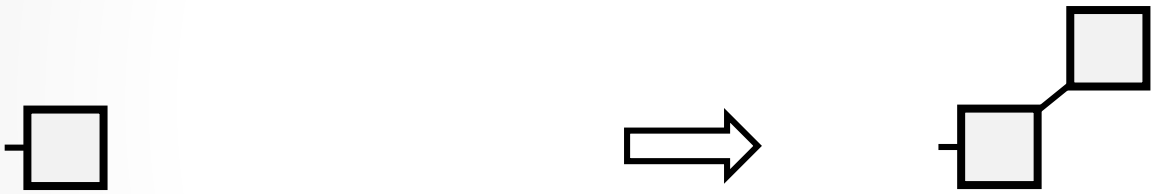

Selfish Mining: Analysis

Selfish Mining – Probabilities

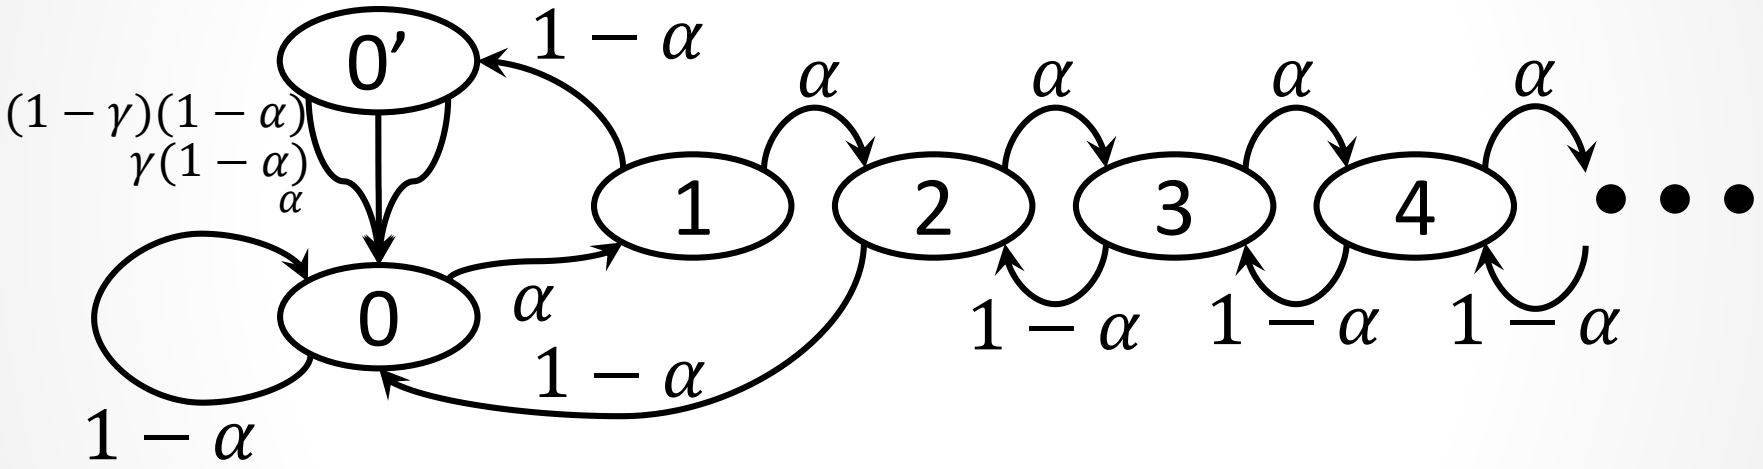

Selfish Mining – Probabilities

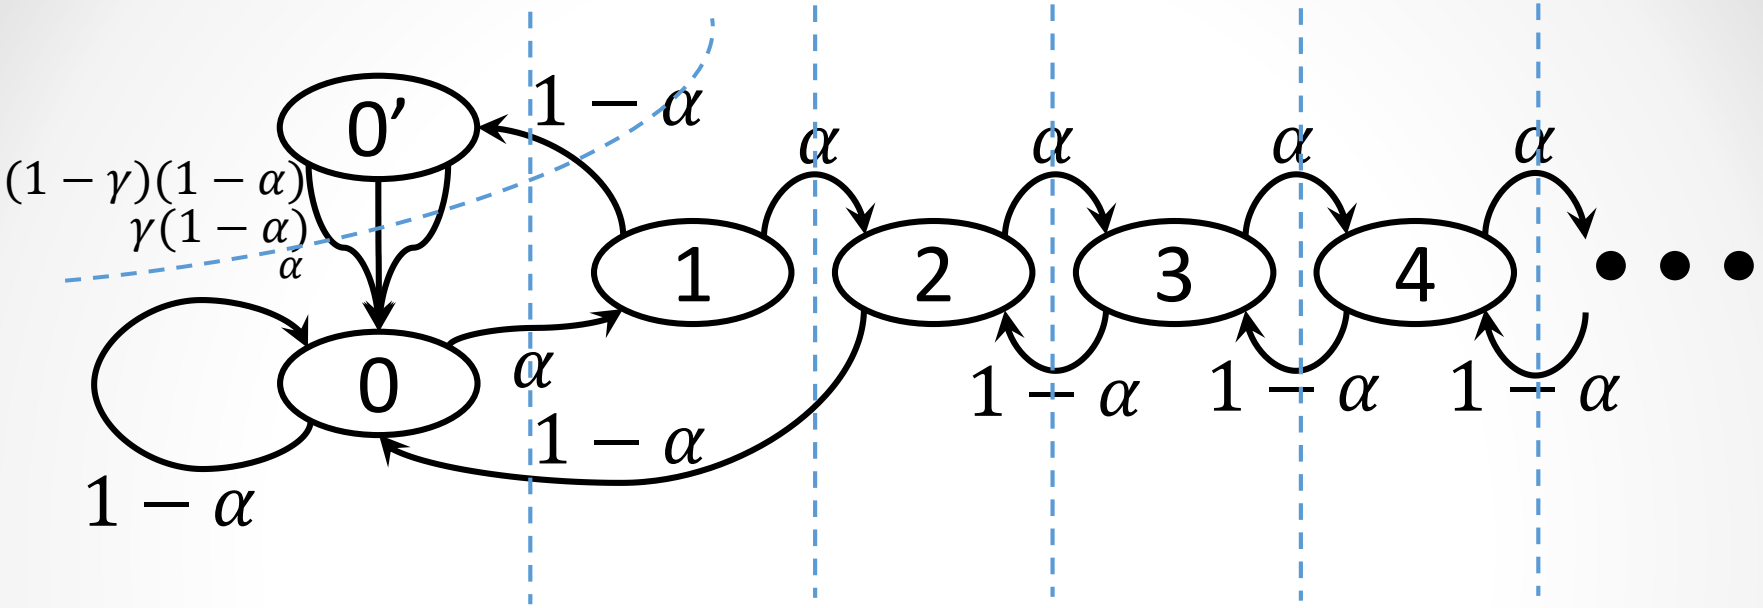

$$
\begin{cases}
\alpha p_0 = (1-\alpha)p_1 + (1-\alpha)p_2 \\
p_{0'} = (1-\alpha)p_1 \\
\alpha p_1 = (1-\alpha)p_2 \\
\forall k \geq 2 : \alpha p_k = (1-\alpha)p_{k+1} \\
\sum_{k=0}^{\infty} p_k + p_{0'} = 1
\end{cases}
$$

Selfish Mining – Revenue

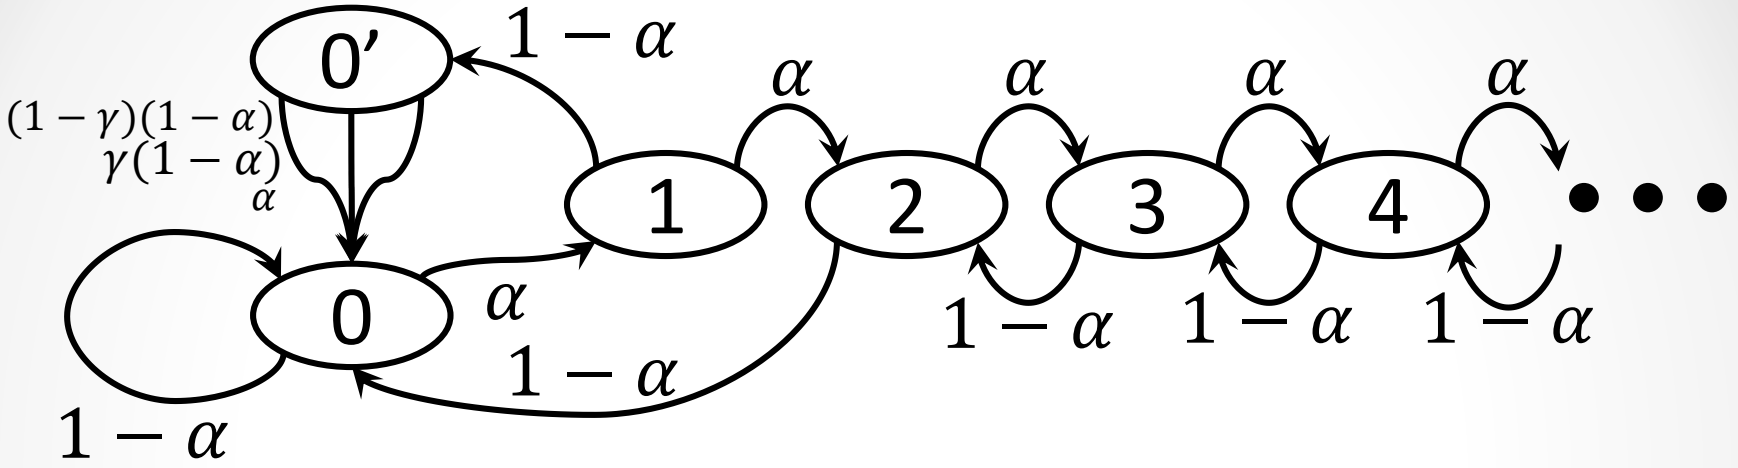

$$r_{\text{others}} = \overbrace{p_{0'} \cdot \gamma(1-\alpha) \cdot 1}^{\text{Case (c)}} + \overbrace{p_{0'} \cdot (1-\gamma)(1-\alpha) \cdot 2}^{\text{Case (d)}} + \overbrace{p_0 \cdot (1-\alpha) \cdot 1}^{\text{Case (e)}}$$

$$r_{\text{pool}} = \overbrace{p_{0'} \cdot \alpha \cdot 2}^{\text{Case (b)}} + \overbrace{p_{0'} \cdot \gamma(1-\alpha) \cdot 1}^{\text{Case (c)}} + \overbrace{p_2 \cdot (1-\alpha) \cdot 2}^{\text{Case (g)}} + \overbrace{P[i>2](1-\alpha) \cdot 1}^{\text{Case (h)}}$$

Selfish Mining – Revenue

Auto-adjusting difficulty, so:

$$R_{pool} = \frac{r_{pool}}{r_{pool} + r_{others}}$$

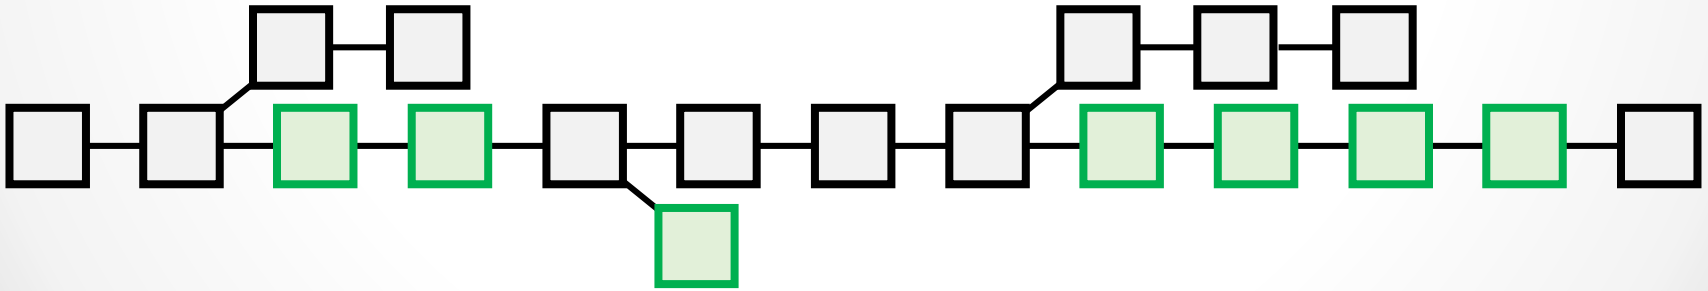

Selfish Mining – Analysis

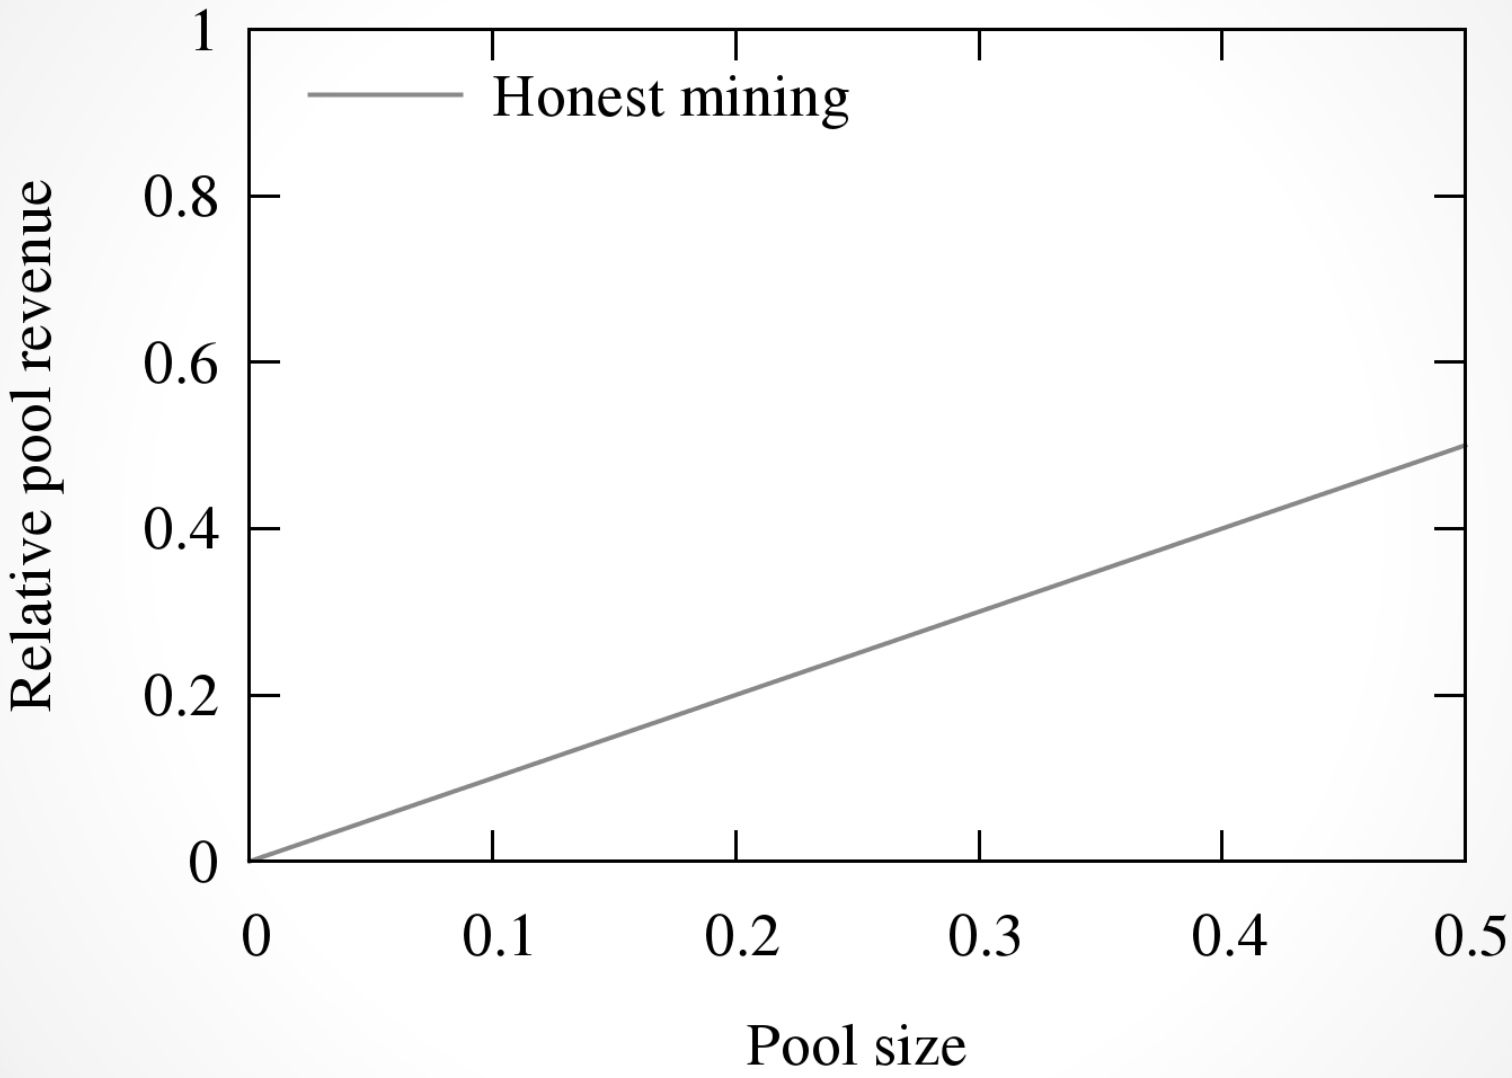

Selfish Mining – Analysis

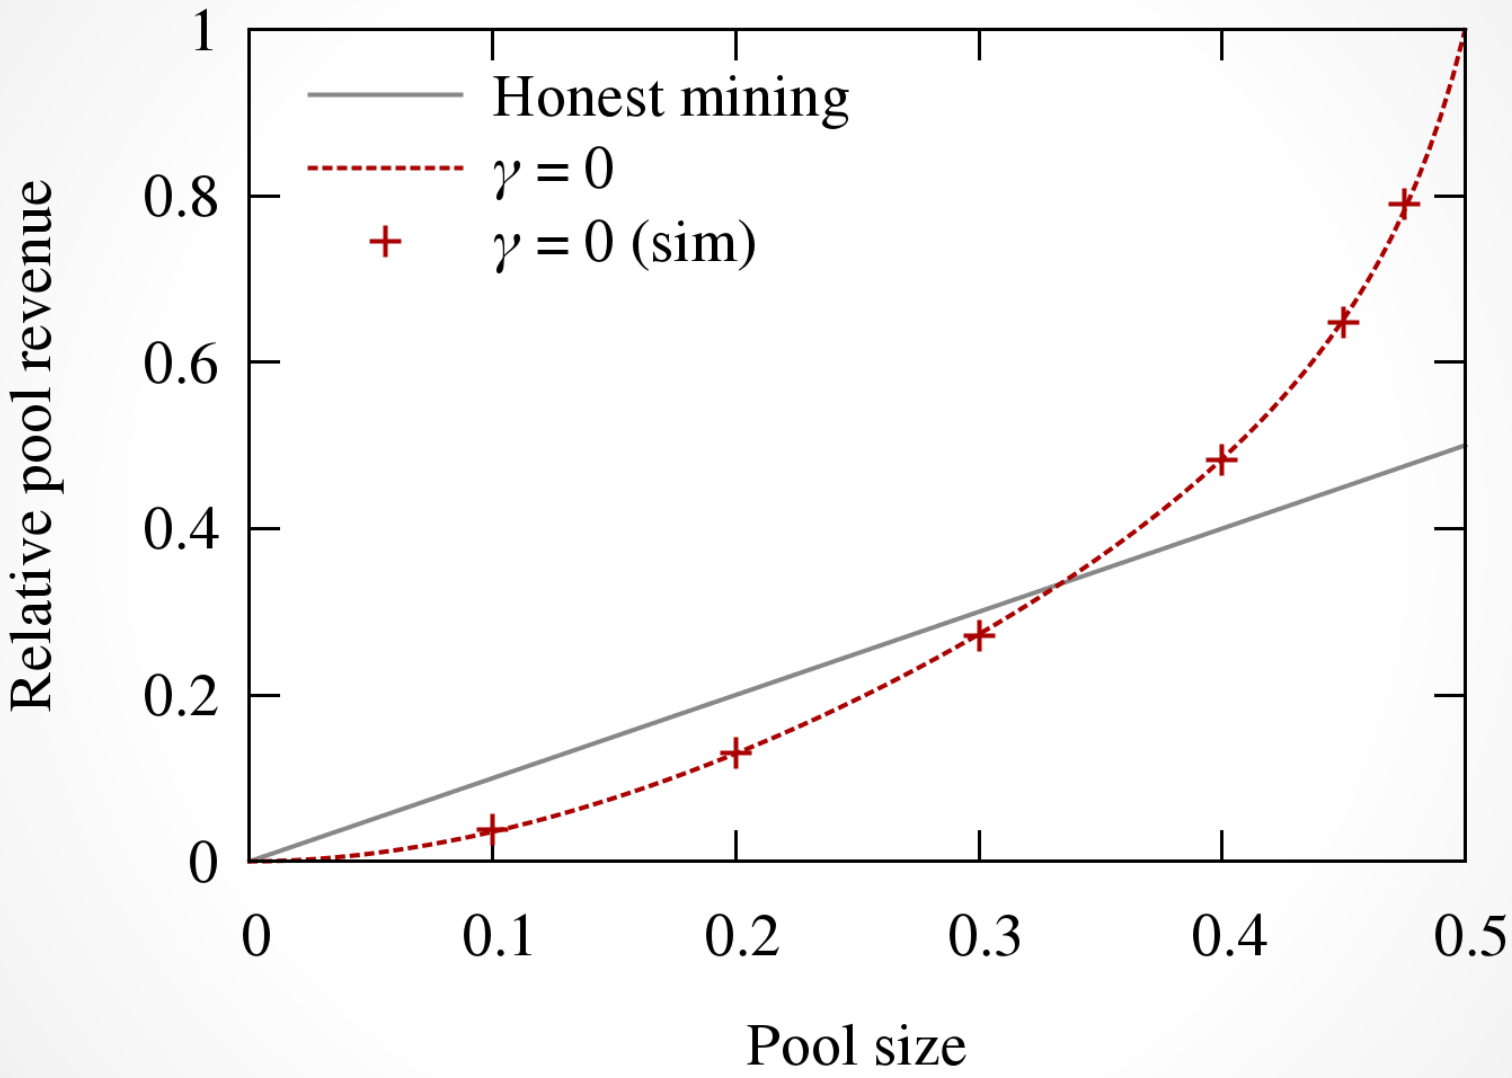

Selfish Mining – Analysis

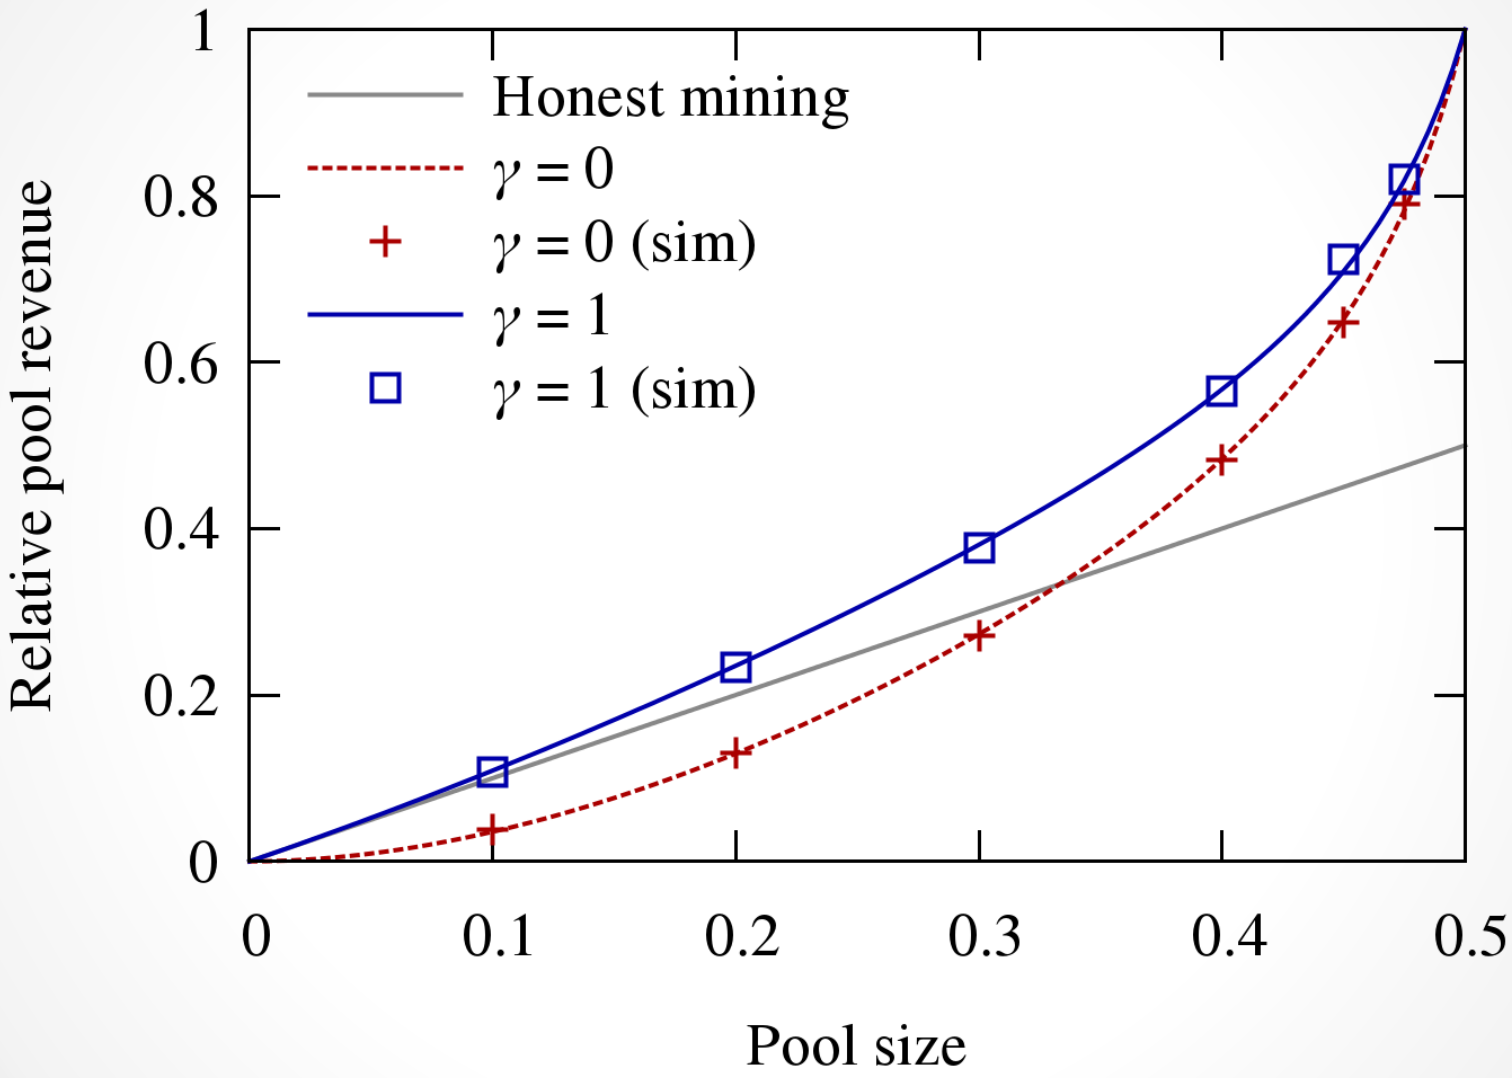

Selfish Mining – Analysis

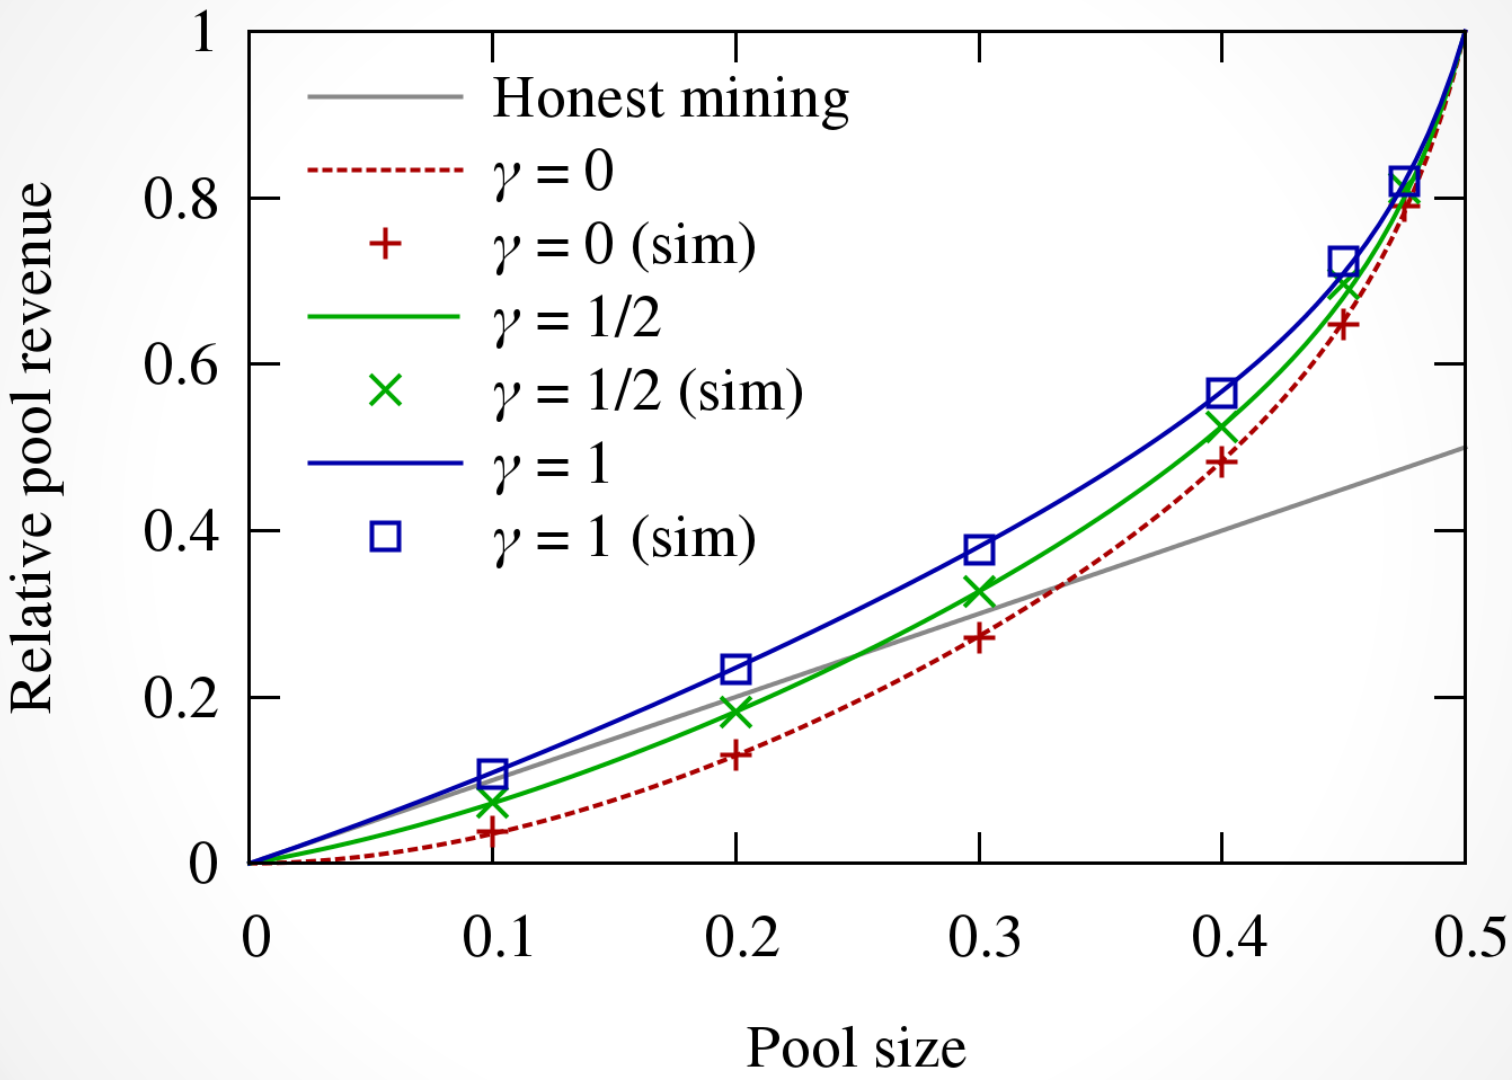

Selfish Mining: Implications

Attack Feasible

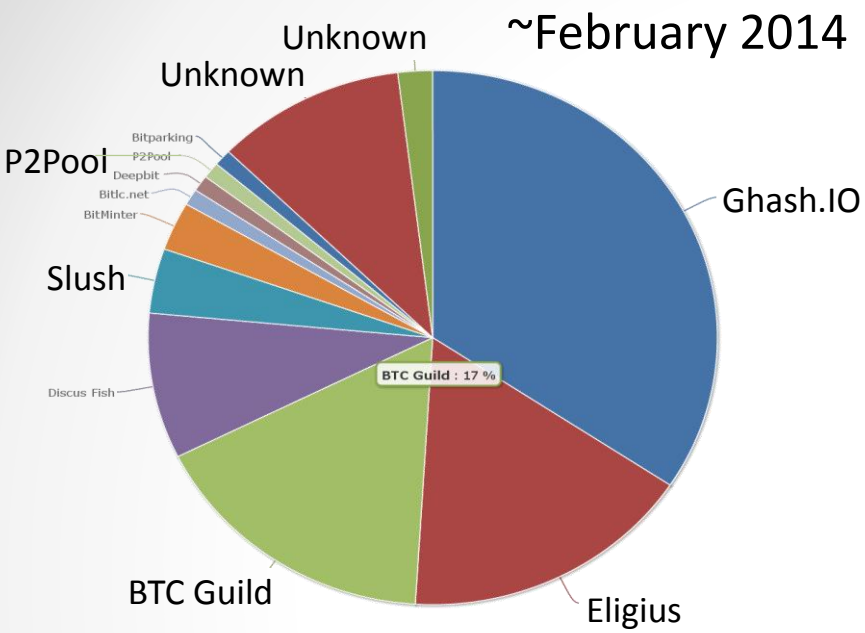

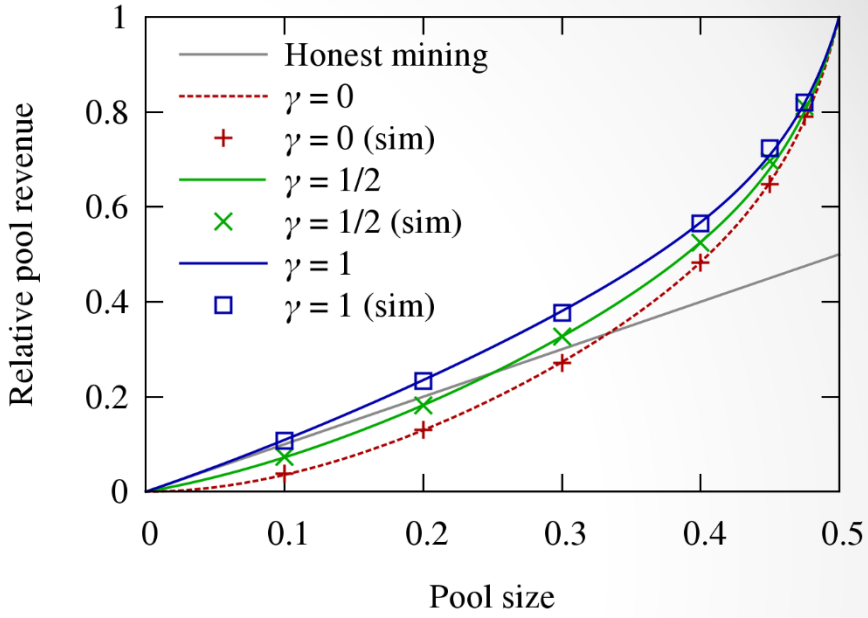

Catastrophe Scenario

After threshold:

Rational miners want to join selfish pool.

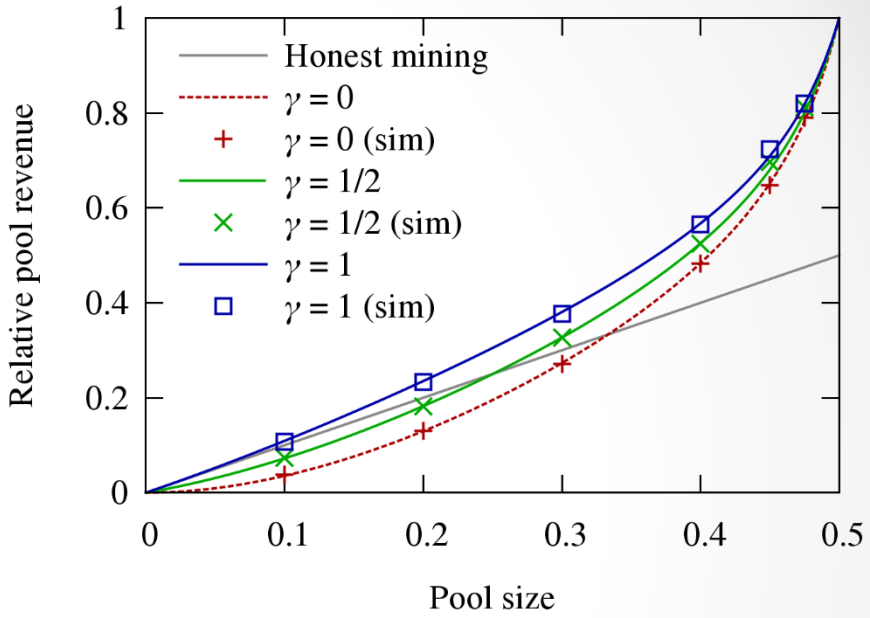

Catastrophe Scenario

Superlinear growth:

Selfish pool wants to grow.

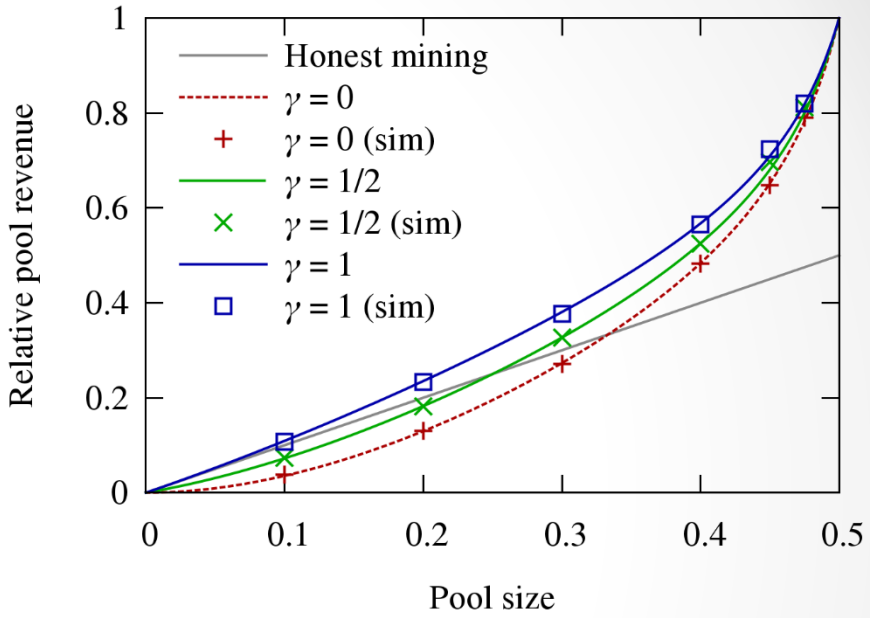

Catastrophe Scenario

Rational miners want to join selfish pool.

+

Selfish pool wants to grow.

=

Selfish pool may grow towards 50%

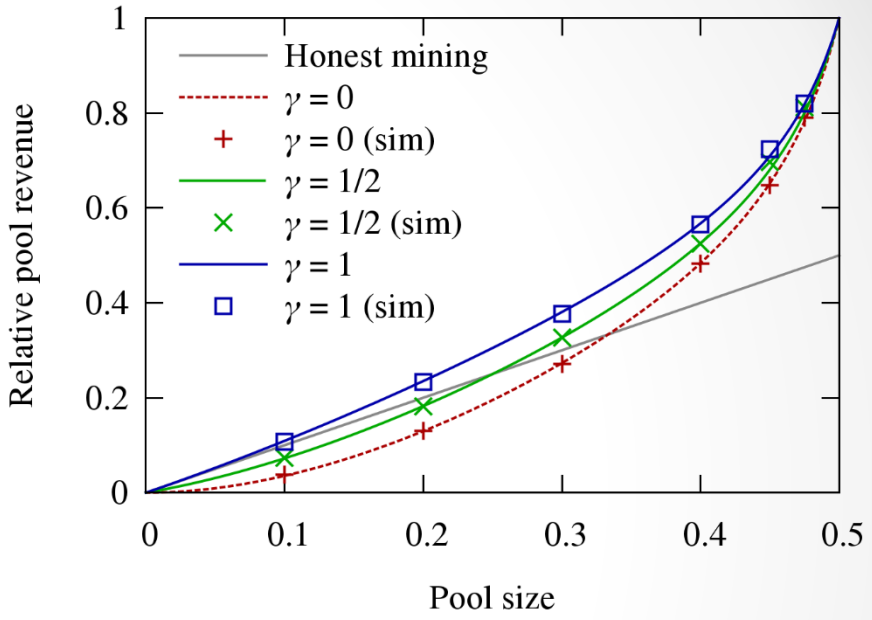

NOT GOOD.

Attack Happening Now?

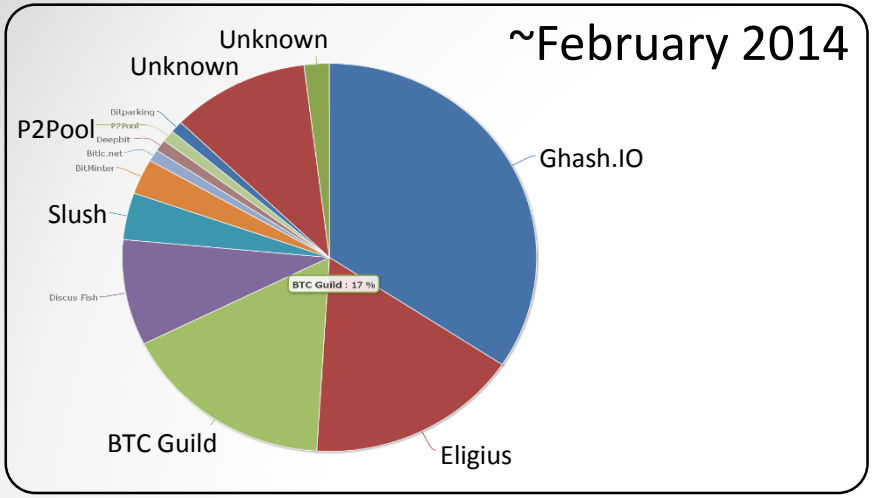

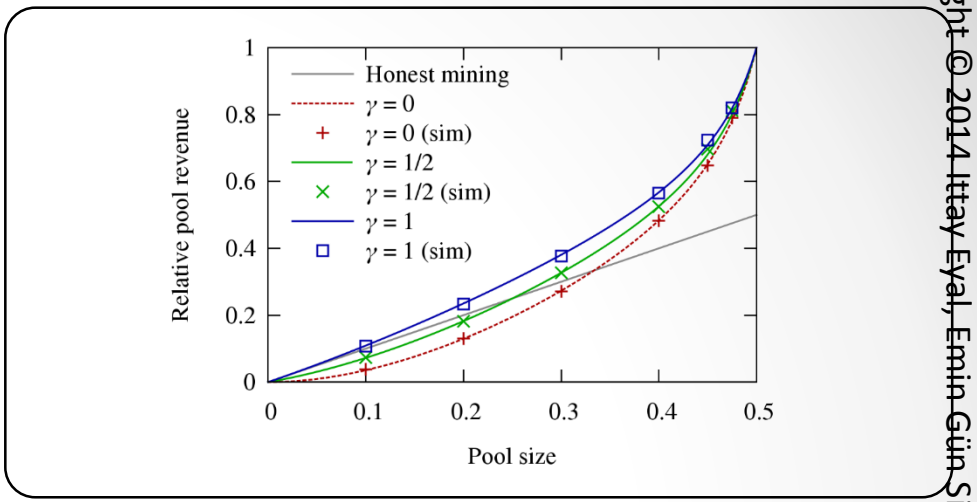

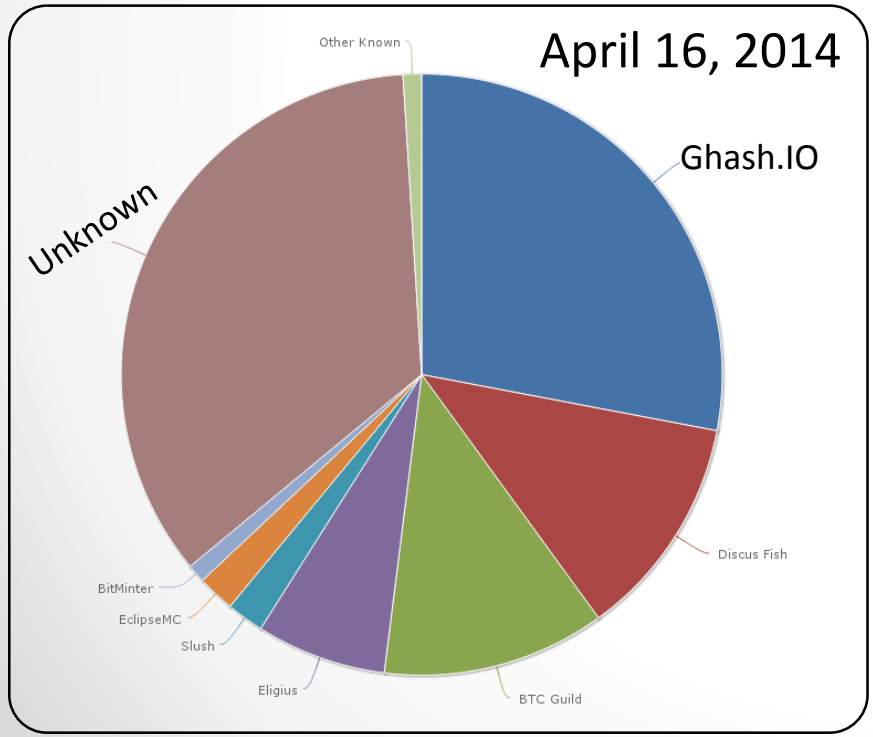

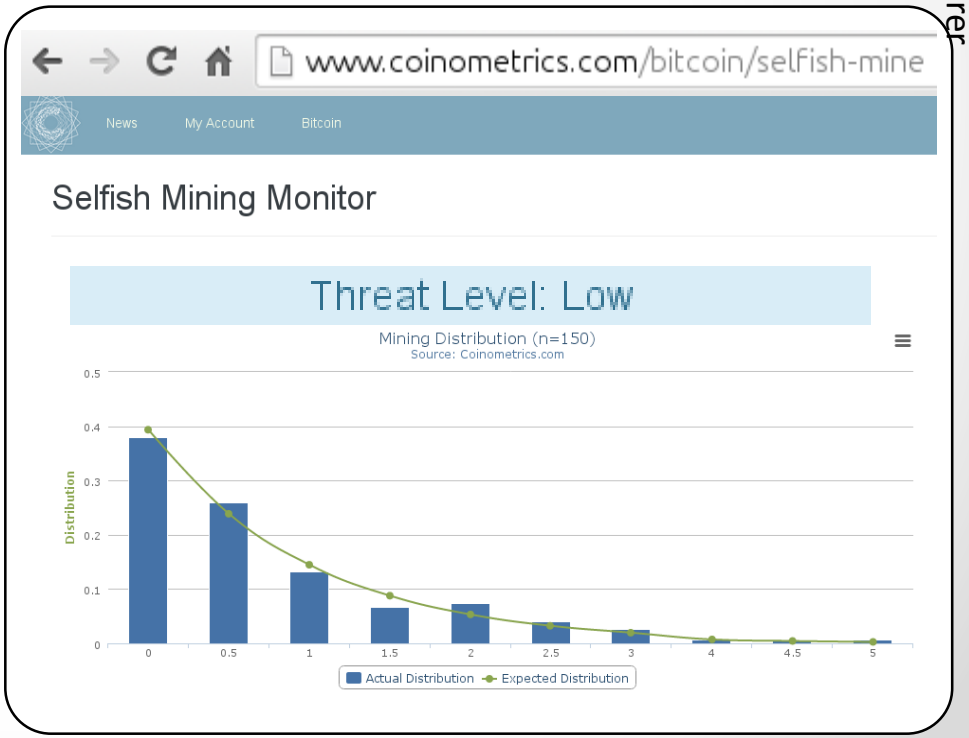

Hardening the protocol

Algorithm change:

- Propagate all blocks of longest chain.
- Choose one at random to mine on.

Hardening the protocol

Algorithm change:

- Propagate all blocks of longest chain.
- Choose one at random to mine on.

Benefits:

- Proved threshold
- Backward compatible
- Progressive
- Simple

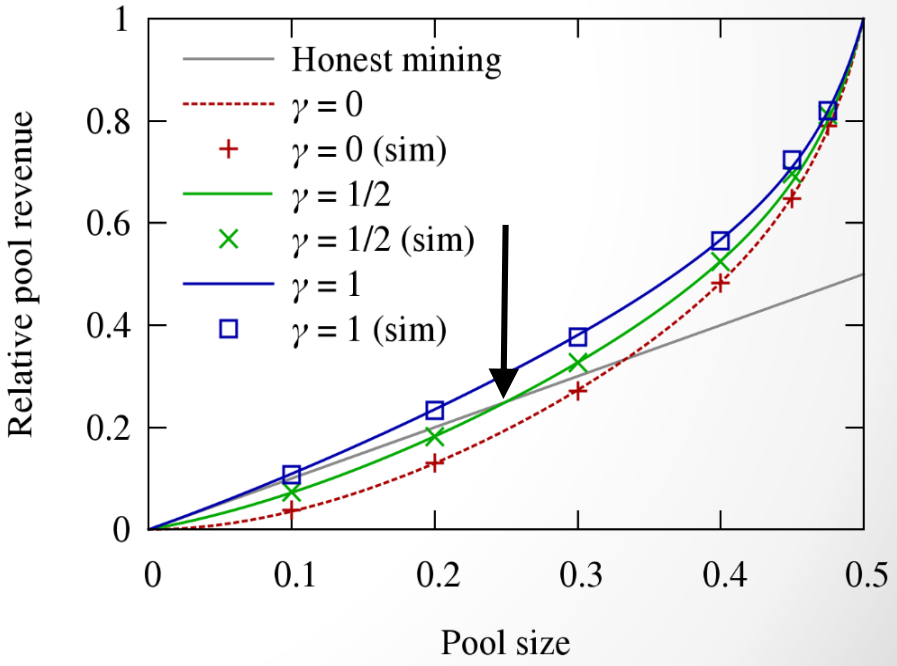

Reducing Pool Sizes

P2Pool [1]

A peer to peer distributed pool.

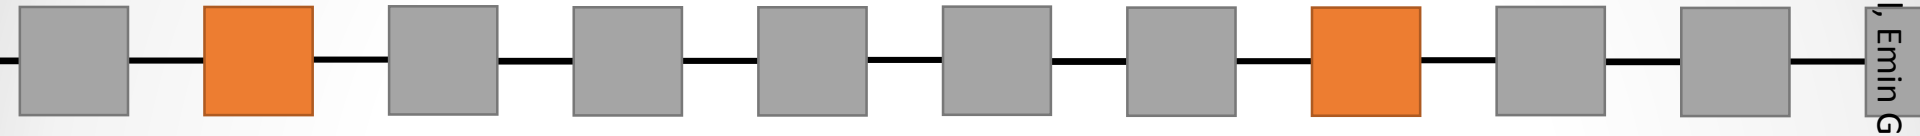

- A separate blockchain with Easy PoW
- Blocks distribute potential revenue among miners.
- Actual revenue on full PoW.

[1] Forrest Voight, *p2pool: Decentralized, DoS-resistant, Hop-Proof pool*, Bitcoin Forum, June 2011

Pool Limiting

- Non-outsourcable PoW [1]
  Cryptographic technique: A miner can steal from the pool when it finds a block.
  - Pool cannot outsource differently.
  - Block does not reveal secret.

- Permacoin [2]
  Proof of storage rather than work.
  Storage should not be outsourceable.

[1] Miller, Shi, Kosba, and Katz. Nonoutsourceable Scratch-Off Puzzles to Discourage Bitcoin Mining Coalitions. TR
[2] Miller, Juels, Shi, Parno and Katz. Permacoin: Repurposing Bitcoin Work for Data Preservation. TR

2-Phase Proof of Work

Split the proof of work.

- Phase 1: Standard Bitcoin, but easier.
- Phase 2: Requires coinbase secret key.

Benefits:

- Existing infrastructure **controlled** phase-out.
  HW, datacenters.
- Pool must trust miners to outsource phase 2.
  Miner could try and steal the coinbase.

User-side security

User-side Security

Client must **keep** private keys **secret**.

High availability vs. **security**

Individual and large organizations security differs only in scale.

Unprecedented security requirements from commodity systems.

Individuals

Tools:

- Standard client
- Software wallets (for phone)
- Online wallets
- Brain wallets
- Hardware wallets

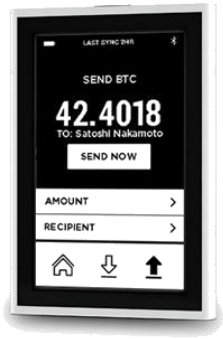

Practice:

- Limited amount on phone
- Cold storage – replicated
- Use correct cryptography [1]

[1] Bos, Halderman, Heninger, Moore, Naehrig and Wustrow: *Elliptic Curve Cryptography in Practice*, FC'14

Large services

Tools:

- Plenty of firewalls
- Bullet proof front-end systems
- Bullet proof back-end systems

Practice:

- Cold storage
- Auditing

Transaction Malleability

Transaction hash used to track transactions.

But it's possible to change a transaction:

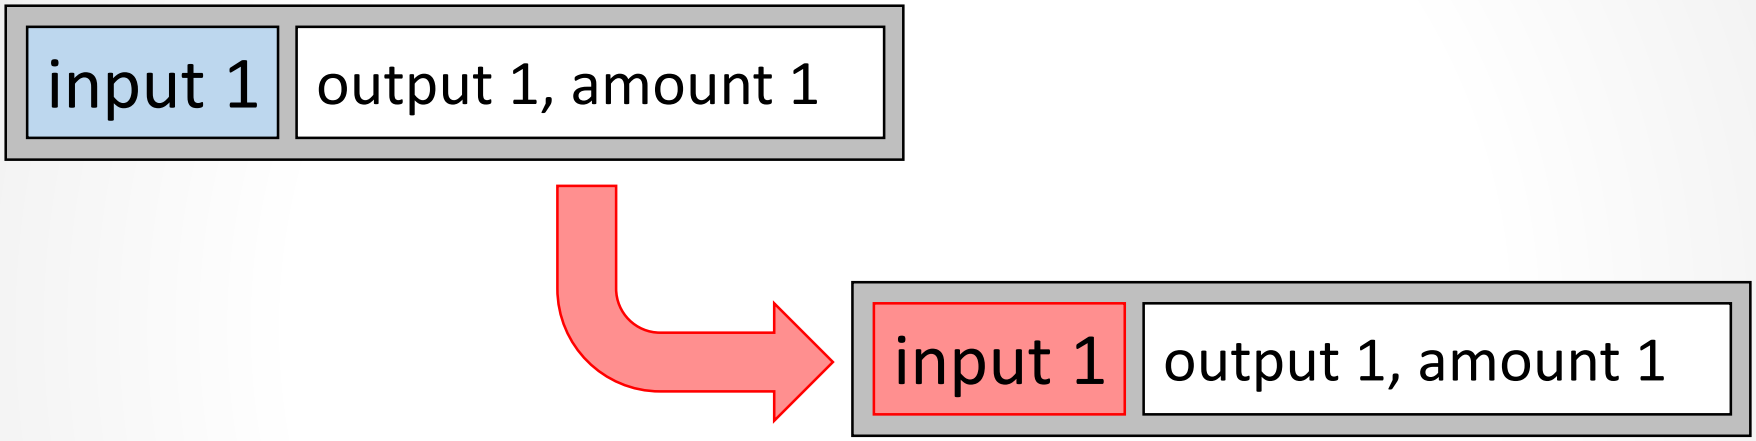

Change scriptSig:

Still valid, for same content, different bits.

1. Change signature. (Crypto trick)
2. Change script. (Protocol trick)

Transaction Malleability

The MtGox con:

Transaction Malleability

The MtGox con:

1. Issue withdraw command.

2. Generate malformed txn, place in public buffer.

3. Change txn and publish it; get the money.

4. Call Mt.Gox to complain.

5. Pay again with new txn.

6. Get money again.

Miners and Pools

- The BGP attack

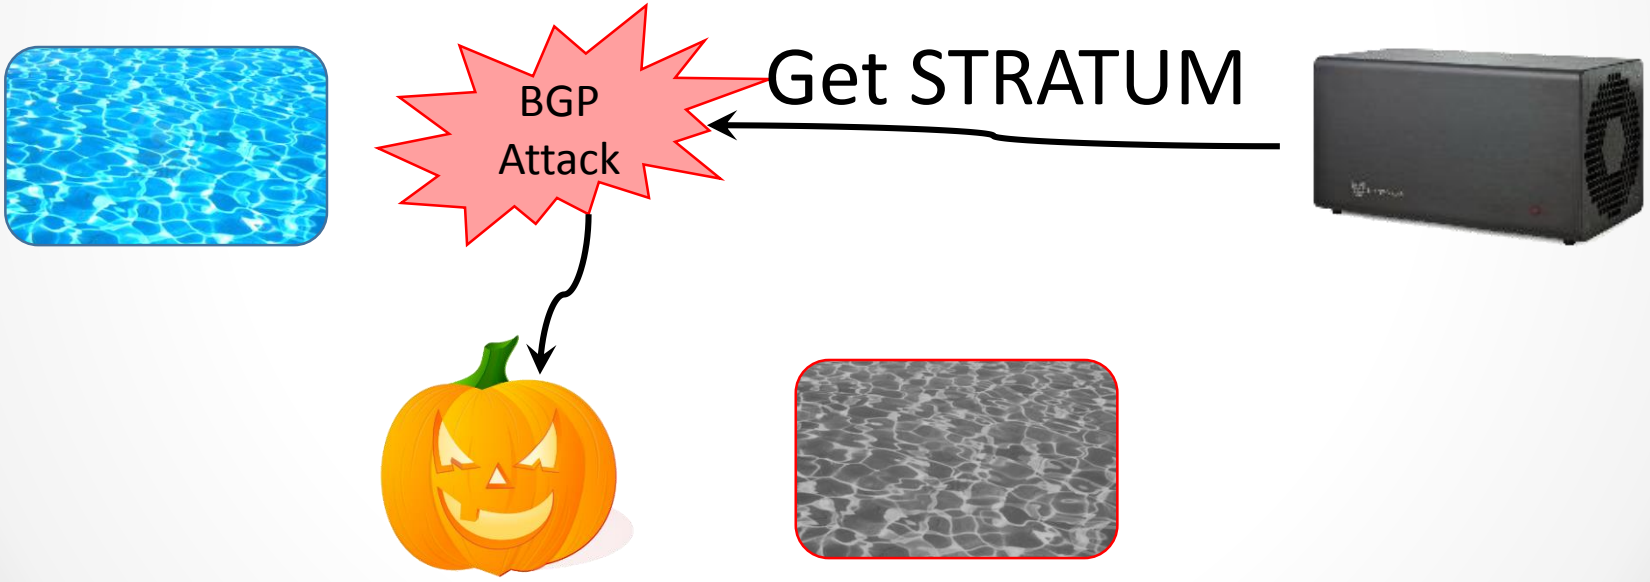

Miners and Pools

- The BGP attack

Miners and Pools

- The BGP attack

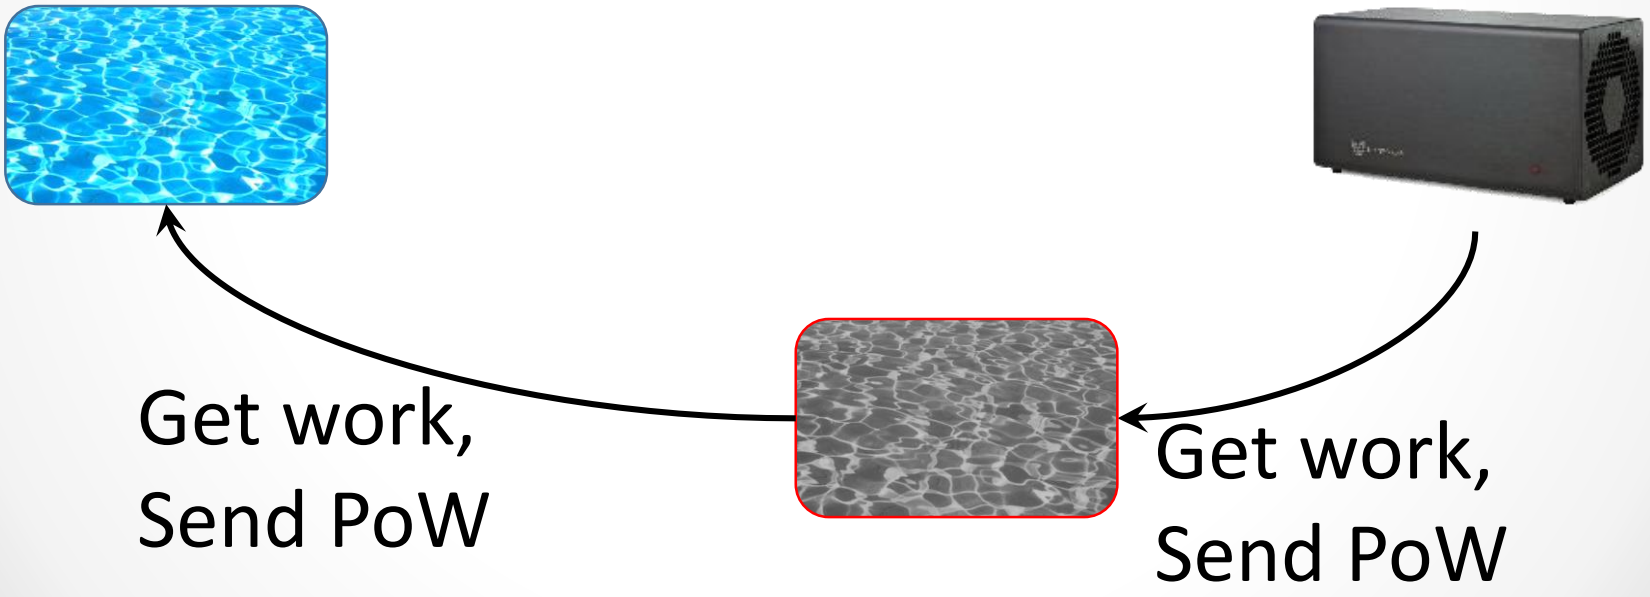

Miners and Pools

- The BGP attack
- Block Withholding
  - Miner sends pool PoW
  - Unless it's an actual solution

Bitcoin:
Concepts, Practice, and Research Directions

Part III
Other Research

Ittay Eyal, Emin Gün Sirer

Computer Science, Cornell University

DISC Bitcoin Tutorial, October 2014

Part 2 – Other Research

- Alt-coins
- Extensions
- Privacy
- Contemporary issues

Alt-coins & Extensions

Parameter changing

- Block frequency
  - Faster confirmation
  - More forks
- PoW choice
  - More green? (no)
  - More fair? (no)
- Difficulty adjustment rate
  - Defense against flash miners

Proof of stake [1]

Goal:

- Save some trees.
- Power to the users! (rather than miners)

Method:

- Proof of Stake (PoS) instead of Proof of Work: Lock coins to create block.

[1] King and Nadal. *PPCoin: Peer-to-Peer Crypto-Currency with Proof of Stake*, August '12

Proof of stake [1]

Goal:

- Save some trees.
- Power to the users! (rather than miners)

Method:

- Proof of Stake (PoS) instead of Proof of Work: Lock coins to create block.

But nothing is at stake!

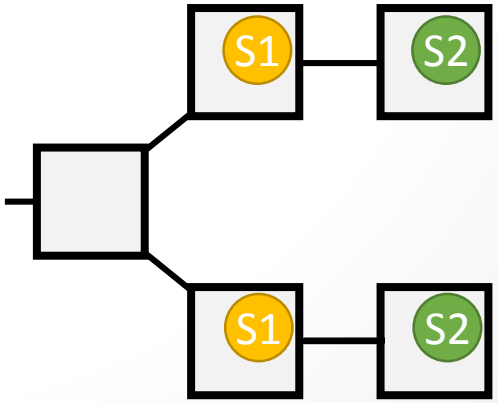

[1] King and Nadal. *PPCoin: Peer-to-Peer Crypto-Currency with Proof of Stake*, August '12

Merged mining

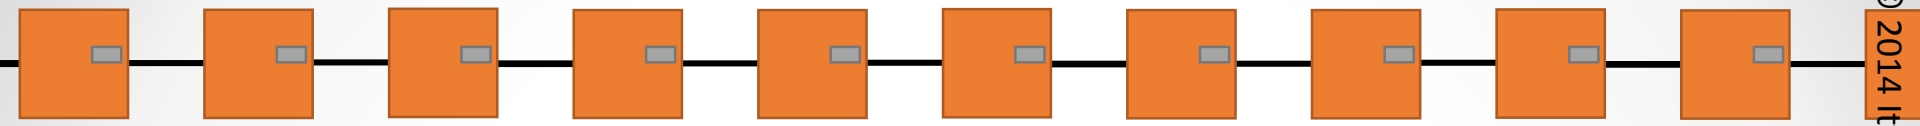

Bitcoin PoW contains:

- Useless transaction (alt-coin header hash *x*)

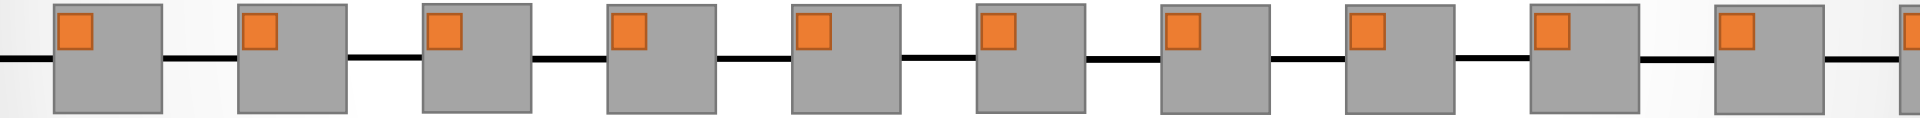

Alt-coin PoW contains:

- Alt-coin header with hash *x*
- Bitcoin header with transaction *x*

Merged mining

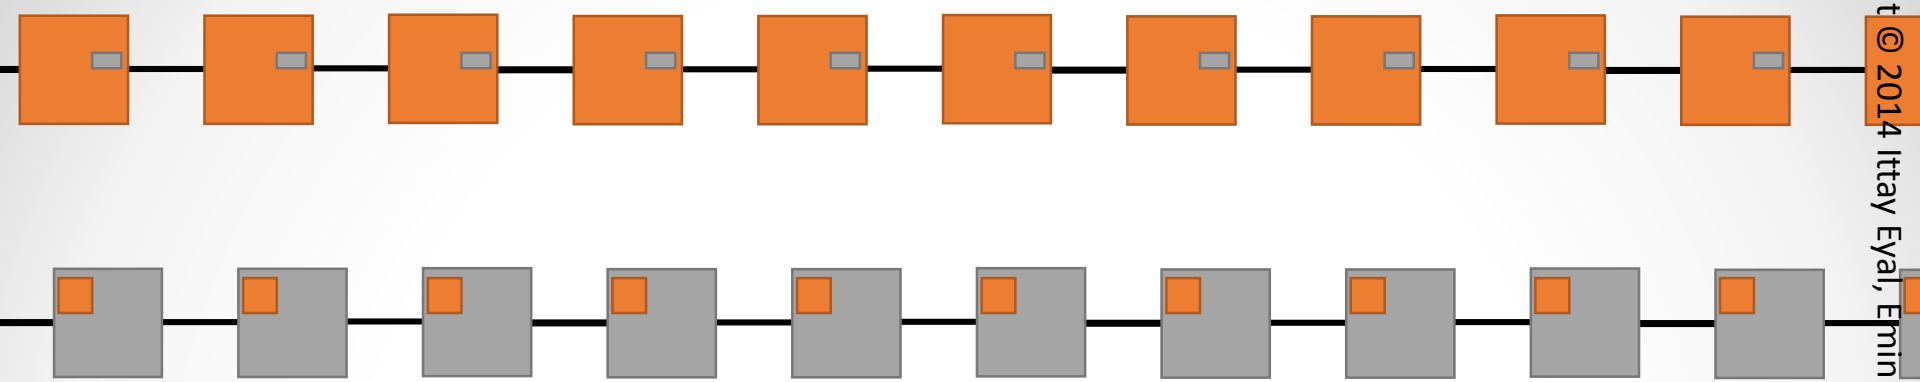

Miners benefit from mining both chains together.
So they do.

Alt-coin gets mining power from day one.

Smart Contracts

- Smart Contracts:
  - m out of n signatures.
  - Time-locked transactions:
    - Time to place in blockchain.
    - Time to use outputs.

- Ethereum: outsource distributed computing (got 31k BTC, at $18 million)
  - Transactions generate transactions.
  - Transactions activate one another.

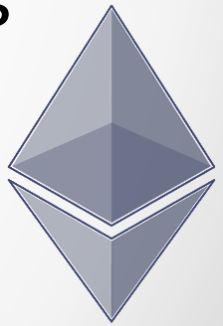

Extensions

- Colored coins:
  Associate assets to individual Bitcoins.

- Side chains:
  - Faster
  - backed by main blockchain
  - less secure

Privacy

Transaction Tracking

All transactions remain in Blockchain forever.

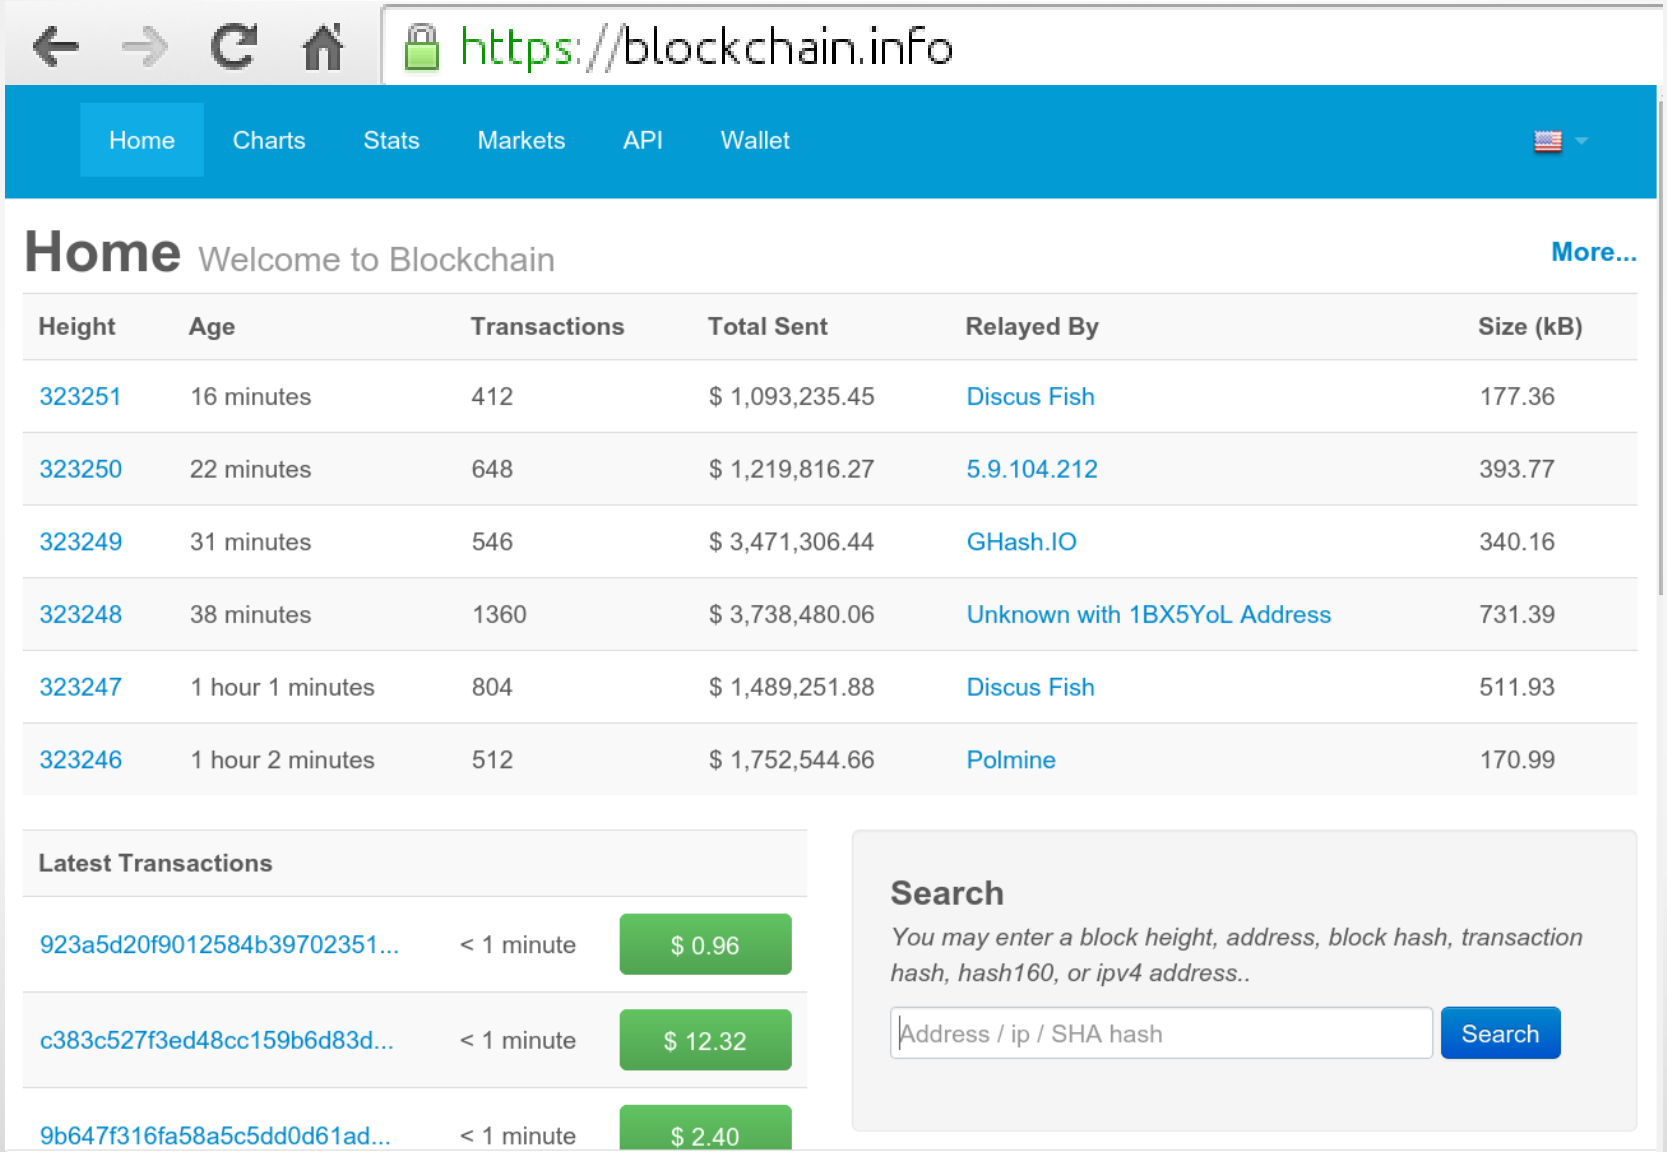

Transaction Tracking

All transactions remain in Blockchain forever.

One can associate addresses by **detective work**.

[1]

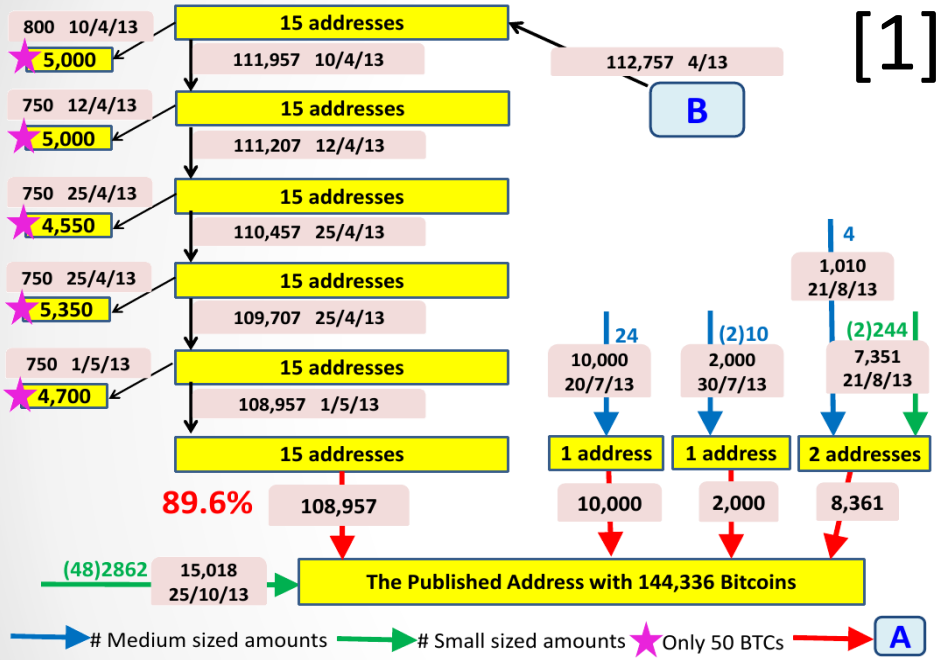

[2]

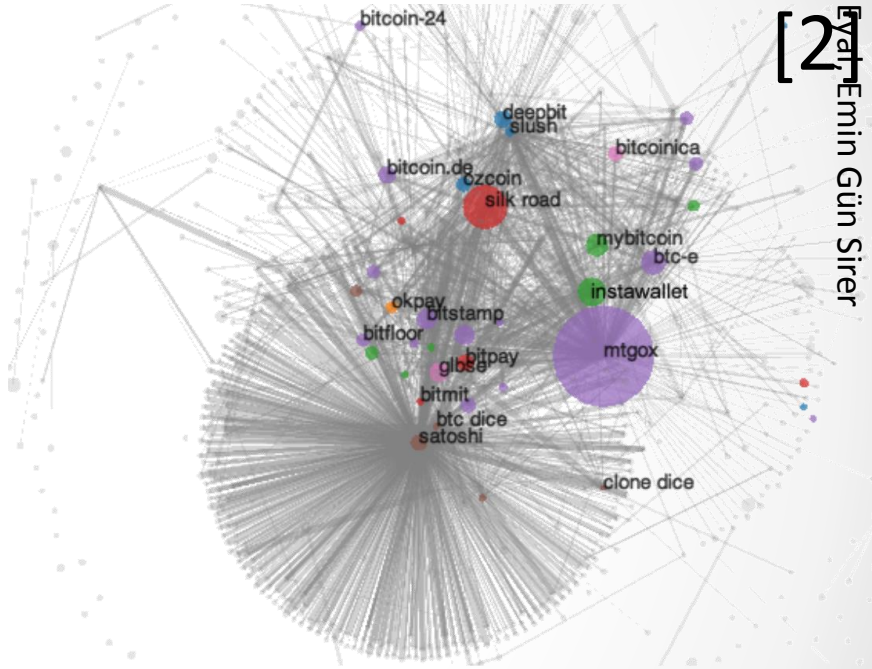

- For large scale crime? Not great.
- For somewhat secret activity? Pretty good.

[1] Ron and Shamir, FC'14

[2] Meiklejohn, Pomarole, Jordan, Levchenko, McCoy, Voelker, Savage. IMC'13

Zerocoin and Zerocash

Zerocash [2]:
Privacy preserving alt-coin on top of Bitcoin.
(preceded by Zerocoin [1])

[1] Miers et al., IEEE S&P, 2013
[2] Ben-Sasson et al., TR, 2014

Zerocoin and Zerocash

The key:
To move funds: prove* that
"I know the secret for moving certain coins".

Without revealing the sources or the value.
But still preventing double-spending.

*Zero-knowledge Succinct Non-interactive Arguments of Knowledge (zk-SNARKs)

[1] Miers et al., IEEE S&P, 2013
[2] Ben-Sasson et al., TR, 2014

Stealth Addresses

The goal:
Untrackable transactions to public address.

The method:
1. Bob publishes address template x.
2. Alice sends Bitcoin to augmented address x'.
3. Bob finds x' and controls it.

- No one but Alice and Bob know x'.
  Need either Alice's secrets or Bob's.
- Only Alice controls x'.

Contemporary Issues

Scalability

Initialization:

- Blockchain over 22GB. Linear growth.
- Long time for bootstrapping

Running (at 7 txn/sec) :

CPU: Insignificant

Memory: ~100MB

Network: ~30Kb/sec

[1] Gavin Andresen, *A Scalability Roadmap*, BTC Foundation Blog, October 2014

Scalability

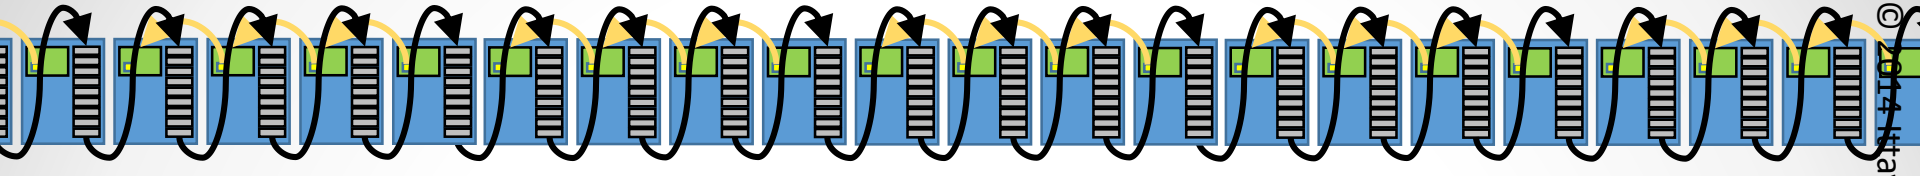

Lightweight clients: **S**imple **P**roof **V**erification

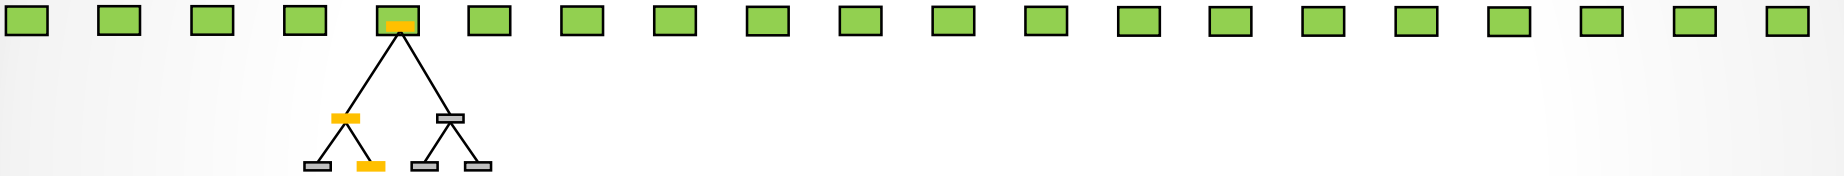

Initialization speedup: [1]

- Headers first
- UTXO first

[1] Gavin Andresen, *A Scalability Roadmap*, BTC Foundation Blog, October 2014

UTXO and Mempool Maintenance

UTXO set becoming large.

Miners can choose to skip transaction verification.

Mempool becoming large

Miners can publish empty blocks.

Block Propagation Time

Block propagation time:

- Too long.
- Depends on block size.

Suggested solutions:

- Transaction set reconciliation.
- Header first.

BIP 70 – Payment Protocol

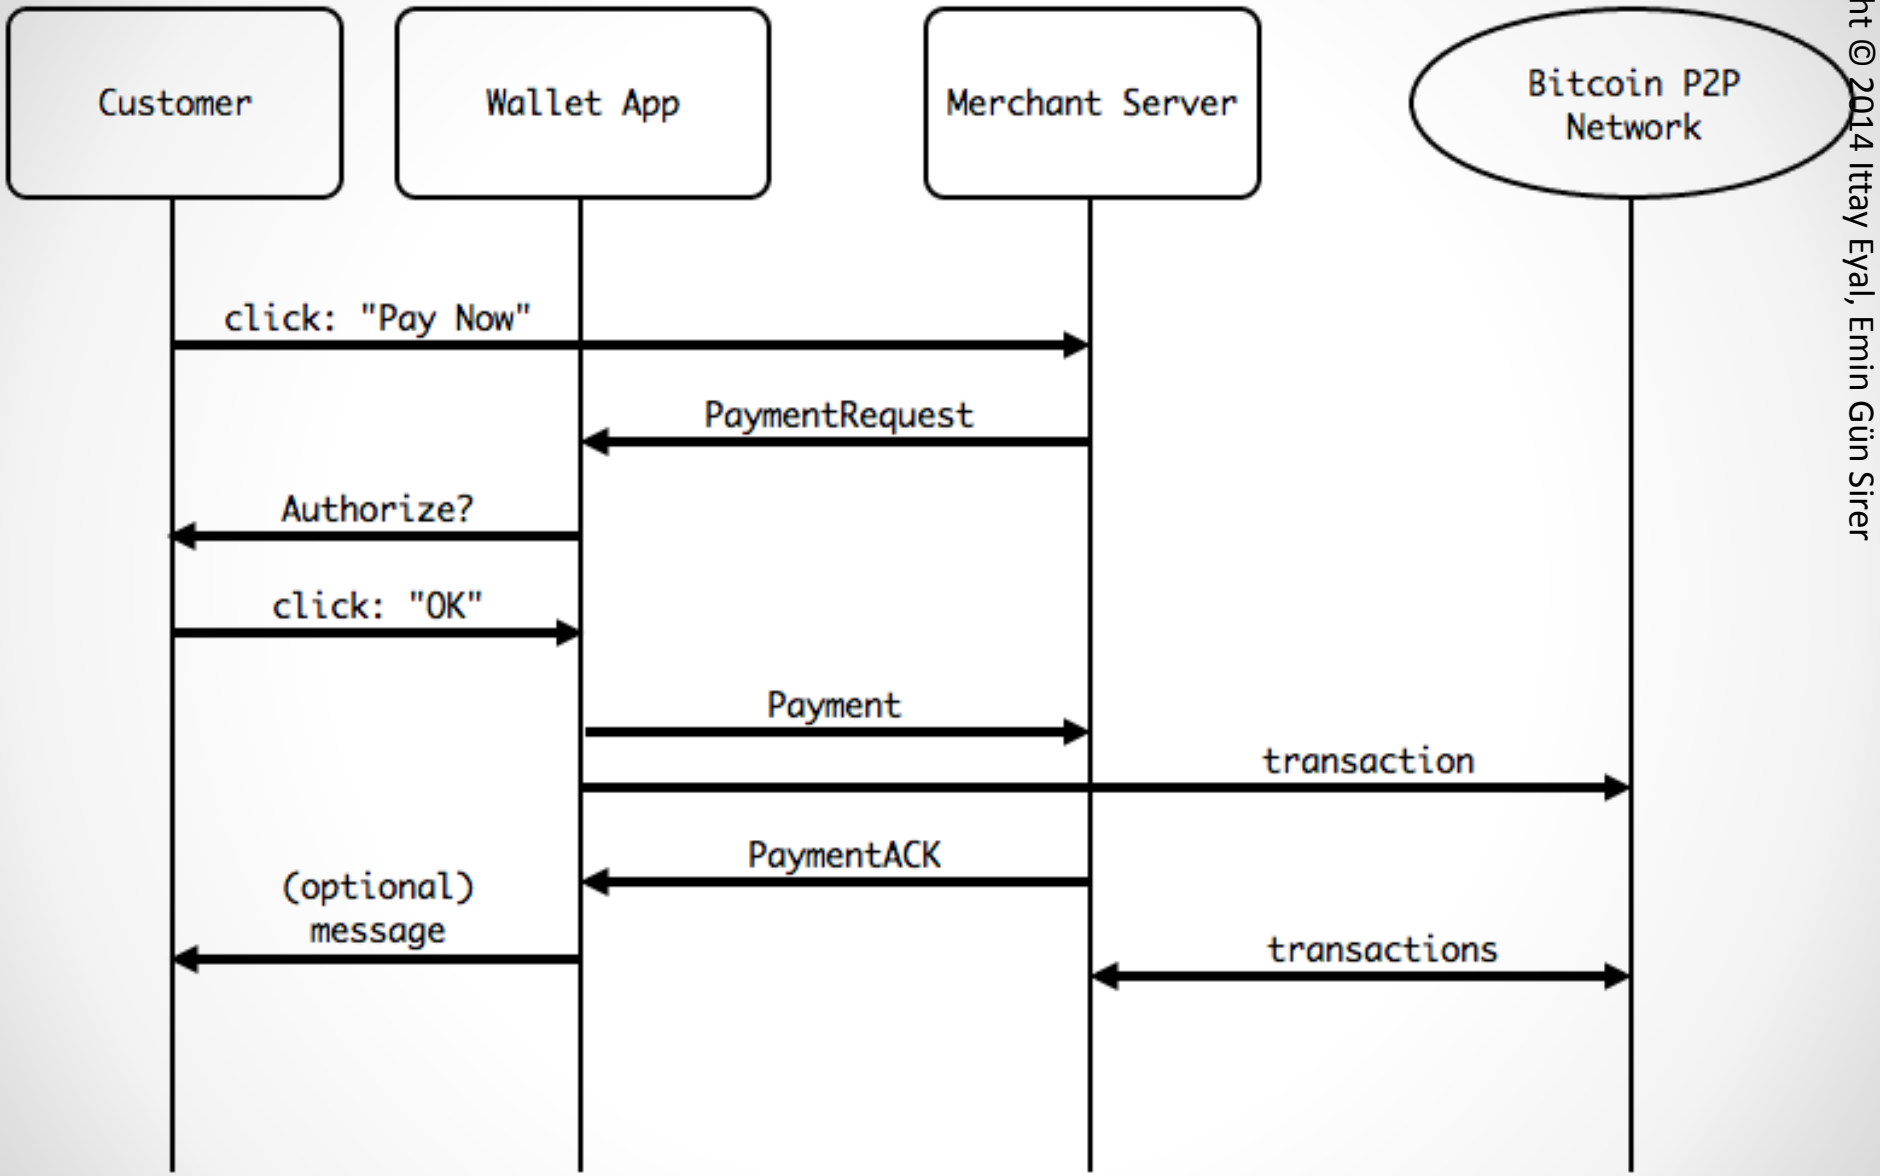

Bitcoin:
Concepts, Practice, and Research Directions

Part IV
Non-technical

Ittay Eyal, Emin Gün Sirer

Computer Science, Cornell University

DISC Bitcoin Tutorial, October 2014

Economy

- Deflationary (21 million total)

- What is it?
  - Store of value?
  - Method to transact USD?

- So what's the potential value (USD/BTC)?
  - Function of mining cost? No! Rate is set.
  - Ratio of world economy?
  - Ratio of world transactions, and a function of the time it needs to store value?

Reasons for Volatility

- Regulation
  - Anti Money Laundering (US/Europe)
  - Adoption / rejection (China, Russia)
  - Fiat regulation (Cyprus, greece)
- Adoption
  - Large companies (Dell, PayPal)
  - Illegal (Silk Road)
- Security
  - Mt. Gox
- Technical
  - Not really

Tax

Commodity or currency? Something else?

Revenue in Bitcoin

Exchange

Mining?

Legal

- Payment for illicit goods.

Shop by category:
Cannabis(203)
Ecstasy(35)
Psychedelics(127)
Opioids(39)
Stimulants(68)
Dissociatives(9)
Other(197)
Benzos(43)

1 hit of LSD (blotter)
฿0.58

1/8 oz high quality cannabis
฿2.05

1 g pure MDMA (white)
฿1.28

Step-by-step:
1. Get **anonymous money**
2. Buy something here
3. Enjoy it when it arrives!

Vacation mode. Important info for **sellers**...

- Money laundering
  - Tumblers
  - w/ pool fees
  - Bitcoin ATMs

Community

Diverse – a lot of players

- Community health
  - Maturing
- Governance
  Mostly the Bitcoin Foundation
  - Protocol changes
  - Interaction with state regulation
  - Bitcoin central bank?
- Large service auditing

Conclusion

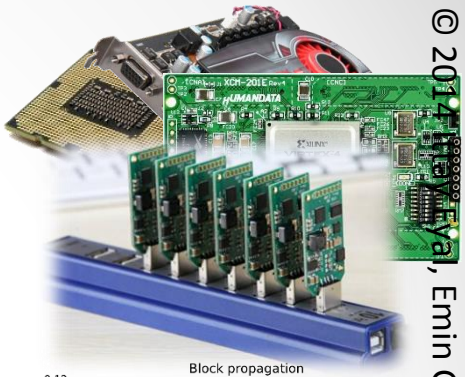

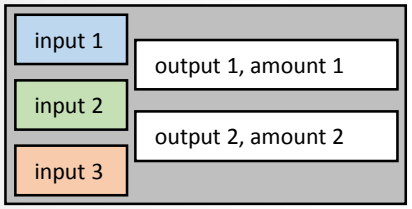

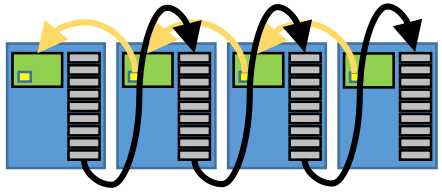

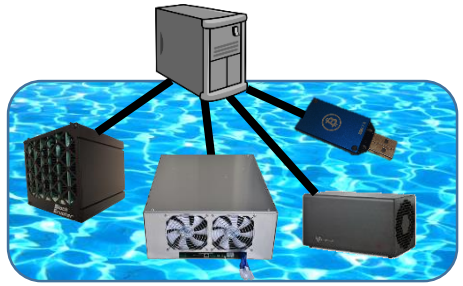

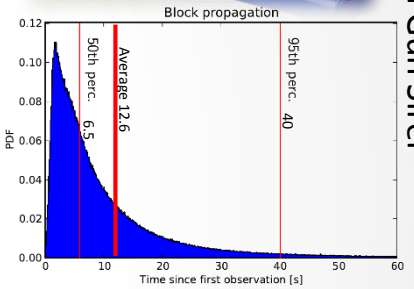

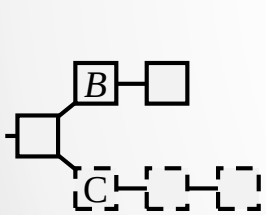

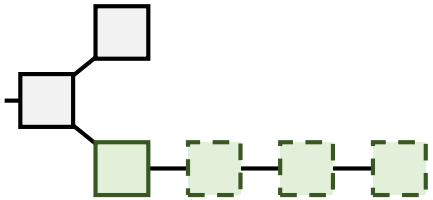

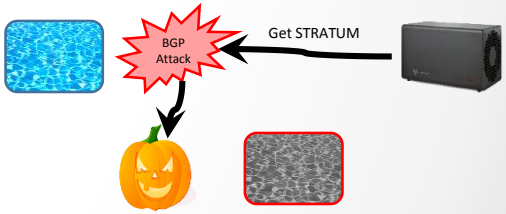

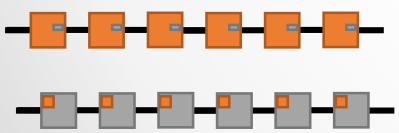

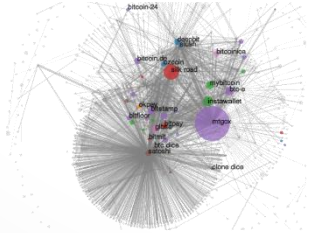

Conclusion