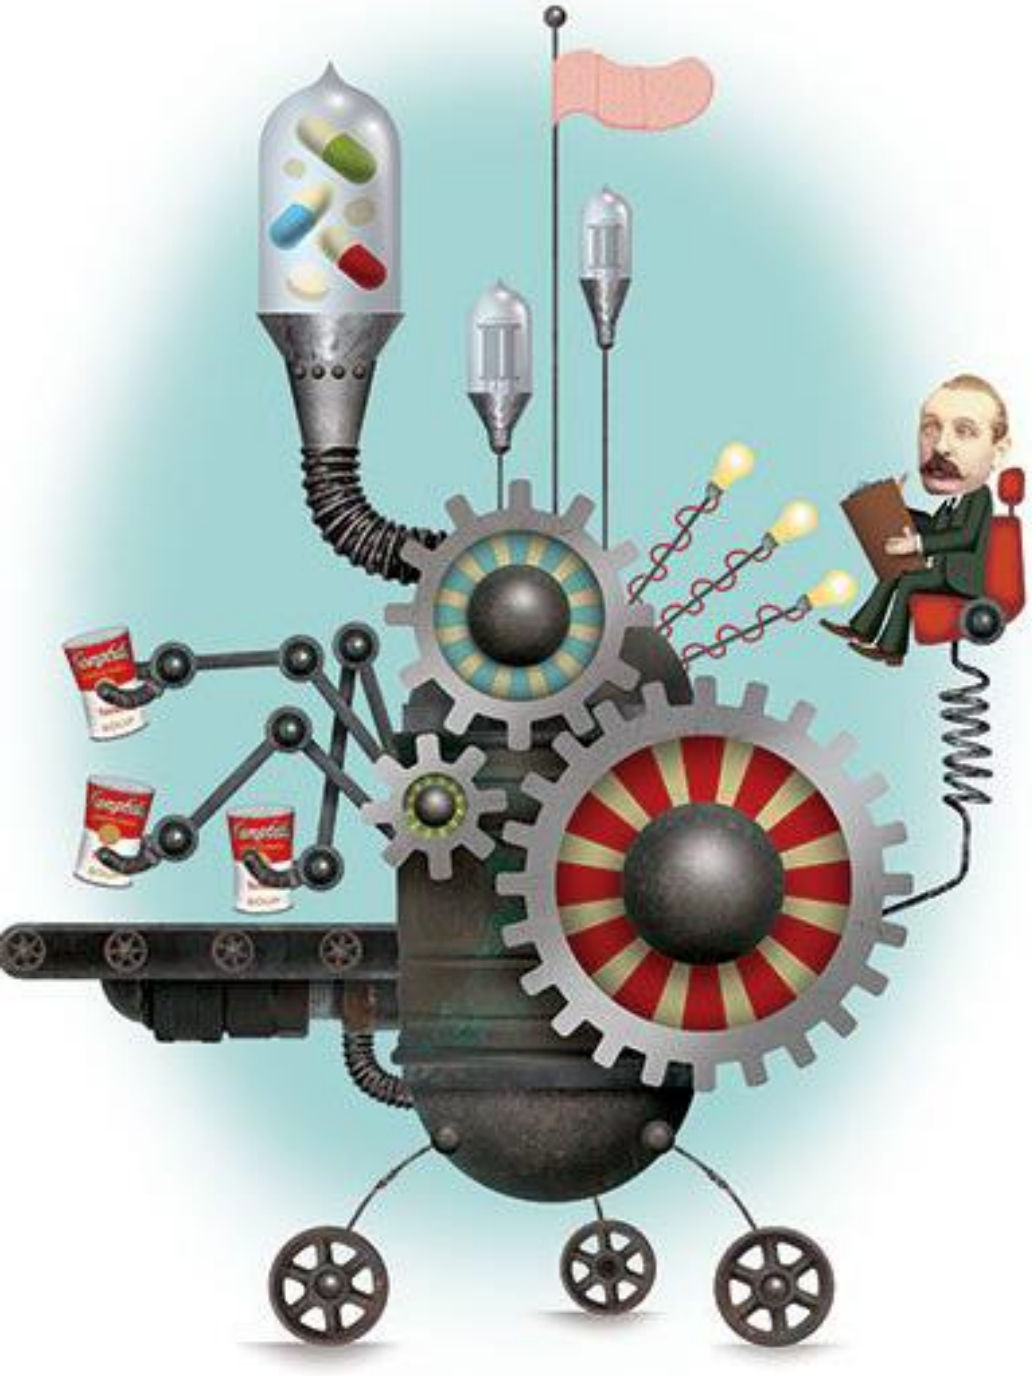

Reinventing the blockchain

Bill Roscoe

University College Oxford Blockchain Research Centre

What do I mean by reinvent?

- Build on the basic idea of a blockchain to create a more efficient and resilient structure
- Building high integrity angelic behaviour around a decentralised community with demonic members.
- Unafraid to throw away some well established ideas
- And embrace ideas not seen in the usual blockchain architectures.
- Take attacks based on non-participation seriously. **What a lazy devil!**

High level aim

- A blockchain that is efficient and secure, and very flexible. And green.
- Should allow smooth transition between public and coalition. Throw away the usual distinctions via a proper theory of decision making.
- Should be self sufficient in the sense that long term drivers of integrity are within chain itself.

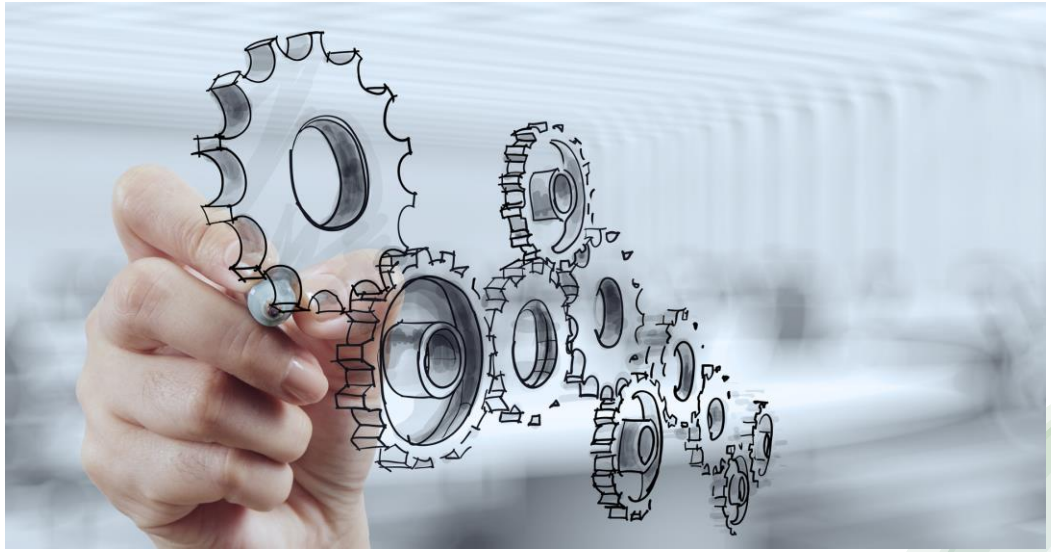

Blockchains are complicated

- Need in general to understand trust, economics, financial regulations, international relations and the nature of society, as well as computer security.
- Impossible to consider all this now.
- We need an understandable interface between the outside world and the technical blockchain:
- Trust model and how to formalise it
- Confidence in the recording and implementation of rules
- And an implementation that is robust under these assumptions from internal and external attacks.
- **Separation of concerns**

The Blockhouse Technology Ltd

Our contributions to the last of these

- Formalising trust models
- Stochastic proof
- Picketing
- Hooks: inserting reverse links to make forking virtually impossible.
- Unbiassable random oracle supporting…….
- Work Your Stake: a carefully designed PoS model.

Formalising trust

- There are many decisions that go to make a blockchain.
- The most important — and summarising most others — are decisions on the validity and finality of blocks.
- We might expect the entire community, or a subset, to make these decisions. The definition of what represents a true decision will be laid down in the chain, as will the groups delegated to make these decisions.
- Relative to a set P of decision makers, there will be a set M of subsets of P such that if any member of M has all its members agreeing on a decision X, then X is officially made.
- It must be impossible for two different members of M all to agree on different decisions.
- Allows for trust models varying, and a continuum between public, where P is everyone, all treated alike, and coalition.

Understanding trust

Trust can be a complex subject. For the permanency of blockchains it needs to take rare events into account……

Independence of risks is important. Is this true of mortgages? Or banks?

Trust can be placed by one person in another; it can be mutual; it can be misplaced since big organisations that were thought to be trustworthy may not be.

Enemy options

- Overt misbehaviour. Taking actions that they expect to be seen soon by good agents and recognised as bat.
- Non-participation: not performing actions expected. May be hard to distinguish from delays afflicting good players.
- Covert misbehaviour. Doing wicked things, such as building a fork, to wheel out later.
- Or combinations.

- Ideally we should disincentivise all but 3 by convincing enemies they will not win.
- **Blockchains that can usually succeed at this can be made much more efficient.**

Statistics

- All public blockchains, and many others, are based on statistical assumptions, such as “at least 51% of mining power is good” and that we can rely on the laws of large numbers.
- The mathematical inferences from those laws do not always make comfortable reading when dealing with strings of, say, 100 blocks.
- It needs to become scientific rather than a belief.
- Define something to be stochastically impossible that is so unlikely we can disregard it.

Stochastic certainty

- Choose a threshold so that in the whole history of a chain, it is most unlikely there will be a counterexample.
- Much less extreme if we can ensure that the number of instances are bounded by a small multiple of the number of blocks, say 10^9, rather than the free for all of PoW.
- So we might set stochastic certainty to be $1\text{-}10^{18}$, set this standard and assume that the likelihood of a block creator being good is at least 50%, we would need 60 pickets. If it were 90% this would be 18. Allowing 2 pickets to be offline/bad would mean that these numbers increase to about 77 and 23 depending on the threat model. All of this is calculated using the binomial distribution: we require the probability that at least p independent choices are all bad to be less than ε, which for us is 10^{-18}. Note that even though some pickets may be the same, the choice of whether the creators of blocks are good or bad remain independent. In statistical terms 8-sigma (an event at least 8 standard deviations from the mean of a normal or binomial distribution.) is close to 10^{-19}.

Decision Thresholds

- Critical to us is SD(N,q) which is the number k of a sample of N independently selected agents, each of which has probability q of being bad, where k being bad is deemed stochastically impossible. Thus if k agents all agree on something one of them must be good. This a safe decision.
- DM(N,q) = max([(N+1)/2,SD(N,q)]) is a dominated, or safe majority, namely a safe decision which is also a majority.
- And SDM(N,q) = [(DM(N,q)+N)/2] is a strong dominated majority, guaranteeing a majority of the good agents agree, as for Byzantine agreement.
- Which of these is appropriate depends on the nature of the decision and the security of communication. Are bulletin boards secure?

Hooks: preventing forks

- Blocks sign statements saying what their N-fold successor is. N>1 to allow for misbehaviour by bad miners.
- Present to eliminate other long term state for justifying choice of next block.
- Reflects recent consensus decisions rather than informing them.
- Allows efficient identification of true blocks of the chain.
- Powerful against forks
- They do the same as hash links
- Only in the other direction
- **Heads down check**

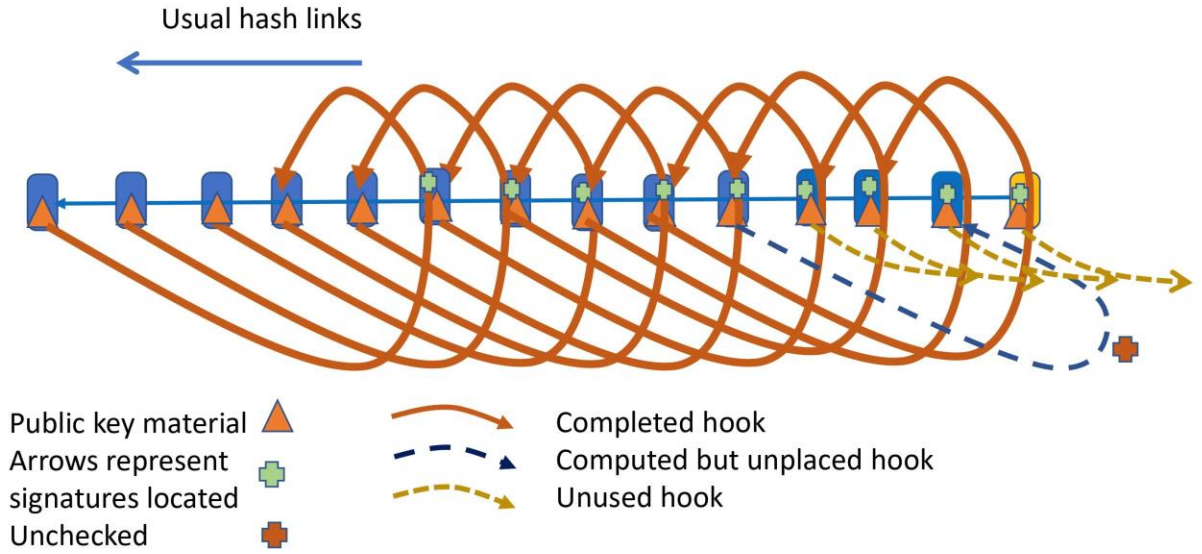

History

- The effect of hooks is that each block B gets N separate endorsements from the creators of earlier ones.
- Once these are in place it will be permitted for blocks to drop analysis of B from their BBs.
- If there are a significant number missing we might want to create a record of this information.
- A detailed analysis of the combinatorics of this, remarkably similar to that of picketing, can be found in the Hooks paper.
- Hooks give us the opportunity to circumvent misbehaviour through non-participation with a back-up mechanism.
- The growing chain maintains the invariant of there being a stochastic proof of the uniqueness and correctness of each block not close to end of the chain embedded within it.

Pickets

- We rely on small groups of agents to control blockchains in the absence of much overt misbehaviour or non-participation. Pick groups that we think are likely big enough for this, and certainly big enough to mean that no decision is made without a (or a majority of) good ones agreeing.
- But overt misbehaviour or non-participation on a moderate scale may thwart progress.
- So we need a back-up mechanism to take over if no decision emerges…..
- All based on the consensus machine…..

The consensus machine

- A group of agents making a decision progress through a defined state machine.
- Synchronising through agreements strong enough for thresholds M.
- When sufficient have signed up to the same state or decision, this represents a state change and others can catch up.
- Aim should be to make a defined high-level decision, such as the next block.
- Primary mechanism must be safe.
- Backup mechanism — larger group of agents — should be sound and complete.
- Getting a safe handover mechanism to work with these decentralised machines was difficult. Much abstraction into CSP, and FDR proof.

The consensus machine

- Allows a clear understanding of decentralised consensus
- And makes it easy to design consensus algorithms.

Determinism

- Because of the concept of proof by dominating majority we need good nodes to agree on whether nodes etc are acceptable or not.
- We have therefore proposed a traffic-light scheme for checking flaws that should lead to unanimity amongst good pickets.
- It means that variance in assessments at this level is potentially disastrous. So we need exceptionally clear specifications of what the result of an assessment should be. Particularly if there are multiple implementations. Presented with two decisions and the same evidence, good agents should produce the same result.

Work your Stake

- A proof of stake model which simulates the economic model of PoW.
- So miners buy mining tokens rather than pay for mining power, electricity.
- They make deposits to ensure good behaviour and the picketing model holds this back for a significant time.
- Like other PoS models, it depends on an unbiassable random oracle for its correctness…..
- We would like the entire blockchain to become a giant proof by induction. That depends on the unbiassed choice of pickets and similar.


How to
stake a
Vampire
Blockchain

Unbiassable random oracle

- We want nodes to contribute so some are surely good
- And the contributors are completely fixed before anyone could have any information about the value.
- They cannot even choose not to contribute— as that would allow them to bias the result.
- We get block creators to embed commitments to shares of later oracles in their blocks, including delay encryption that can be opened without them.
- Delay encryption is here also a multi-party computation using the blockchain.
- Again subject to stochastic analysis.

Verification

- We use verification when we can
- ….and specification in appropriate notations.
- I have used CSP/FDR to develop and prove the handover protocol.
- And a high-level CSP blockchain
- Interested in the automated generation of consensus machines.

Further reading

- See papers at https://tbtl.com/resources/
- Delay and escrow in the blockchain
- The consensus machine……
- Embedding reverse links…..
- The greening of blockchain mining
- And more

- I am organising an academic workshop in Oxford 26/27 June on delay encryption and its applications, sponsored by crypto.com. Contact me for details.