INPUT | OUTPUT

IOHK | BLOG

On the Ouroboros Design: How rigour and engineering are essential for critical infrastructure

JANUARY 11, 2018 AGGELOS KIAYIAS 6 MIN READ

& IOHK Research

A blog post on the Steemit website appeared recently making a number of claims regarding Ouroboros. The article contains several factual inaccuracies. For instance, it is claimed that "DPOS" in the Ouroboros paper stands for "delegated proof of stake", while in fact, DPOS means "dynamic proof of stake", or that the protocol requires a "2/3+" ratio of parties being honest, while in reality it just requires an honest majority, i.e. the stake controlled by parties following the protocol is more than half the total stake.

For the benefit of those that are interested in the Ouroboros protocol and who appreciate its general philosophy, we feel it is appropriate to provide here a response to this article making along the way a few broader points. While pointing out inaccuracies in the blog, we take the opportunity to highlight some of the general approaches followed in the design of Ouroboros and in the related research efforts that are currently underway at IOHK.

Ouroboros is a proof of stake (PoS) protocol that uses delegation in the spirit of the PoS idea as discussed in the Bitcoin forum starting from 2011. The references that influenced its design are listed in our paper. PoS is a powerful concept that has inspired a number of other efforts prior, concurrent and post the first Ouroboros paper. Among all other implemented PoS blockchain systems that carry real assets, Ouroboros is unique

in that it was designed in tandem with a formal security model and a mathematical proof that it implements a robust transaction ledger. This marks a fundamental shift in the methodology of blockchain system design.

Blockchain systems are in a period of transition from curiosities to critical infrastructure; as such, the all too typical software industry approach of releasing a "minimum viable product" as early as possible and then fixing bugs as they appear, is not appropriate. Failures of critical infrastructure have a significant impact on people's lives and thus require rigorous engineering discipline to the highest possible standards. Dependability, rather than maximum performance according to some arbitrarily chosen metric, is the primary goal. Performance is important, of course, but the performance required is a function of the ultimate application domain, and from the point of view of dependability it is the worst-case performance that is important, not the ideal-scenario peak rate.

Like all other protocols in the blockchain space, Ouroboros requires some degree of synchronisation. The block production interval has to be consistent with the likely time to complete the required information exchanges. The 20-second slot time in Ouroboros represents a conservative choice for a block of transactions to traverse the diameter of a peer-to-peer network, where the peers may be significantly geographically distributed, the system is operating at peak transaction load and the interconnection is significantly less than perfect. It is improbable for a block of transactions to consistently traverse a global network much faster than that, and as a result any solution that does significantly better (or claims to do significantly better) is either wrong, or provides a weaker level of decentralisation or security, i.e. it solves an easier problem than Ouroboros. There is a tradeoff between achieving a robust, global, participatory service that delivers sustained effective performance even under an adversarial attack, and creating a high performance, limited participation (in geographical scope or network resource requirement) solution that makes overly optimistic assumptions on network stability.

Irreversibility, the property that transactions persist and are immutable in a blockchain protocol, has to be presented as a function of the level of the adversarial strength. This is true in Nakamoto's Bitcoin paper and also in the Ouroboros paper, see Section 10.1 for the actual time needed for confirmation of transactions. Thus, one should be very wary of statements about irreversibility that do not quantify the level of adversarial power. For instance, Ouroboros will confirm a transaction with 99.9% assurance in just five minutes against an adversary holding 10% of the total stake, which in today's market cap in the Cardano blockchain would amount to more than two billion dollars. Byzantine agreement protocols can provide a more "black and white" irreversibility, in other words the protocol can be guaranteed to be irreversible within a certain time window provided an honest majority or supermajority exists depending on the protocol. Nevertheless, the performance and decentralisation penalty suffered is very high if the level of adversity is allowed to come close to the 1/2 barrier, which is the level of adversity that Ouroboros can withstand.

The issue of possible dominance of the consensus process by a small group of stakeholders holding a large proportion of the stake is important but is not applicable to the current release of the Cardano system (the Byron release). What we have proved for Ouroboros is that it can facilitate a "fair" transaction ledger (where fairness here means that the ledger can fairly record all significant actions that are performed by the protocol participants despite the presence of an adversary). This enabled us to neutralise a number of rational protocol deviations (e.g. the equivalent of selfish mining attacks in the PoS setting) and provide a Nash equilibrium argument showing how the protocol can support many different types of mechanisms for incentivising participant behaviour. Currently, IOHK Research is actively working to finalise the incentive structure that will be incorporated in the Shelley release of Cardano, where stake pools will be supported and delegation behaviour will be properly incentivised so that it offers effective decentralisation of power. The crux of our methodology is the engineering of a novel reward mechanism for rational participants that provides appropriate incentives to partition their delegation rights. The objectives are first, to avoid concentration of power to a small group of participants – as it could happen by a naïve reward mechanism in a Pareto distributed stakeholder population – and second, to provide appropriate incentives to ensure a desired number of delegates. We are very excited about this work; it will be the first of its kind in the area and, as before, we will be disseminating it widely including full technical details, as well as submitting it for peer review.

This brings us to the final distinguishing advantage of the philosophy of Cardano. Scientific peer review has been refined over centuries. The way it is implemented by the International Cryptology Conference (also called Crypto), where Ouroboros was presented, and the other top conferences in the area, strives to remove conflicts of interest and produce the highest level of objectivity. The method of reviewing is known as "double blind", i.e. papers are submitted anonymously and reviewers are experts that also remain anonymous to the authors. The committee of experts that reviews submitted papers each year is formed by two program co-chairs that are appointed by the International Association of Cryptologic Research, the pre-eminent organisation of cryptology research that was founded in 1982.

Being invited to serve in the committee as an expert is an important recognition of an individual's long-term commitment to the area of cryptography (and even a precise count of how many times one has served is maintained). Blockchain protocols fit perfectly within the cryptography scientific literature and thus scientific peer review is to be done by this community. Of course, we welcome reviews from anyone. That is why we make public very detailed whitepapers with precise and specific claims that leave no uncertainty about what is being claimed, and we appreciate any factual discussion about any of these claims. We strongly encourage other projects to submit their work for scientific peer review as well. They will enjoy the benefits of thorough, well-founded and objective critique and they will have the opportunity to showcase any advantages and novelty that their approach possesses.