Linzhi Semiconductors

Linzhi Working Papers
No 8

# ECIP 1043 Revisited.

by Sonia Chen

March 2019

Keywords: LWP8, ECIP, ECIP 1043, Ethereum, DAG

This publication is available on https://linzhi.io
Telegram discussion https://t.me/LinzhiCorp

# ECIP 1043 Revisited

Linzhi ASICs
Mar 9 · 4 min read

ECIP 1043 proposed to set a limit on the maximum size of the DAG, and no longer increase it on epoch schedule.
https://github.com/ethereumproject/ECIPs/pull/91

ECIP 1043 was closed by its author Cody Burns in December 2018, but today we take another look at motivation and specification, from our perspective of being a Shenzhen asicmaker (with upcoming Ethash chip).

Motivation (from the ECIP):

- DAG has failed at ASIC Resistance and now serves as deterrent to broader investment in ASICs
- increasing DAG has negative impact on security and acts as anti-competitive bug preventing investment in mining
- increasing DAG will eventually reach limit of GPU bus speed
- smaller DAG allows more GPUs to mine while lowering entry cost

Specification (from the ECIP):

- *if(blkNumber > forkblock) DAG epoch = 0*

. . .

## ASIC Resistance

We see ASIC Resistance as an economic concept, it increases the cost of ASICs. As such the ASIC Resistance of Ethash has worked.
Maybe that was an unintended consequence, but unintended consequences of major decisions are inevitable and now it seems best to focus on learning instead of reacting with sudden changes that may lead to new unintended consequences.
https://hackernoon.com/asic-resistance-is-nothing-but-a-blockchain-buzzword-b91d3d770366
Keep calm and hash on.

We don’t think the DAG as it is growing in the current protocol (ca. 50 MB/month, 600 MB/year) creates sudden shocks when reaching any GB threshold, even the currently widely used 4 GB (obsolete 2020–11–30).
There is predictable obsolescence for 10 GB and up into the 2030's.
https://investoon.com/tools/dag_size

Ethash has served well since 2015 in creating a stable and predictable mining environment, including the emergence of the Bitmain E3 and Innosilicon A10, as well as upcoming chips like ours or others.

. . .

## Current Mining Landscape

Mining profitability (currently 1.3 cents/MH/day) should be looked at separately from PoW algorithm study. There are many mining business models, welcome ones (using peak wind energy at negative power rates, heaters), unwelcome ones (money laundering), and more.

Our estimates for the current Ethash mining memory distribution are (obsolescence dates for ETC from investoon.com, ETH is comparable):

- 2019–03–17 4% GPU 3GB
- 2020–11–30 10% Bitmain E3 4GB
- 2020–11–30 36% GPU 4GB
- 2024–04–30 36% GPU 6GB
- 2027–09–29 9% GPU 8GB
- 2030+ 5% GPU 10GB and up

. . .

Objectives

- Predictability
  Predictability is connected to fairness, competition, stability and even recyling, and looks like the most important consideration to us. Predictability creates fairness. Shocks and nonlinear benefits should be minimized.

- Cost Advantages
  Decentralization is the result of three goals mainly — to have a system that is trustless, permissionless and resilient.
  From a hardware perspective, avoiding centralization means avoiding persistant cost advantages. Thorough cost analysis must be standard practice.
  What works against this is that you are never sure there is not some unknown way to be faster or cheaper. Calculations are based on one's own assumptions and may not apply to others. We can calculate the limitations of things we know, but we cannot prove that there cannot be something better we don't know.

- Competition
  There are currently 4 vendors for Ethash (Nvidia, AMD, Bitmain, Innosilicon), a 5th one announced (Linzhi). There may be other startups preparing for solutions in the future, but the current economics make investment difficult for everyone.

- Recycling
  We think existing miners should not be bricked frivolously. PoW changes should consider input from chip designers in how the resulting chips can be more useful to scale crypto applications beyond just PoW. Usefulness and security are not mutually exclusive.

- Immutability
  Every hardfork, and every feature in a hardfork, threatens the promise of immutability and should go through a rigorous process — ECIP.

. . .

Technical Considerations

https://github.com/ethereum/wiki/wiki/Ethash#definitions

- seed -> cache -> DAG
- cache is DAG/64
- current DAG is 'relatively' dynamic, the entire 3GB changes every epoch (4–5 days), and grows by approx. 8 MB/epoch
- DAG growth approx. 50 MB/month, 600MB/year
- ROM is 3x cheaper than RAM
- constant data encourages over-specialized development that is technologically a dead-end
- SRAM and eDRAM need to be considered
- maximum chip size is, generally, 26mm x 33mm
- we assume it would be easy for Bitmain to come out with a 6GB E3
- we assume it would be easy for Innosilicon to come out with an 8GB A10

The ECIP 1043 pseudo-code does not clearly specify whether the entire DAG becomes a constant, or whether the cache keeps being modified and the 1GB DAG changes every epoch. We assume that the author did not unintentionally want to create a constant DAG, and an actual implementation would have continued to modify the DAG every epoch, just not grow it anymore. A constant 1GB DAG could allow for single-die ROM miners.

Even a dynamic DAG at 1GB could open the possibility of an ASIC beating a GPU by 10x. 16nm or deeper processes allow integrated 1GB RAM on a single silicon die. On-chip SRAM is >100x faster than GDDR or HBM.
Since Ethash loops are mostly memory accesses and few logic ops, it means the memory bandwidth bottleneck would be completely removed in this case.

. . .

Summary

- We think allowing the DAG to grow at the current approx. 8 MB / epoch (50 MB / month) rates high on predictability, and thus stability, fairness and security. This includes obsolescence of 4 GB GPUs (and Bitmain E3) in December 2020.
- If people believe DAG growth should be limited, we propose 3 GB because that is proven to be stable right now. We also think it's important that the data itself remains dynamic (changes with every epoch), and does not become one large constant. ROM is 3x cheaper than RAM.
- To increase predictability and fairness, we propose a significant interval of one year between hardfork and activation of DAG fixing:
  *if(blkNumber > forkblock + 2200000) DAG epoch = 250*

. . .

If there are no flaws in our analysis or reasoning, something is seriously wrong! If you spot mistakes or just plain disagree, please let us know.

Linzhi Shenzhen
Telegram: https://t.me/LinzhiCorp