Linzhi Semiconductors

Linzhi Working Papers
No 16

Linzhi E1400 – Architecture Overview.

by Werner Almesberger

October 2019

Keywords: LWP16, Ethash, ASIC, Architecture

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

Linzhi E1400 — Architecture Overview

Linzhi ASICs

Oct 8 · 10 min read

[This article is an edited version of the talk given by Werner Almesberger at ETC Summit 2019 in Vancouver, on Oct 4th, 2019.] Video: https://www.youtube.com/watch?v=_enGUG_Jbi4

2019 ETC Summit, E1400 Ethash Architecture Overview

. . .

Good afternoon, my name is Werner Almesberger.

I've worked on the Linux kernel many years ago, and today I talk about what I am working on at the moment. For a change that has nothing to do with software, it's a mining ASIC for ETC, for Ethash. Linzhi is a Shenzhen-based fabless semiconductor company, we announced our Ethash miner at the last ETC Summit, expecting a performance of about 1.4 GH/sec at a power consumption of 1 kW.

. . .

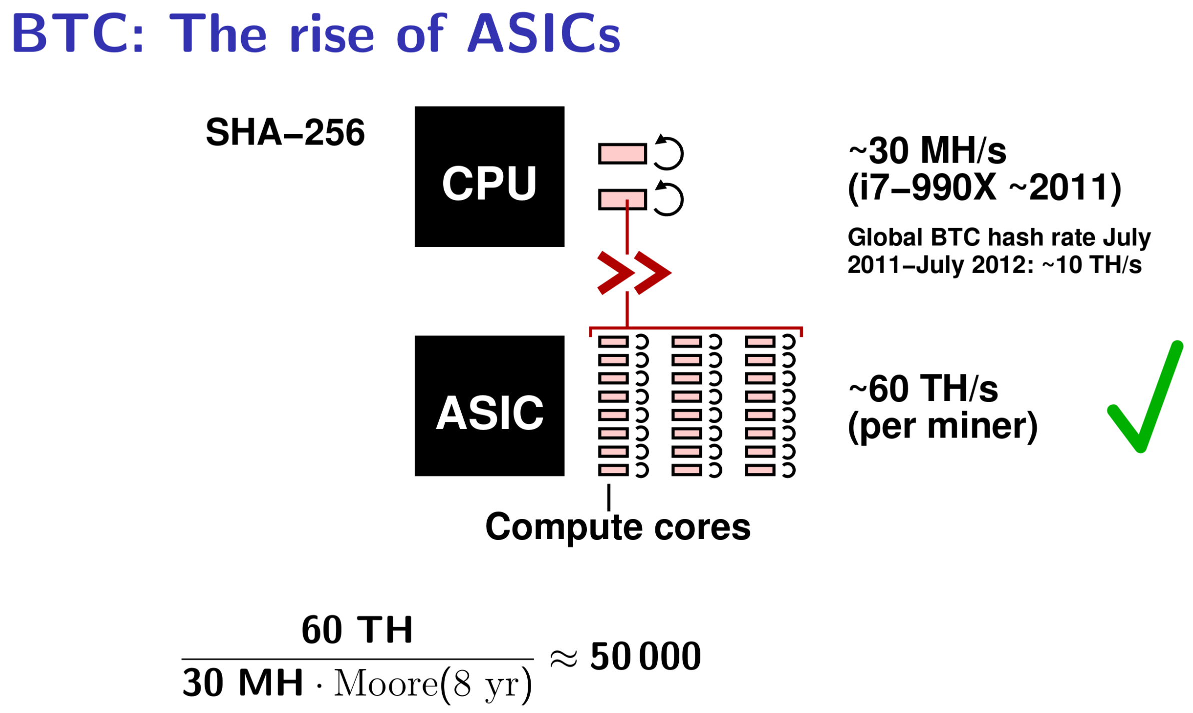

BTC: The rise of ASICs

SHA–256

CPU

~30 MH/s
(i7–990X ~2011)

Global BTC hash rate July 2011–July 2012: ~10 TH/s

ASIC

~60 TH/s
(per miner)

✓

Compute cores

$$\frac{\textbf{60 TH}}{\textbf{30 MH} \cdot \mathrm{Moore}(8\ \mathrm{yr})} \approx \textbf{50 000}$$

Reviewing the beginnings of Bitcoin, there was CPU mining with SHA-256, running on your CPU with one or a small number of compute cores. People quickly realized that you can do better, because the CPU is a big chip with a lot of transistors, you have floating point and other things, even if none of it is used for mining. So why not reduce the functionality, and do only what you actually need to do, namely the hash calculation, and put this on a chip — then you can have a lot of little compute units in the same space, so the chip costs roughly the same. They all work in parallel, and you can do lots of hash operations with the same effort, and gain a huge benefit over the CPU. This is how ASICs were born. The best CPU miners in 2011, when people basically gave up even trying to CPU-mine, did something like 30 MH/sec, while a contemporary ASIC miner today does 60 TH/sec. If you adjust those values for technological progress over time, using Moore's law, then we would get a performance difference between a CPU extrapolated to the present day, and the present day ASIC miner, of about 50,000 — in terms of hashing speed.

You may not be interested in hashing speed per se, but more in power efficiency. There the factor is a little bit smaller, maybe something like 5,000 or a bit more. Still, ASICs have a huge benefit over CPUs. If you had a time machine, sitting around somewhere in your basement, and you would like to put this to some creative use, you could go and shop and buy a state-of-the-art ASIC miner for Bitcoin, hop into your time machine, go back 8 years, get out of your DeLorean, plug in your miner, and suddenly you would have increased the global hashrate from 10 TH to 70 TH. We are not talking about a 51% attack, we are talking about an 85% attack. You basically own the network.

Clearly, when people saw what was happening with the ASICs, they were a bit shocked. And that caused a lot of bad blood with respect to ASICs. Some people started to think about ways to solve the ASIC problem, and this is how they came up with Ethash.

. . .

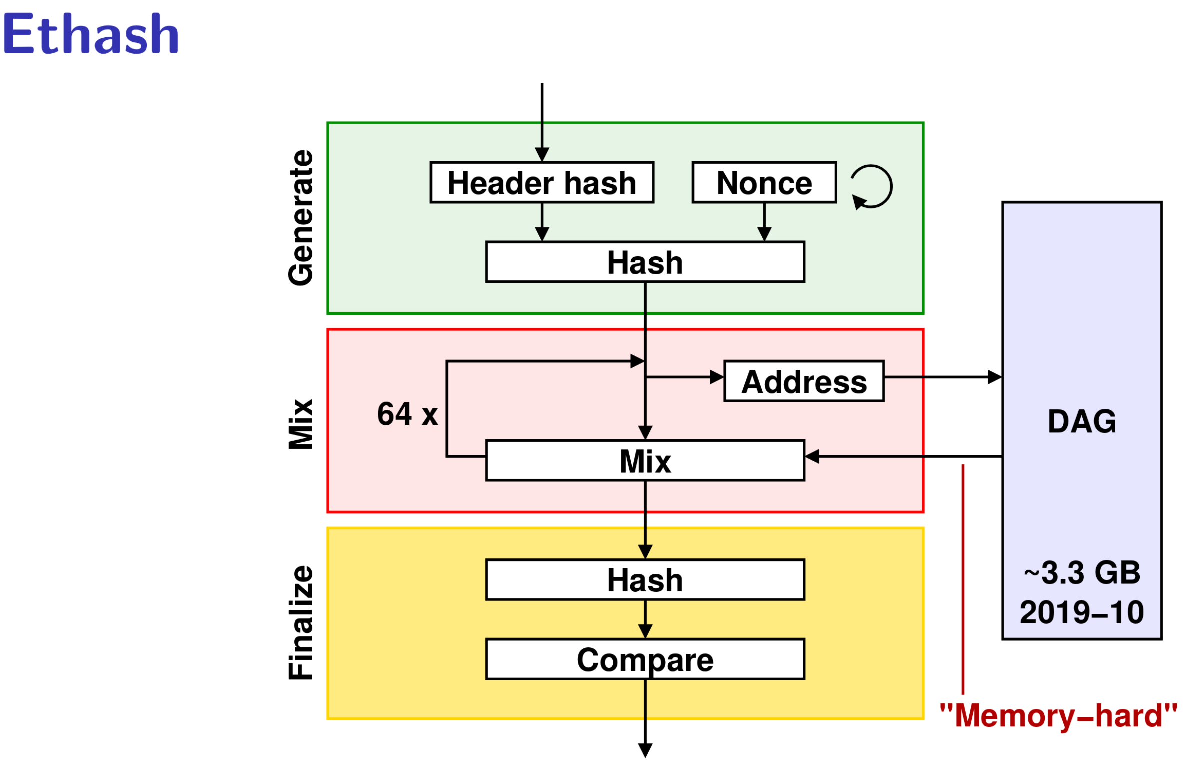

In Ethash what you get is a header hash value, and you have your nonce, which you change on your own for the attempts you make. You hash those two together, and with this you have generated your task. Now you do the mixing, where you first calculate an address from your current value, and lookup this address in the DAG, a big block of memory, and you retrieve the corresponding memory line. You mix it with the result you have so far, and you repeat the whole operation 64 times, then you do some more hashing, and you compare with the difficulty target. The DAG is the new ingredient that should make the thing more difficult for ASICs. It makes it memory hard, and the DAG is very large, so you can't just put it on a little corner on your chip. It's currently about 3

GB and it's growing over time.

Let's look at the comparison:

. . .

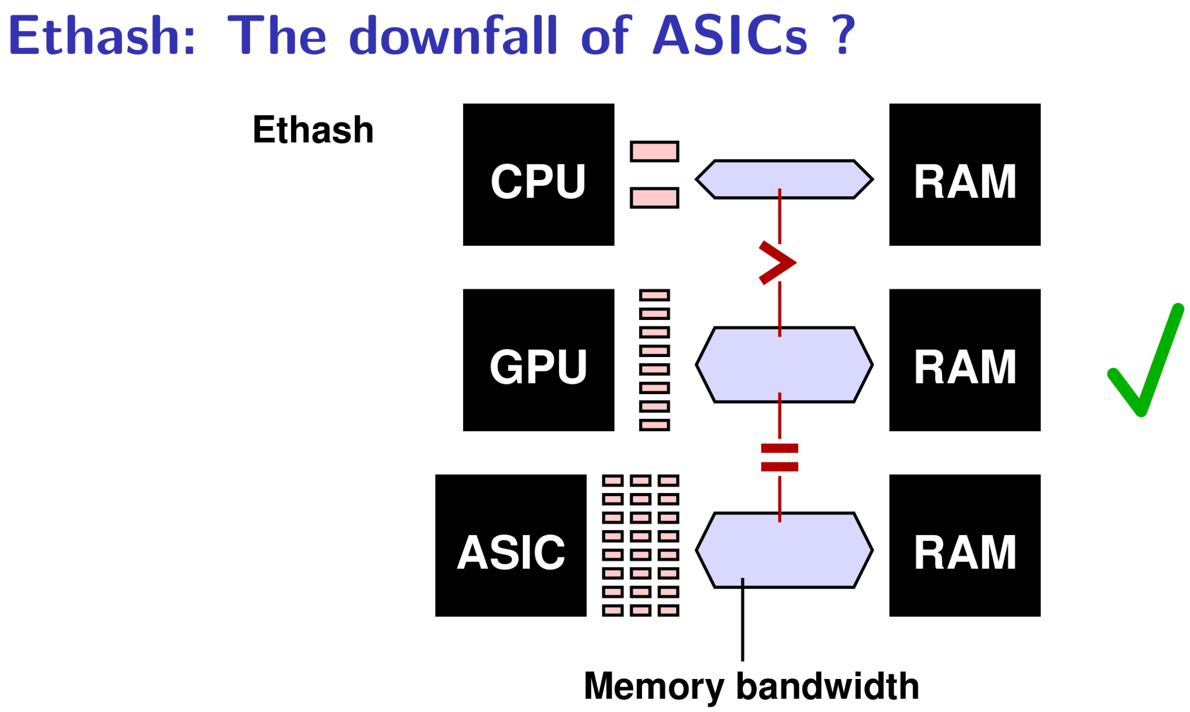

Ethash: The downfall of ASICs ?

Ethash

CPU

RAM

>

GPU

RAM

✓

=

ASIC

RAM

Memory bandwidth

You have your CPU miners, a small number of cores, a reasonable amount of memory bandwidth. You need both, you need to compute, and you need the memory bandwidth.

With GPU miners, you can do better, they are more specialized, so your compute cores get smaller, you can have more of them at the same chip surface, and they are also optimized for memory bandwidth. They have vast memory bandwidth, they have wide buses. So they are better than CPUs for mining.

ASICs could of course still win on the side of the compute units, because the GPUs are still sort-of general purpose. You could make an ASIC that is only Ethash, but you would still have the memory bandwidth as your limiting factor. So unfortunately, you are not really winning. This means that you might be able to more or less pull even with the GPUs, but then GPUs are commodity hardware, they are easy to get, a big market, you have economies of scale on your side with a GPU.

So, ASICs just shouldn't make sense, and this would mean that those evil ASICs would stay away from the precious Ethash. So far the theory.

. . .

[Please briefly review the Ethash diagram above: Green is the generation of the task, red the mixing, yellow the finalizing, and blue is the memory.]

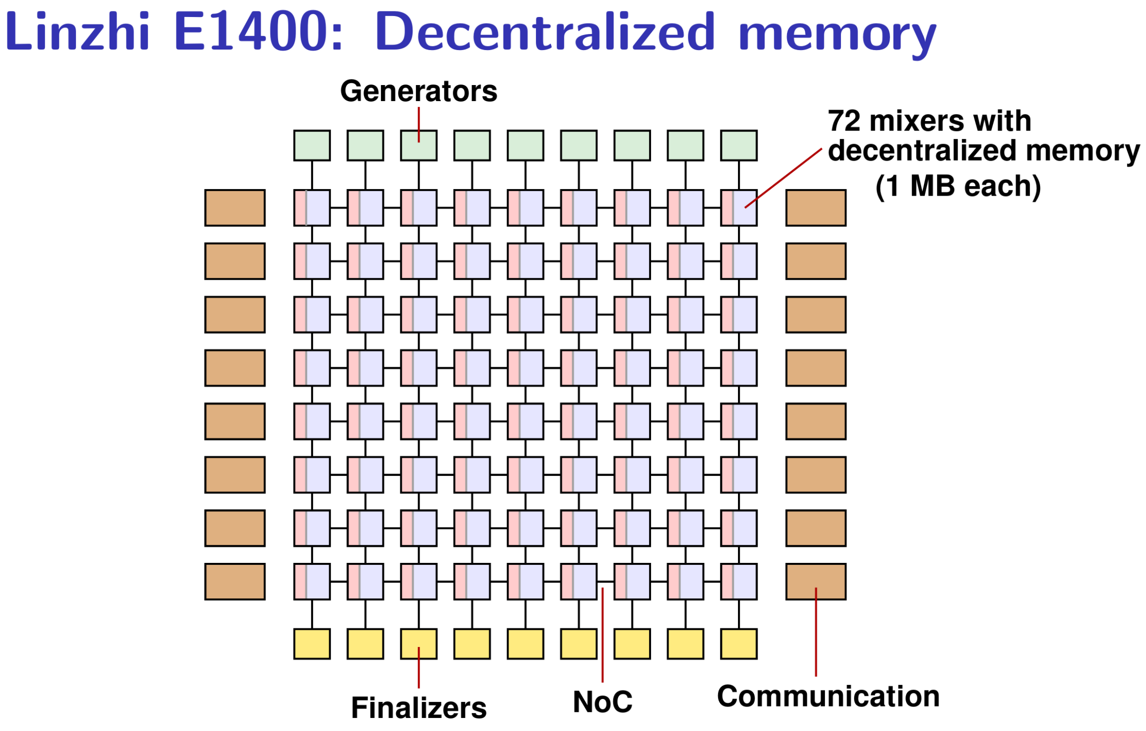
Linzhi E1400: Decentralized memory

In our architecture, this is what our chip more or less looks like. We have those little green boxes which do the generation. We have multiple of them so they don’t have to work too hard. And then, when the task is generated they send it to a so-called mixer node. The mixer node contains a bit of logic to do the computation, and a bit of memory. And there are lots of mixer nodes, 72 of them. They are connected with each other through a network-on-a-chip. And, at the bottom you can see the yellow boxes which are the finalizers which do the final hash and the comparison. Each of the 72 mixing nodes has 1 MB of memory. This means 72 MB of memory per chip. Which is too small for the DAG. We can solve this by using multiple chips, distributing the DAG over those chips, but then we do also need to communicate between chips, and so we add — which is shown on the side — communication interfaces where the chips can talk to each other.

. . .

Hashing with the E1400

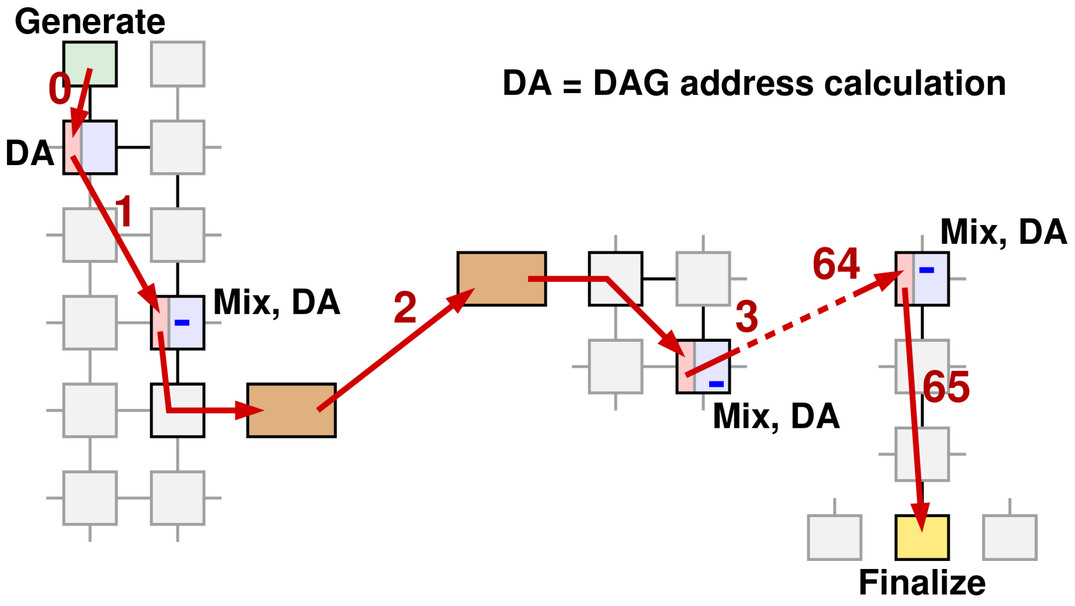

Now let's look at how the hashing operation works. First of all you do the generation, which we call round 0, then you send it to the adjacent mixer node, where only the DAG address is calculated, because we don't know yet where this thing has to go. Then we send it to the node that has the corresponding DAG line in it's memory. The DAG line gets retrieved, the mixing is done, it calculates the next DAG address for the next round. Let's say this DAG address is in the memory on a different chip, so it gets sent to the node where the communication interface connects, then it gets handed over to the the communication interface, arrives at the destination chip, and there it enters on the node where the communication interface is attached. This will then route it to the node which has the memory. There we again retrieve the memory line, do the mixing, calculate the next DAG address, and so on, 64 times in all, until the last round is done. And then it gets sent down to the finalizer on the chip. So generator and finalizer don't have to be on the same chip, it doesn't really matter where things start and where they end.

. . .

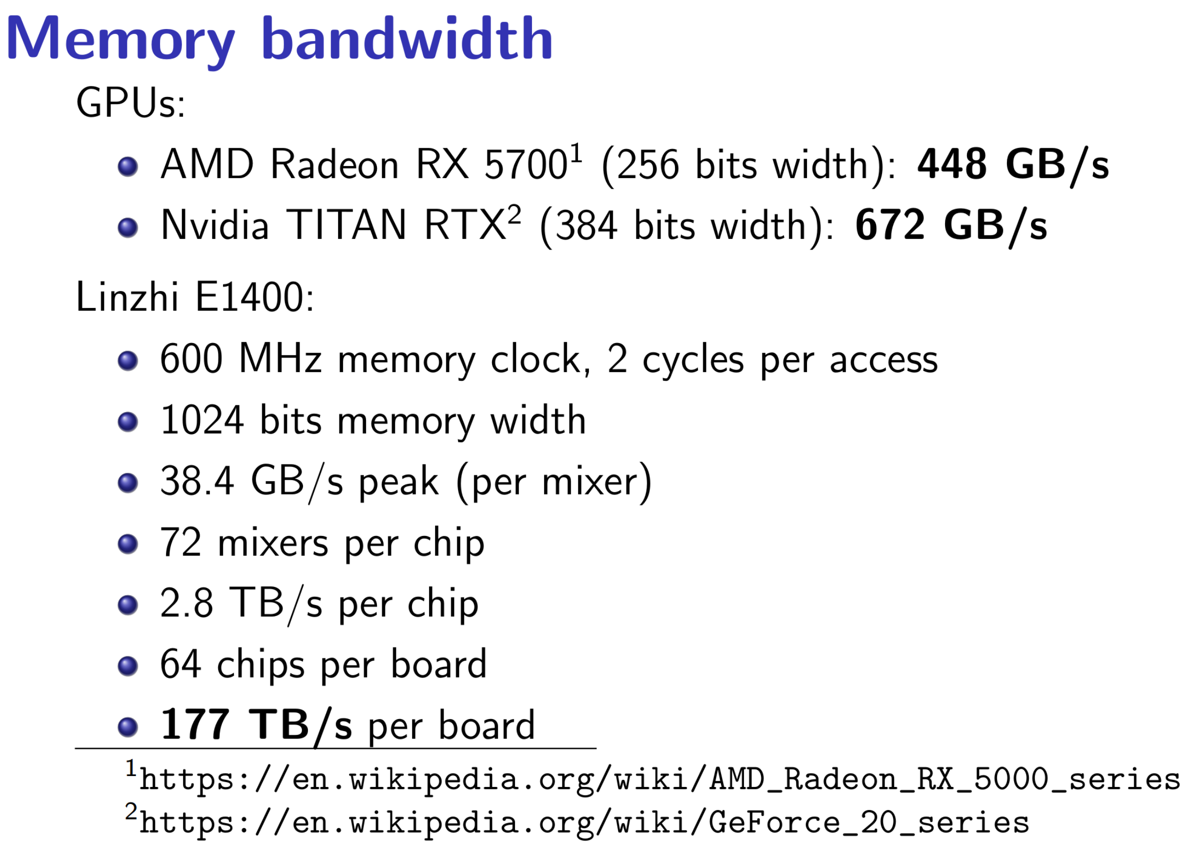

Memory bandwidth

GPUs:

- AMD Radeon RX 5700[1] (256 bits width): **448 GB/s**
- Nvidia TITAN RTX[2] (384 bits width): **672 GB/s**

Linzhi E1400:

- 600 MHz memory clock, 2 cycles per access
- 1024 bits memory width
- 38.4 GB/s peak (per mixer)
- 72 mixers per chip
- 2.8 TB/s per chip
- 64 chips per board
- **177 TB/s** per board

[1] https://en.wikipedia.org/wiki/AMD_Radeon_RX_5000_series
[2] https://en.wikipedia.org/wiki/GeForce_20_series

Let's look at performance — memory bandwidth. There are a few examples for GPUs, contemporary GPUs. I'm not quite sure if they are the best for mining, but they have pretty good, pretty reasonable performance. One would be an AMD Radeon RX 5700, which has a memory bus of 256 bits, and it gets a throughput of about 450 GB/sec. Another one would be the Nvidia Titan RTX, which has a wider memory bus, and does about 670 GB/sec. In our system, memory has a speed of 600 MHz, and it needs 2 cycles for every access. Our memory has a width of 1024 bits, which means that we get throughput of almost 40 GB/sec, per mixer node in the chip. There are 72 mixers in a chip, which means we get to almost 3 TB/sec.

This is already starting to look good in comparison, right?, because the GPUs don't even reach a TB. But then we have 64 of those chips, so our total memory bandwidth in the whole system is almost 180 TB/sec. I think we can confidently say that memory bandwidth is not really a limiting factor here. In fact, we have so much overkill there that we use it extremely inefficiently. We also store a queue in there, so for every access we do to the DAG, we do two accesses for just moving data around. But we have lots of spare bandwidth. However, this isn't our real bottleneck.

. . .

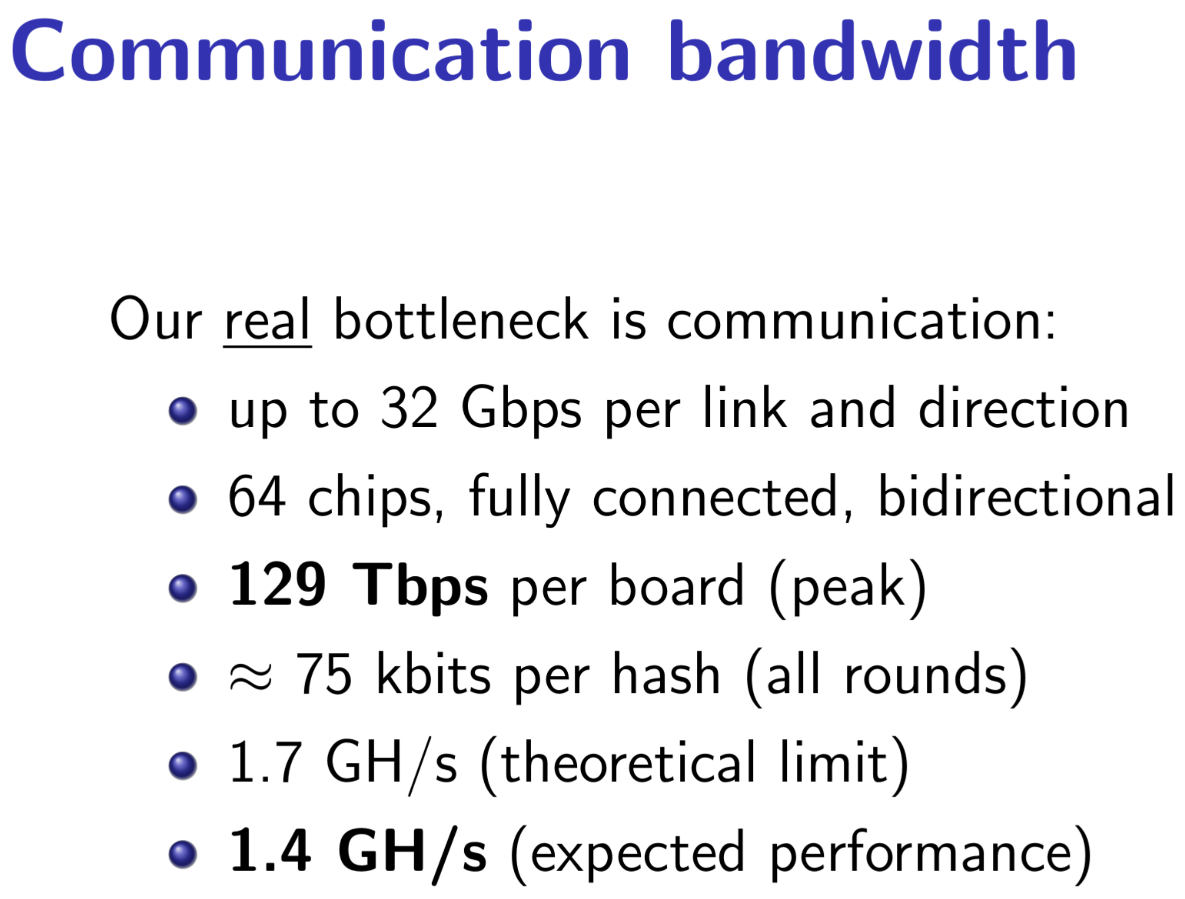

Communication bandwidth

Our real bottleneck is communication:

- up to 32 Gbps per link and direction
- 64 chips, fully connected, bidirectional
- **129 Tbps** per board (peak)
- ≈ 75 kbits per hash (all rounds)
- 1.7 GH/s (theoretical limit)
- **1.4 GH/s** (expected performance)

Our real bottleneck is the communication now, because the chips have to talk to each other. We have up to 42 Gbits/sec, per communication link. All the chips are fully connected with each other, so each chip connects to all 63 others on the board. The links are bidirectional, this means that we have a peak throughput of about 130 Tbits/sec on a board, and for every task, for every complete hash — all 64 rounds, we need to send roughly 75 Kbits, if everything added together. This means that we have a theoretical limit of about 1.7 GH/sec. There are always some factors where we don't use it perfectly, and where things get slowed down, there are some lossages along the way, so we expect a real performance of 1.4 GH/sec.

· · ·

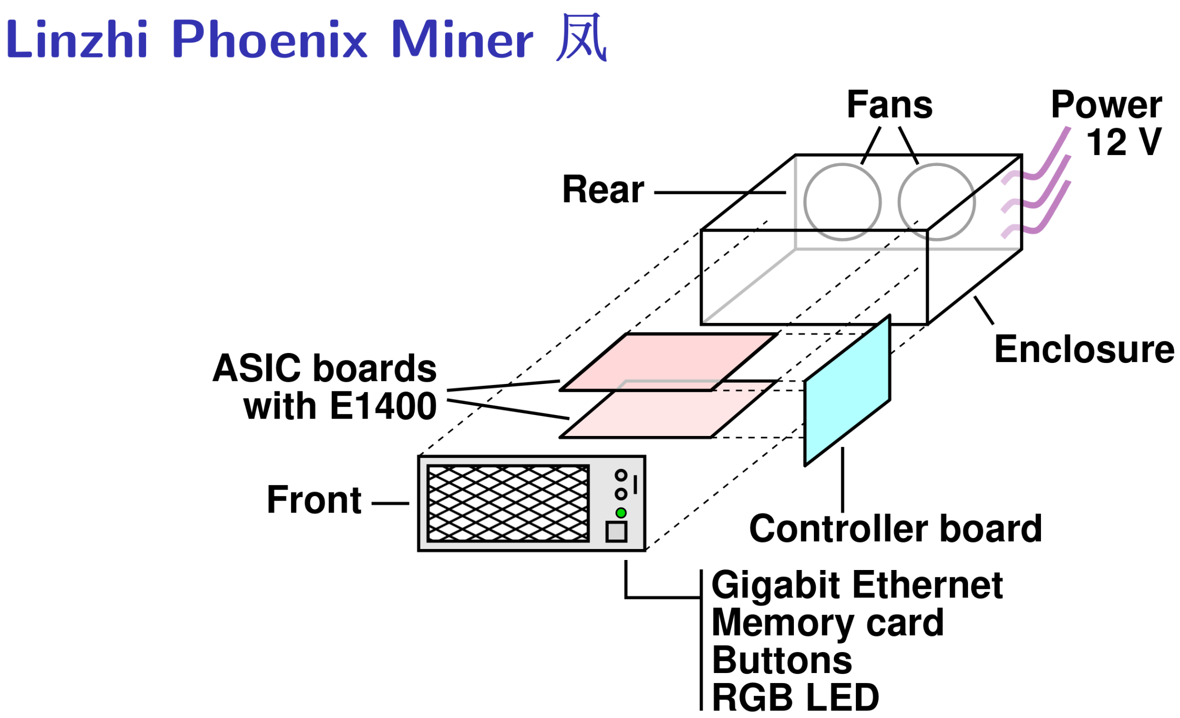
Linzhi Phoenix Miner 凤

There is a little bit more to making a miner. The miner can have one or two ASIC boards. There is a controller board, we have a Xilinx Zynq FPGA, which acts as a microcontroller. We have a few things that we like to do in the FPGA, but it doesn't really do much work. All this goes into an enclosure, which has two fans on the rear. There is also power input, 12V standard mining power. On the front panel we have the Ethernet plug, a reset button, another button for some control functions, plus some blinken lights.

· · ·

We have done our tapeout last month, finally. Tapeout is when you are finished with the design and your simulations and such, and you send the files to the semiconductor factory, and then they make the chip for you. Now we are waiting for those chips to come back, this will take about 2 months, so we expect them to be back around mid to end November, and then we will start with the bringup, which means we will test whether the chip converts power into, perhaps, some calculations, or if it's more interested in converting power into smoke. Then we need to see if we can talk to it, and if it's willing to communicate, and we can maybe try to mine a little bit on a single chip, which is really completely useless, because we don't have the full DAG. But it would show us that the logic is working, we can test this, then combine multiple chips, and so on. Finally integrate everything, check that all the software is working, and we hope that things will be available for sale around first quarter next year.

. . .

By now you may be asking yourself the question: Given that there are now all these ASICs for Ethash — has Ethash failed?
You could say in one sense, if you use a very strict interpretation of what ASIC Resistance means, namely that there could not be any ASICs, they would be economically infeasible, then you could say yes. But, if you think about what happened to Bitcoin, where basically the CPUs got eradicated by ASICs, because they enjoyed such a huge performance difference, and apparently nobody really saw that coming, which is surprising. Then you see that actually Ethash did extremely well. Because now there are ASICs, and they are a bit more efficient than the GPUs, which is to be expected. You

started your ASIC resistant design on using ASICs already. Your CPU, GPU, are just ASICs. So obviously, it can be done on ASICs, and if you are a specialized ASIC it's extremely likely to be a little bit faster, at least. That is pretty much a given. But, you can also see that those ASICs are not a thousand times better than CPUs or GPUs.
That's the good news — Ethash actually worked, it did its job, it prevented the threat of a hostile takeover.

. . .

If you want to know more, you can go to our website linzhi.io, and we are also on Telegram as t.me/LinzhiCorp.
Thank you all.

Werner Almesberger, Buenos Aires/Vancouver