Bitcoin and the Age of Bespoke Silicon

Michael B. Taylor

Associate Professor

University of California, San Diego

This Talk

Introduction

An Overview of the Bitcoin Cryptocurrency

Bitcoin's Computing Evolution

Bespoke Silicon

Interesting Facts about Bitcoin

- The most successful digital currency ever

  Since its deployment in Jan. 2009,

  – **11.7 Million** … Bitcoins (BTC) are in circulation

  – **$142** … is current value of 1 BTC (mtgox.com)

  – **$1.66 Billion** … is the total BTC Market capitalization

- Winklevoss Brothers → bought 1% of BTC supply and are creating a BTC ETF

- You can create (***mine)*** bitcoins with your computer!

What this talk focuses on:

- How Bitcoin mining has raced down the computing specialization hierarchy:
  - distributed CPUs
  - distributed GPUs
  - distributed custom FPGA boards
- And now…
  - Three groups of enthusiasts created three different ***bespoke*** ASICS that have displaced CPU/GPU/FPGAs
  - 80X cheaper/less energy than Intel, AMD, Xilinx…
  - No Venture Capitalists were involved
  - Silicon Valley was not involved
  - They were not backed by any big company
  - *How did they do this in an environment in which new chip startups are almost non-existent? (And can we replicate?)*

This Talk

Introduction

An Overview of the Bitcoin Cryptocurrency

Bitcoin's Computing Evolution

Bespoke Silicon

Bitcoin: User View

- First step: create a bitcoin account
  - run code locally on your computer to create two numbers:
    - ***public key (also known as a BTC address)***: like an email address; people can send BTC to it
      - e.g., 1JVQw1siukrxGFTZykXFDtcf6SExJVuTVE
      - in many cases people publically advertise these
    - **private key**: lets you transfer BTC associated with your public key to somebody else's BTC address
      - 256-bit number
      - keep this in a safe place!
  - no interaction with outside world req'd to **create** account
- Next step: receive, spend and/or mine bitcoin in minimum increments of **1 *satoshi***

  = 1/100,000,000 BTC

Bitcoin Network

- The Bitcoin system:
  - maintains a global, distributed ledger of transactions (e.g. transferring bitcoin) called ***the block chain***, that***:***
    - tracks how many BTC are at each address
    - is replicated across many machines on the internet
    - is maintained via a consensus algorithm by those machines
    - contains a public record of every single transfer of BTC
  - the machines perform an computationally intense operation called ***mining*** that adds new blocks of transactions to the block chain
  - each block contains:
    - a cryptographic hash of the previous block in the chain, maintaining integrity and a total order of blocks in the chain;
    - a merkle hash of all transactions in the block

Bitcoin Network Consensus

- Other nodes will validate new blocks added to the block chain; e.g. no improper creation or destruction of bitcoin
- If the block is valid, nodes will use it as the base for new blocks being added; if they do not, they will use the previous block as the chaining point (a *fork*)
- This is the basis of the consensus algorithm that maintains the integrity of the block chain.
- A block is added roughly every 10 minutes (more on this later).
- Convention: your transaction is legit after 6 blocks have been added to the end of the block chain

Rewards for Bitcoin Mining

- You get 50 BTC *block reward* for adding a block to the chain; paid via a transaction included in the block.
- Reward drops by half every 210K blocks (four years). It has already dropped to 25 and will drop until it reaches a satoshi.
- Total BTC will never exceed 21M; 99% by 2032.
- User often specify a transaction fee (often .0005 BTC) for their transactions to incentivize nodes to add their transaction to the block.
- Miner collects these as well, but only amount to .25% of reward; but this becomes the main incentive when block reward is small.

Bitcoin Mining *Difficulty* (The Catch)

- To add a block, nodes find a ***nonce;*** a value in the block's header, that causes block's double-SHA256 hash to be less than a certain number, maxH.
- Basically, analogous to computing the inverse of a cryptographic hash → hard!
- Brute force (*increment nonce; **hash**; check; repeat*) is the only known method (otherwise SHA256 is easily invertible and a bad hash.)
- maxH is characterized by a number called the **network difficulty:** maxH = (0xFFFF << 208)/difficulty.
- Difficulty is scaled every 2016 blocks to keep network's block creation rate at 1 per 10 minutes.

Bitcoin Mining Profitability

- 6 blocks per hour * 24 hrs = 144 blocks/day
- 144 blocks * 25 BTC = 3600 BTC/day
- 3600 BTC/day = ~$518,000/day
- These mining rewards are spread across the world based on the % of **network hash rate;** basically the world's total brute-force hashing capability.
- Your reward is proportional to your % of the network hash rate.
  - Good CPU: 5-15 MH/s;
  - Good AMD GPU: 600 MH/s; $400 per GPU
  - ASIC is: 300 MH/s; ~$4 per chip

A Brief History of Bitcoin

- Nov ‘08: *Satoshi Nakamoto posts Bitcoin paper*
- Jan ‘09: *System goes live*
- Jul ‘10: *USD/BTC Exchange Created*
- *Sep ‘10: GPUs start being used to mine*
- Apr ‘11: *Satoshi Disappears!*
- *Jun ‘11: FPGAs start being used to mine*
- *Feb ‘13: ASIC hardware appears*

The Value of a Bitcoin

BTC goes from $.05 to $150 in four years: 3000x

Two bubbles: $32 and $266

Early on somebody ordered a pizza for

10,000 BTC = $1.5M!!

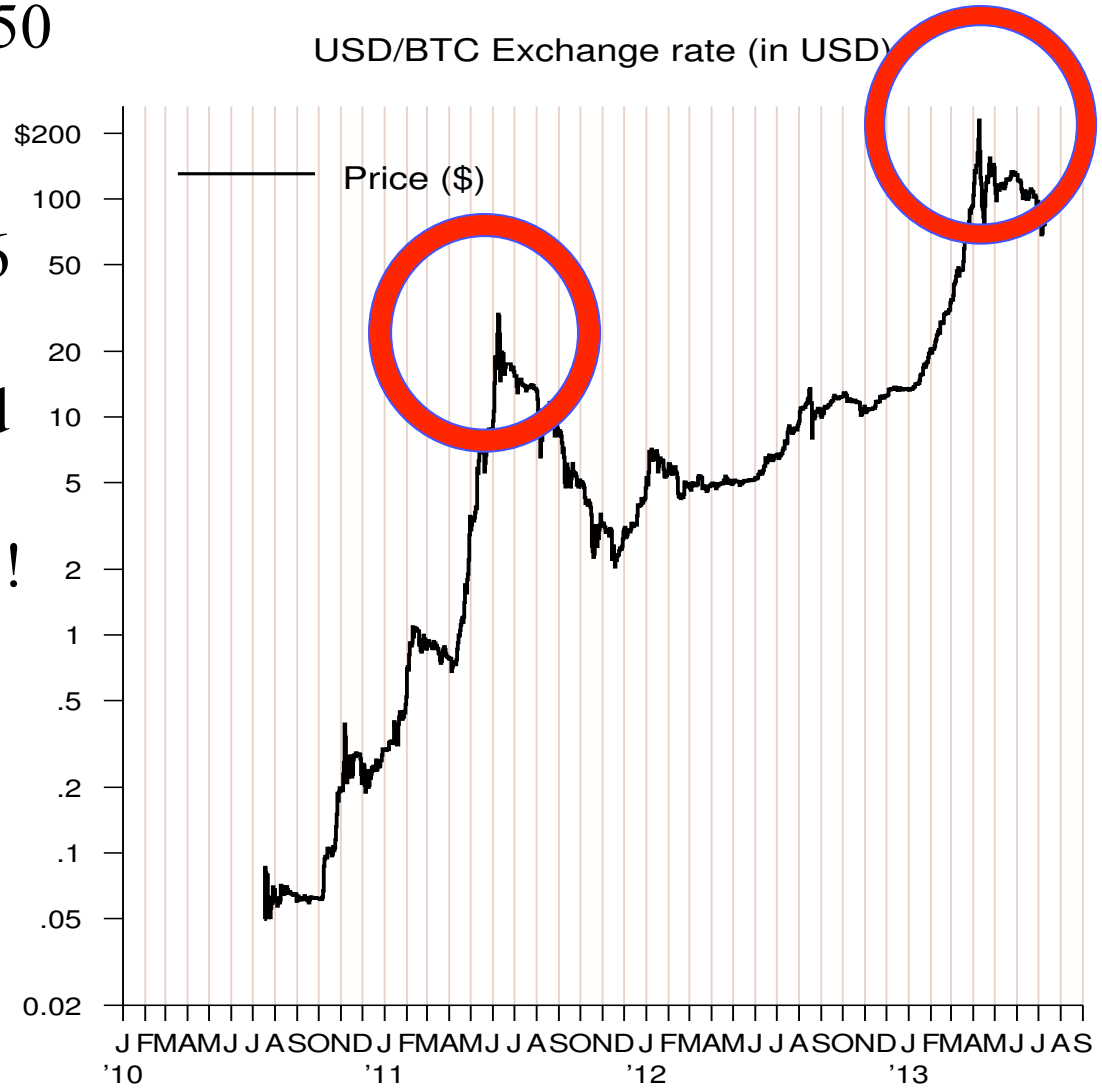

What do people see in Bitcoin?

- Not controlled by any central government
- Fixed money supply; inflation is bounded (21 M BTC)
  - beats gold as a value store
- Pseudo-Anonymous Transfers
  - like Paypal (gov't still can find out who you are)
- Irreversibility
  - no charge-backs like for VISA / Mastercard / Checks
- High portability and physical security
  - better than gold, cash, bank accounts, bearer bonds
  - memorize your private key
- Low transaction fees (5 cents to transfer $1B dollars)

How much will Bitcoin be Worth?

- 21 million BTC total currency supply
- 7.1 billion people
- → 338 *people per BTC*!
- Value of a BTC if it replaces world gold reserves
  - $71,000
- Value of a BTC if it replaces USD as world currency
  - $57,142
- Value of a BTC if BTC reaches VISA/MC mkt cap
  - $9,857

This Talk

Introduction

An Overview of the Bitcoin Cryptocurrency

Bitcoin's Computing Evolution

Bespoke Silicon

BTC Mining Difficulty

- Started at 1.
  - a few CPUs
- Now at 104,000,000
  - 300 million CPUs
  - but actually: 400K ASICs
- Difficulty ramps as:
  - new technologies arrive
  - USD/BTC increases
    - more machines added
- Lines mark dates of introduction of new computing technologies

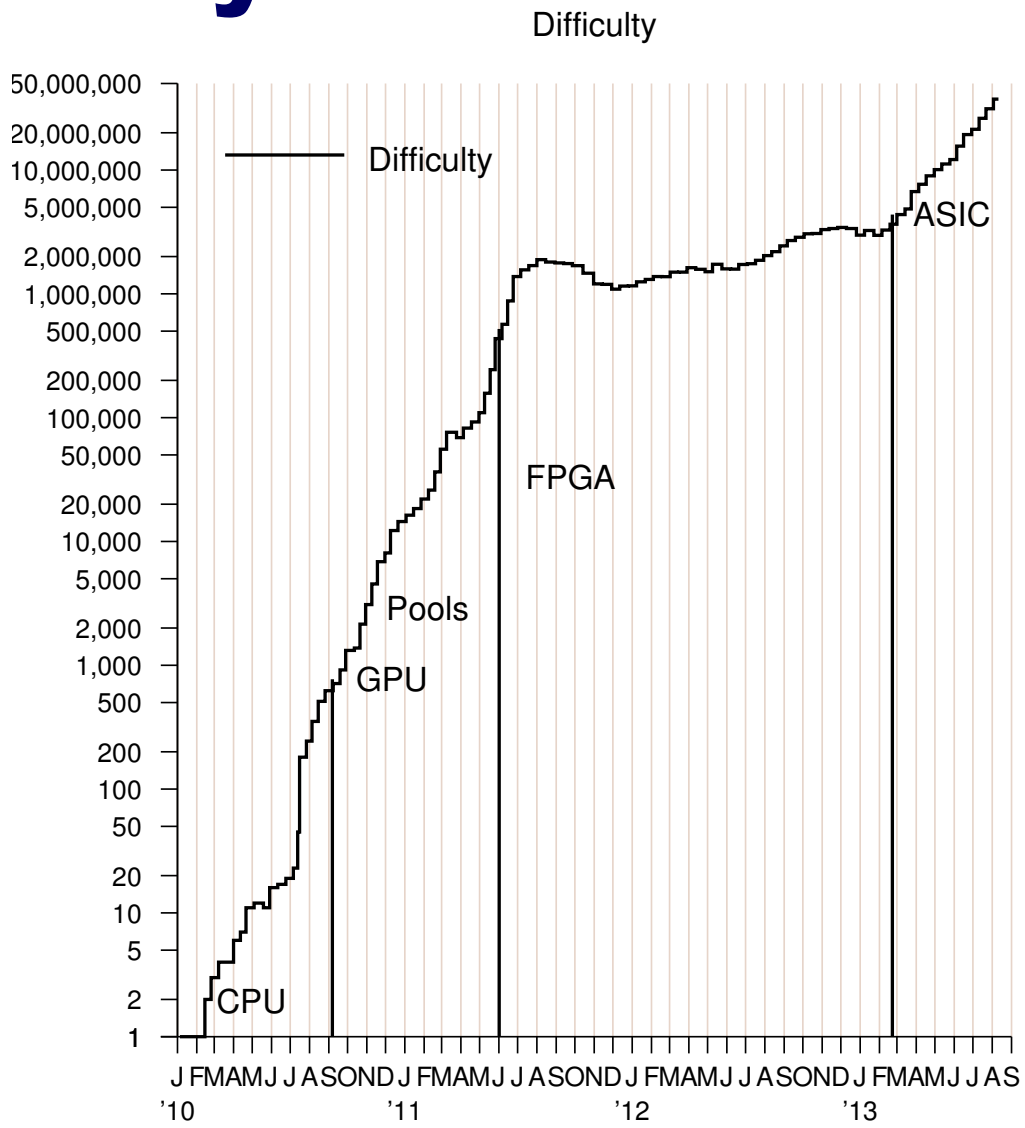

BTC Mining Computing Evolution

- CPU
- GPU
  - Portable OpenCL Imp
  - Completely unrolled double SHA256 hash
  - AMD >> Nvidia
    - instruction set match
    - microarch (VLIW) match
    - higher ALU density
    - memory BW not used
- FPGA
  - verilog
  - “gateway drug to ASIC”: boards, protocols, thermals, verilog
- ASIC

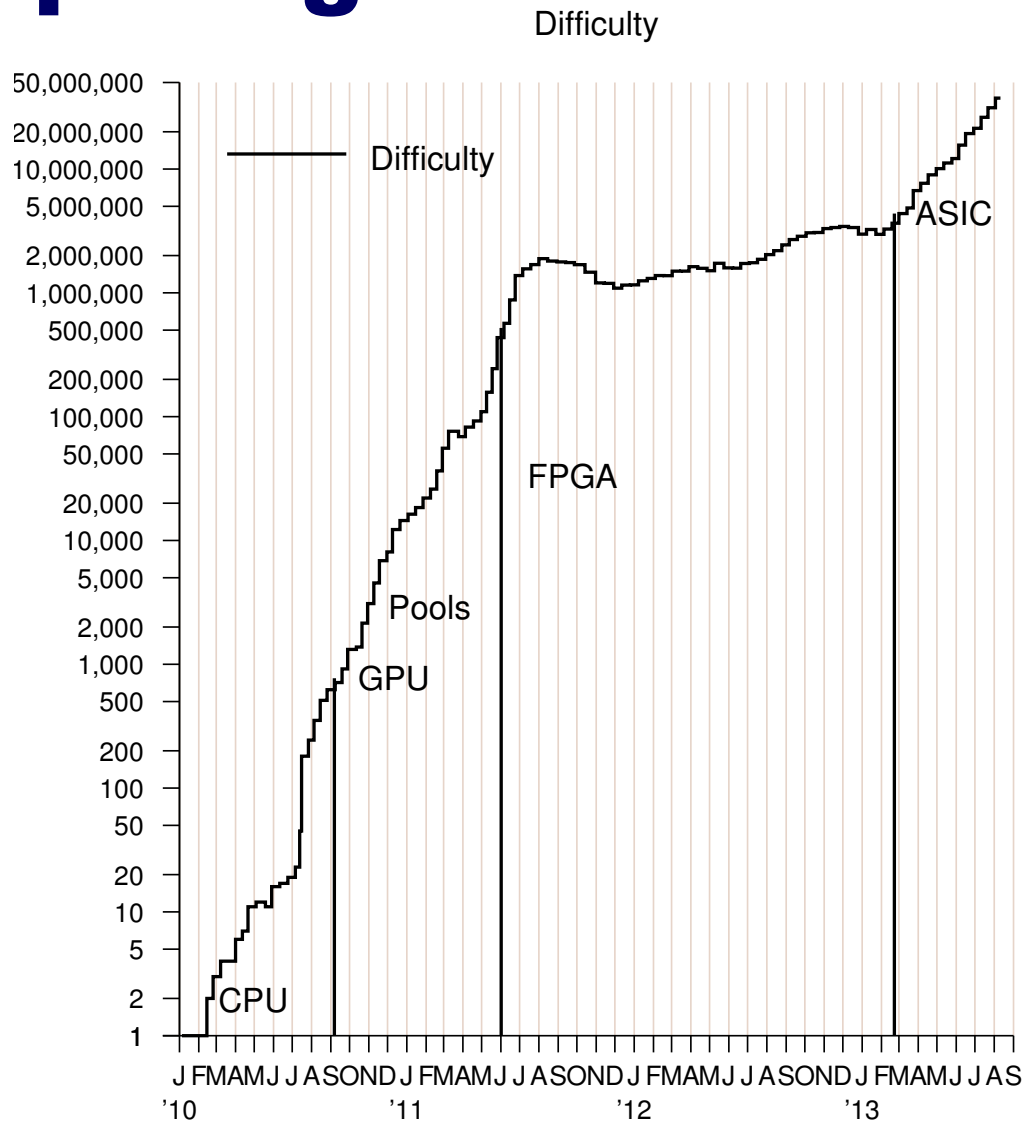

Energy Costs and USD/BTC Say when to unplug/plug HW

- daily $ per Gh/s falls as technology advances and more machines deployed
- daily $/GH/s rises if USD/BTC rises.
- Today, CPUs, GPUs, and even FPGAs do not recoup energy costs
- Rising USD/BTC: old machines get fired up.
- Steady state: cheap energy wins (Iceland?)

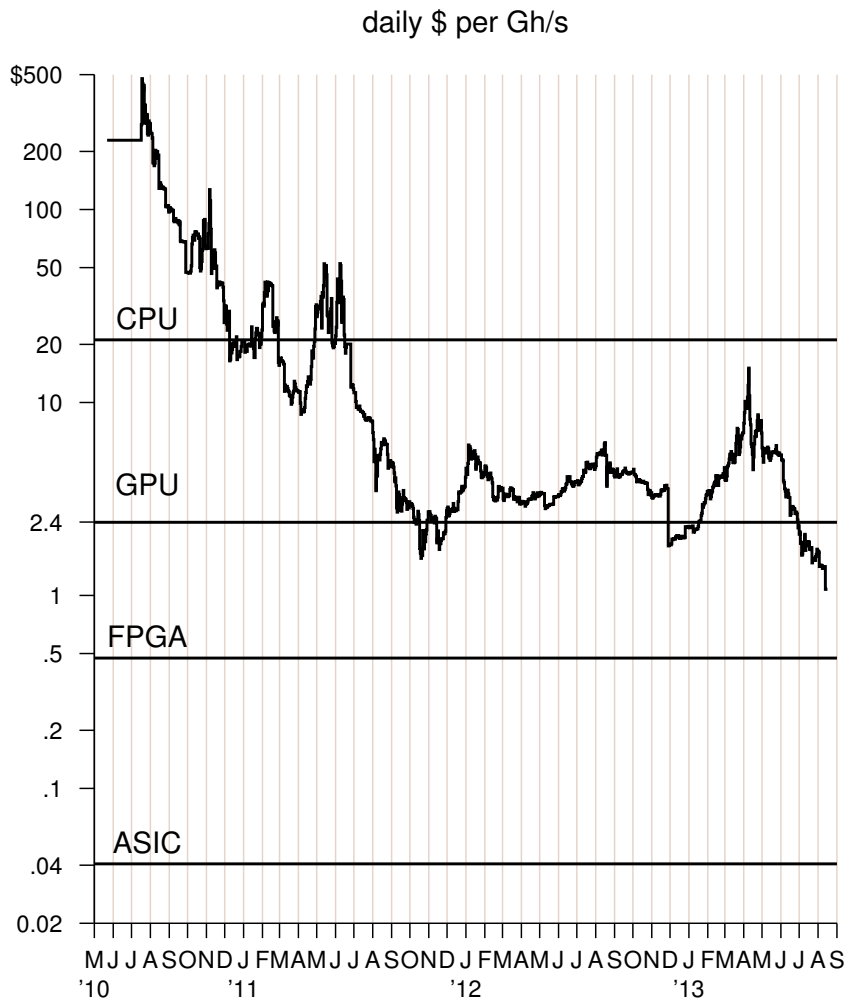

(USD .015 per KWh)

Gen 1: CPU

- CPUs
- Pooled Mining
  - Idea: mining a block takes longer and longer → hours, then day, then months, now years.
  - Solution: have groups of machines work together to reduce variance and uptime.
  - Solution: Divvy up the nonce space among many hosts, and have them get work from a central server, and they get paid for their share of nonce space explored.
  - Problem: Hosts skip work, reporting that they searched and found nothing, and collect BTC.
  - Solution: Hosts return results for blocks that are “close”; server duplicates a small subset of work among different clients.

Gen 2: GPU “Rigs”

- GPUs running OpenCL > 30x GH/s of CPU
- Key challenges in GPU systems
  - wasted $ and energy on CPU / PCB / DRAM
  - power delivery, heat dissipation problems
- Innovative solutions:

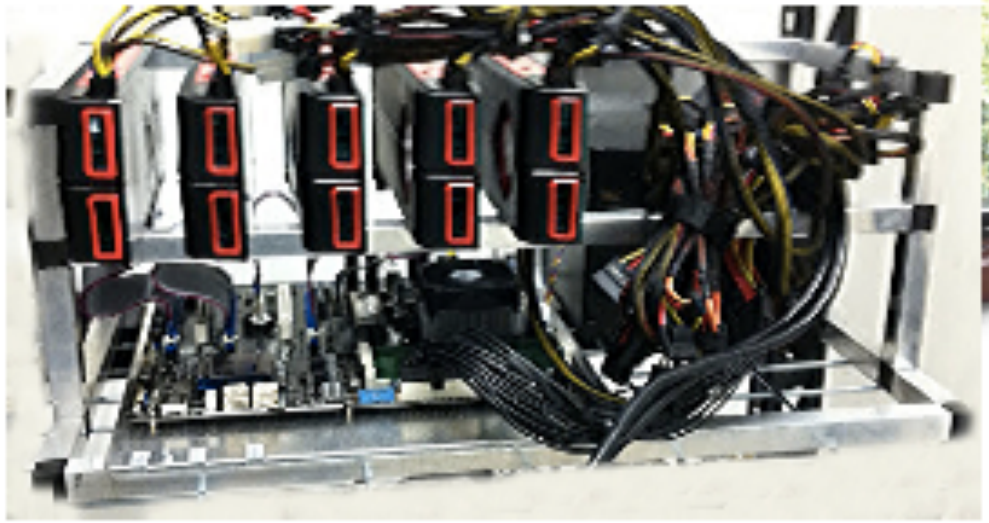

amazon Prime ... "Gets" GPU Mining

Frequently Bought Together

Price for all three: $159.56

Add all three to Cart | Add all three to Wish List

These items are shipped from and sold by different sellers. Show details

- **This item:** Holmes HBF2010A-WM 20 Inch Box Fan, White by Holmes $22.59
- ASRock MB-970EX4 Socket AM3+/ AMD 970/ AMD Quad CrossFireX& nVidia SLI/ SATA3&USB3.0/ A&GbE/ ... $99.99
- AMD Sempron 145 Processor (SDX145HBGMBOX) $36.98

Customers Who Bought This Item Also Bought

ASRock MB-970EX4 Socket AM3+/ AMD 970/ AMD Quad CrossFireX& nVidia SLI/ ...

(39)

$99.99

AMD Sempron 145 Processor (SDX145HBGMBOX)

(35)

$36.98

Seasonic SS-1250XM X-Series ATX PC Power Supply

(13)

$254.99

Sterilite Plastic Storage Crates, Black

(7)

$4.26

PCI-E PCI Express 16X to 1X Riser Card Adapter Extender Flex Flexible Extension Cable

(3)

$4.98

Gen 2.5:
"GPU Datacenter in my Garage"

- Renting data centers was often too expensive
- Roll your own
  - 69 GPUs in one rack
  - Box fans and heat ducts

Gen 3: FPGAs

- 1-2 pipelines per FPGA
- 128-stage pipeline == 1 double SHA hash/cycle
- ~ 216 MHz
- academic FPGA boards: insufficient power; custom boards req'd
- utilized best "consumer" FPGAs (Spartan-150)

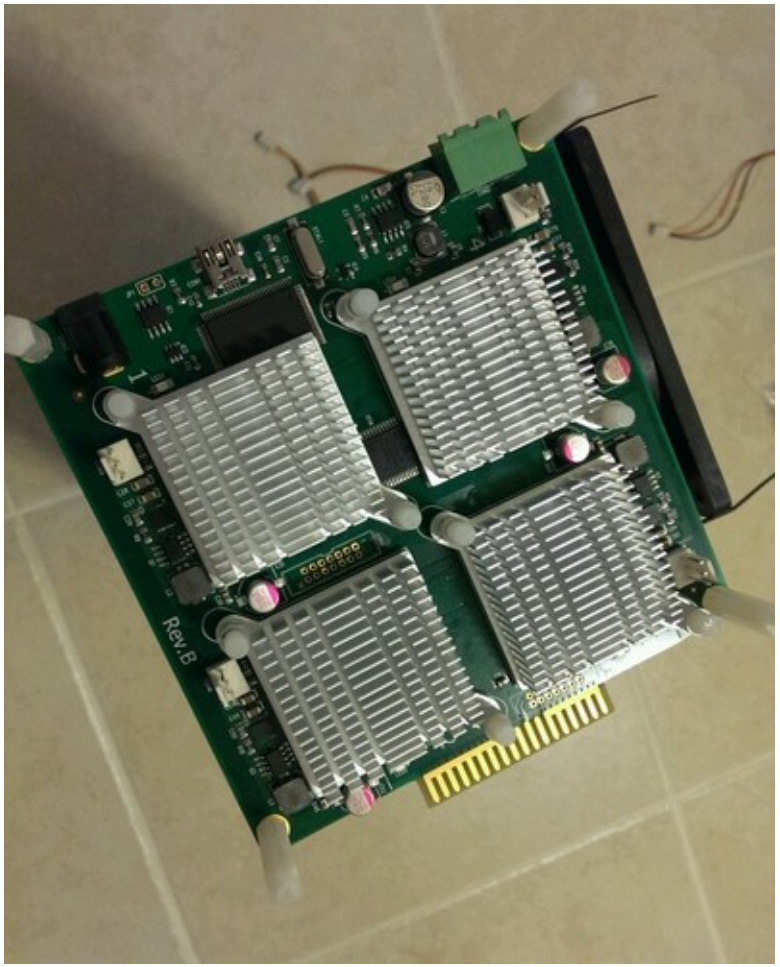

- ~5-8x energy efficiency v. GPU, same cost: but *no resale market! Wins only on* ***Long-term*** *TCO.*
- ASICs came out too quickly after FPGA for FPGAs to obsolete GPUs.

Gen 4: ASICs

- Built and financed by enthusiasts ***on online forums*** – not by existing semiconductor companies or venture capitalists.
- Three parallel efforts:
  - Butterfly Labs (BFL)
  - ASICMINER
  - Avalon
- With the deployment of SoCs and advanced process nodes, new digital chip startups are becoming increasingly rare…
- These efforts are counter to that trend; I refer to them as *interesting cases of* ***bespoke silicon*** – custom built silicon tailored to a particular purpose.

This Talk

Introduction

An Overview of the Bitcoin Cryptocurrency

Bitcoin's Computing Evolution

Bespoke Silicon

Bespoke Silicon: BFL

- Pre-June 2012: Bitcoin forums had constant musings about the potential promise and catastrophe of ASICs, even as GPUs and FPGAs became coming.
- June 2012: Butterfly Labs (BFL), an FPGA miner vendor, announces it is taking pre-orders for:
  - $149 “Jalapeno” at 4.5 GH/s (30x cost/perf over GPU),
  - $1,299 “Single” at 60 GH/s and
  - $30,000 “Minirig” at 1,500 GH/s
  - note: ***entire network was only 12,000 GH/s at the time!***
  - units were sold on a pre-order basis, with delivery promised in November. The funds went towards financing the effort (a la the Kickstarter model.)

Bespoke Silicon: BFL

Bespoke Silicon: BFL

- 65-nm GlobalFoundries process, 7.5 mm x 7.5 mm
- High NRE: $500K - $1M estimated
- 16 lanes of double-SHA256 pipelines @ ~300 MHz
  - like 16 FPGA pipelines in one chip; lanes offer yield control
  - Originally: QFN package; later: 10x10 BGA 144
- Original power estimate: .8W per GH/s
- Actual: 6W per GH/s → Need higher end package
- First HW slips from Nov ‘12 to April ’13
- BFL gave **daily updates** on their progress
  - Tens? of Thousands of investors/customers!
  - Delays have greatly reduced payout to now-angry customers
  - # of units ordered may have made this inevitable!

Bespoke Silicon: ASICMINER

- Early July, after BFL took pre-orders
- Chinese group; connections to Shenzhen
- Raised funding through 100's posts to online forum bitcointalk.org; answering questions about every detail of their operation: business plan, CAD flow, design, qualifications, foundry costs, deployment strategy. Very impressive.
- Did an IPO on a non-SEC regulated online stock exchange **denominated in bitcoin** (!)
- 1/400K share of weekly profits for .1 BTC
- Business plan: build machines and mine 12 TH themselves; then sell hardware.
- *Returned over 40x to investor – in BTC* (~400x in $)

Bespoke Silicon: ASICMINER

- 130-nm 6M China-based process, 6 mm x 6 mm
- MLM process: trade mask costs for stepper costs
- Low NRE (~150K) (one third of BFL's)
- 1 lane of double-SHA256 pipelines @ ~335 MHz
  - just like FPGA version; 1 hash per clock
- 40-pin QFN package; 4.2 W per GH/s est.
- Detailed posts to investors every week
- Unprecedented view into the day-to-day working of a shuttle run --- layer by layer report of shuttle progress!
- First Bitcoin ASIC: Dec 28, 2012
- 2 TH/s deployed: Feb 14, 2012
  - complex ordeal to create a reliable datacenter in Shenzhen

Bespoke Silicon: ASICMINER

- Clever packaging: QFN package has large center ground pad that can transfer heat directly to PCB
- No heatsink or local fan costs
- Heat is spread evenly across large surface

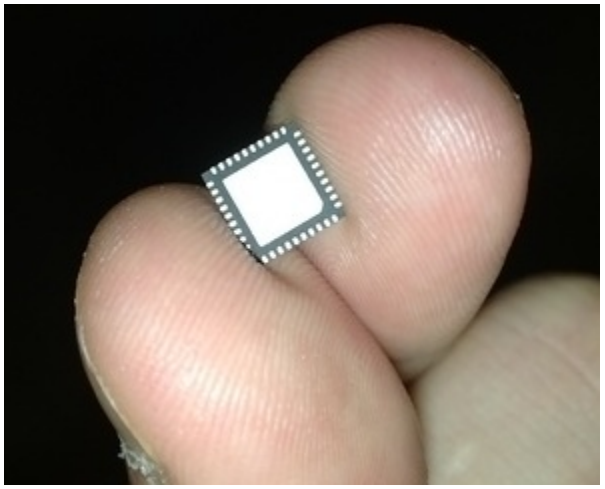

credit: mineforeman.com

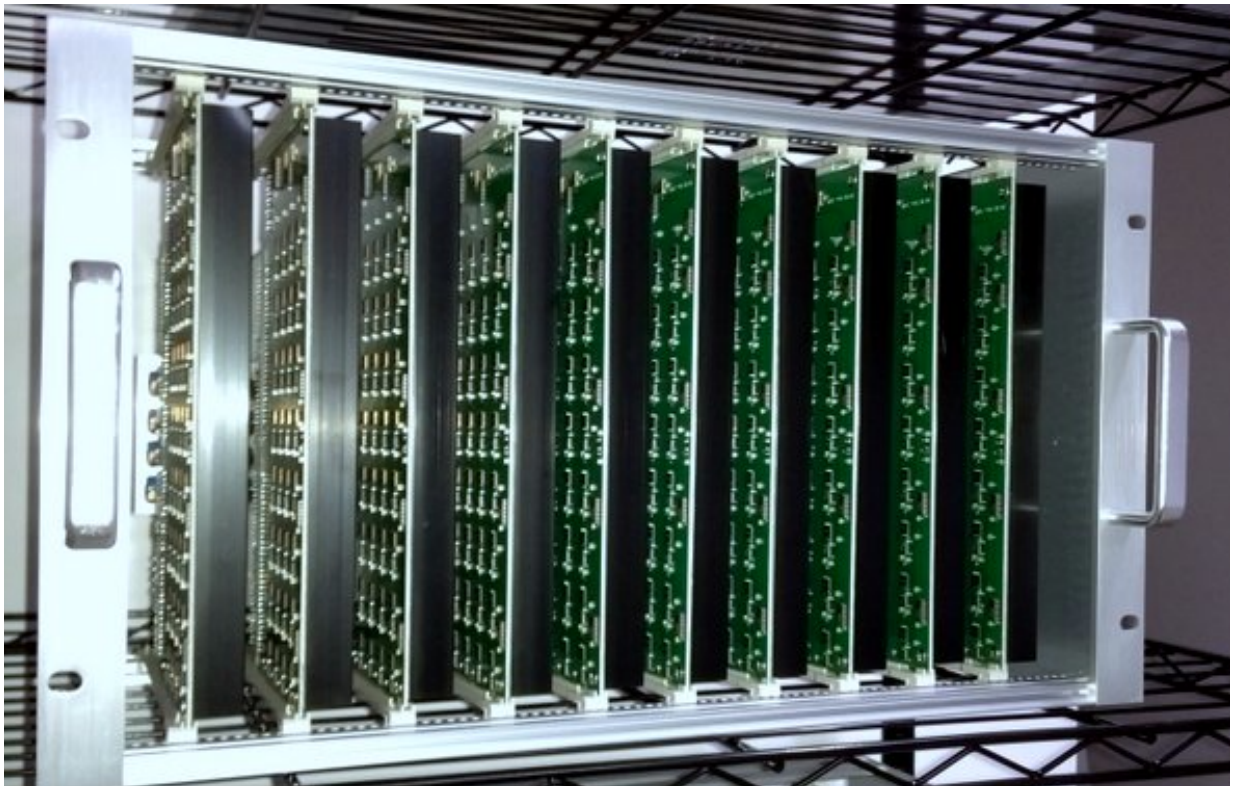

Bespoke Silicon: ASICMINER

- Later, they sold:
  - USB keys containing a single chip; using auctions
  - PCB boards; using auctions
- No reliance on pre-orders; fast Shenzhen supply chain
- Happy customers

Bespoke Silicon: Avalon

- Another Chinese group; connections to Shenzhen
- Jul Preorders: 300 60 GH/s Rigs @ 108 BTC = $1299
- Clever: BTC denomination drives demand for BTC!
- **Delivered first ever ASIC rig to customer Jan 30**. Earned almost 15 BTC in first day!
- After 1st batch, 2nd and 3rd batch of 600 machines each, at 75 BTC, about $7500 at the time
- Then, sold lots of 10K chips for 780 BTC, or $78,000!
- Users banded together to do “group buys” on forum, and to design and procure PCB boards. *(Imagine trusting somebody with $78,000!)*
- Delays in Shenzhen supply chain / shipping → Unhappy customers; but they refunded (unlike BTC)

.

Bespoke Silicon: Avalon

- 110-nm TSMC, 4mm x 4mm, ~220 MHz
- Founder included FPGA miner designer
- Single double-SHA256 pipeline
- **300 chips** across 3 blades in 4U chassis
- Smuggle systems out of China and ship through HK
- Standard PSU
- QFN package with metal platesink.

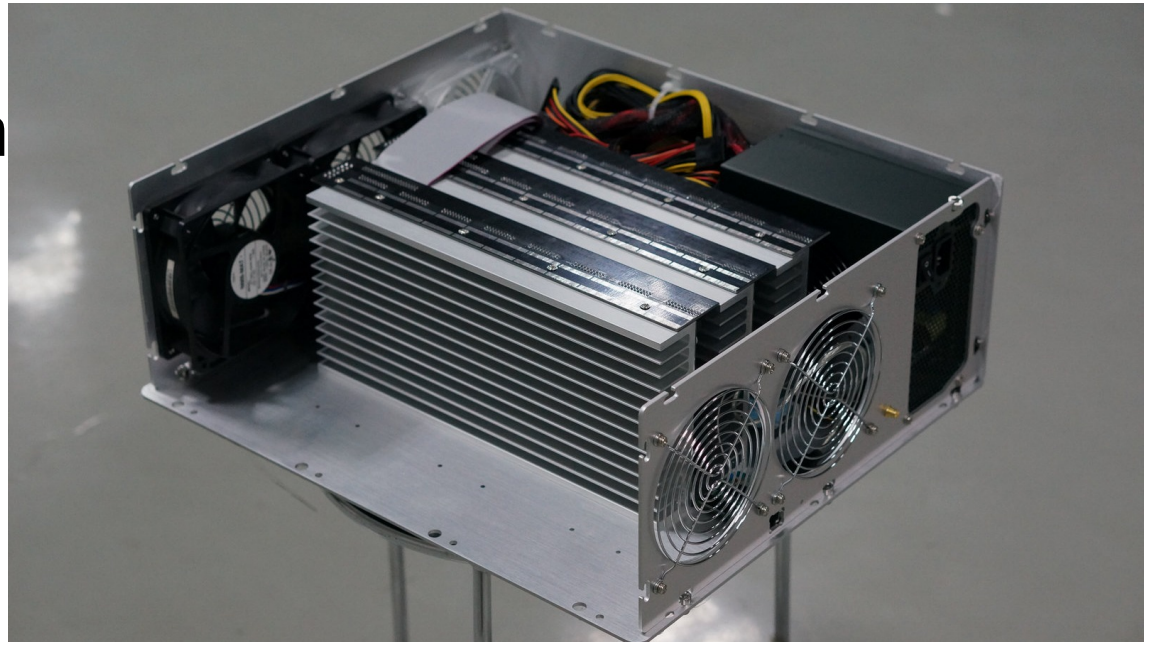

source: gizmodo

Bitcoin Scaling into the Future

- Bitcoin is worst case for dark silicon
  - only linear improvement in throughput and energy per hash due to scaling from 65-nm to 10-nm (6.5x)
- Dark Silicon and Low-Power techniques all apply
  - for instance, near threshold (NTV):
    - no RAMs, little synchronization
  - designs are based on FPGA designs where pipeline registers were free.
  - Next generation will reduce pipelining to bring clock energy under control.
- Probably about 100x left → ~6 W per TH
- Maybe opportunities for specialized circuit design a la DRAMs due to Bitcoin's replicated nature

Observations for Bespoke Silicon

- Specialized devices ***can*** beat general-purpose devices in cost/performance by orders of magnitude
  - if the application benefits from “weak scaling”
- Users are willing to finance when VC’s were not.
  - But they were demanding; even annoying
  - Bitcoin had a (local) linear utility curve for performance
- CPU->GPU->FPGA->ASIC: good progression for new domains: scale up effort as premise is proved
- Old process generations combine well w/ bespoke
  - specialization compensates for “old silicon”
  - low startup costs for trying new ideas
  - time-to-market was inversely correlated with feature width!
  - avoid design complexity issues associated w/ power density
    - e.g. BGA versus QFN; active vs. passive cooling; leakage; power grid

Observations for Bespoke Silicon

- Have we lost the ability to do cheap chips in the US?
  - Two of the three teams were from China.
  - They were the best at executing.
- Academia's fixation on latest process generations does not prepare HW students to do quick startups, unlike their software peers.
  - Are we hamstringing our students and killing innovation?
- Training on million-dollar tools makes it hard to "design cheap" when students exit academia and have to pay
- Technologies like multi-layer masks can bring down the cost of chip startups
- ***Bitcoin is a unique case; but could offer insights into how to build new, bespoke HW for new domains for cheap and revitalize the chip industry!***