

Bitcoin and The Age of Bespoke Silicon

Michael Bedford Taylor

University of California, San Diego

ABSTRACT

Recently, the Bitcoin cryptocurrency has been an international sensation. This paper tells the story of Bitcoin hardware: how a group of early-adopters self-organized and financed the creation of an entire new industry, leading to the development of machines, including ASICs, that had orders of magnitude better performance than what Dell, Intel, NVidia, AMD or Xilinx could provide.

We examine this story for clues as to how we can foster greater innovation in the semiconductor industry and enable this phenomenon to occur more broadly for more application areas, spawning a new age of hardware innovation tailored to emerging application domains—an Age of Bespoke Silicon.

Categories and Subject Descriptors B.7.1 [*Integrated Circuits*]: Types and Design Styles

General Terms Design, Performance, Economics

Keywords Dark Silicon, Bitcoin, Specialization

1. INTRODUCTION

Bitcoin, since its Jan 2009 deployment, has experienced explosive, exponential growth. As of the writing of this paper, there are 11.5 million Bitcoins (BTC, or ₿) in circulation, and the USD/BTC exchange rate is \$104, which means that the market capitalization of Bitcoin is just shy of \$1.2 billion USD. Notably the Winklevoss twins, of *The Social Network* fame, have purchased \$11 million worth of BTC, and have submitted a proposal to the SEC to create an Exchange-Traded Fund (ETF) to allow broader access to investors.

With such rapid growth, Bitcoin is the most successful digital currency, exceeding the next most successful open digital currency, Litecoin, by an order of magnitude. Underpinning Bitcoin's success is a series of technological innovations, spanning from algorithms, to distributed software, and also into hardware. Amazingly, none of this success has been underwritten by a corporate or government entity, but rather emerged through a grass-roots collaboration of enthusiasts.

In this paper, we will introduce the algorithms and software that underpin the Bitcoin system, discuss the turbulent history of Bitcoin so far, and then delve into the fascinating hardware ecosystem that has emerged—from GPUs, to custom FPGA systems to custom ASICs.

The latest round of hardware—dedicated ASICs—have been financed, developed, and deployed by Bitcoin users, which is perhaps an unprecedented event in recent history. One question is whether this model can scale to other application areas

and usher in a new era of *bespoke silicon*—that is, customized silicon that has been developed in small volumes—that leverages specialization to outperform high-volume general-purpose SoCs built by major billion-dollar companies.

2. THE BITCOIN CRYPTOCURRENCY

We overview the bitcoin system below: how it is used, pricing and difficulty trends, and how it is mined.

2.1 How Bitcoin Works: User Perspective

The first step is to create a Bitcoin account. Bitcoin addresses can be created locally on your computer using open-source software, free of charge. The software outputs both a public key and a private key. No interaction with the outside world is necessary. The private key must be kept secret in order to protect the account, and is needed whenever you plan on sending money from the account. If the private key is lost, then the funds are also irrevocably lost. The public key, on the other hand, may be freely distributed to those people who might possibly want to make a payment to your account. To transfer funds to you, they will enter in their own account's private key, and your public key, and a small user-specified transaction fee (typically .0005 ₿, but as low as a single Satoshi¹ or even zero).

The Bitcoin system maintains a global, distributed cryptographic ledger of transactions, called the *block chain*, which is maintained through a consensus algorithm running across a large number of computers distributed across the world. These computers perform a computationally intense function called *mining*, which integrates the transaction into the block chain. The transaction to debit from the sender's account and credit to your account is aggregated with other pending transactions together into a *block* by one of these machines and posted to the head of the block chain. A block also contains a hash of the previous head block of the *block chain*, creating a total order on all blocks in the block chain.

Upon receiving notice of the block being posted to the network, other nodes will verify that the transaction is in order—for instance, not improperly creating or destroying bitcoin, or over-spending from an account—and then use the new block as the head block for blocks that they are trying to post to the block chain. Each such additional block that is posted to the chain is referred to as a *confirmation*. If, by chance, two machines simultaneously append a block to the same link on the block chain, the *fork* will be resolved by picking the branch that has the longest chain of successors—essentially preferring the path that is followed by the majority of mining nodes. As new blocks are posted to the block chain, about every ten minutes, the transaction gets exponentially less likely to be reversed. Up until now, most BTC services are satisfied with six confirmations, but bitcoin creation (see the next section) requires 100 confirmations.

2.2 Bitcoin Mining: Miner's Perspective

¹The atomic unit of Bitcoin, equivalent to .00000001 BTC.

Bitcoin mining is the heart of the distributed consensus algorithm that enforces the consistency of BTC transactions. Bitcoin miners span a wide spectrum of personalities:

1. High school and college students making use of cheap electricity and/or hardware from their parents or universities;
2. Gamers who subsidize their game machines by running GPU bitcoin mining codes on them when not in use;
3. Extreme hobbyists that buy multiple machines (“mining rigs”) until they max out the cooling capacity of their basements (and/or the tolerance of their spouses);
4. Hackers deploying botnets robbing computation from networks of zombie machines;
5. Online collaboratives that raised funding to purchase mining hardware and share in profits, and
6. Companies that raised funding from Bitcoin enthusiasts via an IPO on a BTC-denominated non-SEC-regulated online stock exchange, and are designing ASIC hardware to mine BTC and distribute dividends.

What incentivizes bitcoin miners to perform the *mining* operation that is integral to transaction verification? The answer is that for each block they add to the block chain, the miner receives two rewards:

- First, they are given a *block reward*, which started out at 50 BTC and is halved every 210,000 blocks, about every four years. As of the writing of this article, we are on block 251,660; which means that the block reward is 25 BTC. Due to this halving, the total number of BTC will never exceed 21 million; 54.8% of BTC has already been issued, and 99% of all BTC will be issued by 2032.

- Second, they reap all of the transaction fees that are attached to the transactions in the block. The miner has the option of excluding or including transactions in the block, so the transaction fee creates an incentive for the miner to expend the additional bandwidth, storage or compute required for that transaction. Requests with low transaction fees can take a long time (a day or more) to be added to a block.

Currently transaction fees average around a quarter of a bitcoin per block, while the block reward is 25 BTC. Over time, the block reward will drop, and BTC transaction volume will increase, so we can expect that transaction fees will become increasingly the incentive for miners.

After bitcoin are earned, the user has the option of either selling them on an exchange like Mt. Gox (whether in USD, Euro or other currency; Bitcoin is truly international), or simply retaining them, hoping that they will appreciate.

Since a new block is supposed to be generated at around every ten minutes, how is this enforced? The answer is that in addition to aggregating the transactions into a block, the miners must find a *nonce* value that makes a double SHA-256 hash of the block’s header be less than $(65535 \ll 208)/difficulty$. Since SHA-256 has been designed to be non-invertable, the primary approach is to use brute force. If the difficulty value is twice as large, then it takes twice as many brute-force tries to find the corresponding nonce.

The difficulty is scaled every 2016 blocks, using the world’s collective hashrate, the *network hashrate*, in the preceding period, to target an average block creation time of ten minutes. In practice, the time between generated blocks occurs somewhat randomly, with some blocks being generated within a few seconds of each others, and other consecutive blocks taking over an hour.

Thus, in the typical situation where mining capacity is increasing on the network (i.e. more machines are being put in place to mine), groups of 2016 blocks will be mined more quickly than the targeted two week period, and the

difficulty will be adjusted upwards, but always trailing the ever-growing rate. Each machine, or *rig*, that is in place to mine will get a correspondingly smaller fraction of the current $24 \times 6 \times 25 = 3600$ BTC bounty that is available per day. At some point, enough rigs have been added, and the difficulty increased, that the energy and maintenance cost of mining equals the value of BTC earned. At this point, the network enters steady state. Since the USD/BTC exchange fluctuates, mining profitability has also fluctuated—during dips, less energy efficient rigs are taken off-line, and the difficulty lowers, and the opposite when USD/BTC rises.

2.3 Bitcoin’s “In Plain View” Anonymity

Since all Bitcoin transactions must be posted to the block chain, the Bitcoin system is inherently public because the block chain is public².

Bitcoin does not explicitly require personal identifying information to perform transactions, which makes it highly useful for performing irreversible transactions with third-parties that you may not want to share physical address information with. However, the public record of the transactions can potentially allow identities to be de-anonymized by determined parties, even if users follow the practice of using a different address for every sum of money that is received.

For instance, dispensing with BTC that you receive will often require that you payout sums of money to several different parties. By receiving a transfer from that address, you are now able to identify other addresses that also do business with the same entity. Moreover, many entities publish their addresses so that they may use the blockchain in order to enhance trust—for instance showing that a set of payments have indeed been paid out to stakeholders as promised.

Once an address is identified with an organization, government entities can subpoena the records of a given public owner of an address to discover the corresponding real-life information about the user who owns a particular address.

Improvements in anonymity are attainable through “dark pools”—services that perform transactions among accounts internally and only post the net differences to the block chain.

2.4 A Brief History

Bitcoin resulted from a refinement of ideas in prior digital cryptocurrencies, and first came to light when a user identifying themselves as Satoshi Nakamoto posted a paper on Nov 1, 2008 outlining the cryptocurrency system. On Jan 3, 2009, the system went live, and use grew slowly, then exponentially. Notably, one person paid for a pizza with 10,000 BTC (now worth over \$1 million USD). Nakamoto maintained the software base, communicating with others online, but by April 2011 had transferred responsibility for the code base and disappeared. Even today, Nakamoto’s identity is still a mystery, and the subject of rampant speculation.

BTC Pricing Trends. Figure 1 shows the exchange rate of BTC to USD over time. Starting in 2010, the value of BTC really started to take off, rising from 5 cents in July 2010 to \$105 in August 2013, a difference of 2100×. In that time, there were two bubbles, one in June 2011, peaking at \$31.50 and one in April 2013, peaking at \$266. Large drops in BTC value tend to follow periods of intense media attention that direct large numbers of people to speculate on BTC. Occasionally, there are also scares—based on rumors of weakness in the protocols or electronic-breakins in the institutions that

²E.g., the complete transaction history for address **1JVQw1siukrxGFTZykXFDtcf6SExJVuTVE** is visible at <https://blockchain.info/address/1JVQw1siukrxGFTZykXFDtcf6SExJVuTVE>.



Figure 1: BTC to USD Exchange Rate (\$ per BTC) on the Mt Gox exchange. Since July 2010, the value of a BTC has increased two thousand-fold. Data from <http://bitcoincharts.com>.

facilitate BTC/fiat currency conversion—that cause a massive rush to sell, overwhelming the BTC exchanges. One of the key aspects that makes BTC attractive to speculators is the upper limit of 21 million bitcoin that will ever exist. If BTC were to replace gold as a value store, then the 1.5 trillion USD equivalent currently in world-wide gold reserves, when allocated to bitcoin, would make a single bitcoin worth \$71,000—significantly above its current value.

BTC Difficulty Trends. Figure 2 shows the trend of difficulty over time. The difficulty started as 1.0, and has scaled up to 50 million. This is notable because the initial difficulty corresponded to 4-8 general-purpose cores running the nonce-search algorithm, trying out $\sim 7\text{M}$ double-SHA hashes per second, and now the collective hashing rate of the network is 50 million times that, trying out over 350 TeraHash/sec!

Two factors increase difficulty. First, due to the rising USD/BTC rate, mining can cover the expense of more rigs. Second, continual improvements have been made in both the software and hardware for bitcoin mining. Dips in BTC difficulty can be noted to line up with bubble bursts in the BTC price; in these cases, the value of the BTC did not justify the costs of running some of the more inefficient miners in pool, and their operators pulled them off-line.

Timeline of Innovations in Mining Hardware and Software. Innovation has been amazingly fast. The first publicly available CUDA miner was released in Sept 2010, with the first OpenCL miner following in Oct 2010. Shortly afterwards, in Nov 2010, a new innovation was released—pooled mining—where groups of computers could work together and split up the nonce-space. Participants were rewarded according to the fraction of the explored nonce space they contributed before the correct nonce was found. These mining pools, which rapidly scaled to thousands of members, allowed




Figure 2: BTC Mining Difficulty has increased by $50\text{ M} \times$. Data from <http://blockchain.info>. Lines indicate the introduction dates of new technologies. GPUs and Pool overlapped.

users to get incremental payouts every day as opposed to a large, 50 or 25 BTC payout every several months—by this time, mining a block was equivalent to several months of computation for a single high-end consumer GPU, and the amount of time could vary widely. One of the key innovations was figuring out how to make sure that the end-machines actually did the work that they claimed to have done; and also to make sure they did not “run off” with the winning nonce. Unfortunately, pools concentrate the distributed nature of bitcoin, resulting in potential integrity threats to the majority-based confirmation process.

Shortly afterwards, the first open-source FPGA miner code was released in June 2011. And then—the first ASIC miner came out in Jan 2013, and other efforts rapidly followed afterwards. Figure 2 shows the debut dates of these technologies. **Advances in Performance and Energy-Efficiency.** High-end, overclocked six-core CPUs like Core i7 990x eventually reached 33 megahash (MH)/s when using SIMD extensions. NVidia high-end consumer-grade GPUs like the GTX570 reached 155 MH/s rates, while \$450 AMD GPUs like the 7970 performed even better, reaching 675 MH/s³. The next evolutionary step were FPGA-based miners, which emerged in June 2011. Open-source versions used four cost-effective Xilinx parts (\$ per LUT), Spartan 150s, falling short of the \$/MH/s of AMD GPUs, but on a 60 Watt power budget instead of 200 W. A commercial company, Butterfly Labs (BFL), began to market and sell a range of FPGA miners. FPGAs would have supplanted GPUs due to energy costs; however, ASICs came out, providing orders of magnitude cost reduction, driving up network hash rates, and inexorably driving GPU and then FPGA profits negative.

³ See https://en.bitcoin.it/wiki/Mining_hardware_comparison for a host of statistics.

2.5 Miner Strategy

An important question that Bitcoin miners need to consider is whether the investment of USD in a new piece of hardware will pay off, versus simply buying the BTC on an exchange. Many custom BTC mining rigs (or shares in companies that maintain them on your behalf) are denominated in BTC⁴, so it's embarrassing to buy such a rig and never recoup the original BTC cost in its mining profits, especially since maintaining the rigs requires round-the-clock monitoring and considerable energy bills. A simple solution is to evaluate the return of the mining operation in terms of BTC.

With Bitcoin's exponential increase in hashing difficulty, a rig's ability to generate BTC drops exponentially over time. At the average of $1.199\times$ growth of difficulty rating per 14-day period (see Figure 2), more than 66% of a rig's lifetime BTC earnings comes in the first quarter, 22% will come in Q2, 7% in Q3, and 4% in Q4-Q ∞ . The lifetime earnings in BTC top out at $\sim 84\times$ the initial daily earnings. Practically speaking, you will unplug the rig in two cases: first, when the daily earnings in USD is less than the cost of the energy bill, and second, when you need to clear space for your newly purchased set of much faster hashing hardware which has begun its rapid depreciation cycle.

The rig's value is the sum of these exponentially declining expected payments, minus operating costs, plus a final payment, which is the salvage value of reselling the hardware at the end of its life cycle. From this, we can compute the ROI, with P =price, X =exchange rate, ME =maintenance and energy costs, and DR_i initial daily revenue⁵:

$$ROI = \frac{\frac{P_{resale}}{X_{resale}} + 80 * DR_i - ME}{\frac{P_{purchase}}{X_{purchase}}} - 1$$

Making a decision about mining hardware requires you to estimate several of these parameters. Some are known at the time of purchase; for instance, the purchase price and purchase exchange rate, and the cost of maintenance. For GPUs, it is easy to estimate the DR_i , since ship dates are easy to estimate, and because the online forums will frequently contain postings with the hashing rate (in gigahash per second or Gh/s) of the hardware. The ME are also easy to estimate from your electric bill costs and from the power specs of the GPU. GPU resale price can be estimated from E-bay sales of prior GPU hardware. The primary risk is the USD/BTC exchange rate when you resell the hardware. For general-purpose hardware like GPUs, you might recover a significant fraction of the salvage value in USD. However, the currency risk is considerable in this conversion—if BTC appreciates significantly, then salvage value, denominated in BTC, will be very low, and you will hold onto the hardware because the energy costs are proportionally low. On the other hand, if USD/BTC rate stays steady or declines in value, then selling the hardware early will greatly improve your ROI. Inspection of E-bay shows relatively low depreciation on year-old AMD GPUs. On the other hand, the story for custom hardware is the opposite. Since its only purpose is for mining, everybody will be dumping it on the market at the same time. Custom hardware like FPGA boards and ASICs has much more significant risks that focus around the delivery date. Every single effort has slipped to date, with grass-roots efforts tending to be optimistic in how quickly they can assemble and ship the hardware. Managing this delivery risk is a major part of

⁴Smart, as it creates demand for BTC, driving up its value!

⁵We assume the rig will operate for 3 quarters, producing $.95*84 = 80\times$ initial daily revenues; if daily ME costs are high relative to DR, a shorter time frame would be apropos.

bitcoin mining. First, you must decide which of several competing efforts for a new technology (whether ASIC or FPGA) is most likely to deliver first. Then, within that effort, you need to get yourself early on the wait list relative to other customers. Otherwise, although you picked the right technology, the difficulty of the mining pool will have already ramped to meet the new technology, and you will lose the most valuable, early profits of the technology. For example, a Bitcoin software developer who was selected to receive the first available Avalon ASIC rig cost spent 108 BTC, earning over 15 BTC on its first day of operation in Feb 2013, while in August 2013, there are back-ordered rigs of the same type even though the rig only pulls in 0.9 BTC per day. Using the formula above, we can see that the revenue difference is 1200 BTC versus 72 BTC, a jaw-dropping difference for the same physical hardware.

2.6 Retiring a rig

Figure 3 graphs the daily revenue per Gh/s that the Bitcoin network paid out since 2010. This graph combines historical hashing difficulty data with the historical USD/BTC exchange rate. The drop in late Nov 2012 corresponded to the transition from 50 BTC to 25 BTC payouts per block. The horizontal lines show the daily energy cost per Gh/s of CPUs (Core i5), GPU (AMD 7970), FPGA (Bitforce SHA256), and 110-nm ASIC (Avalon Batch 1) at 20 cents per kilowatt energy cost. When the revenue per Gh/s of mining drops below these costs, profits turn negative and the rig should be turned off. The network is currently experiencing a large-scale build out of ASIC capacity, which will drive daily \$ per Gh/s below the FPGA line and ultimately below the 110-nm ASIC line. Downward voltage scaling can possibly provide a few extra months of life. Since difficulty increases largely exponentially, flat or upward regions in daily \$ per Gh/s are typically the result of appreciation of BTC relative to USD.

3. BITCOIN'S COMPUTING EVOLUTION

In this section, we examine some notable challenges and developments in the evolution of “bespoke” customized compute systems intended for Bitcoin mining.

3.1 CPU: First Generation Mining

The bitcoin miner source code can be found on github, and is surprisingly simple (see <https://github.com/bitcoin/bitcoin/blob/master/src/miner.cpp>). The basic computation,

```
while (1)
  HDR[kNoncePos]++;
  IF (SHA256(SHA256(HDR)) < (65535 << 208) / DIFFICULTY)
    return;
```

can leverage existing high-performance libraries for implementing the SHA256 hash. One simple optimization that is used is the use of a mid-state buffer, which contains the beginning portion of the block header that precedes the nonce and has a constant intermediate hash value.

The SHA256 computation takes in 512 bits blocks and performs 64 rounds of a basic encryption operation involving several long chains of 32-bit additions and rotates, as well as bit-wise functions including xors, majority, and mux functions. An array of 64 32-bit constants is used as well. Each round is dependent on the last round, creating a chain of dependencies between operations. Although separate rounds of a SHA256 computation cannot be parallelized, each separate nonce trial can be performed in parallel in a classic Eureka-style computation, making this very amenable to parallelization. Furthermore, some of the operations inside a round are parallelizable. However, typical multicore machines have extra

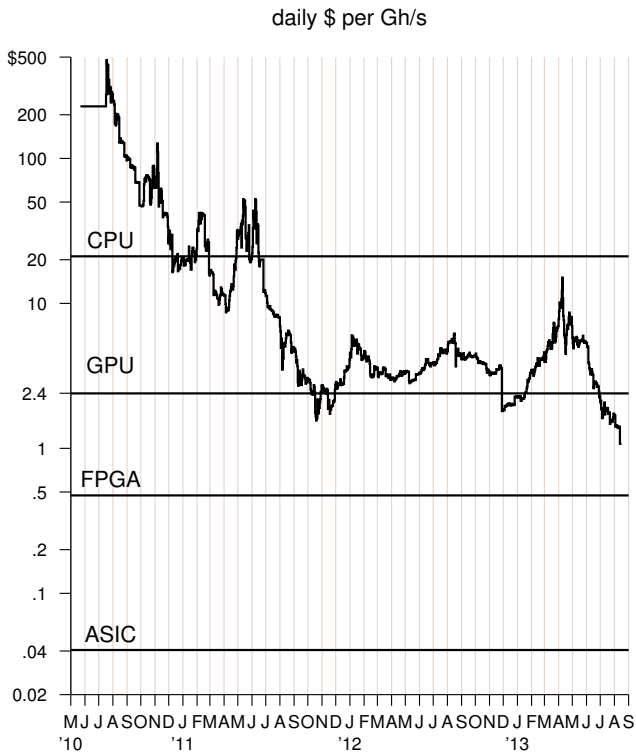


Figure 3: Daily \$ per Gh/s. The figure graphs the daily revenue per Gh/s that the Bitcoin network paid out since 2010. The horizontal lines show the daily energy costs per Gh/s of CPUs, GPU, FPGA, and 110-nm ASIC at 20 cents/kWh energy cost. When revenue per Gh/s drops below these costs, profits turn negative and the rig should be turned off.

hardware optimized for less regular computations, resulting in wasted performance and energy efficiency.

3.2 GPU: Second Generation Mining

In October 2010, an open-source OpenCL miner was released on the web, and it was rapidly optimized and adapted by several open-source efforts. Typically these miners would implement the Bitcoin protocol in another language such as Java or Python, and the core nonce-search algorithm as a single OpenCL file⁶ that was compiled down by installed runtimes into the hidden native ISA of the GPU.

Different variants of the OpenCL file emerged as coders attempted to coax the compilers to improve code quality. The non-OpenCL code is also responsible for invoking an OpenCL API to use the GPU, for double-checking answers, and for controlling GPU parameters in response to temperature and user-specified tuning parameters.

Since these rigs will be left to mine for many months at time, users aggressively tweaked the voltages (lower to reduce mining costs, or higher, with frequency, to increase Gh/s) and operating frequencies of video ram (lower to save energy, since memory is unused) and the GPU core itself, as well as parameters of the code such as the number of threads that are enqueued at a time, so as to maximize throughput within reasonable bounds of stability and temperature. Since the Bitcoin computation does not exercise the memory system or floating point units, many of the critical paths and bottlenecks in the GPU are not exercised, which means that the system can be pushed beyond the normal bounds of reliability.

⁶<https://github.com/Diablo-D3/DiabloMiner/blob/master/src/main/resources/DiabloMiner.cl>.

ity. Over time, it would often become necessary to retune the parameters as fans and power delivery system wear eventually caused a critical path to run too slowly.

Mainstream AMD GPUs tended to outperform NVidia GPUs in terms of Gh/s per \$, in part due to an instruction set well-suited for the bit-level nature of the SHA256 algorithm, and also because the AMD VLIW ISA provides for more parallel ALUs running at slightly lower frequency than NVidia products. In particular, rotate operations and bit-wise choose operations could be implemented with single instructions in the AMD ISAs.

The core code itself was specified in OpenCL rather than machine or assembly code, and in some cases, patched after the binary was generated to make use of special instructions that were not directly supported by OpenCL. The code is scheduled via the AMD software into the VLIW4 or VLIW5 instruction sets, which allows some of the operations in each round to execute in parallel. The OpenCL implementations are consistently a single linear code region which at start selects a nonce based on the thread work item id, and which perform both chained 64 SHA256 hash rounds in a single unrolled loop. External memory is not accessed in steady state. Successful nonces are flagged at the end of the OpenCL routine, to be acted on by the driver code.

A Datacenter In My Garage. After shelling out \$300-600 on a GPU-based mining rig that is literally minting cash, and spending considerable time tweaking its parameters, the natural inclination is to scale it up. Buy the same GPU again, and reuse the settings, and you double your money. In fact, the BTC are coming in so fast, and growing so quickly in value, maybe it makes sense to buy ten or twenty GPUs! Although this could have led to catastrophe, with group behavior leading to a massive collective drop in profitability due to a sky-rocketing difficulty rating, it turned out that BTC appreciated so quickly that these gung-ho BTC minters did not grow regret their decision.

GPUs tended to be much more accessible than FPGAs for end users, requiring PC-building skills and avid forum-reading but no formal training in parallel programming or FPGA tools. The goal of scaling BTC hash rate through GPUs pushed the limits of consumer computing in amazing and novel ways. GPUs had a few key limitations:

1. The GPUs are not standalone. Each GPU had to be plugged into a PCI-E 8x or 16x slot, of which there are relatively few on commercial motherboards.
2. The motherboard, processor, hard-drive and RAM are all but unused in GPU mining, and increase the \$ per Gh/s cost of your mining operation. Typical users had a single computer lying around to host the GPU, but didn't have more computers to host follow-on GPUs.
3. The GPUs would typically require 200-300W of additional power, per GPU, quickly exceeding typical power supply (PSU) ratings and requiring an upgrade.
4. Cases are typically not designed to provide the airflow required for multiple GPUs.
5. The power consumption of multiple GPUS rapidly exceeded natural cooling, power and noise limits of typical residential spaces.
6. OpenCL required a display to be attached to the GPU.
7. GPUs typically take two slots in a case or motherboard, preventing you from attaching many GPUs to a system.

The solution that evolved, addressed these issues as follows. First, because Bitcoin running on the GPUs did not make use of the bandwidth to the motherboard, a 1x slot actually had sufficient bandwidth, and in fact the GPU func-




Figure 4: Two open-air GPU mining rigs. In both cases, five GPUs are suspended above the motherboards, with riser cables connecting the PCI-E connector of the GPU to the motherboard below, and a single high-wattage power supply powering both. Note that the second rig is blowing exhaust heat out of the opened window. Left photo credit: James Gibson (gigavps). Right photo credit: Sophokles.

tions fine with only a 1x connection. A simple \$8 PCIe riser cable converts from the 16x GPU connector to the 1x motherboard slot. However, this meant the card could not actually be plugged into the case—which led to hackers to get rid of the case and just create rail sets that suspend the GPUs over the motherboard. With an appropriate motherboard with many cheap 1x slots, this problem was solved. The use of rails allowed an open design with more surface area for dissipating heat. A resistor was inserted into the GPU’s DVI adapter to fake out the presence of a monitor for OpenCL. These approaches enabled a low-cost motherboard, CPU, and DRAM combination to be amortized across 5 or 6 GPUs, improving capital efficiency. Figure 4 shows a few examples of well-designed open-air multi-GPU rigs.

Interestingly, some of these systems would work for a few months but then would develop stability issues. The GPUs would pull too much 12V current through the PCI-E slots, overloading the current-carrying capacity of the power connector on the ATX motherboard. The solution was to breakout the 12V wire from the riser cable and connect it directly to the power supply via molex, bypassing the motherboard.

After the technical issues of reducing per-GPU overhead, the next scaling challenge is in dealing with the prodigious power and cooling requirements of maintaining many GPUs. With each GPU consuming 200 Watts or more, the power density is comparable or in excess of that of many high-density data centers. Data-centers were almost never used because of high costs and data-center imposed requirements for FCC certification of the hardware. Few residential homes are equipped to deal with these demands and in states like California, residential energy prices allow only one or two GPUs before the price per KWh is jacked up to 35 cents per KWh. The most successful Bitcoin mining operations typically relocated to warehouse space with a large volume of air for cooling and cheap industrial power rates. Figure 5 shows a homebrew data center consisting of 69-GPU rack that is cooled by an array of 12 box fans and an airduct.

3.3 FPGA: Third Generation Mining

June 2011 brought the first open-source FPGA bitcoin miner implementations. FPGA are inherently good at both rotate-by-constant operations, and at bit-level operations both used by SHA256, but not so good at SHA256’s 32-bit adds.

An interesting challenge of the open-source efforts for FPGA

miners was providing a design that scaled to a variety of levels of FPGAs, from high-end to low-end. The resulting design⁷ addressed this challenge very elegantly by replicating a single SHA-256 module that had a parameter that specified an unroll factor. With full unrolling, the module would create different hardware for each of the 64 rounds of the hash, each of which was separated by pipeline registers. These registers would contain the running hash digest as well as a full copy of the 512 bit block being hashed. The state for a given nonce trial would proceed down the pipeline, one stage per cycle, allowing for a throughput of one nonce trial (hash) per cycle. Lesser unroll factors could be specified that would recycle values in the pipeline, and calculate a hash every N cycles. If the FPGA were large enough, several such unrolled pipelines could be instantiated, and trade-offs could be attempted between unrolling and duplication of pipelines.

This unrolled approach resulted in relatively high numbers of registers being allocated, due to all of the mostly-redundant copies of the 512 bit block, but in many FPGAs, each logic LUTs is paired with registers, reducing their cost.

The key challenge encountered with BTC FPGA miners is that the power consumption was much higher than typical for FPGA—essentially the activity factor of LUTs was extremely high with the pipelined design. As a result, the majority of pre-made boards, such as educational boards readily available to student hackers, could neither supply enough current nor dissipate enough heat to sustain usage. The problem was doubly so for higher end Xilinx parts that had more resources. As a result, hackers developed custom boards that minimized unnecessary cost due to parts like RAM and I/O and focused on providing sufficient power and cooling. These boards attained 215 MH/s rates with Spartan XC6SLX150 parts, and quad-chip boards were developed to reduce board fabrication, assembly and bill-of-materials costs⁸, reaching 860 MH/s at 216 MHz and 39 W, and costing \$1060.

Another manufacturer, Butterfly Labs, based in Kansas, offered a non-open-source version that cost \$599 with similar 830 MH/s performance. They also offered a higher-end Goliath FPGA-based machine, the BFL Mini-rig, which cost upwards of \$15K and reached 25.2 Gh/s, and was based on

⁷For instance, <https://github.com/fpgaminer/Open-Source-FPGA-Bitcoin-Miner/tree/master/src>.

⁸See <http://www.ztex.de/btcminer/> for schematics and <http://www.opencores.org/project,btcminer> for Verilog.



Figure 5: Two pictures of a homebrew 69-GPU Bitcoin mining data center. Note the ample power delivery on the photos on the left, and the cooling system, consisting of the box fans and air duct on the photos on the right. The GPUs are arranged in racks as shown in Figure 4. Photos Credit: James Gibson (gigavps).

higher-end Altera FPGAs that individually reached 650–750 MH/s per chip. Four of these rigs are shown in Figure 7. BFL was by all accounts the most successful commercial closed-source Bitcoin company to date.

Unfortunately, FPGAs had trouble competing on cost per Gh/s with high-volume GPUs that would go on-sale on Newegg; often costing 30% more with less potential for resale. It did not help that FPGA-based systems trailed GPUs in reaching the latest, most energy efficient process generation; Spartan-6 was in 45-nm, and GPUs had reached 28-nm. The main benefit of FPGA was the reduction of energy consumption to one-fifth, breaking-even on total cost of ownership (TCO) after a year or two, holding resale equal.

The reign of FPGAs was brief, because little time passed before the next generation of hardware, ASICs, provided an order of magnitude cost and energy-efficiency advantage. However, FPGA development efforts were not wasted; instead they served as a quick stepping-stone to ASIC. The ASIC Verilog designs used were remarkably similar to the FPGA Verilog implementations that preceded them, and the board, packaging and distribution infrastructure and expertise could be re-applied to the ASIC generation. We examine the first three ASICs that came to market.

3.4 Butterfly Labs (BFL)

BFL was the first to announce an ASIC product line, confident from their prior success in their FPGA product line. BFL took pre-orders in June 2012 for three types of machines; \$149 Jalapenos rated at 4.5 Gh/s, \$1,299 SC Singles rated at 60 Gh/s and \$30K SC MiniRigs rated at 1,500 Gh/s. At these prices, the machines could generate 20-50× more bitcoins per dollar invested versus GPUs. The funds from these pre-orders, which exceeded \$250K just in the first day, and lasted many days after that, presumably covered the considerable NRE mask costs for BFL's 65-nm GLOB-ALFOUNDRIES process⁹, with a speculated cost of \$500K¹⁰

⁹Some rumors state that BFL took investments.

¹⁰A post by Friedcat, lead representative of the ASICMINER outfit in July 2012 indicated that NRE costs in China run 150K for 130nm, and 500K for 65nm, and even less with MLM. BFL may have paid more since it used premium outfit

and also ASIC design service costs incurred by Butterfly.

The BFL chip used in all three products contains 16 lanes of double-SHA256 hash pipelines, essentially integrating 16 Spartan-6 FPGAs into one ASIC. The die size was 7.5 mm on a side, and it was placed into a 10x10 mm BGA 144 package. **Surprises.** BFL initially targeted the first half of Nov for shipping their product, however the schedule experienced repeated slips after setbacks and delays from the fab, packaging and BFL itself. The targeted power consumption of the chip was 0.8W per Gh/s, but a month or so before the chips were expected to roll off line, BFL revised its target to 1.2W, and switched from QFN package to a flip-chip BGA package, after tape-out, in anticipation of potential power problems. The energy-efficiency of the device ended up being a major setback; it ended up consuming 4–8× as much power than expected, which required that they underclock the chips from 500 MHz to 250 MHz. These factors required a redesign of all of the systems using the chips. For example, the Jalapenos, which were supposed to use 1 chip, were shipped with two chips to meet the 4.5 Gh/s rate, and they typically operated at 30 Watts, closer to 6W per Gh/s. The MiniRigs, shown in Figure 7, would ship as three separate 500 Gh/s boxes.

While the chips were intended to operate at 500 MHz, instead they would be clocked at ~250 MHz, in order to keep the thermal dissipation within chip limits. The three designs inherently leverage the redundancy of the lanes to control yield; BFL reports that 60% of the chips have 16 fully-functional hash lanes, 20% have 15, 15% have 14 lanes, and 5% have 12-13 lanes functional.

Dynamics of Customer-Financed Hardware. Considerable drama is recorded on the online-forums as customers, who essentially financed the company with millions of dollars, posted on the forums demanding answers for the delays. These setbacks, combined with refinements necessary to their packaging, motherboards, and cases resulted in long waits. BFL started shipping to customers in April 2013, five months after initial estimates, and almost a year after customers had paid for their units in Bitcoin. A large order backlog still exists in Aug 2013, however BFL has shipped

GLOBALFOUNDRIES.



Figure 6: A USB hub hosting an array of ASICMINER Block Erupter USB-stick style bitcoin miners, and a USB-powered cooling fan. Each USB-stick uses a 130-nm ASIC that hashes at 330 MH/s, or about half the performance of \$450 28-nm AMD 7970 GPU. Photo Credit: DennisD7.



Figure 7: Newly arrived \$22,484 65-nm ASIC-based BFL 500 Gh/s MiniRig SC, at center, with 4 surrounding last-generation BFL FPGA MiniRig, and a bunch of smaller mining rigs. Note the two hefty power cables; the rig consumes 2500 to 2700 watts. Photo Credit: James Gibson (gigavps).

many units, by order of ship date, within each category.

Although BFL's customers were understandably concerned that their purchase was rapidly depreciating before they even received it, BFL's initial expectations about the amount of time it took to bring up the chip and ship it at scale as a product were wildly optimistic, and in the end the delays were probably not atypical, especially for a company's first ASIC product. Anecdotally, companies like Intel can often take a year from first silicon to shipping products.

What was perhaps most atypical was the high level of transparency that BFL offered, most likely unparalleled in almost any chip product, presumably brought on by their pre-order based model which, while raising ample capital, also pushed them to over-promise. Nonetheless, this resulted in frustration and animosity from many enthusiastic customers anxiously looking at the rising BTC difficulty curve and wondering if they had bet on the slow horse.

3.5 ASICMINER

The ASICMINER (www.asicminer.co) Bitcoin effort started in early July, after BFL had started taking pre-orders for their machines, and consisted of three Chinese-national founders. One of the motivations was to prevent BFL from being the sole purveyor of Bitcoin mining hardware. Their approach was quite different than BFL's, since they did not have the credibility that BFL had from an existing product line.

Remarkable, the entire process of raising funding was performed exclusively through online forums, namely bitcointalk.org, and also some Chinese-language forums. Using these forums, they outlined carefully their plan for developing an ASIC, and responded to hundreds of questions, many of them very technical, by the online community, regarding their business model, their technical decisions, and their financial trustworthiness. We summarize the openly posted developments.

By July 18, the ASICMINER team had registered a company in Shenzhen, China, and signed contracts with the IC manufacturer and received the files required for starting the

chip design process. By July 29, they had completed an initial IC design, which targeted 1.25 Gh/s per chip in 130-nm, and used $17.5mm^2$ of silicon, at 13.3 W of power. 130-nm was explicitly chosen because the NRE was low—on the order of 150K for a design based on a multi-layer mask set produced by a fab in Shanghai, China. According to their posts, they used an industry-standard flow: Verilog, VCS-based simulation, Verdi-based debugging Design Compiler for synthesis, IC Compiler for Place and Route, Calibre for design rule checks (DRC) and layout-versus-synthesis (LVS), Virtuoso for merging the layout, StarRCXT-based extraction, PrimeTime-based static timing analysis and Formality for verification. This tool suite would be quite expensive in the US, but they cite that EDA license costs are cheap in China, as are labor costs.

In early August, after completing an initial place-and-route, they proceeded to raise funds through an IPO on an online stock exchange, GLBSE, in which the securities were Bitcoin-related and further denominated in bitcoin. They proposed to sell a 1/400K share of the company for 0.1 BTC, with up to 200K shares going to shareholders. Their business plan was to start out by mining shares with 12Th/s of their own hardware, and then later sell hardware or chips directly to customers. The profits would be split as weekly dividends equally across all shares. The forum posting contained a professional-looking prospectus, including risk factors, and detailed payout schedule including preferential payments for shareholders to recoup their investment before payout to the founders, and estimates for profits based on Bitcoin difficulty trends and projections about competitors. Shareholder votes were used to attain guidance on key issues, such as whether to convert 8,000 BTC raised in the IPO to fiat currency (535K RMB) in order to hedge currency risks and ensure payment to the fab. The IPO closed Aug 27, selling 163,962 shares, roughly equivalent to 160K USD.

By Sept 22, they finalized the chip's spec: 1.05V, 335 MHz, 6 mm x 6 mm QFN40 package, 4.2 W per Gh/s, 6-metal

130-nm process, with a simple memory mapped interface for writing, midstate, data and nonces into the part. The part would contain a single double SHA256 hash unit, essentially replicating the Spartan-6 design at a higher frequency, lower power, and much, much smaller cost. The final design focused on reducing power consumption so that the QFN package could be used to reduce packaging and cooling expenses. A tapeout followed shortly afterwards.

On Oct 6, 2012, the plot thickened. The GLBSE exchange was shut down due to a security breach and disagreement among its founders. Since the shares were held in an anonymous system by GLBSE, ASICMINER did not know who its shareholders were, and would not be able to distribute dividends. To make matters worse, some fraction of ASICMINER funds from the IPO was trapped in GLBSE's accounts. Over time, the ASICMINER founders relied on emails and documentation from the shareholders to prove out the ownership of > 150,000 shares. Finally, GLBSE delivered the list after two months of anxiety.

By Oct 14, the mask-set was generated and in the wafer-process queue. PC board design had commenced and the foundry received payment for the masks and first wafer set.

By Oct 31, the initial set of wafers was in the metal layer processing stage, the last stage before the wafers are sliced and shipped to be packaged. However, in Nov 7, it became clear that business at the fab had picked up, delaying further production of ASICMINER chips in favor of bigger players. For the next month, the founders posted to say how many layers of processing remained left, with 1.4-1.5 day/layer for "hot runs" and 1.1-1.2 day/layer for "bullet" runs. By Dec 5, 12 out of 29 layers remained. On Dec 22, half of the wafers had finished via the bullet run, with the other half waiting pre-metalization for potential bug fixes.

On Dec 28, 2012, ASICMINER posted on the forum with chip carrier pictures—the first Bitcoin mining ASICs ever.

By Jan 31, 2013, ASICMINER had boards in hand, each with 64 chips, and was aiming to deploy 800 boards, and mounting them into 10-board backplanes, by early Feb.

By Feb 14, they had 2Th/s deployed and hashing. Officially the ASIC Bitcoin movement was in full force!

Over time, ASICMiner continued to deploy units, however they encountered some significant stalls after their initially stellar rollout. They needed to train workers to assemble the units, and acquire a warehouse space that had sufficiently stable power and cooling to host the machines. Finally, their fab, encountering a busy period, insisted that they respin the masks to avoid MLM technology which reduced fab throughput. After a period, encountering problems scaling out their mining datacenter, they switched to a second phase where they sold their hardware directly to consumers. They auctioned off 60 individual 83W 10.7 Gh/s blades on the forums, for prices as high as 50-75 bitcoin (roughly \$5K-7.5K), to end customers, and then developed a USB miner stick, the Block Erupter, containing a single ASIC, which sold initially for 2 bitcoin in large lots to be resold by others, and rapidly dropped in price. They are currently available for purchase on Amazon.com for \$44. Figure 6 shows a USB hub hosting an array of ASICMiner Block Erupter USB-stick style bitcoin miners, and a USB-powered cooling fan. Each USB-stick users a 130-nm ASIC that hashed at 330 MH/s at 1.05 V and 2.5 W, but can be overclocked to 392 MH/s at 1.15V. The ASIC performs one hash per clock cycle, mirroring earlier FPGA designs. The ASIC is 40× more energy efficient than the 28-nm AMD 7970 GPU, and 4.4× cheaper per Gh/s.

ASICMINER shares now sell for 4 BTC each, signifying a 40× return to the initial investors. Of the three efforts, clearly ASICMINER was the most innovative in trying out

new products and business models for their chips.

3.6 Avalon

The Avalon company was another grass-roots effort that secured funding by direct Internet pre-sales of units via an online store. A key founder, ngzhang, had established a reputation his design of a top Bitcoin FPGA board, Icarus.

They focused on an 110-nm TSMC implementation of a single double-SHA256 pipeline, measuring 4 mm on a side, and packaged 300 chips across 3 blades inside a 4U-ish machine. Like ASICMiner, they were based in Shenzhen, China, which provided significant challenges in shipping the rigs internationally—they essentially had to transport the rigs through a shaky customs process to Hong Kong, where the units could be mailed out. They ran pre-order sales for 300 rigs, each selling for \$1299 each, or 108 bitcoin at the time, and hashing at 66 Gh/s on 600W.

They taped out slightly after ASICMiner, encountering delays due to higher-priority TSMC customers, with a target date of Jan 10. On Jan 30, 2013, Jeff Garzik, a Bitcoin developer, was the first customer in history to receive a Bitcoin ASIC mining rig, which earned ~15 BTC the first day.

Subsequently, Avalon sold off new machine batches, a 2nd batch of 600 rigs for 75 BTC (\$1599) on Feb 2, and a third batch of 600 rigs, also for 75 BTC (\$5500) on Mar 25. They sold out almost immediately. Avalon followed up with direct chip sales, selling over 100 batches of 10,000 chips for 780 BTC per batch, or about \$78,000. Groups of users banded together to perform "group buys", ensuring security by nominating well-known online users to perform escrow. Other groups banded together to design boards for the new chips, including USB sticks and multi-chip boards. In response, BFL has also started to sell their chips in bulk. Meanwhile, Avalon has started work on a 55-nm chip.

3.7 Bitcoin Hardware Scaling

Already, pre-orders have been placed for 28-nm units from a new upstart company. This leads us to the question of how well Bitcoin chips will scale. Due to the dark silicon problem, improvements in performance due to process technology are gated by improvements in energy efficiency at $\sim 1.4\times$ per process generation [3], i.e., the ratio of the two feature widths. In fact, Bitcoin logic is close to the worst case for dark silicon, much worse than multicore [1] or GPUs, because of high duty cycles and lack of low power-density SRAMs.

Thus, if we suppose 10-nm is the terminal process generation for CMOS scaling, there is only a $6.5\times$ improvement that we can expect in performance/\$ due to process generation improvements versus 65-nm. Thus, the additional transistor count and frequency gained by these advanced processes will not pay the dividends one might expect, due to power limitations. BFL, with its advanced 65-nm chip, already bumped up against these dark silicon related limits when it had to scale back performance due to power consumption.

However, unlike in the "race to ASIC" days, the cost/performance difference of future generations of hardware will not be great enough to quickly obsolete the last generation. Rather, it will be energy costs that are likely to dictate which ASIC will be the most profitable. This is especially true in the case where there is a supply glut of chips of a given generation, such as is likely to happen in the next year, as the NREs have been paid, and the three groups are simply paying wafer costs now. One can imagine Bitcoin users dumping their chips, and groups with access to cheap energy buying them for almost free and putting them back to use for mining.

Of course, there are two factors that dictate energy costs—the cost of energy, and the energy consumption of the part.

The parties with the greatest advantage will be those that have cheaper access to large quantities of energy and already have their mining hardware paid off when returns on hashing were higher. Cheaper energy allows these parties to pay off their newly acquired hardware over longer cycles, and to continue to operate even when \$ per Gh/s, as shown in Figure 3, drops precipitously low. Others may have an advantage because they have more energy efficient hardware designs.

Optimizing Energy Efficiency. BFL’s 65-nm part hashes at 5.5 W per Gh/s, while Avalon’s 110 nm part is 9W, and ASICMINER’s 130-nm is 8W. Post-Dennard Scaling [2] predicts that a 14-nm process could allow energy efficiency to improve another $65/14 = 4.6\times$ to around 1 W per Gh/s.

Since the first round of ASIC parts was essentially a race to ASIC, there is likely ample room for optimization of the underlying circuits, including relatively off-the-shelf optimizations that improve energy efficiency without decreasing performance or increasing area, such as reducing energy cost by replacing flip-flops with latches, using multi- V_t flows, or using dual-edge triggered flip-flops, or even using self-timed logic to reduce clock energy. Additionally, there are system-level power distribution and cooling overheads (especially for BFL) that can be reduced. I would estimate at least a factor of 4 energy efficiency to be gained from these approaches.

Beyond this initial level of optimization, Bitcoin hash engines are very friendly targets for dark-silicon types of optimization [2]. In particular near-threshold voltage (NTV) operation is a great fit for Bitcoin mining because hash units can hash with almost no communication, and there are no SRAMs to limit V_{dd} scaling. We could expect that near-threshold could offer an additional $5\times$ in energy efficiency.

Because SHA256 circuits are relatively simple, we could imagine very specialized fabrication processes emerging for them, along the lines of DRAMs today, that take advantage of the limited diversity in the circuit.

4. CONCLUSIONS FOR BESPOKE SILICON

In this paper, we examined the Bitcoin hardware movement, which led to the development of customized silicon ASICs without the support of any major company. The users self-organized and self-financed the hardware and software development, bore the risks and fiduciary issues, evaluated business plans, and braved the task of developing expensive chips on extremely low budgets. This is unheard of in modern times, where last-generation chip efforts are said to cost \$100 million or more, and the # of ASIC starts drop yearly.

What lessons can learn from this? Under what conditions is bespoke silicon truly possible? Some thoughts:

- Bespoke silicon is most competitive against high volume silicon when it passes the “concentration test”: the benefit of customized silicon is contained almost entirely in the part itself, and not in other parts of the system. In the case of Bitcoin, the profitability is a direct function of the silicon, with few other factors except access to electricity and cooling.

- As in the case of Bitcoin, it makes sense that the initial steps towards realizing a bespoke implementation of an algorithm would start with successively harder levels of programming (i.e. cloud, then GPU), and then use FPGA as a gateway to a low-cost ASIC. As was the case with Bitcoin, if the computation exhibits “weak scaling”, where the data size can be scaled up arbitrarily to provide additional benefit, and the specialized implementation is much smaller than the equivalent general purpose or FPU code, then it will be a good fit for a cheap ASIC, making the jump.

- Surprisingly, university research played a limited role in the development process. There are good reasons for this. For one, university research focuses on the latest fabrication

processes, which inherently are unsuitable for bootstrapping. Furthermore, the university has free access to tens of millions of dollars worth of CAD tools that are licensed for non-commercial use only. This means that recent graduates don’t know how to “do hardware on the cheap”.

- The arrival time to market was in direct inverse order of the process node targeted: 130-nm came first, and 65-nm came last. In situations where ASIC really makes a difference, what’s important is that you get the ASIC working and financed, not what generation it is, or how optimized it is.

- Venture capital appeared not to play a significant role. VC conventional wisdom is that hardware startups are too costly, and take too long. Here, user bases were able to self-finance through Bitcoin-denominated stock sales, through online forums, and through pre-order websites, even in cases where the order price ran in the thousands. Kickstarter, another option, surprisingly did not play a role in the contest, although it helped Adapteva. Essentially, the model is that new technology comes into being when crowdsourced early adopters believe in it enough to risk their money.

- Bitcoin machines had a very strong value proposition from the outset: you buy the machine, and it makes you lots of money. Furthermore, users had already tested this value proposition with prior generations of miners. It’s possibly the easiest product to motivate people to risk capital for.

- Two of the teams were from China, and were operating in Shenzhen. Although Silicon Valley is known for its cutting-edge design, much of the work done on a tight budget is performed in Asia. This gives them an inherent advantage in bootstrapping. Cheap access to labor and CAD tools also played an essential role. The CAD flow mentioned by ASICMINER would run \$400K+ for a single seat in the US.

- Ultimately, for innovation in the hardware space, we need lots of new ideas to be tried out for cheap. However, the semiconductor model has increasingly moved away from this direction to expensive chips. As a result, chip startups are largely non-existent and there are few markets in which high-risk, innovative ideas can be examined. At the same time, demand for hardware engineers is dropping and fresh hardware talent is being diverted away from hardware companies to software companies that offer higher salaries. This creates an increasingly unhealthy death spiral where fewer new ideas are being tried and the top talent is leaving the field.

We need to think about strategic ways to enable cheaper chips for new ideas—perhaps via open-source CAD tools, through new technologies to reduce chip costs like MLM, through more fluid financing methods that spread risk better, and through better education and training—in order to enter the **Age of Bespoke Silicon**.

This work was partially supported by NSF Awards 0846152, 1018850, 0811794, and 1228992, Nokia and AMD gifts, and by STARnet, an SRC program sponsored by MARCO and DARPA. We thank James Gibson, Sophokles and DennisD7 for providing photographs, and the Bitcoin online forums for information. Thanks to Saman Amarasinghe and Krste Asanovic for feedback.

- [1] Goulding et al. “GreenDroid: A mobile application processor for a future of dark silicon.” In *HOTCHIPS*, 2010.
- [2] Taylor. “Is dark silicon useful? harnessing the four horsemen of the coming dark silicon apocalypse.” In *Design Automation Conference (DAC)*, 2012.
- [3] Venkatesh et al. “Conservation cores: Reducing the energy of mature computations.” In *ASPLOS*, 2010.