

Quantum attacks on Bitcoin, and how to protect against them

Divesh Aggarwal,^{1,2} Gavin K. Brennen,³ Troy Lee,^{4,2} Miklos Santha,^{5,2} and Marco Tomamichel⁶

¹*National University of Singapore, Singapore*

²*Centre for Quantum Technologies, National University of Singapore, Singapore*

³*Macquarie University, Sydney, Australia*

⁴*SPMS, Nanyang Technological University, Singapore*

⁵*CNRS, IRIF, Université Paris Diderot, France*

⁶*University of Technology Sydney, Australia*

The key cryptographic protocols used to secure the internet and financial transactions of today are all susceptible to attack by the development of a sufficiently large quantum computer. One particular area at risk are cryptocurrencies, a market currently worth over 150 billion USD. We investigate the risk of Bitcoin, and other cryptocurrencies, to attacks by quantum computers. We find that the proof-of-work used by Bitcoin is relatively resistant to substantial speedup by quantum computers in the next 10 years, mainly because specialized ASIC miners are extremely fast compared to the estimated clock speed of near-term quantum computers. On the other hand, the elliptic curve signature scheme used by Bitcoin is much more at risk, and could be completely broken by a quantum computer as early as 2027, by the most optimistic estimates.

We analyze an alternative proof-of-work called Momentum, based on finding collisions in a hash function, that is even more resistant to speedup by a quantum computer. We also review the available post-quantum signature schemes to see which one would best meet the security and efficiency requirements of blockchain applications.

I. INTRODUCTION

Bitcoin is a decentralized digital currency secured by cryptography. Since its development by Satoshi Nakamoto in 2008 [Nak09], Bitcoin has proven to be a remarkably successful and secure system and has inspired the development of hundreds of other cryptocurrencies and blockchain technologies in a market currently worth over 150 billion USD.

The security of Bitcoin derives from several different features of its protocol. The first is the proof-of-work that is required to write transactions to the Bitcoin digital ledger. The work required to do this safeguards against malicious parties who possess less than 50% of the computational power of the network from creating an alternative history of transactions. The second is the cryptographic signature that is used to authorize transactions. Bitcoin currently uses a signature scheme based on elliptic curves.

The coming development of quantum computers poses a serious threat to almost all of the cryptography currently used to secure the internet and financial transactions, and also to Bitcoin. In this paper, we comprehensively analyse the vulnerability of Bitcoin to quantum attacks. We find that the proof-of-work used by Bitcoin is relatively resistant to substantial speedup by quantum computers in the next 10 years, mainly because specialized ASIC miners are extremely fast compared to the estimated clock speed of near-term quantum

computers. This means that transactions, once on the blockchain, would still be relatively protected even in the presence of a quantum computer.

The elliptic curve signature scheme used by Bitcoin is well-known to be broken by Shor’s algorithm [Sho99] for computing discrete logarithms. We analyse exactly *how long* it might take to derive the secret key from a published public key on a future quantum computer. This is critical in the context of Bitcoin as the main window for this attack is from the time a transaction is broadcast until the transaction is processed into a block on the blockchain with several blocks after it. By our most optimistic estimates, as early as 2027 a quantum computer could exist that can break the elliptic curve signature scheme in less than 10 minutes, the block time used in Bitcoin.

We also suggest some countermeasures that can be taken to secure Bitcoin against quantum attacks. We analyse an alternative proof-of-work scheme called Momentum [Lar14], based on finding collisions in a hash function, and show that it admits even less of a quantum speedup than the proof-of-work used by Bitcoin. We also review alternative signature schemes that are believed to be quantum safe.

II. BLOCKCHAIN BASICS

In this section we give a basic overview of how bitcoin works, so that we can refer to specific parts of the protocol when we describe possible quantum attacks. We will keep this discussion at an abstract level, as many of the principles apply equally to other cryptocurrencies with the same basic structure as bitcoin.

All Bitcoin transactions are stored in a public ledger called the *blockchain*. Individual transactions are bundled into blocks, and all transactions in a block are considered to have occurred at the same time. A time ordering is placed on these transactions by placing them in a chain. Each block in the chain (except the very first, or *genesis* block) has a pointer to the block before it in the form of the hash of the previous block’s header.

Blocks are added to the chain by *miners*. Miners can bundle unprocessed transactions into a block and add them to the chain by doing a *proof-of-work* (PoW). Bitcoin, and many other coins, use a PoW developed by Adam Back called Hashcash [Bac02]. The hashcash PoW is to find a well-formed *block header* such that $h(\text{header}) \leq t$, where h is a cryptographically secure hash function and header is the block header. A well-formed header contains summary information of a block such as a hash derived from transactions in the block ¹, a hash of the previous block header, a time stamp, as well as a so-called *nonce*, a 32-bit register that can be freely chosen. An illustration of a block can be found in Table I. The parameter t is a target value that can be changed to adjust the difficulty of the PoW. In Bitcoin, this parameter is dynamically adjusted every 2016 blocks such that the network takes about 10 minutes on average to solve the PoW.

In Bitcoin the hash function chosen for the proof of work is two sequential applications of the SHA256 : $\{0, 1\}^* \rightarrow \{0, 1\}^{256}$ hash function, i.e. $h(\cdot) = \text{SHA256}(\text{SHA256}(\cdot))$. As the size of the range of h is then 2^{256} , the expected number of hashes that need to be tried to accomplish the hashcash proof of work with parameter t is $2^{256}/t$. Rather than t , the Bitcoin proof-of-work is usually specified in terms of the *difficulty* D where $D = 2^{224}/t$. This is the expected number of hashes needed to complete the proof of work divided by 2^{32} , the number of available nonces. In other words, the difficulty is the expected number of variations of

¹ Specifically the root of a Merkle tree of hashes of the transactions

Version	0x20000012
Previous block header hash	00...0dfff7669865430b...
Merkle Root	730d68233e25bec2...
Timestamp	2017-08-07 02:12:18
Difficulty	860,221,984,436.22
Nonce	941660394
Transaction 1	
Transaction 2	
⋮	

TABLE I. Illustration of a block. The data in the top constitutes the block header.

transactions and time stamps that need to be tried when hashing block headers, when for each fixing of the transactions and time stamp all nonces are tried.

Miners can bundle unprocessed transactions into a block however they like, and are awarded a number of bitcoins for succeeding in the PoW task. The “generation” transaction paying the mining reward is also a transaction included in the block, ensuring that different miners will be searching over disjoint block headers for a good hash pre-image.

Once a miner finds a header satisfying $h(\text{header}) \leq t$, they announce this to the network and the block is added to the chain. Note that it is easy to verify that a claimed header satisfies the PoW condition—it simply requires one evaluation of the hash function.

The purpose of the PoW is so that one party cannot unilaterally manipulate the blockchain in order to, for example, double spend. It is possible for the blockchain to fork, but at any one time the protocol dictates that miners should work on the fork that is currently the longest. Once a block has k many blocks following it in the longest chain, a party who wants to create a longest chain not including this block would have to win a PoW race starting k blocks behind. If the party controls much less than half of the computing power of the network, this becomes very unlikely as k grows. In Bitcoin, a transaction is usually considered safe once it has 6 blocks following it.

The first question we will look at in Section III A is what advantage a quantum computer would have in performing the hashcash PoW, and if it could unilaterally “come from behind” to manipulate the blockchain.

The second aspect of Bitcoin that is important for us is the form that transactions take. When Bob wants to send bitcoin to Alice, Alice first creates (an ideally fresh) private-public key pair. The public key is hashed to create an *address*. This address is what Alice provides to Bob as the destination to send the bitcoin. Bitcoin uses the hash of the public key as the address instead of the public key² not for security reasons but simply to save space. As we see later, this design choice does have an impact on the *quantum* security.

To send bitcoin to Alice, Bob must also point to transactions on the blockchain where bitcoin was sent to addresses that he controls. The sum of bitcoin received to these referenced transactions must add up to at least the amount of bitcoin Bob wishes to send to Alice. Bob proves that he owns these addresses by stating the public key corresponding to each address and using his private key corresponding to this address to sign the message saying he is giving these bitcoins to Alice.

² In early versions of the Bitcoin protocol the public key could be used as an address.

III. QUANTUM ATTACKS ON BITCOIN

A. Attacks on the Bitcoin proof-of-work

In this section, we investigate the advantage a quantum computer would have in performing the hashcash PoW used by Bitcoin. Our findings can be summarized as follows: Using Grover search [Gro96], a quantum computer can perform the hashcash PoW by performing quadratically fewer hashes than is needed by a classical computer. However, the extreme speed of current specialized ASIC hardware for performing the hashcash PoW, coupled with much slower projected gate speeds for current quantum architectures, essentially negates this quadratic speedup, at the current difficulty level, giving quantum computers no advantage. Future improvements to quantum technology allowing gate speeds up to 100GHz could allow quantum computers to solve the PoW about 100 times faster than current technology. However, such a development is unlikely in the next decade, at which point classical hardware may be much faster, and quantum technology might be so widespread that no single quantum enabled agent could dominate the PoW problem.

We now go over these results in detail. Recall that the Bitcoin PoW task is to find a valid block header such that $h(\text{header}) \leq t$, where $h(\cdot) = \text{SHA256}(\text{SHA256}(\cdot))$. The security of the blockchain depends on no agent being able to solve the PoW task first with probability greater than 50%. We will investigate the amount of classical computing power that would be needed to match one quantum computer in performing this task.

We will work in the random oracle model [BR93], and in particular assume that $\Pr[h(\text{header}) \leq t] = t/2^{256}$ where the probability is taken uniformly over all well-formed block headers that can be created with transactions available in the pool at any given time (such well-formed block headers can be found by varying the nonce, the transactions included in the block as well as the least significant bits of the timestamp of the header). On a classical computer, the expected number of block headers and nonces which need to be hashed in order to find one whose hash value is at most t is $D \times 2^{32}$ where D is the hashing difficulty defined by $D = 2^{224}/t$ ³.

For quantum computers in the random oracle model we can restrict our attention to the generic quantum approach to solving the PoW task using Grover's algorithm [Gro96]. By Grover's algorithm, searching a database of N items for a marked item can be done with $O(\sqrt{N})$ many queries to the database (whereas any classical computer would require $\Omega(N)$ queries to complete the same task).

Let $N = 2^{256}$ be the size of the range of h for the following. By our assumptions, with probability at least 0.9999 a random set of $10 \cdot N/t$ many block headers will contain at least one element whose hash is at most t . We can fix some deterministic function g mapping $S = \{0, 1\}^{\lceil \log(10 \cdot N/t) \rceil}$ to distinct well-formed block headers. We also define a function f which determines if a block header is "good" or not

$$f(x) = \begin{cases} 0 & \text{if } h(g(x)) > t \\ 1 & \text{if } h(g(x)) \leq t \end{cases}.$$

A quantum computer can compute f on a superposition of inputs, i.e. perform the mapping

$$\sum_{x \in S} \alpha_x |x\rangle \rightarrow \sum_{x \in S} (-1)^{f(x)} \alpha_x |x\rangle.$$

³ According to blockchain.info, on August 8, 2017, the hashing difficulty was $D = 860 \cdot 10^9$ and target was $t = 2^{184.4}$

Each application of this operation is termed an *oracle call*. Using Grover’s algorithm a quantum algorithm can search through S to find a good block header by computing $\#\mathcal{O} = \frac{\pi}{4}\sqrt{10 \cdot N/t} = \pi 2^{14}\sqrt{10 \cdot D}$ oracle calls. The Grover algorithm can be adapted to run with this scaling even if the number of solutions is not known beforehand, and even if no solutions exist [BBHT98].

While the number of oracle calls determines the number of hashes that need to be performed, additional overhead will be incurred to compute each hash, construct the appropriate block header, and to do quantum error correction. We now analyze these factors to determine a more realistic estimate of the running time in two ways. First, we estimate the running time based on a well studied model for universal quantum computers with error correction.

On a classical computer, a hash function such as SHA256 uses basic boolean gate operations, whereas on a quantum computer, these elementary boolean gates are translated into reversible logical quantum gates which introduces some overhead. There are a total of 64 rounds of hashing in the SHA256 protocol and each round can be done using an optimized circuit with 683 Toffoli quantum gates [PRS15]. Most quantum error correction codes use T gates rather than Toffoli gates as the representative time consuming gate, and a careful analysis of the cost to perform the SHA256 function call as well as the inversion about the mean used in the Grover algorithm finds a total T gate count of 474168 for one oracle call [SFL⁺13]. In that circuit decomposition, the T gates can only be parallelized by roughly a factor of three.

There is additional overhead needed by quantum computers to perform error correction. In deciding on a good quantum error correction code there are a variety of tradeoffs to consider: tolerance to a particular physical error model, versatility for subroutines, number of qubits used, logical gate complexity, and the amount of classical processing of error syndromes and feedback. Adopting the surface code, which has advantages of a relatively high fault tolerance error threshold and local syndrome measurements, we can adapt the analysis in Ref. [SFL⁺13] to estimate the total run time of the quantum algorithm. The time needed to run the Grover algorithm and successfully mine a block is

$$\tau = \#\mathcal{O} \times \#G/s = \pi 2^{14}\sqrt{10 \cdot D} \times \#G/s,$$

where $\#G$ is the number of cycles needed for one oracle call, and s is the quantum computer clock speed. Using a surface code, where the dominant time cost is in distilling magic states to implement T gates, one finds

$$\#G = 297784 \times c_\tau(D, p_g),$$

where the first factor includes the logical T gate depth for calling the SHA256 function twice as required by bitcoin PoW, and twice again to make the circuit reversible, as well as the inversion about the mean. The second factor, c_τ , is the overhead factor in time needed for quantum error correction. It counts the number of clock cycles per logical T gate and is a function of difficulty and the physical gate error rate p_g . For a fixed gate error rate, the overhead factor c_τ is bounded above by the cost to invert a 256 bit hash (maximum difficulty).

Because the quantum algorithm runs the hashing in superposition, there is no direct translation of quantum computing power into a hashing rate. However, we can define an effective hash rate, h_{QC} , as the expected number of calls on a classical machine divided by

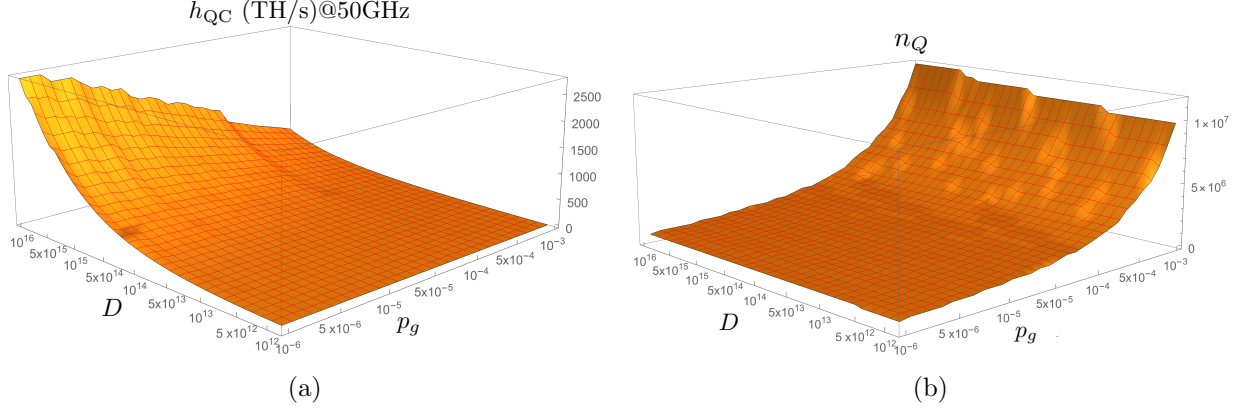


FIG. 1. Performance of a single quantum computer for blockchain attacks as a function of physical gate error rate p_g , which is an internal machine specification, and mining Difficulty D , which is set by the blockchain protocol. (a) Effective hash rate h_{QC} for a quantum computer operating at 50GHz clock speed which is at the optimistic limit of foreseeable clock speeds. The hash rate increases as the square root of difficulty (note the log scale). For d quantum computers acting in parallel the effective hash rate increases by a factor of $2.56 \times \sqrt{d}$. (b) Number of physical qubits n_Q used by the quantum computer.

the expected time to find a solution on the quantum computer, viz.

$$h_{QC} \equiv \frac{N/t}{\tau} = \frac{0.28 \times s\sqrt{D}}{c_\tau(D, p_g)}.$$

Because the time overhead is bounded, asymptotically the effective hashing rate improves as the square root of the difficulty, reflecting the quadratic advantage obtainable from quantum processors.

The Grover algorithm can be parallelized over d quantum processors. In the optimal strategy, each processor is dedicated to search over the entire space of potential solutions, and the expected number of oracle calls needed to find a solution is $\#\mathcal{O} = 0.39 \times \#\mathcal{O}/\sqrt{d}$ [GWC00]. This implies an expected time to find a solution is

$$\tau_{||} = 0.39 \times \tau/\sqrt{d},$$

and the effective hash rate using d quantum processors in parallel is

$$h_{QC,||} = 2.56 \times h_{QC}\sqrt{d}.$$

The number of logical qubits needed in the Grover algorithm is fixed at 2402, independent of the difficulty. The number of physical qubits needed is

$$n_Q = 2402 \times c_{n_Q}(D, p_g),$$

where c_{n_Q} is the overhead in space, i.e. physical qubits, incurred due to quantum error correction, and is also a function of difficulty and gate error rate.

In appendix A we show how to calculate the overheads in time and space incurred by error correction. The results showing the performance of a quantum computer for blockchain




FIG. 2. This plot shows two estimates of the hashing power (in hashes per second) of the bitcoin network (blue striped curves) vs. a single quantum computer (red striped curves) as a function of time for the next 25 years. We give more and less optimistic estimates and uncertainty regions (blue and orange area). The model is described in detail in Appendices B and C. Prior to 2028 (in the more optimistic estimate) there will not be any quantum computer with sufficiently many qubits to implement the Grover algorithm. For comparison, the black dotted line shows the hash rate of a single ASIC device today.

attacks are given in Fig. 1. To connect these results to achievable specifications, we focus on superconducting circuits which as of now have the fastest quantum gate speeds among candidate quantum technologies and offer a promising path forward to scalability. Assuming maximum gate speeds attainable on current devices of $s = 66.7\text{MHz}$ [KKL⁺17] and assuming an experimentally challenging, but not implausible, physical gate error rate of $p_g = 5 \times 10^{-4}$, and close to current difficulty $D = 10^{12}$, the overheads are $c_\tau = 538.6$ and $c_{n_Q} = 1810.7$, implying an effective hash rate of $h_{QC} = 13.8\text{GH/s}$ using $n_Q = 4.4 \times 10^6$ physical qubits. This is more than one thousand times slower than off the shelf ASIC devices which achieve hash rates of 14TH/s ⁴; the reason being the slow quantum gate speed and delays for fault tolerant T gate construction.

Quantum technologies are posed to scale up significantly in the next decades with a quantum version of Moore’s law likely to take over for quantum clock speeds, gate fidelities, and qubit number. Guided by current improvements in superconducting quantum circuit technology, forecasts for such improvements are given in Appendices B and C. This allows us to estimate of the power of a quantum computer as a function of time as shown in Figure 2. Evidently, it will be some time before quantum computers outcompete classical machines for this task and when they do, a single quantum computer will not have majority hashing power. However, even given modest improvement in power over competing miners’ classical

⁴ Using e.g. the AntMiner S9 available at <https://asicminermarket.com>

machines, certain attacks become more profitable such as attacks on mining pools that use smart contracts to reward pool miners to withhold their blocks [VTL17]. For example, given the optimistic assumptions outlined in Appendix A where the effective hash rate scales like $h_{\text{QC}} = 0.04 \times s\sqrt{D}$, at $D = 10^{13}$ and $s = 50\text{GHz}$, the fractional hashing power of a group of 20 quantum machines running in parallel relative to total hashing power is 0.1%. This would allow for a quantum attack on pool mining with minimal bribing cost that would reduce pool revenue by 10%.

B. Attacks on signatures

Signatures in bitcoin are made using the Elliptic Curve Digital Signature Algorithm based on the secp256k1 curve. The security of this system is based on the hardness of the Elliptic Curve Discrete Log Problem (ECDLP). While this problem is still believed to be hard classically, an efficient quantum algorithm to solve this problem was given by Shor [Sho99]. This algorithm means that a sufficiently large universal quantum computer can efficiently compute the private key associated with a given public key rendering this scheme completely insecure. The implications for bitcoin are the following:

1. (Reusing addresses) To spend bitcoin from an address the public key associated with that address must be revealed. Once the public key is revealed in the presence of a quantum computer the address is no longer safe and thus should never be used again. While always using fresh addresses is already the suggested practice in Bitcoin, in practice this is not always followed. Any address that has bitcoin and for which the public key has been revealed is completely insecure.
2. (Processed transactions) If a transaction is made from an address which has not been spent from before, and this transaction is placed on the blockchain with several blocks following it, then this transaction is reasonably secure against quantum attacks. The private key could be derived from the published public key, but as the address has already been spent this would have to be combined with out-hashing the network to perform a double spending attack. As we have seen in Section III A, even with a quantum computer a double spending attack is unlikely once the transaction has many blocks following it.
3. (Unprocessed transactions) After a transaction has been broadcast to the network, but before it is placed on the blockchain it is at risk from a quantum attack. If the secret key can be derived from the broadcast public key before the transaction is placed on the blockchain, then an attacker could use this secret key to broadcast a new transaction from the same address to his own address. If the attacker then ensures that this new transaction is placed on the blockchain first, then he can effectively steal all the bitcoin behind the original address.

We view item (3) as the most serious attack. To determine the seriousness of this attack it is important to precisely estimate how much time it would take a quantum computer to compute the ECDLP, and if this could be done in a time close to the block interval.

For an instance with a n bit prime field, a recently optimized analysis shows a quantum computer can solve the problem using $9n + 2\lceil\log_2(n)\rceil + 10$ logical qubits and $(448\log_2(n) + 4090)n^3$ Toffoli gates [RNSL17]. Bitcoin uses $n = 256$ bit signatures so the number of Toffoli




FIG. 3. Performance of a quantum computer operating at 10GHz clock speed for attacks on digital signatures using the elliptic curve digital signature algorithm. (a) Time in minutes to break the signature as a function of physical gate error rate p_g . (b) Number of physical qubits used by the quantum computer.

gates is 1.28×10^{11} , which can be slightly parallelized to depth 1.16×10^{11} . Each Toffoli can be realized using a small circuit of T gate depth one acting on 7 qubits in parallel (including 4 ancilla qubits) [Sel13].

Following the analysis of Sec. III A, we can estimate the resources needed for a quantum attack on the digital signatures. As with block mining, the dominant time is consumed by distilling magic states for the logical T gates. The time to solve the ECDLP on a quantum processor is

$$\tau = 1.28 \times 10^{11} \times c_\tau(p_g)/s,$$

where the time overhead c_τ now only depends on gate error rate, and s is again the clock speed. The number of physics qubits needed is

$$n_Q = 2334 \times c_{n_Q}(p_g),$$

where the first factor is the number of logical qubits including 4 logical ancilla qubits, and c_{n_Q} is the space overhead.

The performance of a quantum computer to attack digital signatures is given in Fig. 3. Using a surface code with a physical gate error rate of $p_g = 5 \times 10^{-4}$, the overhead factors are $c_\tau = 291.7$ and $c_{n_Q} = 735.3$, and the time to solve the problem at 66.6 MHz clock speed is 6.49 days using 1.7×10^6 physical qubits. Looking forward to performance improvements, for 10GHz clock speed and error rate of 10^{-5} , the signature is cracked in 30 minutes using 485550 qubits. The latter makes the attack in item (3) quite possible and would render the current Bitcoin system highly insecure. An estimate of the time required for a quantum computer to break the signature scheme as a function of time is given in Figure 4, based on the model described in Appendices B and C.

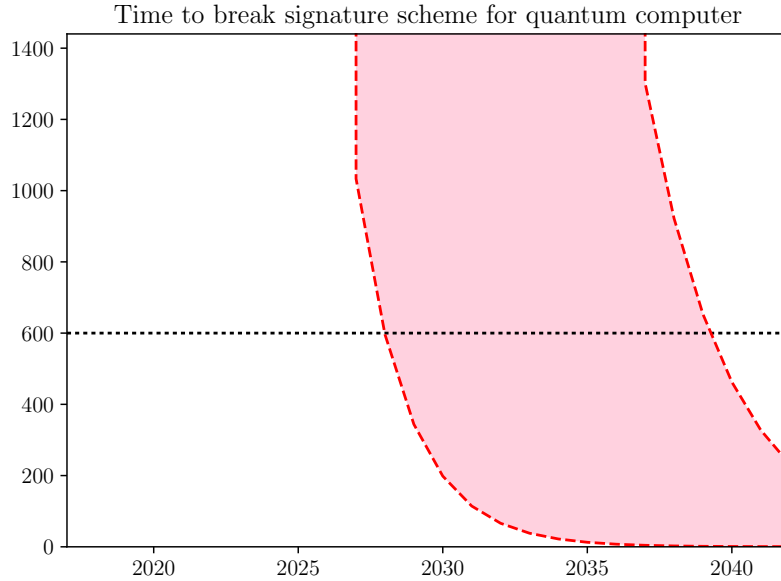


FIG. 4. This plot shows two estimates of the time (in seconds) required for a quantum computer to break the signature scheme (red curves) as a function of time for the next 25 years. We give more and less optimistic estimates (red striped lines). The models are described in detail in Appendix C. According to this estimate, the signature scheme can be broken in less than 10 minutes (600 seconds, black dotted line) as early as 2027.

C. Future enhancements of quantum attacks

We have described attacks on the Bitcoin protocol using known quantum algorithms and error correction schemes. While some of the estimates for quantum computing speed and scaling may appear optimistic, it is important to keep in mind that there are several avenues for improved performance of quantum computers to solve the aforementioned problems.

First, the assumed error correction code here is the surface code which needs significant classical computational overhead for state distillation, error syndrome extraction, and correction. Other codes which afford transversal Clifford *and* non-Clifford gates could overcome the need for slow state distillation [PR13]. In fact the slow down from classical processing for syndrome extraction and correction could be removed entirely using a measurement free protocol, see e.g. [PSBT10] which in a recent analysis show error thresholds [CJS16] only about 5 times worse than the measurement heavy surface code. This could potentially dramatically improve overall speed of error correction.

Second, reductions in logical gate counts of the quantum circuits are possible as more efficient advanced quantum-computation techniques are being developed. For example, using a particular large-size example problem (including oracle implementations) that was analyzed in a previous work [SVM⁺17], a direct comparison of the concrete gate counts, obtained by the software package Quipper, has been achieved between the old [HHL09] and the new [CKS15] linear-systems solving quantum algorithms, showing an improvement of

several orders of magnitude⁵. Given that the quantum Shor and Grover algorithms have been well studied and highly optimized, one would not expect such a dramatic improvement, nonetheless it is likely some improvement is possible.

Third, different quantum algorithms might provide relative speedups. Recent work by Kaliski [Kal17], presents a quantum algorithm for the Discrete Logarithm Problem: find m given $b = a^m$, where b is a known target value and a is a known base, using queries to a so called “magic box” subroutine which computes the most significant bit of m . By repeating queries using judiciously chosen powers of the target value, all bits of m can be calculated and the problem solved. Because different bits are solved one by one the problem can be distributed to multiple quantum processors. Each processor requires a number of logical qubits comparable to solving the entire problem, but the overall time would be reduced by the parallelization. Furthermore, the overhead for quantum error correction is likely reduced as the phases in the quantum fourier transform part of the circuit need not be as accurate as in the original Shor algorithm.

IV. COUNTERMEASURES

A. Alternative proofs-of-work

As we have seen in the last section, a quantum computer can use Grover search to perform the bitcoin proof-of-work using quadratically fewer hashes than are needed classically. In this section we investigate alternative proofs-of-work that might offer less of a quantum advantage. The basic properties we want from a proof-of-work are:

1. (Difficulty) The difficulty of the problem can be adjusted in accordance with the computing power available in the network.
2. (Asymmetry) It is much easier to verify the proof-of-work has been successfully completed than to perform the proof-of-work.
3. (No quantum advantage) The proof-of-work cannot be accomplished significantly faster with a quantum computer than with a classical computer.

The bitcoin proof-of-work accomplishes items (1),(2), but we would like to find an alternative proof of work that does better on (3).

Similar considerations have been investigated by authors trying to find a proof-of-work that, instead of (3) look for proofs-of-work that cannot be accelerated by ASICs. An approach to doing this is by looking at memory intensive proofs of work. Several interesting candidates have been suggested for this such as Momentum [Lar14], based on finding collisions in a hash function, Cuckoo Cycle [Tro15], based on finding constant sized subgraphs in a random graph, and Equihash [BK17], based on the generalized birthday problem. These are also good candidates for a more quantum resistant proof-of-work.

These schemes all build on the hashcash-style proof-of-work and use the following template. Let $h_1 : \{0,1\}^* \rightarrow \{0,1\}^n$ be a cryptographically secure hash function and $H = h_1(\text{header})$ be the hash of the block header. The goal is then to find a nonce x such that

$$h_1(H \parallel x) \leq t \text{ and } P(H, x) ,$$

⁵ A. Scherer, personal communication

for some predicate P . The fact that the header and nonce have to satisfy the predicate P means that the best algorithm will no longer simply iterate through nonces x in succession. Having a proof-of-work of this form also ensures that the parameter t can still be chosen to vary the difficulty.

In what follows, we will analyse this template for the Momentum proof-of-work, as this can be related to known quantum lower bounds. For the momentum proof of work, let $h_2 : \{0, 1\}^* \rightarrow \{0, 1\}^\ell$ be another hash function with $n \leq \ell$. In the original Momentum proposal h_1 can be taken as SHA-256 and h_2 as a memory intensive hash function, but this is less important for our discussion. The proof-of-work is to find H, a, b such that

$$h_1(H \parallel a \parallel b) \leq t \text{ and } h_2(H \parallel a) = h_2(H \parallel b) \text{ and } a, b \leq 2^\ell. \quad (1)$$

First let's investigate the running time in order to solve this proof-of-work, assuming that the hash functions h_1, h_2 can be evaluated in unit time. Taking a subset $S \subset \{0, 1\}^\ell$ and evaluating $h_2(H \parallel a)$ for all $a \in S$, we expect to find about $|S|^2/2^\ell$ many collisions. Notice that by using an appropriate data structure, these collisions can be found in time about $|S|$.

One algorithm is then as follows. For each H , we evaluate h_2 on a subset S and find about $|S|^2/2^\ell$ many pairs a, b such that $h_2(H \parallel a) = h_2(H \parallel b)$. For each collision we then test $h_1(H \parallel a \parallel b) \leq t$. In expectation, we will have to perform this second test $2^n/t$ many times. Thus the number of H 's we will have to try is about $m = \max\{1, \frac{2^{n+\ell}}{t|S|^2}\}$, since we have to try at least one H . As for each H we spend time $|S|$, the total running time is $m|S|$. We see that it is the smallest when $|S| = \sqrt{\frac{2^{n+\ell}}{t}}$, that is when $m = 1$, and we just try one H . This optimal running time is then $T = \sqrt{\frac{2^{n+\ell}}{t}}$, and to achieve it we have to use a memory of equal size to the running time, which might be prohibitive. For some smaller memory $|S| < \sqrt{\frac{2^{n+\ell}}{t}}$ the running time will be $\frac{2^{n+\ell}+1}{t|S|}$.

Now let us look at the running time on a quantum computer. On a quantum computer we can do the following. Call H *good* if there exists $a, b \in S$ such that $h_1(H \parallel a \parallel b) \leq t$ and $h_2(H \parallel a) = h_2(H \parallel b)$. Testing if an H is good requires finding a collision, and therefore necessitates at least $|S|^{2/3}$ time by the quantum query lower bound of Aaronson and Shi [AS04]. Note that this lower bound is tight as finding such a collision can also be done in roughly $|S|^{2/3}$ time using Ambainis' element distinctness algorithm [Amb07]. We have argued above that a set of size $m = \max\{1, \frac{2^{n+\ell}}{t|S|^2}\}$ is needed to find at least one good H . By the optimality of Grover search [BBBV97] we know that we have to perform at least \sqrt{m} many tests to find a good H . As testing if an H is good requires time $|S|^{2/3}$, the total running time is at least $\sqrt{m}|S|^{2/3}$. As the classical running time is $m|S|$, we see that unlike for the current proof of work in Bitcoin, with this proposal a quantum computer would not be able to achieve a quadratic advantage as soon as S is more than constant size. In particular, since $\sqrt{m}|S|^{2/3}$ is minimized also when $S = \sqrt{\frac{2^{n+\ell}}{t}}$, the running time of even the fastest quantum algorithm is at least $T^{2/3}$, which is substantially larger than $T^{1/2}$.

B. Review of post-quantum signature schemes

Many presumably quantum-safe public-key signature schemes have been proposed in the literature. Some examples of these are hash-based signature schemes (LMS [LM95], XMSS [BDH11], SPHINCS [BHH⁺15], NSW [NSW05]), code-based schemes (CFS [CFS01],

Type	Name	security level (bits)	PK length (kb)	Sig. length (kb)	PK + Sig. lengths (kb)
I.1	GPV	100	300	240	540
I.2	LYU	100	65	103	168
I.3	BLISS	128	7	5	12
I.4	DILITHIUM	138	11.8	21.6	33.4
II.1	RAINBOW	160	305	0.244	305
III.1	LMS	128	0.448	20	20.5
III.2	XMSS	128	0.544	20	20.5
III.3	SPHINCS	128	8	328	336
III.4	NSW	128	0.256	36	36
IV.1	CFS	83	9216	0.1	9216
IV.2	QUARTZ	80	568	0.128	568

TABLE II. Comparison of the public key (PK) and signature lengths of post-quantum signature schemes in kilobits (kb). The security level given is against classical attacks. Type I are lattice based, type II based on multivariate polynomials, type III hashing based, and type IV code based.

QUARTZ [PCG01]), schemes based on multivariate polynomials (RAINBOW [DS05]), and lattice-based schemes (GPV [GPV08], LYU [Lyu12], BLISS [DDLL13], DILITHIUM [DLL⁺17], NTRU [MBDG14]). Each of these cryptosystems have varying degree of efficiency. For a comparison in terms of signature size and key size, see Table II.

In the blockchain context the most important parameters of a signature scheme are the signature and public key lengths, as these must be stored in some capacity to fully verify transactions, and the time to verify the signature. Looking at Table II, with respect the sum of signature and public key lengths, the only reasonable options are hash and lattice based schemes.

Hash based schemes like XMSS have the advantage of having provable security, at least assuming the chosen hash function behaves like a random oracle. The generic quantum attack against these schemes is to use Grover’s algorithm which means that their quantum security level is half of the classical security level. In contrast, the best known quantum attack against DILITHIUM at 138 bit classical security level requires time 2^{125} . Thus at the same level of *quantum* security, lattice based schemes have some advantage in signature plus public key length.

Although the lattice based scheme BLISS has the shortest sum of signature and public key lengths of all the schemes in Table II, there are some reasons not to choose BLISS in practice. The security of BLISS relies on hardness of the NTRU problem and the assumption that solving this problem is equivalent to finding a short vector in a so-called NTRU lattice. It has been shown recently [KF17] that this assumption might be too optimistic, at least for large parameters. Moreover, there is a history of attacks on prior NTRU-based signature schemes [NR06, DN12]. Perhaps most fatally, BLISS is difficult to implement in a secure way as it is very susceptible to side channel attacks. The production grade strongSwan implementation of BLISS has been attacked in this way by Pessl et al. [PBY17], who showed that the signing key could be recovered after observing about 6000 signature generations.

```

function CALCULATEFACTORYRESOURCES( $p_g, n_T$ )      ▷ iterates layers of error correction in
factory
   $p_{\text{tol}} \leftarrow \frac{1}{n_T}$                                 ▷ (uncorrected) error tolerance
   $i \leftarrow 0$ 
  while  $p_{\text{tol}} < 10p_g$  do
     $i \leftarrow i + 1$                                 ▷ add layer
     $d_i \leftarrow \min \left\{ d \in \mathbb{N} : 192d \cdot (100p_g)^{\frac{d+1}{2}} \geq \frac{p_{\text{tol}}}{2} \right\}$   ▷ code distance in this layer
     $p_{\text{tol}} \leftarrow \left( \frac{p_{\text{tol}}}{70} \right)^{\frac{1}{3}}$                 ▷ increased error tolerance
  end while
  layers  $\leftarrow i$ 
   $\tau \leftarrow n_T \cdot 10 \sum_{i=1}^{\text{layers}} d_i$                 ▷ total clock cycles (only counts  $T$  gates)
   $Q_{\text{factory}} \leftarrow 50(d_{\text{layers}})^2 \cdot 15^{\text{layers}-1}$     ▷ total physical qubits for factory
  return ( $\tau, Q_{\text{factory}}$ )
end function

function CALCULATECIRCUITRESOURCES( $p_g, n_C, n_L$ )
   $d_C \leftarrow \min \left\{ d \in \mathbb{N} : (80p_g)^{\frac{d+1}{2}} \geq \frac{1}{n_C} \right\}$   ▷ code distance for circuit (single layer)
  return  $Q_{\text{circuit}} \leftarrow 3.125n_L d_C$                 ▷ total physical qubits for circuit
end function

```

TABLE III. Algorithms to compute space and time resources for quantum attacks. The inputs are p_g , the physical gate error rate; n_C , the total number of Clifford gates in the logical circuit; n_T , the total number of T gates in the logical circuit; and n_L , the number of logical qubits. The outputs are τ , the time cost in number of clock cycles; and $n_Q = Q_{\text{circuit}} + Q_{\text{factory}}$, the number of physical qubits used for the computation including state distillation.

Appendix A: Estimating error correction resource overheads for quantum attacks

Here we describe how the overhead factors for quantum error correction are calculated in order to obtain resource costs for quantum attacks on blockchains and digital signatures. The method follows the analysis given in Refs. [FMMC12, MDMG⁺16]. We first determine n_T and n_C , the number of T gates and Clifford gates respectively needed in the algorithm. The pseudo-code to compute the overhead is given in Table III. For the blockchain attack on $n_L = 2402$ qubits, these values are

$$n_T = 297784 \times \pi 2^{14} \sqrt{10 \cdot D}, \quad n_C = 29.4 \times n_T.$$

For the Digital Signature attack on $n_L = 2334$ qubits, the values are ⁶

$$n_T = 1.28 \times 10^{11}, \quad n_C = 20 \times n_T.$$

If we look some years into the future we can speculate as to plausible improvements in quantum computer technology. If we assume a quantum error correction code that supports transversal Clifford and non-Clifford gates so there is no distillation slow down and that it is done in a measurement free manner so that no classical error syndrome processing

⁶ The factor of 20 for the number of Clifford gates per T gate is based on the construction of T gate depth one representations of the Toffoli gate in [Sel13].

is necessary, then the number of cycles needed for one oracle call is determined solely by the circuit depth which is 2142094. This is based on an overall circuit depth calculated as follows. The oracle invokes two calls to the SHA256 hash function, and this is done twice, once to compute it and again to uncompute it. Each hash has a reversible circuit depth of 528768. Similarly, there are two multi-controlled phase gates used, one for inversion about the mean and one for the function call, each having a circuit depth 13511, for a total depth $4 \times 528768 + 2 \times 13511 = 2142094$ (these numbers are from [SFL⁺13] but could be further optimized). Then accepting potential overhead in space and physical qubit number, but assuming no time penalty for error correction or non Clifford gate distillation, this implies an improved effective hashing rate of

$$h_{\text{QC}} = 0.04 \times s\sqrt{D}.$$

which is substantially faster. For superconducting circuits, ultrafast geometric phase gates are possible at ~ 50 GHz, essentially limited by the microwave resonator frequency [RBW⁺12]. Using the above very optimistic assumptions, at difficulty $D = 10^{12}$ the effective hash rate would be $h_{\text{QC}} = 2.0 \times 10^3 \text{TH/s}$.

Appendix B: Modeling the development of bitcoin network hash rates and difficulty

The total number of hashes per second in the whole bitcoin network are taken from blockchain.info. The data points in Figure 5(a) are the hash rates for the first of January (2012–2015) and first of January and July (2016–2017). The two dotted curves correspond to optimistic and less optimistic assumptions for the extrapolations. The optimistic extrapolation assumes that the present growth continues exponentially for five years and then saturates into a linear growth as the market gets saturated with fully optimized ASIC bitcoin miners. The less optimistic assumption assume linear growth at the present rate.

From the extrapolation of the bitcoin network hashrate we can determine the difficulty as a function of time. The expected number of hashes required to find a block in 10 minutes (600 seconds) is given by $\text{rate}(t) \cdot 600$, where $\text{rate}(t)$ is the total hash rate displayed in Figure 5(a). Thus the bitcoin hashing difficulty is calculated as $D(t) = \text{rate}(t) \cdot 600 \cdot 2^{-32}$ for the two scenarios discussed above. In Figure 5(b) we compare this with values from blockchain.info for the first of January of 2015–2017.

Appendix C: Modeling the development of quantum computers

There are several aspects of the development of quantum technologies that we must model. Since only few data points are available at this early stage of the development there is necessarily a lot of uncertainty in our estimates. We therefore give two different estimates, one that is optimistic about the pace of the development and another one that is considerably more pessimistic. Nonetheless, these predictions should be considered as a very rough estimate and might need to be adapted in the future.

First, we need to make an assumption on the number of qubits available at any point of time. As we focus only on solid state superconducting implementations there are only a few data points available. We assume that the number of available qubits will grow exponentially in time in the near future. The optimistic assumption is that the number will double every 10 months whereas the less optimistic assumption assumes the number doubles every 20

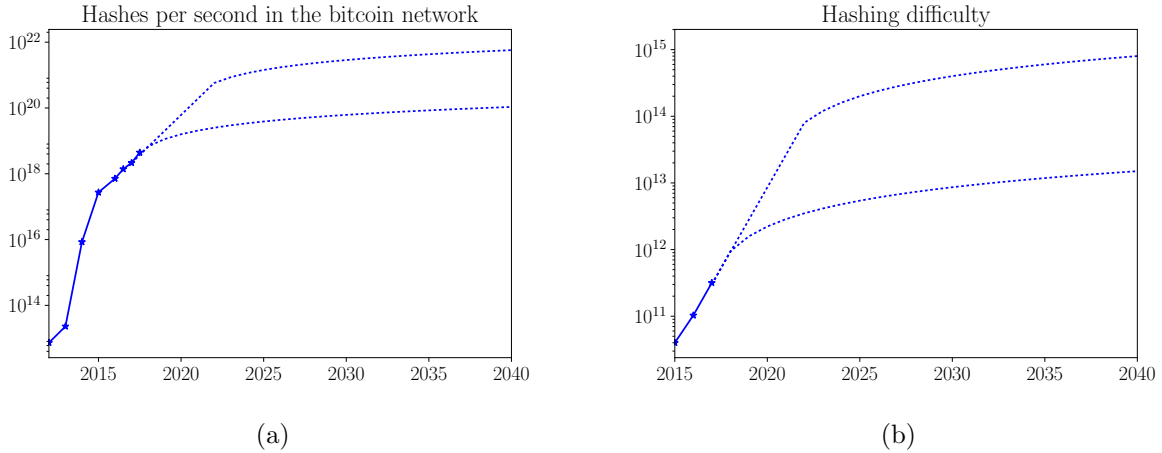


FIG. 5. Prediction of the hash rate of the bitcoin network (in number of hashes per second) and the hashing difficulty as a function of time.

months. These two extrapolations are plotted in Figure 6(a). The data points are taken from the following table:

number of qubits	year	reference
2	2013	[CGC ⁺ 13]
5	2014	[BKM ⁺ 14]
3	2014	[CGM ⁺ 14]
5	2016	IBM
16	2017	IBM
20	2017	[Goo]
49	2018	[Goo]

We predict that the quantum gate frequency grows exponentially for the next years. This assumes that the classical control circuits will be sufficiently fast to control quantum gates at this frequencies. After a couple of years the growth slows down considerably because faster classical control circuits are necessary to further accelerate the quantum gates. We cap the quantum gate frequency at 50 GHz (for the optimistic case) or 5 GHz (for the less optimistic case), respectively, mostly because we expect that classical control circuits will not be able to control the quantum gates at higher frequencies. (See, e.g., [HHOI11] for progress in this direction.) This is shown in Figure 6(b). The data points are taken from the following table:

gate time	year	reference
420ns	2013	[CGC ⁺ 13]
433ns	2015	[CMS ⁺ 15]
160ns	2016	[SMCG16]
42ns	2017	[DBE17]
25ns	2018	Google, projected for end of 2017

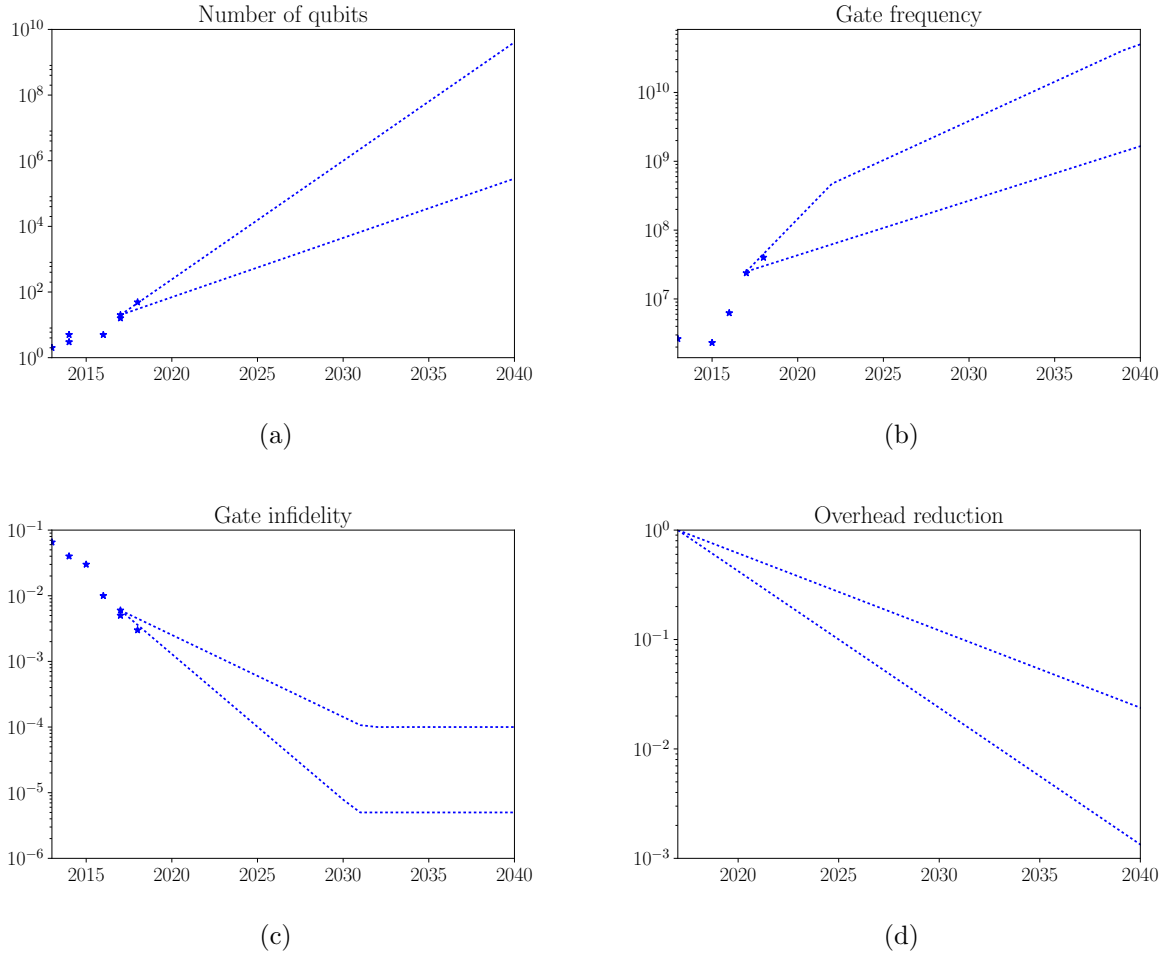


FIG. 6. Prediction of the number of qubits, the quantum gate frequency (in gate operations per second) and the quantum gate infidelity as a function of time. The forth plot models a reduction of the overhead due to theoretical advances.

The predicted development of the gate infidelity is shown in Figure 6(c). We assume that the gate infidelity will continue to drop exponentially but that this development will stall at an infidelity of $5 \cdot 10^{-6}$ (optimistic case) or $5 \cdot 10^{-5}$ (less optimistic case). For the optimistic case we expect that the gate infidelity will continue to follow DeVincenzo's law which predicts a reduction of the infidelity by a factor of 2 per year. The data points are taken from the following table:

gate fidelity	year	reference
0.9347	2013	[CGC+13]
0.96	2014	[CGM+14]
0.97	2015	[CMS+15]
0.99	2016	[SMCG16]
0.995	2017	[Goo]
0.997	2018	[Goo]

Finally, we assume that the number of qubits and time steps required by any algorithm will be reduced over time for two reasons. First, the gate fidelity will increase over time and thus allow for more efficient fault-tolerant schemes to be used. Second, theoretical advances will allow to decrease the number of qubits and gates required to implement the algorithm and fault-tolerant schemes. We expect that this factor will be $\text{overhead}(t) = \beta^{t-2017}$ where $\beta \in \{0.75, 0.85\}$ for optimistic and less optimistic assumptions, respectively.

ACKNOWLEDGEMENTS

MT and GB would like to thank Michael Bremner for initial discussions. TL would like to thank John Tromp for helpful comments and discussions about proof-of-work and Ronald de Wolf for conversations about parallel quantum search.

-
- [Amb07] A. Ambainis. Quantum walk algorithm for element distinctness. *SIAM Journal on Computing*, 37:210–239, 2007.
 - [AS04] S. Aaronson and Y. Shi. Quantum lower bounds for the collision and the element distinctness problems. *Journal of the ACM*, 51(4):595–605, 2004.
 - [Bac02] Adam Back. Hashcash—a denial of service counter-measure, 2002. Available at: <http://www.hashcash.org/papers/hashcash.pdf>.
 - [BBBV97] C.H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani. Strengths and weaknesses of quantum computing. *SIAM Journal on Computing*, 26:1510–1523, 1997.
 - [BBHT98] Michel Boyer, Gilles Brassard, Peter Høyer, and Alain Tapp. Tight bounds on quantum searching. *Fortschritte der Physik*, 46(4-5):493–505, 1998.
 - [BDH11] Johannes Buchmann, Erik Dahmen, and Andreas Hülsing. XMSS—a practical forward secure signature scheme based on minimal security assumptions. *Post-Quantum Cryptography*, pages 117–129, 2011.
 - [BHH⁺15] Daniel J Bernstein, Daira Hopwood, Andreas Hülsing, Tanja Lange, Ruben Niederhagen, Louiza Papachristodoulou, Michael Schneider, Peter Schwabe, and Zooko Wilcox-O’Hearn. SPHINCS: practical stateless hash-based signatures. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 368–397. Springer, 2015.
 - [BK17] Alex Biryukov and Dmitry Khovratovich. Equihash: Asymmetric proof-of-work based on the generalized birthday problem. *Ledger*, 2:1–30, 2017.
 - [BKM⁺14] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O’Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and John M. Martinis. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, 508(7497):500–503, apr 2014.
 - [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *ACM Conference on Computer and Communications Security*, pages 62–73, 1993.
 - [CFS01] Nicolas Courtois, Matthieu Finiasz, and Nicolas Sendrier. How to achieve a mceliece-based digital signature scheme. In *Asiacrypt*, volume 2248, pages 157–174. Springer, 2001.

- [CGC⁺13] A. D. Córcoles, Jay M. Gambetta, Jerry M. Chow, John A. Smolin, Matthew Ware, Joel Strand, B. L. T. Plourde, and M. Steffen. Process verification of two-qubit quantum gates by randomized benchmarking. *Physical Review A*, 87(3):030301–, 03 2013.
- [CGM⁺14] Jerry M. Chow, Jay M. Gambetta, Easwar Magesan, David W. Abraham, Andrew W. Cross, B R Johnson, Nicholas A. Masluk, Colm A. Ryan, John A. Smolin, Srikanth J. Srinivasan, and M Steffen. Implementing a strand of a scalable fault-tolerant quantum computing fabric. *Nature Communications*, 5, jun 2014.
- [CJS16] Daniel Crow, Robert Joynt, and M. Saffman. Improved error thresholds for measurement-free error correction. *Physical Review Letters*, 117(13):130503–, 09 2016.
- [CKS15] Andrew M Childs, Robin Kothari, and Rolando Somma. Quantum linear systems algorithm with exponentially improved dependence on precision. *arXiv:1511.02306*, 2015.
- [CMS⁺15] A.D. Córcoles, Easwar Magesan, Srikanth J. Srinivasan, Andrew W. Cross, M. Steffen, Jay M. Gambetta, and Jerry M. Chow. Demonstration of a quantum error detection code using a square lattice of four superconducting qubits. *Nature Communications*, 6:6979, apr 2015.
- [DBE17] Xiu-Hao Deng, Edwin Barnes, and Sophia E. Economou. Robustness of error-suppressing entangling gates in cavity-coupled transmon qubits. *Physical Review B*, 96(3):035441–, 07 2017.
- [DDLL13] Léo Ducas, Alain Durmus, Tancrede Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal gaussians. In *Advances in Cryptology–CRYPTO 2013*, pages 40–56. Springer, 2013.
- [DLL⁺17] Léo Ducas, Tancrede Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. Crystals–dilithium: Digital signatures from module lattices. Technical report, IACR Cryptology ePrint Archive, 2017: 633, 2017.
- [DN12] Léo Ducas and Phong Q Nguyen. Learning a zonotope and more: Cryptanalysis of ntrusign countermeasures. In *ASIACRYPT*, volume 7658, pages 433–450. Springer, 2012.
- [DS05] Jintai Ding and Dieter Schmidt. Rainbow, a new multivariable polynomial signature scheme. In *ACNS*, volume 5, pages 164–175. Springer, 2005.
- [FMMC12] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A*, 86:032324, Sep 2012.
- [GLP12] Tim Güneysu, Vadim Lyubashevsky, and Thomas Pöppelmann. Practical lattice-based cryptography: A signature scheme for embedded systems. In *CHES*, volume 7428, pages 530–547. Springer, 2012.
- [Goo] <https://www.newscientist.com/article/2138373-google-on-track-for-quantum-computer-breakthrough-by-end-of-2017/>.
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 197–206. ACM, 2008.
- [Gro96] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proc. ACM STOC 1996*, pages 212–219. ACM, may 1996.
- [GWC00] Robert M. Gingrich, Colin P. Williams, and Nicolas J. Cerf. Generalized quantum search with parallelism. *Physical Review A*, 61(5):052313–, 04 2000.
- [HHL09] Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Physical Review Letters*, 103(15):150502–, 10 2009.

- [HHOI11] Quentin P. Herr, Anna Y. Herr, Oliver T. Oberg, and Alexander G. Ioannidis. Ultra-low-power superconductor logic. *Journal of Applied Physics*, 109(10):103903, may 2011.
- [Kal17] B. S. Jr. Kaliski. A quantum “magic box” for the discrete logarithm problem. *eprint.iacr.org/2017/745*, 2017.
- [KF17] Paul Kirchner and Pierre-Alain Fouque. Revisiting lattice attacks on overstretched NTRU parameters. In *Advances in Cryptology - EUROCRYPT 2017 - 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 - May 4, 2017, Proceedings, Part I*, pages 3–26, 2017.
- [KKL⁺17] S Kirchhoff, T Kessler, PJ Liebermann, E Assemat, S Machnes, F Motzoi, and FK Wilhelm. Optimized cross-resonance gate for coupled transmon systems. *arXiv:1701.01841*, 2017.
- [Lar14] Daniel Larimer, 2014. Available at: <http://www.hashcash.org/papers/momentum.pdf>.
- [LM95] Frank T Leighton and Silvio Micali. Large provably fast and secure digital signature schemes based on secure hash functions, July 11 1995. US Patent 5,432,852.
- [Lyu12] Vadim Lyubashevsky. Lattice signatures without trapdoors. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 738–755. Springer, 2012.
- [MBDG14] Carlos Aguilar Melchor, Xavier Boyen, Jean-Christophe Deneuville, and Philippe Gaborit. Sealing the leak on classical ntru signatures. In *International Workshop on Post-Quantum Cryptography*, pages 1–21. Springer, 2014.
- [MDMG⁺16] Amy Matthew, Olivia Di Matteo, Vlad Gheorghiu, Michele Mosca, Alex Parent, and John Schanck. Estimating the cost of generic quantum pre-image attacks on SHA-2 and SHA-3. *Cryptology ePrint Archive, Report 2016/992*, 2016.
- [Nak09] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2009. Available at: <http://www.bitcoin.org/pdf>.
- [NR06] Phong Q Nguyen and Oded Regev. Learning a parallelepiped: Cryptanalysis of ggh and ntru signatures. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 271–288. Springer, 2006.
- [NSW05] Dalit Naor, Amir Shenhav, and Avishai Wool. One-time signatures revisited: Have they become practical? *IACR Cryptology ePrint Archive*, 2005:442, 2005.
- [PBY17] Peter Pessl, Leon Groot Bruinderink, and Yuval Yarom. To BLISS-B or not to be—attacking strongswan’s implementation of post-quantum signatures. In *ACM Conference on Computer and Communications Security (CCS)*, 2017.
- [PCG01] Jacques Patarin, Nicolas Courtois, and Louis Goubin. Quartz, 128-bit long digital signatures. In *Cryptographers? Track at the RSA Conference*, pages 282–297. Springer, 2001.
- [PR13] Adam Paetznick and Ben W. Reichardt. Universal fault-tolerant quantum computation with only transversal gates and error correction. *Physical Review Letters*, 111(9):090505–, 08 2013.
- [PRS15] Alex Parent, Martin Rötteler, and Krysta Marie Svore. Reversible circuit compilation with space constraints. *CoRR*, abs/1510.00377, 2015.
- [PSBT10] Gerardo A. Paz-Silva, Gavin K. Brennen, and Jason Twamley. Fault tolerance with noisy and slow measurements and preparation. *Physical Review Letters*, 105(10):100501–, 08 2010.
- [RBW⁺12] G. Romero, D. Ballester, Y. M. Wang, V. Scarani, and E. Solano. Ultrafast quantum gates in circuit qed. *Physical Review Letters*, 108(12):120501–, 03 2012.
- [RNSL17] M Roetteler, M Naehrig, K.M. Svore, and K Lauter. Quantum resource estimates for computing elliptic curve discrete logarithms. *arXiv:1706.06752*, 2017.

- [Sel13] Peter Selinger. Quantum circuits of $\$t\$$ -depth one. *Physical Review A*, 87(4):042302–, 04 2013.
- [SFL⁺13] Martin Suchara, Arvin Faruque, Ching-Yi Lai, Gerardo Paz, Frederic Chong, and John D. Kubiatowicz. Estimating the resources for quantum computation with the qure toolbox. Technical Report UCB/EECS-2013-119, EECS Department, University of California, Berkeley, <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2013/EECS-2013-119.html>, May 2013.
- [Sho99] Peter W. Shor. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM Review*, 41(2):303–332, jan 1999.
- [SMCG16] Sarah Sheldon, Easwar Magesan, Jerry M. Chow, and Jay M. Gambetta. Procedure for systematically tuning up cross-talk in the cross-resonance gate. *Physical Review A*, 93(6):060302, jun 2016.
- [SVM⁺17] Artur Scherer, Benoît Valiron, Siun-Chuon Mau, Scott Alexander, Eric van den Berg, and Thomas E. Chapuran. Concrete resource analysis of the quantum linear-system algorithm used to compute the electromagnetic scattering cross section of a 2d target. *Quantum Information Processing*, 16(3):60, 2017.
- [Tro15] John Tromp. Cuckoo cycle: a memory bound graph-theoretic proof-of-work. *Financial Cryptography and Data Security: BITCOIN*, pages 49–62, 2015.
- [VTL17] Y Velner, J Teutsch, and L Luu. Smart contracts make bitcoin mining pools vulnerable. *Cryptology ePrint Archive*, page <http://eprint.iacr.org/2017/230>, 2017.